Elastic Cloud Server

User Guide

Issue 42

Date 2023-07-31

HUAWEI

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks 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 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 Instances

1.1 Selecting an ECS Billing Mode

1.1.1 Yearly/Monthly Billing

Concept

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

For more billing information, see **Billing**.

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.

Resources Supporting Yearly/Monthly Billing

Resources billed in yearly/monthly mode include:

- ECSs (vCPUs and memory)
- Images, including prepaid Marketplace images
- EVS disks purchased together with a yearly/monthly ECS
- Bandwidth purchased together with a yearly/monthly ECS

  EIP and dedicated bandwidth are billed together. For details, see the pricing for dedicated bandwidths.

When you purchase a yearly/monthly ECS, the configuration price covers the above resources.

For details about ECS prices, see **Product Pricing Details**.

1.1.2 Pay-per-Use Billing

Concept

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.

For more billing information, see **Billing**.

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.

Billing Examples

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 x 30 x 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).

Resources Supporting Pay-per-Use Billing

Resources billed on a pay-per-use basis include:

- ECSs (vCPUs and memory)
- Images, including Marketplace images as well as shared or customized images based on Marketplace images
- EVS disks purchased together with a yearly/monthly ECS
- Bandwidth purchased with a pay-per-use ECS

  For details about ECS prices, see **Product Pricing Details**.

Billing for Stopped ECSs

Common ECSs refer to ECSs without local disks or FPGAs attached. After a common ECS is stopped, it is billed as follows:

- ECS basic resources (vCPUs, memory, and image) are released and no longer generate costs. Its associated resources such as its EVS disks, EIPs, and bandwidth will continue to be billed.
- When you try to start the ECS the next time, the system will allocate vCPUs and memory again, but if resources are insufficient, the startup may fail. In

this case, you can try again later or resize the ECS specifications first before trying to start it.

Special pay-per-use ECSs will continue to be billed after being stopped and its resources such as vCPUs and memory are still retained.

NOTE

Special ECSs include:

- BMSs
- ECSs attached with local disks, such as disk-intensive ECSs and ultra-high I/O ECSs
- FPGA-based ECSs

To stop billing for special ECSs, delete them and their associated resources.

1.1.3 Spot Pricing

1.1.3.1 Spot Pricing ECSs

Concept

HUAWEI CLOUD sells available computing resources at a discount. The price changes in real time depending on market demands. This is the spot pricing billing mode.

An ECS billed in spot pricing billing mode is a spot ECS.

In spot pricing billing mode, you can purchase and use ECSs at a discount price. A spot ECS performs as well as the ECSs with the same specifications in other billing modes. However, when inventory resources are insufficient, or the market price increases and exceeds your expected price, the system will automatically release your ECS resources and reclaim the ECSs. Compared with pay-per-use and yearly/monthly ECSs, spot ECSs offer the same level of performance while at lower costs.

Working Rules

The market price for the ECSs of a certain flavor fluctuates due to supply-and-demand changes. You can purchase and use spot ECSs at a low market price to reduce computing costs.

When purchasing a spot ECS, you are required to set the maximum price you are willing to pay for a specified flavor. A higher price ensures a greater success rate for you to purchase such an ECS.

- If the maximum price is greater than or equal to the market price and the inventory resources are sufficient, the spot ECS can be purchased and will be billed at the market price.
- If the maximum price is less than the market price, the spot ECS cannot be purchased.

After purchasing a spot ECS, you can use it like using the ECSs in other billing modes. However, the system will periodically compare the maximum price with the market price and check the inventory resources.

- If the maximum price is greater than or equal to the market price and the inventory resources are sufficient, you can continue using the ECS.
- If the maximum price is less than the market price or the inventory resources are insufficient, the system notifies you of releasing the ECS resources (notifications enabled) and automatically deletes the ECS in about 5 minutes.

Figure 1-1 Lifecycle of a spot ECS

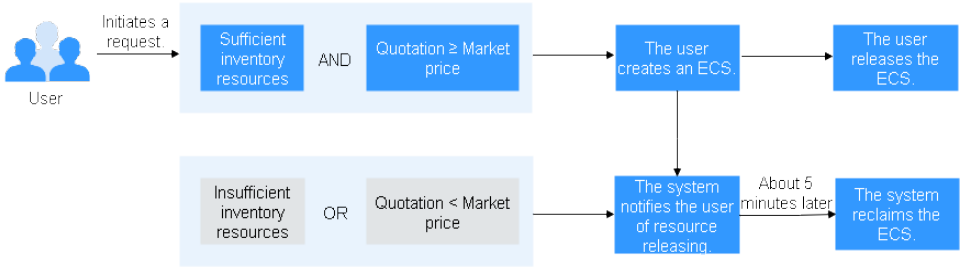

Application Scenarios

- Suitable workloads

  Spot ECSs are suitable for image rendering, stateless web service, gene sequencing, offline analysis, function calculation, batch calculation, sample analysis, CI/CD, and test.

  NOTE

  When the market price is higher than the maximum price you are willing to pay or the inventory resources are insufficient, the spot ECSs will be reclaimed. Therefore, back up data when using such ECSs.

- Unsuitable workloads

  To prevent ECS reclamation from interrupting services, do not use spot ECSs to run workloads requiring long-time operations or high stability.

Notes

- Only KVM ECSs support spot pricing payments. For details about supported ECS flavors, see the information displayed on the management console.
- The market prices of the ECSs of the same flavor may vary depending on AZs.
- Spot ECSs do not support OS change.
- Spot ECSs do not support automatic recovery.
- Spot ECSs do not support specifications modification.
- Spot ECSs cannot be created using a Marketplace image.
- Spot ECSs cannot be switched to yearly/monthly ECSs.
- When a spot ECS is being reclaimed,
  - It cannot be used to create system disk images and full-ECS images. However, data disks of the ECS can be used to create data disk images.
  - It cannot be deleted.
- By default, the data disks and EIP of a spot ECS will not be released after it is reclaimed. If you want to be notified when a spot ECS is reclaimed so that you can determine whether to manually release data disks and EIP, set a reclaim

notification. For details, see "Enabling Reclaim Notifications" in **Purchasing a Spot ECS**.

Billing Rules

For details, see **Billing**.

Billing Examples

- **If the market price is higher than the maximum price you set, the spot ECS is released. The spot ECS is billed based on the market price. Example:**

  At 08:30, the market price is $0.02 USD/hour, and the maximum price is $0.04 USD/hour. Then, the ECS is billed at $0.02 USD/hour.

  At 09:00, the market price is $0.03 USD/hour.

  At 10:00, the market price is $0.04 USD/hour.

  At 10:30, the market price is $0.05 USD/hour, which is higher than the maximum price. Then, the system notifies the user of ECS releasing.

  This ECS is billed in three billing periods.

  During 08:30-09:00, the ECS had been running for 30 minutes and it is billed by the second: 0.02/3600 x 30 x 60 = $0.01 USD

  During 09:00-10:00, the ECS had been running for 1 hour and it is billed at the market price at 09:00, which is $0.03 USD ($0.03 USD/hour x 1 hour = $0.03 USD).

  During 10:00-10:30, the ECS had been running for 30 minutes and it is billed by the second: 0.04/3600 x 30 x 60 = $0.02 USD

  The total price is $0.06 USD for the running duration of 2 hours.

- **If inventory resources are insufficient, the system releases a price ECS and bills it based on the market price. Example:**

  At 08:30, the market price is $0.02 USD/hour, and the maximum price is $0.06 USD/hour. Then, the ECS is billed at $0.02 USD/hour.

  At 09:00, the market price is $0.03 USD/hour.

  At 10:00, the market price is $0.04 USD/hour.

  At 10:30, the market price is $0.05 USD/hour. Although the market price is lower than the maximum price, the system releases this ECS due to insufficient inventory resources.

  This ECS is billed in three billing periods.

  During 08:30-09:00, the ECS had been running for 30 minutes and it is billed by the second: 0.02/3600 x 30 x 60 = $0.01 USD

  During 09:00-10:00, the ECS had been running for 1 hour and it is billed at the market price at 09:00, which is $0.03 USD ($0.03 USD/hour x 1 hour = $0.03 USD).

  During 10:00-10:30, the ECS had been running for 30 minutes and it is billed by the second: 0.04/3600 x 30 x 60 = $0.02 USD

  The total price is $0.06 USD for the running duration of 2 hours.

Purchasing a Spot ECS

You can purchase a spot ECS on the management console or by calling APIs.

- For instructions about how to purchase a spot ECS on the management console, see **Purchasing a Spot ECS**.
- For instructions about how to purchase a spot ECS by calling APIs, see **Creating an ECS**.

Reclaiming an ECS

HUAWEI CLOUD may reclaim and terminate your spot ECS at any time. A spot ECS that is being reclaimed cannot be used to create images.

An ECS may be reclaimed due to:

- Higher market price than the maximum price you are willing to pay
- Insufficient inventory resources

  NOTE

  - If a spot ECS is reclaimed within the first hour after it is provisioned, the spot ECS is not billed.
  - In the first settlement period (in hours) of a spot ECS, the spot ECS is billed, regardless of whether it is started or not.
  - It takes 5 minutes to reclaim a spot ECS. If during that 5 minutes, the spot pricing hour is exceeded, any time in excess of that hour will be billed at the new market price.
  - During the running of a spot ECS, its price is updated once an hour. After a spot ECS is restarted, or it is stopped and then started, it is billed at the market price when the ECS starts.

Back up data on spot ECSs. Before the system reclaims your spot ECSs, it will notify you of the release if notifications are enabled. To enable notifications, see **Purchasing a Spot ECS**.

FAQs

See **Spot ECSs**.

1.1.3.2 Purchasing a Spot ECS

Scenarios

A spot ECS is billed in spot pricing mode. You can purchase and use such ECSs at a discount price. A spot ECS performs as well as the ECSs with the same specifications in other billing modes. However, when inventory resources are insufficient, or the market price increases and exceeds your expected price, the system will automatically release your ECS resources and reclaim the ECSs.

Compared with pay-per-use and yearly/monthly ECSs, spot ECSs offer the same level of performance while at lower costs. For more information about the spot pricing payments, see **Spot Pricing**.

Purchasing a Spot ECS

Follow the instructions provided in **Purchasing an ECS** and **Logging In to an ECS** to buy and log in to spot ECSs. Pay attention to the following settings:

When purchasing a spot ECS:

- Set **Billing Mode** to **Spot pricing**.

  In **Spot pricing** billing mode, your purchased ECS is billed based on the service duration at a lower price than that of a pay-per-use ECS with the same specifications. However, a spot ECS may be reclaimed at any time based on the market price or changes in supply and demand.
- Set **Maximum Price**, which can be **Automatic** or **Manual**.
  - **Automatic** is recommended, which uses the pay-per-use price as the highest price you are willing to pay for a spot ECS.
  - **Manual** requires you to set the upper price limit for a spot ECS. The maximum price must be greater than or equal to the market price and less than or equal to the pay-per-use price.
- Click **Next**, confirm that the specifications and price are correct, agree to the service agreement, and click **Submit**.

  NOTE

  A spot ECS may be reclaimed by the system. Therefore, back up your data.

Constraints

- Only KVM ECSs support spot pricing payments. For details about supported ECS flavors, see the information displayed on the management console.
- The market prices of the ECSs of the same flavor may vary depending on AZs.
- Spot ECSs do not support OS change.
- Spot ECSs do not support automatic recovery.
- Spot ECSs do not support specifications modification.
- Spot ECSs cannot be created using a Marketplace image.
- Spot ECSs cannot be switched to yearly/monthly ECSs.
- When a spot ECS is being reclaimed,
  - It cannot be used to create system disk images and full-ECS images. However, data disks of the ECS can be used to create data disk images.
  - It cannot be deleted.
- By default, the data disks and EIP of a spot ECS will not be released after it is reclaimed. If you want to be notified when a spot ECS is reclaimed so that you can determine whether to manually release data disks and EIP, set a reclaim notification. For details, see "Enabling Reclaim Notifications" in **Purchasing a Spot ECS**.

(Optional) Enabling Reclaim Notifications

After purchasing a spot ECS, you can use it like using the ECSs in other billing modes. However, a spot ECS may be reclaimed at any time based on the market price or changes in supply and demand.

You can enable reclaim notifications to be notified ahead of about 5 minutes before the system starts to release your spot ECS if the maximum price you are willing to pay is lower than the market price or the inventory resources are insufficient.

Use Cloud Trace Service (CTS) and Simple Message Notification (SMN) to enable notifications. For details, see **Cloud Trace Service User Guide**.

Step 1 Enable CTS. For details, see **Enabling CTS**.

Once CTS is enabled, the system automatically identifies the cloud services enabled on the cloud platform, obtains key operations on the services, and reports traces of these operations to CTS.

Step 2 Configure reclaim notifications.

You can configure key event notifications on CTS so that SMN can send messages to notify you of key operations. This function is triggered by CTS, but notifications are sent by SMN.

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Under **Management & Governance**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Key Event Notifications**.
5. Click **Create Key Event Notification** in the upper right corner of the page and set parameters listed in **Table 1-1**.

Table 1-1 Parameters for configuring key event notifications

Type	Parameter	Configuration
Basic Information	Notification Name	The value is user-defined, for example, **spottest**.
Operation	Operation Type	Select **Custom**.
	Operation List	Choose **ECS** > **server** > **interruptServer** and click **Add**.
User	Specified Users	If you do not specify users, CTS notifies all users when key operations are initiated.
Topic	Send Notifications	Select **Yes**.

Type	Parameter	Configuration
	SMN Topic	Select a topic from the drop-down list. If there are no proper SMN topics, create one. 1. Click **Topic** to switch to the **Topics** page. 2. On the SMN console, choose **Topic Management** > **Topics**. Then, click **Create Topic** and set parameters as required. For details, see **Creating a Topic**. 3. Locate the newly added topic and click **Add Subscription** in the **Operation** column. Then, you can receive notifications sent for the topic. For details, see **Adding a Subscription to a Topic**.

After the configuration is complete, you will receive a notification 5 minutes before the system deletes your spot ECS.

Step 3 (Optional) View reclaimed spot ECSs.

1. Under **Management & Governance**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in **Table 1-2** and search for traces as needed.

Table 1-2 Setting filter criteria to search for reclaimed ECSs

Parameter	Configuration
Trace Source	ECS
Resource Type	server
Search By	**Trace name** > **interruptServer**
Operator	All operators
Trace Status	All trace statuses

4. Locate the target trace and expand the trace details.
5. Click **View Trace** in the **Operation** column for details.

----End

1.1.4 Reserved Instances

1.1.4.1 Reserved Instance Overview

Concept

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to the use of pay-per-use ECSs in your account. When the attributes of

your pay-per-use ECSs **match** those of an RI, the RI billing benefit automatically applies to your ECSs. The combination of RIs and pay-per-use billing fully utilizes the flexibility of pay-per-use resources at lower costs.

NOTE

- A purchased RI is billed, regardless of whether it is used or not.
- RIs cannot be used on the ECSs running Microsoft SQL Server.

Table 1-3 ECS billing modes

Billing Mode	What It Is	How to Use
RI	A billing discount applied to pay-per-use ECSs.	When the attributes of your pay-per-use ECSs **match** those of an RI, the RI billing benefit automatically applies to your ECSs.
Pay-per-use	ECS billed based on usage frequency and duration. Such an ECS can be created or deleted at any time.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.
Yearly/ Monthly	ECS billed based on the service duration. This mode is ideal when the duration of ECS usage is predictable.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.
Spot pricing	ECS billed in spot pricing billing mode.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.

- For instructions about how to purchase an RI, see **Enabling and Purchasing a Reserved Instance**.
- For instructions about how to modify an RI, see **Modifying RI Attributes**.

What Is Attribute Mapping Between an RI and a Pay-per-Use ECS?

A regional RI is purchased within a region and without an AZ specified. A zonal RI is purchased within an AZ.

- Attribute mapping of a regional RI: indicates whether the region, OS type, ECS series, and vCPU/memory ratio of a pay-per-use ECS are the same as those specified in a regional RI.
- Attribute mapping of a zonal RI: indicates whether the AZ, OS type, flavor of a pay-per-use ECS are the same as those specified in a zonal RI.

Application Scenarios

If your ECSs are to be used in a short term, it is a good practice to use pay-per-use rates. If you plan to use ECSs for one or three years, it is a good practice to use RIs. RIs offer discounts for pay-per-use ECSs with matched attributes.

For example, after you purchase two s3.2xlarge Linux RIs with a one-year term in AZ 1, the billing benefit of the RIs is immediately applied to up to two pay-per-use s3.2xlarge Linux ECSs running in AZ 1.

Working Rules

For example, you have a running pay-per-use ECS in your account. After you purchase an RI that matches the attributes of this ECS, the billing benefit of the RI is automatically applied to your ECS when the RI takes effect. A purchased RI takes effect at the next hour.

Table 1-4 lists RI attributes. You can purchase your desired RIs based on these attributes.

Table 1-4 RI attributes

Parameter	Description
Region or AZ	• Regional RI: indicates an RI purchased in a region, without an AZ specified. Capacity reservations are not supported for regional RIs. • Zonal RI: indicates an RI purchased with an AZ specified. Capacity reservations are supported for zonal RIs.
Flavor	• When purchasing a regional RI, ensure that the ECS series and vCPU/memory ratio specified in the RI are the same as those of the target pay-per-use ECS. • When purchasing a zonal ECS, ensure that the flavor specified in the RI is the same as that of the target pay-per-use ECS.
OS	The OS of the ECS to be bought, which must match the OS specified in your RI. For example, if you want to use a Linux RI, select a Linux public or private image when purchasing an ECS.
Term	The service duration of an RI. A year is defined as 31,536,000 seconds (365 days).
Offering Class	Standard: Certain attributes, such as the ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

Zonal RIs

A zonal RI, which is purchased within a specified AZ, offers a billing discount for the ECSs with the same OS and flavor as the RI in that AZ.

For example, after you purchase two c3.xlarge.2 Linux RIs with a one-year term in AZ 1, the billing benefit of the RIs is immediately applied to up to two pay-per-use c3.xlarge.2 Linux ECSs running in AZ 1.

Regional RIs

A regional RI, which is purchased within a specified region, has the following characteristics:

- AZ flexibility: The RI discount applies to pay-per-use ECS usage in any AZ within a region.
- ECS size flexibility: The RI discount applies to pay-per-use ECS usage when the ECS OS, ECS series, and vCPU/memory ratio of the target ECS are the same as those specified in the regional RI. ECS size flexibility is determined based on the normalization factor of the ECS size. ECS size flexibility does not apply to zonal RIs.

ECS size flexibility is applied from the smallest to the largest ECS size within the ECS series based on the normalization factor. **Table 1-5** describes ECS size within an ECS type and corresponding normalization factor per hour.

NOTE

An ECS automatically benefits from the billing discount offered by a regional RI only when the ECS series and vCPU/memory ratio are the same as those specified in the RI.

For example, a regional c3.large.4 RI cannot be used on a c3.large.2 ECS because their vCPU/memory ratios are different.

Table 1-5 Normalization factors

ECS Size	Normalization Factor
small	1
medium	1
large	2
xlarge	4
2xlarge	8
4xlarge	16
6xlarge	24
7xlarge	28
8xlarge	32
9xlarge	36
12xlarge	48

ECS Size	Normalization Factor
14xlarge	56
15xlarge	60
16xlarge	64
26xlarge	104
52xlarge	208
nxlarge	n x 4

For example, an s3.large.2 ECS has a normalization factor of 2. You purchase an s3.large.2 Linux RI for the CN-Hong Kong region of HUAWEI CLOUD with a one-year term.

- If you have two running s3.medium.2 pay-per-use Linux ECSs in this region, the billing benefit is fully applied to both ECSs.

  Figure 1-2 Example RI 1

  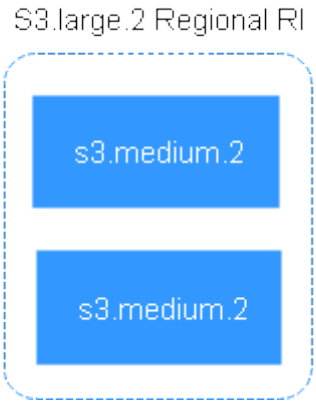

- If you have one running s3.xlarge.2 pay-per-use Linux ECS with a normalization factor of 4 in this region, the billing benefit is applied to 50% of the usage of the ECS.

  Figure 1-3 Example RI 2

  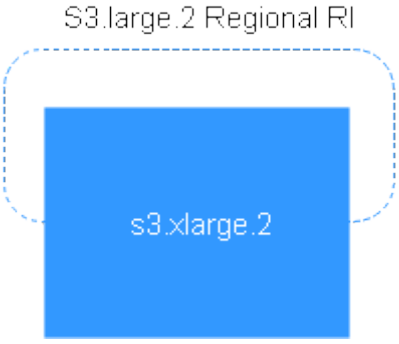

Table 1-6 Comparison between regional and zonal RIs

RI Type	AZ Flexibility	ECS Size Flexibility	Capacity Reservation
Regional RI	Supported A regional RI applies to any AZ in the region.	Supported A regional RI applies when the ECS series and vCPU/memory ratio of the target ECS are the same as those specified in the RI.	Not supported Resources are not reserved so ECS creation may fail when resources are insufficient.
Zonal RI	Not supported A zonal RI applies only in a specified AZ.	Not supported A zonal RI applies only when the flavor of the target ECS is the same as that specified in the RI.	Supported Desired resources can be reserved for creating a pay-per-use ECS.

Examples

If you have the following pay-per-use ECSs in region A:

- Five s3.large.2 Windows ECSs in AZ 1
- Three m3.xlarge.2 Windows ECSs in AZ 2
- One c3.xlarge.2 Windows ECS in AZ 3

You purchase the following RIs in the same region (region A):

- Five s3.large.2 Windows RIs with a one-year term in AZ 1
- Six m3.large.2 Windows RIs with a one-year term in region A
- One c3.large.2 Windows RI with a one-year term in region A

The RI benefits are applied as follows:

- The discount of the five s3.large.2 zonal RIs is used by the five s3.large.2 ECSs because the attributes (AZ, OS, and ECS type) between the RIs and ECSs match.
- The m3.large.2 regional RIs offer AZ flexibility and ECS size flexibility.

  An m3.large.2 RI is equivalent to two normalization factors. The six m3.large.2 regional RIs are equal to 12 normalization factors (6 x 2). In account A, there are three running m3.xlarge.2 ECSs, which are equivalent to 12 normalization factors (3 x 4). In this case, the six m3.large.2 regional RIs are equivalent to three m3.xlarge.2 ECSs.
- The c3.large.2 regional RI offers AZ flexibility and ECS size flexibility and can be applied to c3.xlarge.2 ECSs.

  A c3.large.2 RI is equivalent to two normalization factors (1 x 2). A c3.xlarge.2 ECS requires an RI with four normalization factors (1 x 4). Therefore, the c3.large.2 RI billing discount applies to 50% of c3.xlarge.2 usage. The remaining c3.xlarge.2 usage is billed at the pay-per-use rate.

1.1.4.2 Enabling and Purchasing a Reserved Instance

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to pay-per-use ECSs in your account. When the attributes of your pay-per-use ECSs match those of an RI, the RI's discount rate automatically applies to your ECSs.

RIs are suitable for scenarios where the resource usage duration can be predicted. Billing automatically applies your RI's discounted rate when attributes of your ECS usage match attributes of an RI.

- For more information about RIs, see **Reserved Instance Overview**.
- For instructions about how to modify an RI, see **Modifying RI Attributes**.

Quota Constraints

- The quota for the number of RIs that you can purchase in the current region is displayed in the upper left area of the **Reserved Instance** page. The quota for the number of RIs that can be purchased by a user in each region is 20.
- The quota for the number of RIs is automatically reset every month.
- The remaining quota for the number of RIs (Remaining quota = Total quota – Used quota) is reduced only after more RIs are purchased. It will not be changed if RIs are modified, split, combined, or unsubscribed.

Enabling RIs

Before purchasing an RI, contact customer service to apply for the required permissions.

Purchasing an RI

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Reserved Instance**.
5. Click **Buy RI**.

   The **Buy RI** page is displayed.
6. Confirm the region.

   If the RIs in the selected region do not meet your requirements, select another region.
7. (Optional) Select **Show offerings that reserve capacity** to view the AZs that support capacity reservations.
   - Zonal RIs offer capacity reservation.
   - Regional RIs offer capacity reservation.
8. (Optional) Select an AZ to purchase a zonal RI for capacity reservation.

   Perform this operation only when you purchase RIs for a specified AZ.
9. Select an RI type.

The cloud platform provides various RI types for you to choose from based on your application scenarios.

10. Filter for RI specifications.

    Set flavor, OS, term, offering class, and payment option to search for the target RI specifications.

    Table 1-7 shows specifications parameters.

    Table 1-7 RI attributes

Parameter	Description
Region or AZ	• Regional RI: indicates an RI purchased in a region, without an AZ specified. Capacity reservations are not supported for regional RIs. • Zonal RI: indicates an RI purchased with an AZ specified. Capacity reservations are supported for zonal RIs.
Flavor	• When purchasing a regional RI, ensure that the ECS series and vCPU/memory ratio specified in the RI are the same as those of the target pay-per-use ECS. • When purchasing a zonal ECS, ensure that the flavor specified in the RI is the same as that of the target pay-per-use ECS.
OS	The OS of the ECS to be bought, which must match the OS specified in your RI. For example, if you want to use a Linux RI, select a Linux public or private image when purchasing an ECS.
Term	The service duration of an RI. A year is defined as 31,536,000 seconds (365 days).
Offering Class	Standard: Certain attributes, such as the ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

11. Select specifications.

    The cloud platform provides various RI types for you to choose from based on your application scenarios. On the **Buy RI** page, view released RI types and specifications.

    Effective Rate: amortized hourly costs of the RI, which is equivalent to the total cost (including any upfront payment) of the RI over the entire term divided by the total number of hours over the entire term. (Effective rate = Total cost of the RI/Entire term of the RI)

    Upfront Price: fee that needs to be paid before you purchase an RI.

    Hourly Rate: amortized hourly costs of the RI, which is equivalent to the difference between the total cost of the RI and the upfront payment divided

by the total number of hours over the entire term (Hourly rate = Total cost of the RI – Upfront payment/Entire term of the RI)

12. Specify an RI name.

    The name can be customized. It can contain 1 to 128 characters, which can only be letters, digits, underscores (_), and hyphens (-).

13. Set the number of RIs to be purchased.
    - **Quantity**: The system displays the number of RIs that you can purchase.
    - **Total Normalization Factors**: measures the ECS size flexibility. The value is determined based on the specifications of the RI to be purchased.
    - **Total Upfront Price + Pay-per-use Price**: The price to be paid for consists of the total upfront price and the pay-per-use price. The total upfront price is a product of the upfront fee per RI and the number of RIs. The pay-per-use price is a product of the pay-per-use fee per RI and the number of RIs.

      For details, click **Pricing details**.

14. Click **Next**.

    To learn more about the price, click **Pricing details**.

15. On the page for you to confirm RI specifications, view details and submit the request.

    After verifying the configurations and price, click **Submit** and pay for the order as prompted.

16. Return to the RI list as prompted and view the purchased RI.

Follow-up Operations

- **Purchase a pay-per-use ECS that matches an RI.**

  Locate the target RI and click **Buy ECS** in the **Operation** column. The system automatically switches to the page for purchasing ECSs, and the specifications of the ECSs selected by default are the same as those specified in the RI.

  NOTE

  - If the OS of the target ECS does not match the OS specified in the RI, or the target ECS is not billed on a pay-per-use basis, the RI cannot be used. When the attributes of the ECS match those of the RI, including the ECS series and vCPU/memory ratio, the ECS automatically benefits from the billing discount offered by the RI.

- **Check the usage of RIs.**

  On the **Reserved Instance** page, click the name of the target RI. In the lower part of the RI details page, view the usage of the RI.

  As shown in **Figure 1-4**, the horizontal coordinate indicates the number of in-service days, and the vertical coordinate indicates the usage of RI on the current day. Click the chart to view the RI usage in the selected time period on the current day.

Figure 1-4 Viewing RI usage

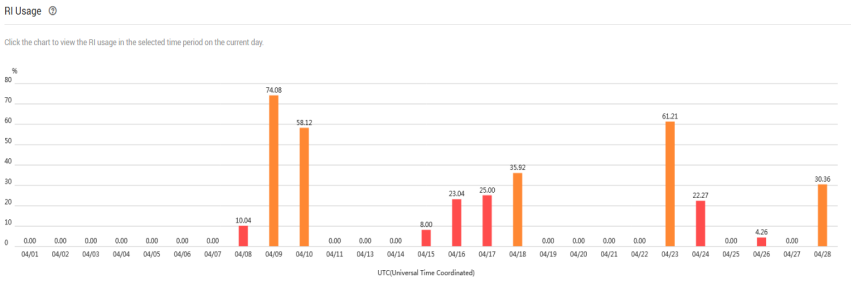

1.1.4.3 Modifying RI Attributes

Scenarios

If an RI type cannot meet your computing requirements, you can modify the RI attributes and then apply it to your pay-per-use ECSs.

You can modify the scope, AZ, and ECS size of a standard RI.

- For more information about RIs, see **Reserved Instance Overview**.
- For instructions about how to purchase an RI, see **Enabling and Purchasing a Reserved Instance**.

Constraints

- RIs can be combined only when their attributes, including the OS, payment option, offering class, term, expiration time, region, ECS series, vCPU/memory ratio, and discount are the same.
- The total normalization factors must be the same before and after the modification.
- A maximum of five RIs can be modified in a batch.
- One RI can be split to multiple ones, and multiple RIs can only be combined into one.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**. On the displayed console, choose **Reserved Instance** from the left navigation pane.
4. On the **Reserved Instance** page, select the target RI and click **Modify RI** in the upper left corner of the list.
5. Modify the RI attributes as required.

Table 1-8 Common operations for modifying an RI

Allowed Operation	Description
Splitting an RI or combining RIs	For example, there are six s3.xlarge.2 RIs in an account, and an s3.xlarge.2 RI has a normalization factor of 4. Then, the six s3.xlarge.2 RIs are equivalent to 24 normalization factors. Then, these RIs can be combined into three s3.2xlarge.2 RIs or split to 24 s3.medium.2 RIs. Just ensure that the splitting or combination matches to the total normalization factor.
Changing a regional RI to a zonal one	A regional RI can be changed to a zonal RI.

NOTICE

Total normalization factors are the number of RIs multiplied by the normalization factor of such an RI. The total normalization factors must be the same before and after the modification.

For example, there are six s3.large.4 RIs with the total normalization factors of 12 (6 x 2) before the modification. These RIs can be split to two s3.xlarge.4 RIs and four s3.medium.4 RIs. After the modification, the total normalization factors are still 12 (2 x 4 + 4 x 1).

6. Verify the modified RI attributes and click **Submit**.

1.1.5 Changing Pay-per-Use to Yearly/Monthly

Scenarios

- **Pay-per-use**: a postpaid billing mode, in which an ECS is billed by usage duration. You can provision or delete such an ECS at any time.
- **Yearly/Monthly**: a prepaid billing mode, in which an ECS is billed based on the purchased duration. This mode is more cost-effective than the pay-per-use mode and is suitable for predictable usage.

If you need to use an ECS for a long time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs.

NOTE

For certain associated resources, their billing modes can be changed together with the ECS to yearly/monthly.

For associated resources whose billing modes cannot be changed together with the ECS to yearly/monthly, they will retain their original billing modes. For details, see **Billing Mode Change Rules for Associated Resources**.

Billing Mode Change Rules for Associated Resources

Resources associated with ECSs include disks and EIPs. **Table 1-10** shows the billing mode change rules for these associated resources.

Table 1-9 Billing mode change rules for disks attached to an ECS

Disk Type	Billing Mode	Shared	Changed Together with ECS to Yearly/ Monthly	Measure
Local disks	N/A	No	N/A	None
DSS/ DESS disks	Yearly/Monthly (the same as the storage pool billing mode)	No	N/A	None
EVS disks	Pay-per-use	No	Yes (not supported for extreme SSD V2 disks)	None
EVS disks	Pay-per-use	Yes	No	On the EVS console, change the billing mode of EVS disks from pay-per-use to yearly/monthly. For details, see **Billing for Disks**.
EVS disks	Yearly/Monthly	No	No	The billing mode is already yearly/monthly. No actions are required.
EVS disks	Yearly/Monthly	Yes	No	The billing mode is already yearly/monthly. No actions are required.

Table 1-10 Billing mode change rules for EIPs bound to an ECS

Resource	Billing Mode	Billed By	Bandwidth Type	Changed Together with ECS to Yearly/ Monthly	Measure
EIP	Pay-per-use	Bandwidth	Dedicated	Yes	None
EIP	Pay-per-use	Traffic	Dedicated	No	On the EIP console page, change the billing mode from billing by traffic (pay-per-use) to billing by bandwidth (pay-per-use) first and then to yearly/monthly. For details, see **Changing EIP Billing Mode**.
EIP	Pay-per-use	Billed by bandwidth	Shared	No	On the EIP console, change the billing mode from pay-per-use to yearly/monthly. For details, see **Changing EIP Billing Mode**.
EIP	Yearly/ Monthly	Bandwidth	Dedicated or shared	No	The billing mode is already yearly/monthly. No actions are required.

Prerequisites

The selected ECS is billed on a pay-per-use basis.

Procedure

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, select the target ECS.
5. Choose **More** > **Change Billing Mode** in the **Operation** column.

NOTE

You can batch change the billing modes of multiple ECSs. To do so, perform the following operations:

1. Select the target ECSs.
2. Choose **More** > **Change Billing Mode** above the ECS list.

6. Confirm the ECS details, specify the usage duration, and pay for the order.

1.1.6 Changing Yearly/Monthly to Pay-per-Use

Scenarios

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 usage frequency and duration, you can change the billing mode from yearly/monthly to pay-per-use.

NOTE

After the billing mode is changed from yearly/monthly to pay-per-use, the new billing mode takes effect only after the yearly/monthly subscription has expired.

Prerequisites

- 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.
- A yearly/monthly subscription can be changed to pay-per-use before its expiration date. However, the change takes effect only after the subscription has expired.
- The billing modes of products in a solution portfolio cannot be changed from yearly/monthly to pay-per-use.

Procedure

1. Log in to the management console.
2. On the top navigation bar, choose **Billing & Costs** > **Renewal**.

   The **Renewals** page is displayed.
3. Customize search criteria.
   - On the **Pay-per-Use After Expiration** tab, you can search for the ECSs with the billing mode changed to pay-per-use.
   - On the **Manual Renewals**, **Auto Renewals**, and **Renewals Canceled** tabs, you can also change the billing mode of the ECSs to pay-per-use (taking effect after the subscription expires).

Figure 1-5 Renewals

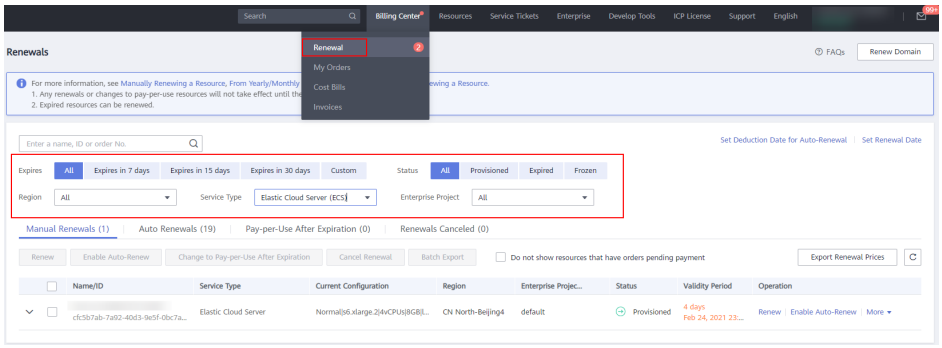

4. Change the ECS billing mode to pay-per-use after the yearly/monthly subscription expires.
   - Single ECS: Select the ECS for which you want to change the billing mode, and choose **More** > **Change to Pay-per-Use After Expiration** in the **Operation** column.

     Figure 1-6 Changing the billing mode of a single ECS from yearly/monthly to pay-per-use after expiration

     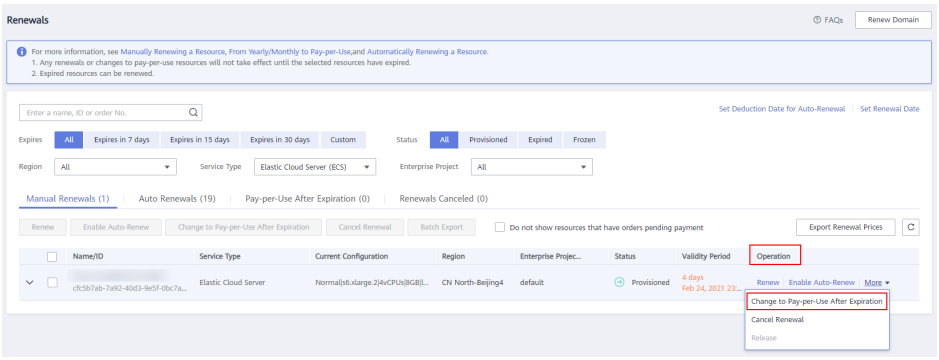

   - Multiple ECSs: Select the ECSs for which you want to change the billing mode, and click **Change to Pay-per-Use After Expiration** above the ECS list.

     Figure 1-7 Batch changing the billing mode of ECSs from yearly/monthly to pay-per-use after expiration

     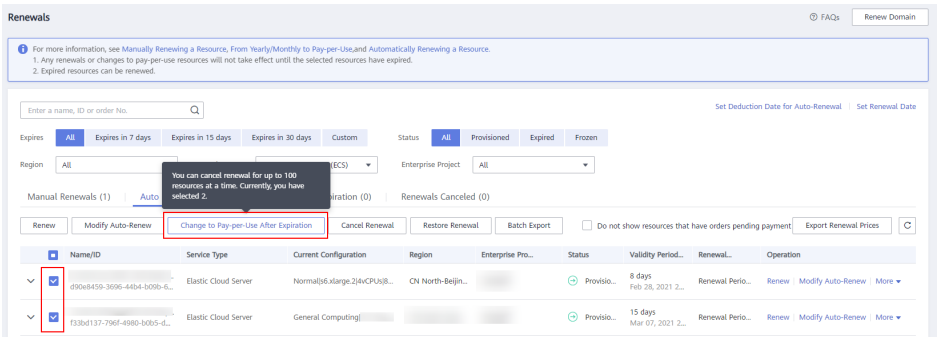

5. Confirm the change details and click **Change to Pay-per-Use**.

1.2 Purchasing an ECS

1.2.1 Purchasing the Same ECS

Scenarios

If you have bought an ECS and want to buy new ones with the same configuration, it is a good practice to use "Buy Same ECS" to rapidly buy the new ones.

Notes

Large-memory ECSs do not support "Buy Same ECS".

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click and choose **Compute** > **Elastic Cloud Server**.
4. Select the target ECS, click **More** in the **Operation** column, and select **Buy Same ECS**.

   Figure 1-8 Buy Same ECS

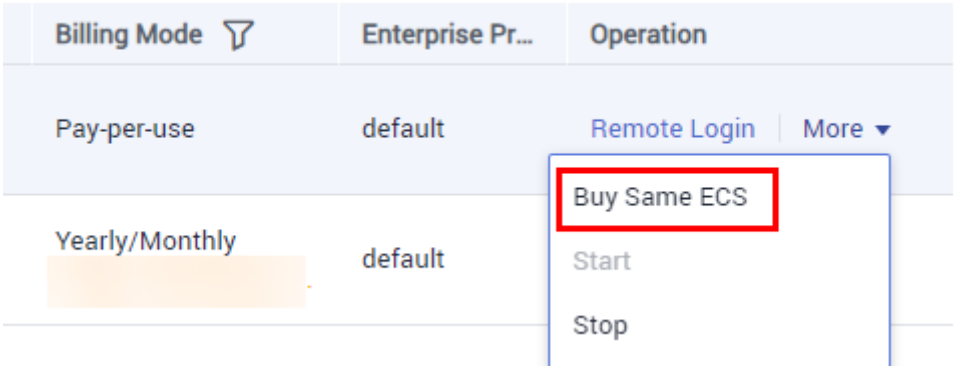

5. The system switches to the ECS purchase page and automatically copies the parameter settings of the selected ECS. Adjust the settings of the new ECSs as needed, confirm the configuration, and click **Submit**.

NOTE

For security purposes, you must manually configure some of the settings for the new ECSs, including:

- Manually add data disks if the quantity of data disks needed exceeds 10.
- Manually add NICs if the quantity of NICs needed exceeds 5.
- Manually add security groups if the quantity of security groups needed exceeds 5.
- Select a new data disk image if the disks of the source ECS are created using a data disk image.
- If the source ECS is created from a full-ECS image, only the disks included in this image are displayed. Add disks if necessary.
- Select **Encryption** if the disks of the source ECS have been encrypted.
- Configure the functions in **Advanced Options**.
- Configure **EIP** if required because it is set to **Not required** by default.

1.3 Viewing ECS Information

1.3.1 Viewing ECS Creation Statuses

Scenarios

After submitting the request for creating an ECS, you can view the creation status. This section describes how to view the creation status of an ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. View the ECS status in the **Status** column after purchasing an ECS.

   NOTE

   - An ECS that is being created is in one of the following states:
     - **Creating**: The ECS is being created.
     - **Faulty**: Creating the ECS failed. In such a case, the system automatically rolls back the task and displays an error code on the GUI, for example, **Ecs.0013 Insufficient EIP quota**.
     - **Running**: The request of creating the ECS has been processed, and the ECS is running properly. An ECS in this state can provide services for you.
   - If you find that the task status area shows an ECS creation failure but the ECS list displays the created ECS, see **Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?**

1.3.2 Viewing Failed Tasks

Scenarios

You can view the details of failed task (if any) in the **Failures** area, including the task names and statuses. This section describes how to view failures.

Failure Types

Table 1-11 lists the types of failures that can be recorded in the **Failures** area.

Table 1-11 Failure types

Failure Type	Description
Creation failures	A task failed. For a failed task, the system rolls back the task and displays an error code, for example, **Ecs.0013 Insufficient EIP quota**.
Operation failures	• Modifying ECS specifications If an ECS specifications modification failed, this operation is recorded in **Failures**.

Procedure

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. View **Failures** on the right side of common operations.

   Figure 1-9 Failures

   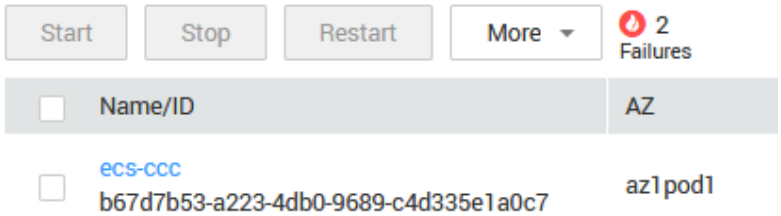

5. Click the number displayed in the **Failures** area to view task details.
   - **Creation Failures**: show the failed ECS creation tasks.
   - **Operation Failures**: show the tasks with failed operations and error codes, which help you troubleshoot the faults.

1.3.3 Viewing ECS Details (List View)

Scenarios

After obtaining ECSs, you can view and manage them on the management console. This section describes how to view detailed ECS configurations, including its name, image, system disk, data disks, VPC, NIC, security group, and EIP.

To view the private IP address of an ECS, view it on the **Elastic Cloud Server** page.

Procedure

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**.

   The **Elastic Cloud Server** page is displayed. On this page, you can view your ECSs and the basic information about the ECSs, such as their specifications, images, and IP addresses.
4. In the search box above the ECS list, select a filter (such as ECS name, ID, or private IP address), enter the corresponding information, and click  .
5. Click the name of the target ECS.

   The page providing details about the ECS is displayed.
6. View the ECS details.

   You can click the tabs and perform operations. For details, see **Changing a Security Group**, **Attaching a Network Interface**, **Adding Tags**, and **Binding an EIP**.

1.3.4 Exporting ECS Information

Scenarios

The information of all ECSs under your account can be exported in an XLSX file to a local directory. The file includes the IDs, private IP addresses, and EIPs of your ECSs.

Procedure

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. In the upper right corner above the ECS list, click  .

   The system will automatically export all ECSs in the current region under your account to a local directory.

NOTE

To export certain ECSs, select the target ECSs and click  in the upper right corner of the page.

5. In the default download path, view the exported ECS information.

1.4 Logging In to a Windows ECS

1.4.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Windows ECS is **Administrator**.
- If the login password is forgotten, reset the password on the ECS console.

  To reset a password, locate the row containing the target ECS, and choose **More** > **Reset Password** in the **Operation** column. For details, see **Resetting the Password for Logging In to an ECS on the Management Console**.
- If an ECS uses key pair authentication, use the password obtaining function available on the management console to decrypt the private key used during ECS creation to obtain a password.
- Certain G series of ECSs do not support remote login provided by the cloud platform. If you need to remotely log in to the ECSs, install the VNC server on them. For details, see **GPU-accelerated ECSs**. You are suggested to log in to the ECSs using MSTSC.
- If you log in to a GPU-accelerated ECS using MSTSC, GPU acceleration will fail. This is because MSTSC replaces the WDDM GPU driver with a non-accelerated remote desktop display driver. In such a case, you must log in to the ECS using other methods, such as VNC. If the remote login function available on the management console fails to meet your service requirements, you must install a suitable remote login tool, such as TightVNC, on the ECS.

  To download TightVNC, log in at **https://www.tightvnc.com/download.php**.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 1-12 Windows login modes

ECS OS	Local OS	Connection Method	Requirement
Windows	Windows	Use MSTSC. Click **Start** on the local computer. In the **Search programs and files** text box, enter **mstsc** to open the **Remote Desktop Connection** dialog box. For details, see **Login Using MSTSC**.	The target ECS has had an EIP bound. (If you log in to an ECS through an intranet, for example, through VPN or Direct Connect, the ECS does not require an EIP.)
	Linux	Install a remote connection tool, for example, rdesktop. For details, see **Logging In to a Windows ECS from a Linux Computer**.	
	macOS	Install a remote connection tool, for example, Microsoft Remote Desktop on the macOS. For details, see **Logging In to a Windows ECS from a Mac**.	
	Mobile terminal	Install a remote connection tool, for example, Microsoft Remote Desktop. For details, see **Logging In to a Windows ECS from a Mobile Terminal**.	
	Windows	Through the management console. For details, see **Login Using VNC**.	No EIP is required.

Helpful Links

- **Login Password Resetting**
- **Multi-User Logins**
- **Remote Logins**

1.4.2 Login Using VNC

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS.

Prerequisites

If an ECS uses key pair authentication, make sure that the key file has been used to resolve the login password before logging in to the ECS. For details, see **Obtaining the Password for Logging In to a Windows ECS**.

Logging In to a Windows ECS

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Obtain the password for logging in to the ECS.

   Before logging in to the ECS, you must have the login password.

   - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
   - If your ECS uses key pair authentication, obtain the password by following the instructions provided in **Obtaining the Password for Logging In to a Windows ECS**.
5. In the **Operation** column of the target ECS, click **Remote Login**.

   Figure 1-10 Remote Login

   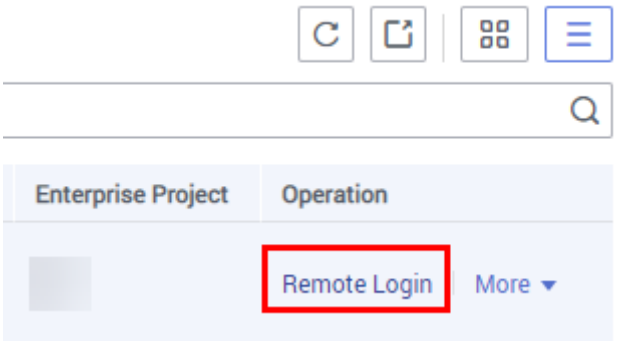

6. In the **Logging In to a Windows ECS** dialog box, click **Log In** in the **Other Login Modes** area.
7. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Ctrl+Alt+Del** in the upper part of the remote login page to log in to the ECS.

   Figure 1-11 Ctrl+Alt+Del

   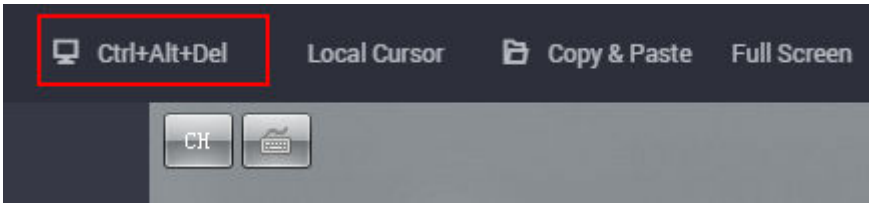

8. Enter the ECS password as prompted.

Helpful Links

- **Login Password Resetting**
- **Multi-User Logins**
- **Remote Logins**

1.4.3 Login Using MSTSC

Scenarios

This section describes how to use the remote login tool MSTSC to log in to a Windows ECS from a local computer.

Prerequisites

- The target ECS is running.
- If your ECS uses key pair authentication, you have obtained the password for logging in to the Windows ECS. For details, see **Obtaining the Password for Logging In to a Windows ECS**.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.

  An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see **Configuring Security Group Rules**.
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- Remote Desktop Protocol (RDP) needs to be enabled on the target ECS. For ECSs created using public images, RDP has been enabled by default. For instructions about how to enable RDP, see **Enabling RDP**.

Logging In to a Windows ECS Using MSTSC

If your local server runs Windows, you can use the remote desktop connection tool MSTSC delivered with the Windows OS to log in to a Windows ECS.

The following uses Windows Server 2012 ECS as an example.

Figure 1-12 Logging in to an ECS using MSTSC

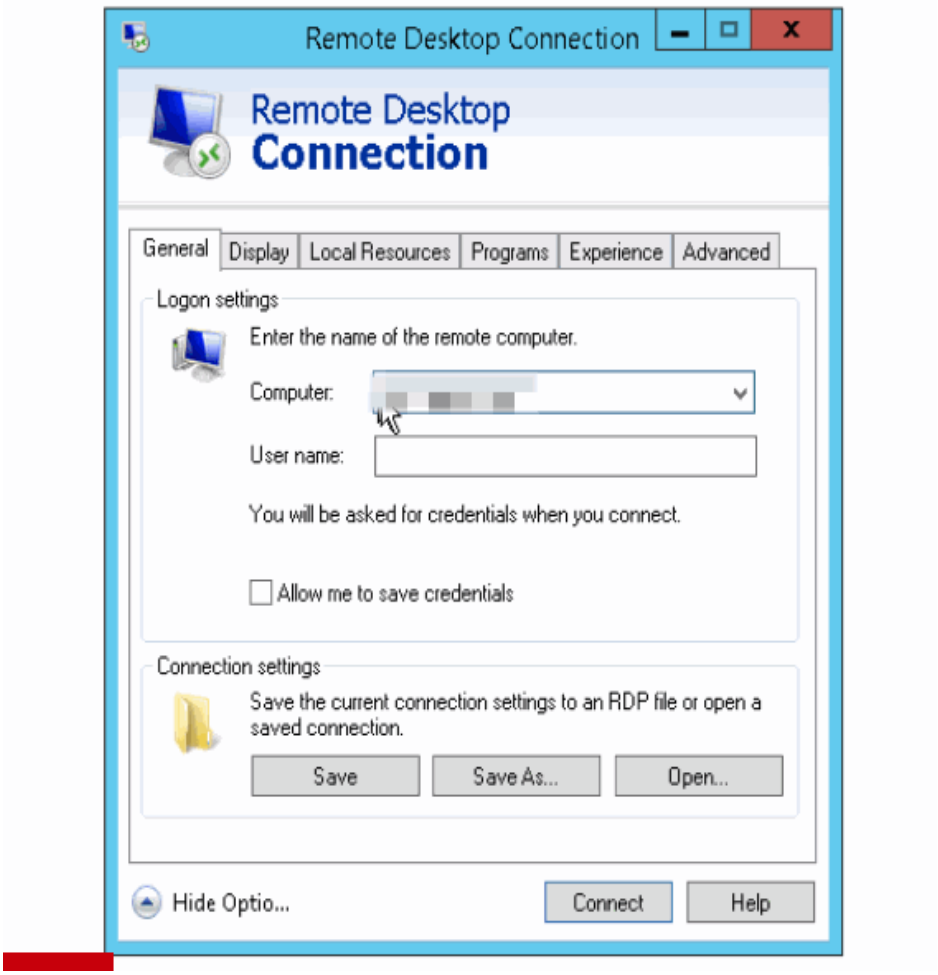

For details, see the following procedure:

1. Click the start menu on the local server.
2. In the **Search programs and files** text box, enter **mstsc**.
3. In the **Remote Desktop Connection** dialog box, click **Show Options**.

Figure 1-13 Show Options

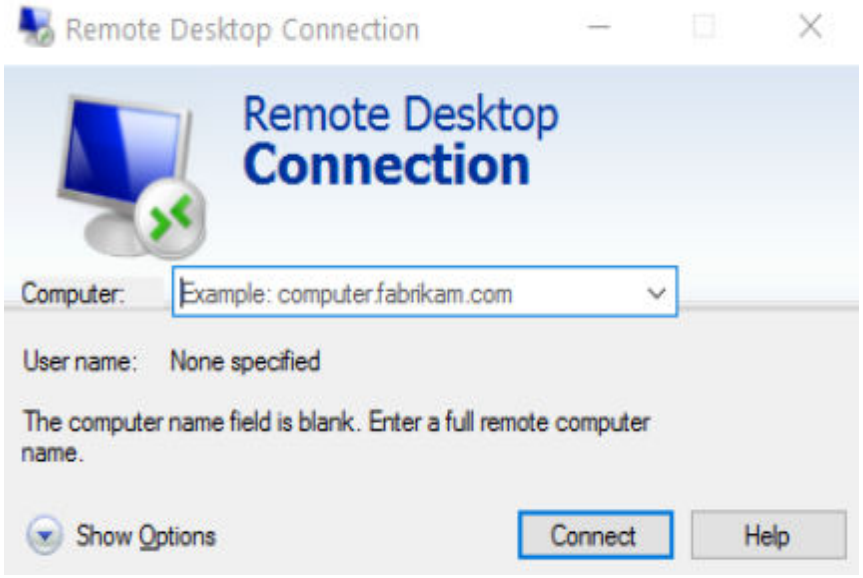

4. Enter the EIP and username (**Administrator** by default) of the target ECS.

   NOTE

   If you do not want to enter the username and password in follow-up logins, select **Allow me to save credentials**.

Figure 1-14 Remote Desktop Connection

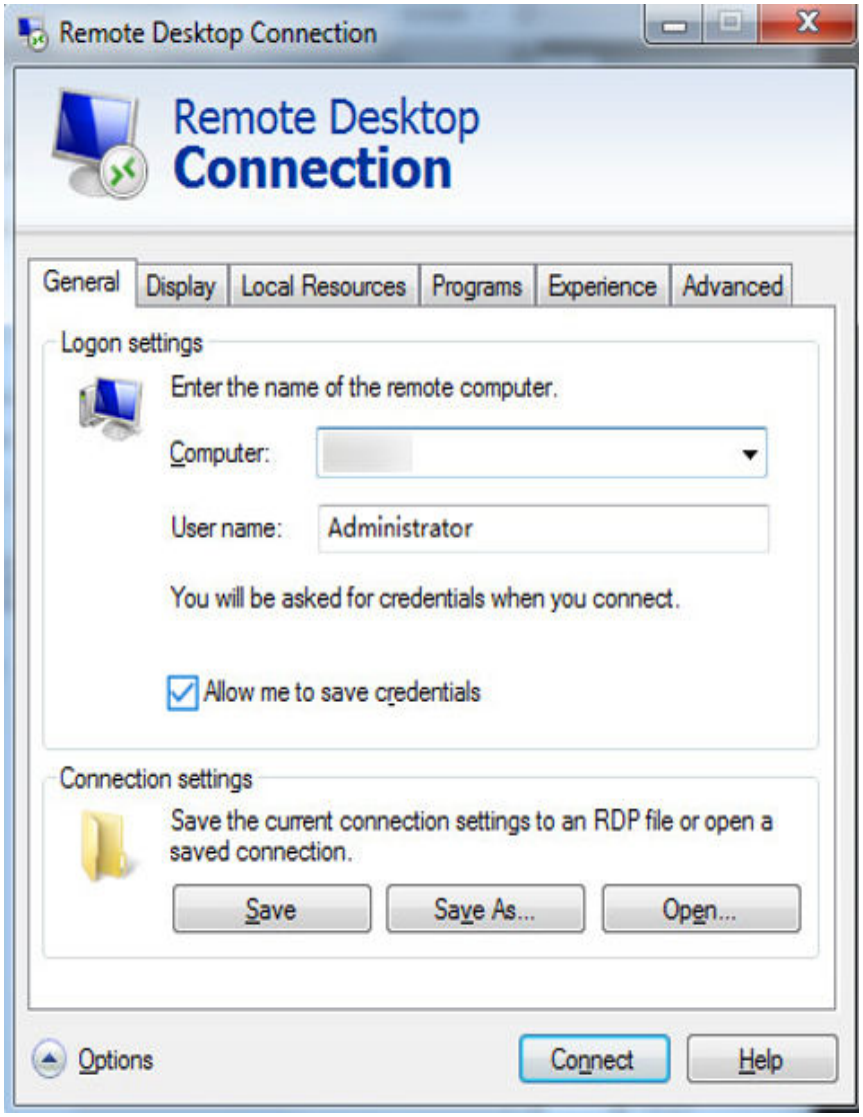

5. (Optional) To use local server resources in a remote session, configure parameters on the **Local Resources** tab.

   To copy data from the local server to your ECS, select **Clipboard**.

Figure 1-15 Clipboard

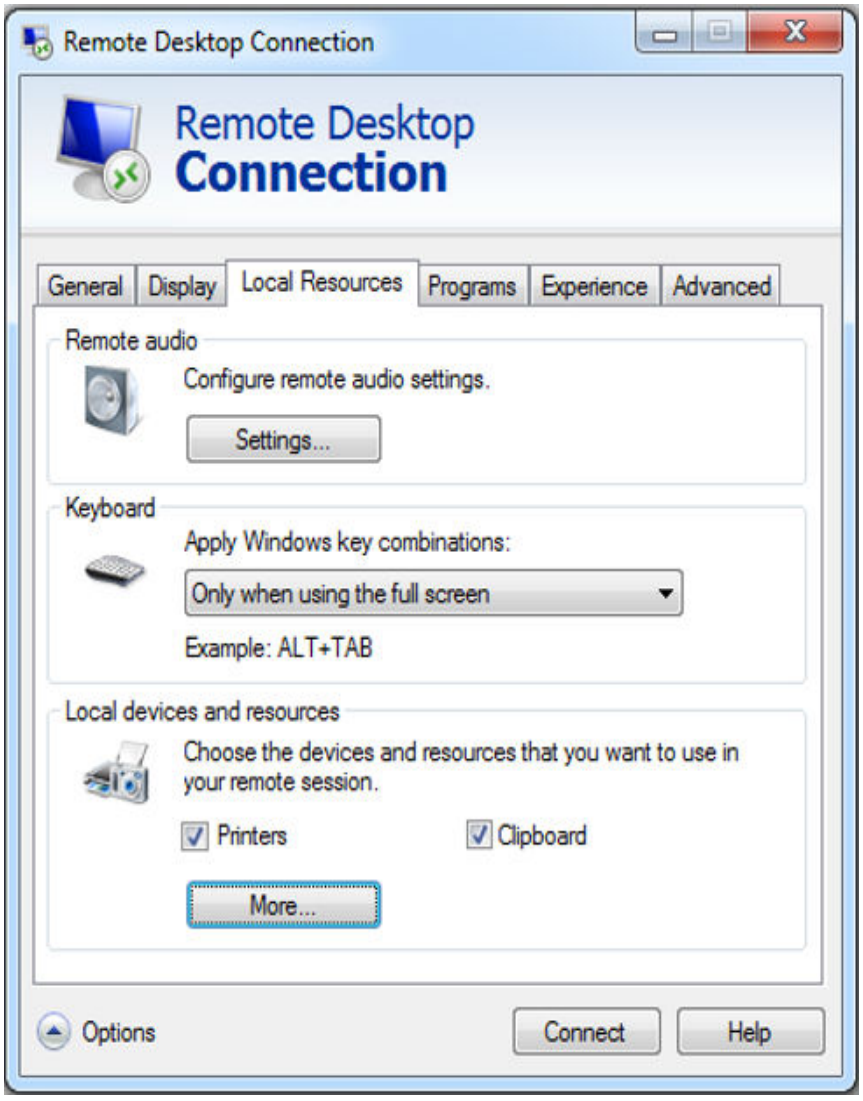

To copy files from the local server to your ECS, click **More** and select your desired disks.

Figure 1-16 Drives

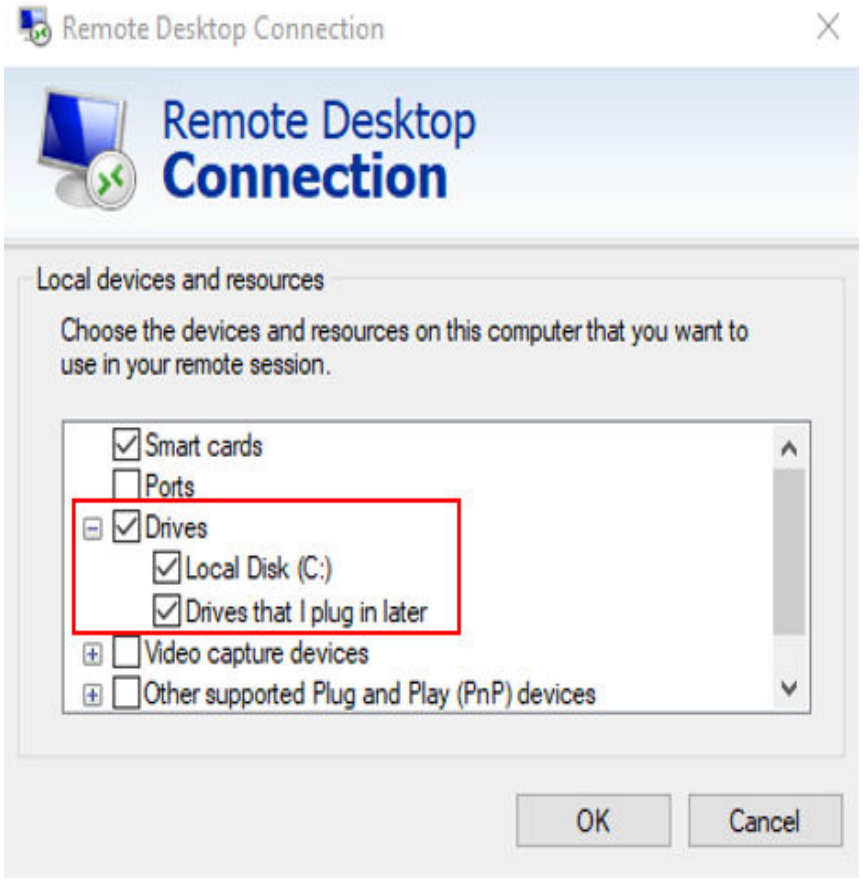

6. (Optional) Click the **Display** tab and then adjust the size of the remote desktop.

Figure 1-17 Adjusting the size of the desktop

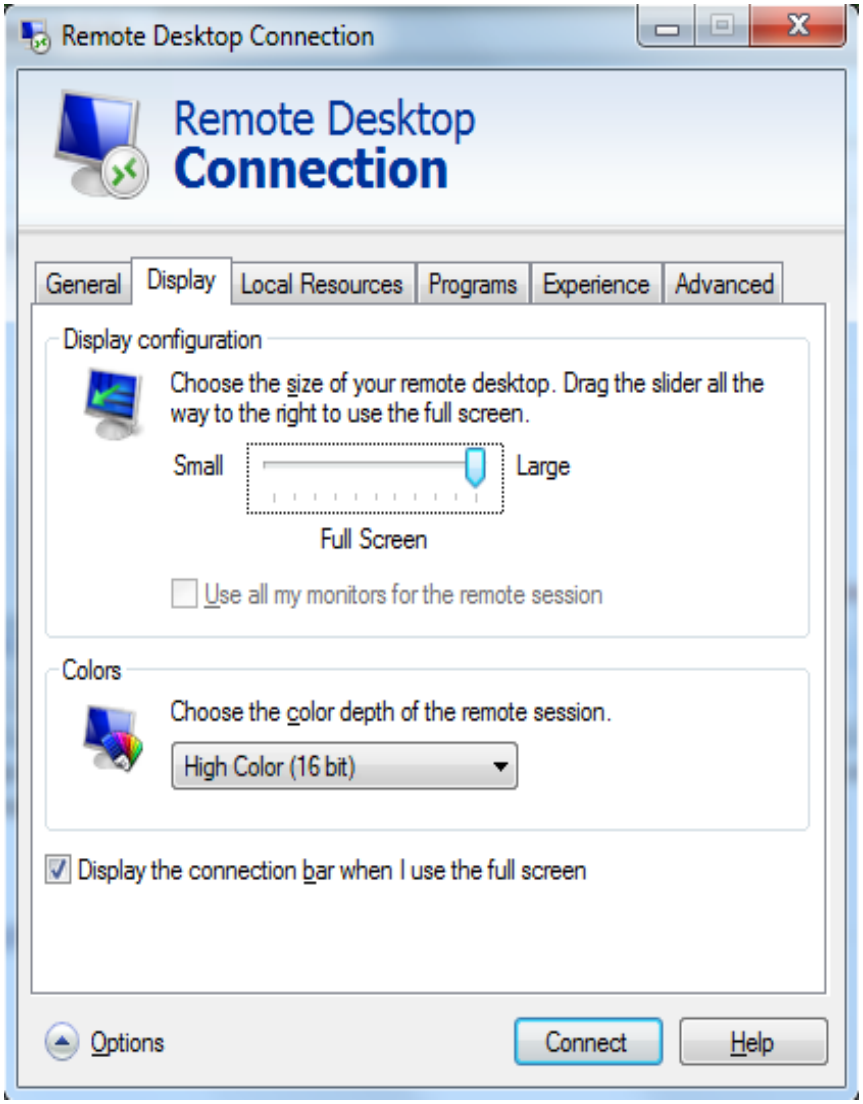

7. Click **Connect** and enter the login password as prompted to log in to the ECS.

   To ensure system security, change the login password after you log in to the ECS for the first time.

8. (Optional) Copy local files to the Windows ECS using clipboard. If the file size is greater than 2 GB, an error will occur.

   To resolve this issue, see **troubleshooting cases**.

Enabling RDP

For your first login, use VNC to log in and enable RDP for your ECS. Then, use MSTSC to log in.

NOTE

By default, RDP has been enabled on the ECSs created using a public image.

1. Log in to the Windows ECS using VNC.

   For details, see **Login Using VNC**.

2. Click **Start** in the task bar and choose **Control Panel** > **System and Security** > **System** > **Remote settings**.

   The **System Properties** dialog box is displayed.

   Figure 1-18 System Properties

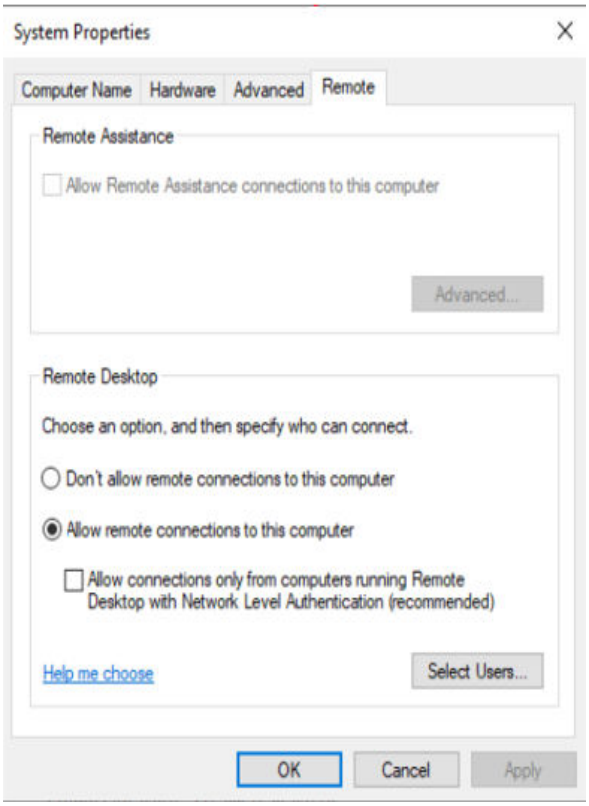

3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

Helpful Links

- **Login Password Resetting**
- **Multi-User Logins**
- **Remote Logins**

1.4.4 Logging In to a Windows ECS from a Linux Computer

Scenarios

This section describes how to log in to a Windows ECS from a Linux computer.

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS.

  An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.

- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs.
- Data can be exchanged between the login tool and the target ECS. For example, the default port 3389 is not blocked by the firewall.
- RDP has been enabled on the target ECS. By default, RDP has been enabled on the ECSs created using a public image. For instructions about how to enable RDP, see **Enabling RDP**.

Procedure

To log in to a Windows ECS from a local Linux computer, use a remote access tool, such as rdesktop.

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

   rdesktop

   If the message "command not found" is displayed, rdesktop is not installed. In such a case, obtain the rdesktop installation package at the **official rdesktop website**.

2. Run the following command to log in to the ECS:

   rdesktop -u *Username* **-p** *Password* **-g** *Resolution EIP*

   For example, run **rdesktop -u administrator -p password -g 1024*720 121.xx.xx.xx**.

   Table 1-13 Parameters in the remote login command

Parameter	Description
-u	Username, which defaults to **Administrator** for Windows ECSs
-p	Password for logging in to the Windows ECS
-f	Full screen by default, which can be switched using **Ctrl+Alt +Enter**
-g	Resolution, which uses an asterisk (*) to separate numbers. This parameter is optional. If it is not specified, the remote desktop is displayed in full screen by default, for example, **1024*720**.
EIP	EIP of the Windows ECS to be remotely logged in. Replace it with the EIP bound to your Windows ECS.

Enabling RDP

For your first login, use VNC to log in and enable RDP for your ECS. Then, use MSTSC to log in.

NOTE

By default, RDP has been enabled on the ECSs created using a public image.

1. Log in to the Windows ECS using VNC.
   For details, see **Login Using VNC**.
2. Click **Start** in the task bar and choose **Control Panel** > **System and Security** > **System** > **Remote settings**.
   The **System Properties** dialog box is displayed.

   Figure 1-19 System Properties

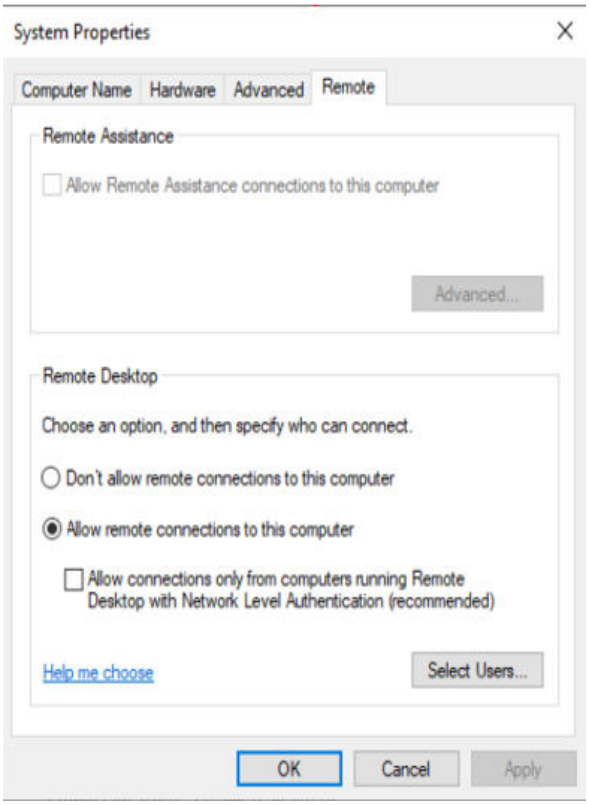

3. Click the **Remote** tab and select **Allow remote connections to this computer**.
4. Click **OK**.

1.4.5 Logging In to a Windows ECS from a Mobile Terminal

Scenarios

This section describes how to log in to an ECS running Windows Server 2012 R2 DataCenter 64bit from a mobile terminal via the Microsoft Remote Desktop client.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If the password is forgotten, reset the password by referring to **Resetting the Password for Logging In to an ECS on the Management Console**.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see **Configuring Security Group Rules**.

- **Microsoft Remote Desktop** has been installed on the mobile terminal.

Procedure

1. Start the Microsoft Remote Desktop client.
2. In the upper right corner of the **Remote Desktop** page, tap and select **Desktop**.

   Figure 1-20 Remote Desktop

   

3. On the **Add desktop** page, set login information and tap **SAVE**.
   - **PC name**: Enter the EIP bound to the target Windows ECS.
   - Perform the following operations to set **User name**:
     i. Tap **User name** and select **Add user account** from the drop-down list.

        The **Add user account** dialog box is displayed.

     ii. Enter username **administrator** and password for logging in to the Windows ECS and tap **SAVE**.

Figure 1-21 Setting the login information

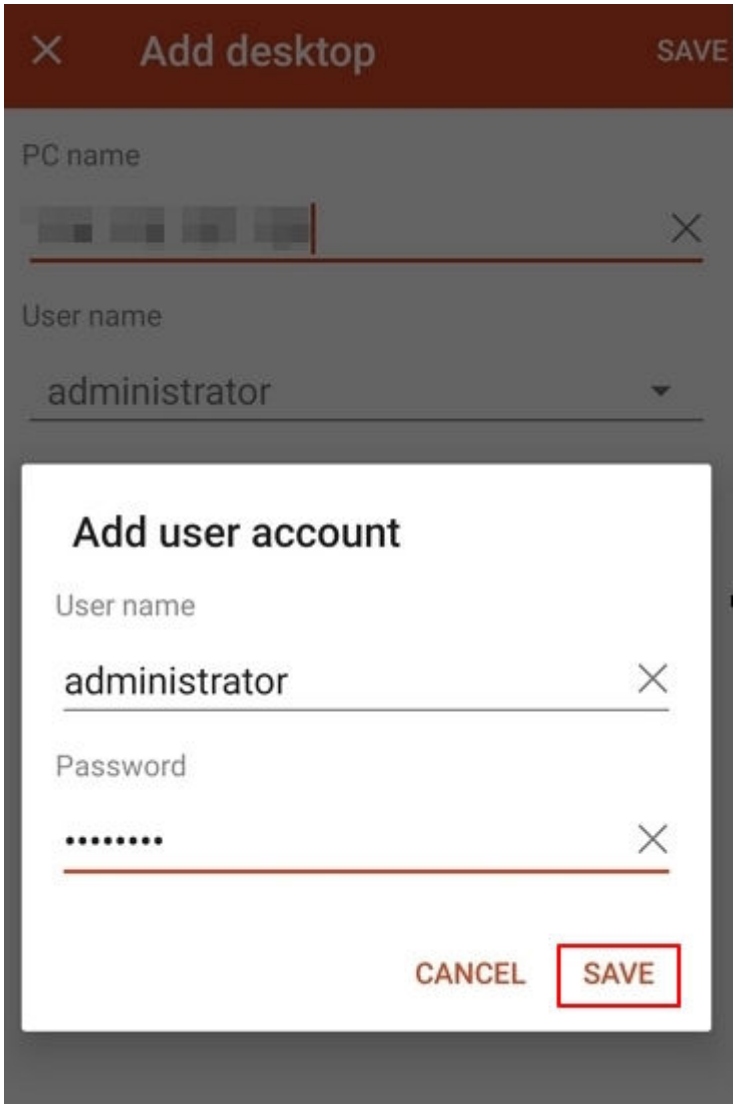

Figure 1-22 Saving the settings

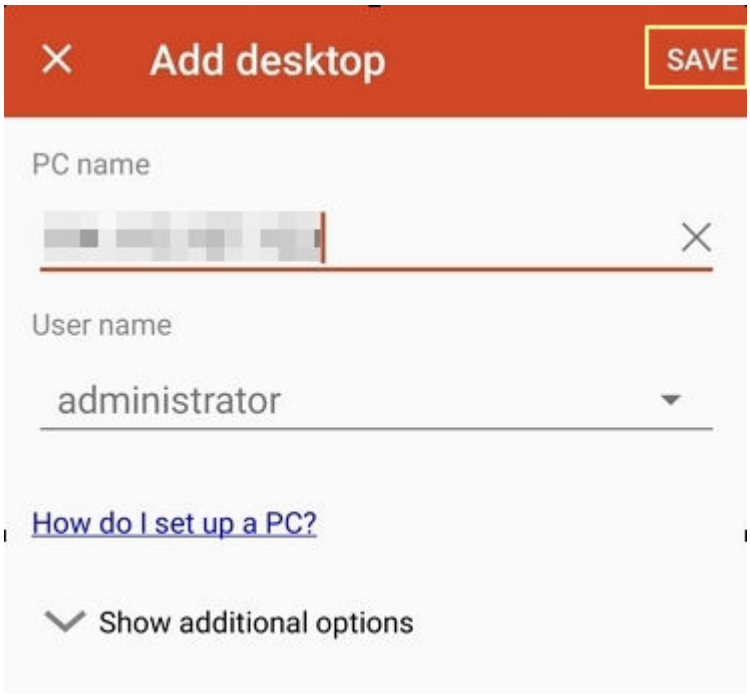

4. On the **Remote Desktop** page, tap the icon of the target Windows ECS.

Figure 1-23 Logging in to the Windows ECS

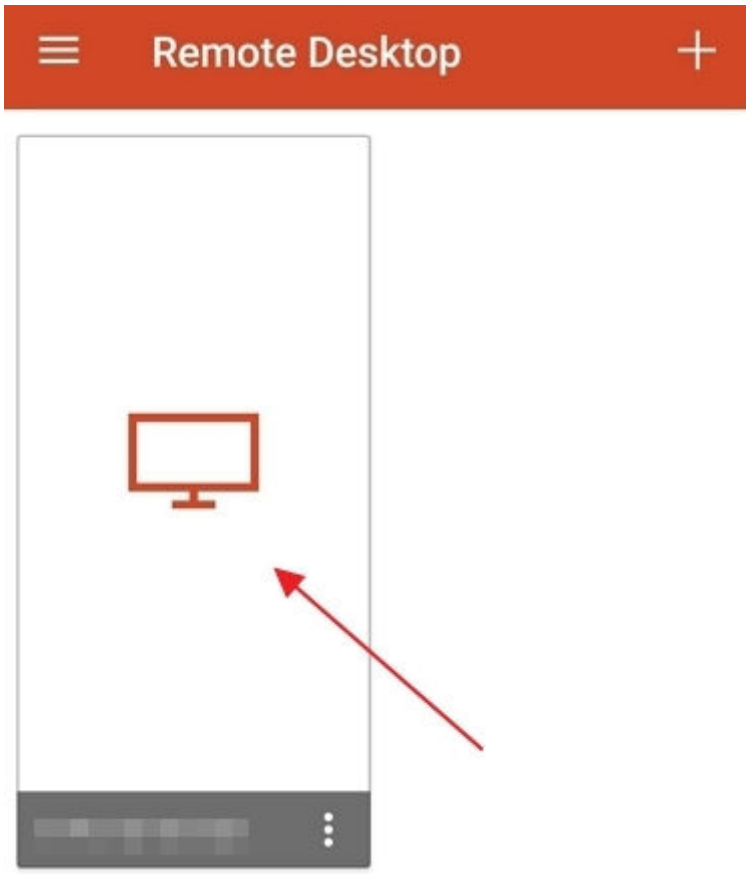

5. Confirm the information and tap **CONNECT**.

Figure 1-24 CONNECT

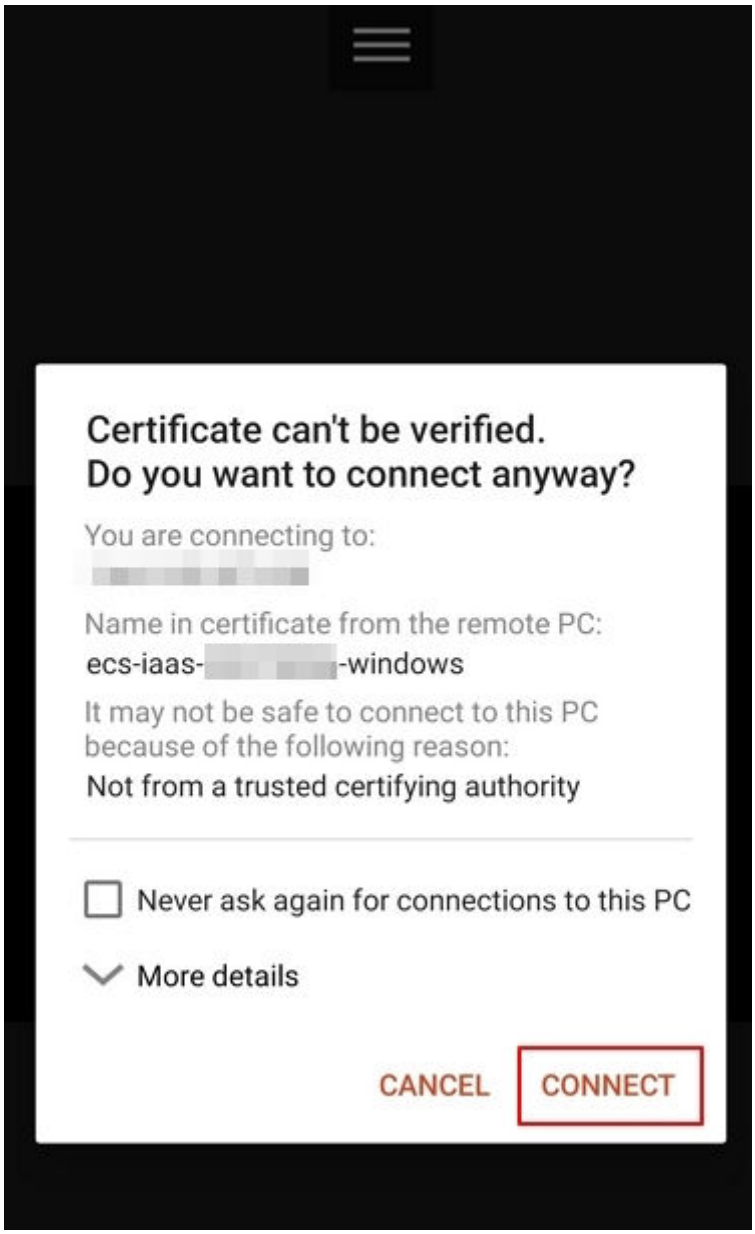

You have logged in to the Windows ECS.

Figure 1-25 Successful login

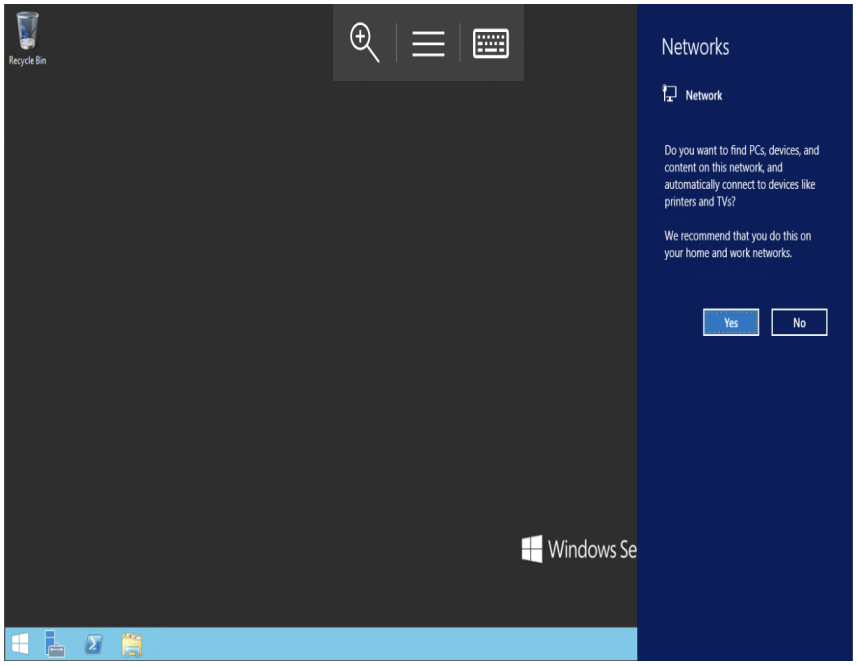

1.4.6 Logging In to a Windows ECS from a Mac

Scenarios

This section describes how to use a remote login tool to log in to a Windows ECS from a Mac. In this section, the remote login tool Microsoft Remote Desktop for Mac and the ECS running Windows Server 2012 R2 Data Center 64bit are used as an example.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If the password is forgotten, reset the password by referring to **Resetting the Password for Logging In to an ECS on the Management Console**.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see **Configuring Security Group Rules**.
- The remote access tool supported by Mac, such as Microsoft Remote Desktop for Mac has been installed. For details, see **Download Microsoft Remote Desktop for Mac**.

  Microsoft stopped providing the link for downloading the Remote Desktop client. You can download the beta version by visiting **Microsoft Remote Desktop Beta**.

Procedure

1. Start Microsoft Remote Desktop.

2. Click **Add Desktop**.

Figure 1-26 Add Desktop

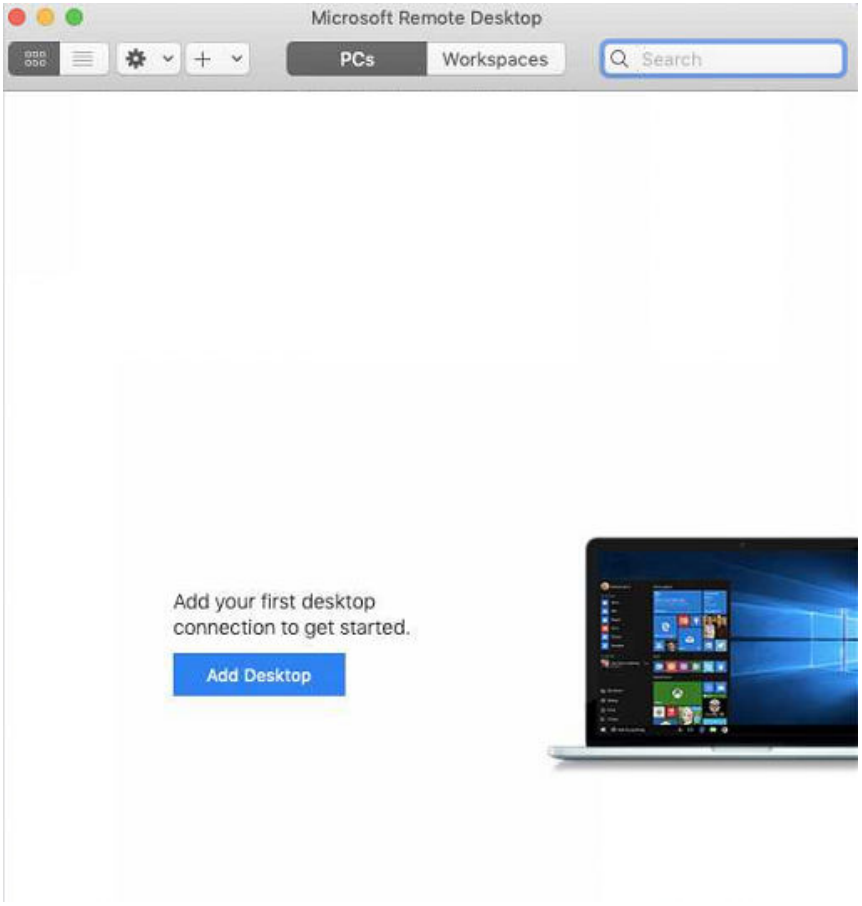

3. On the **Add PC** page, set login information.
   - **PC name**: Enter the EIP bound to the target Windows ECS.
   - **User account**: Select **Add user account** from the drop-down list.

     The **Add user account** dialog box is displayed.

     i. Enter username **administrator** and password for logging in to the Windows ECS and click **Add**.

Figure 1-27 Add user account

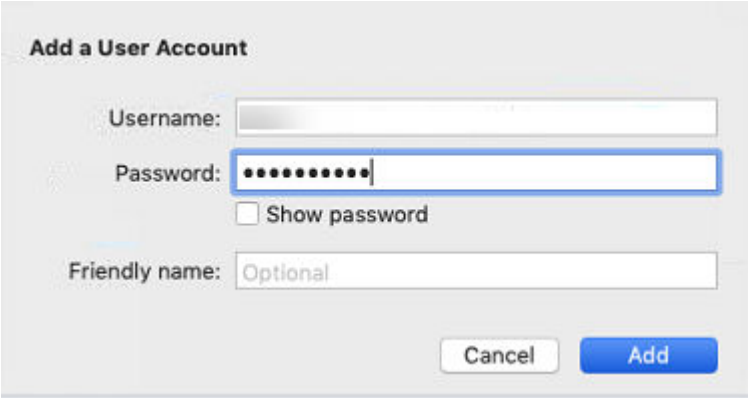

Figure 1-28 Add PC

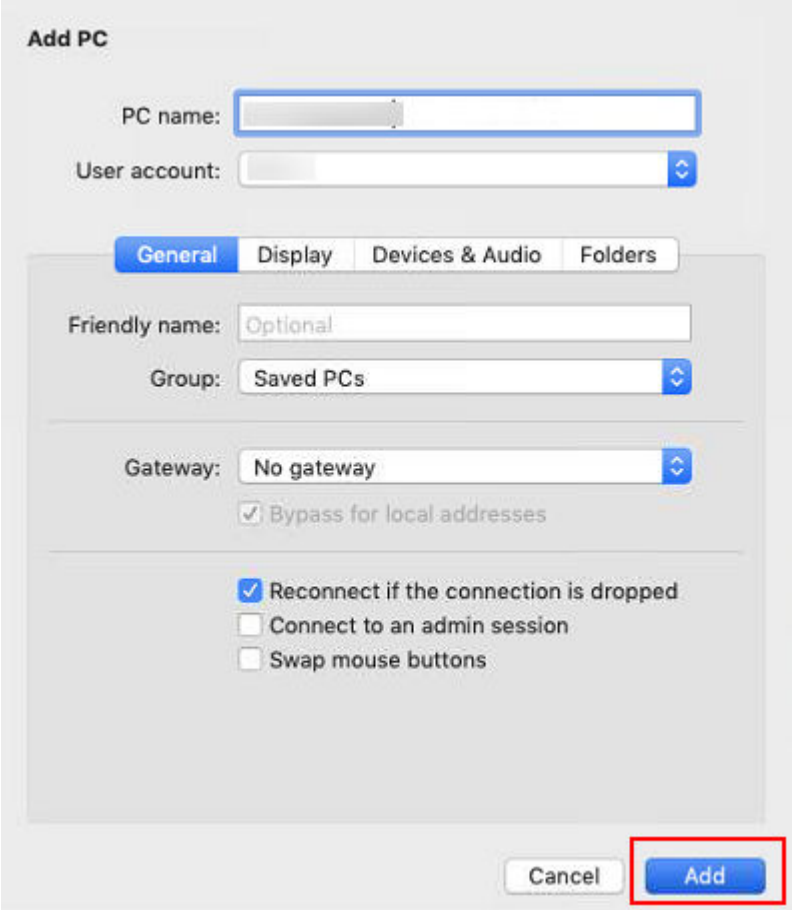

4. On the **Remote Desktop** page, double-click the icon of the target Windows ECS.

Figure 1-29 Double-click for login

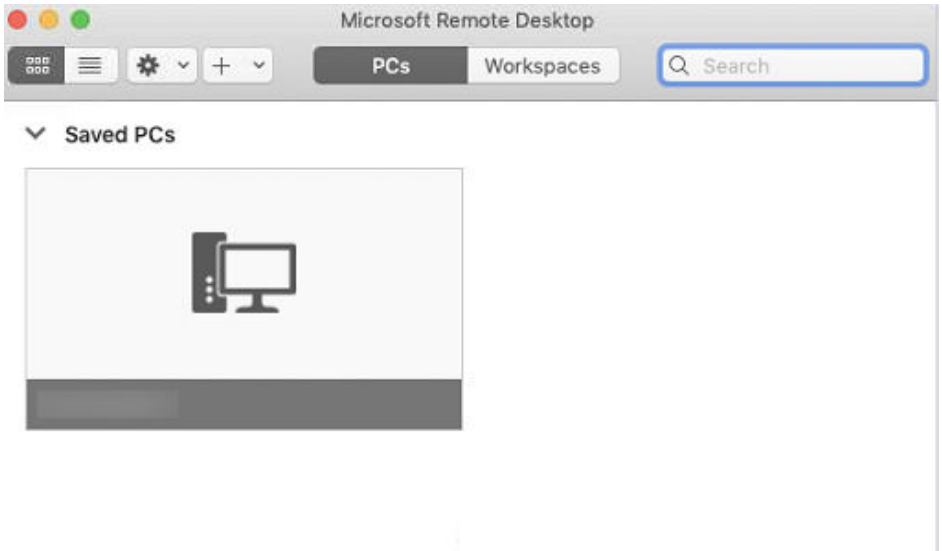

5. Confirm the information and click **Continue**.

   You have logged in to the Windows ECS.

Figure 1-30 Successful login

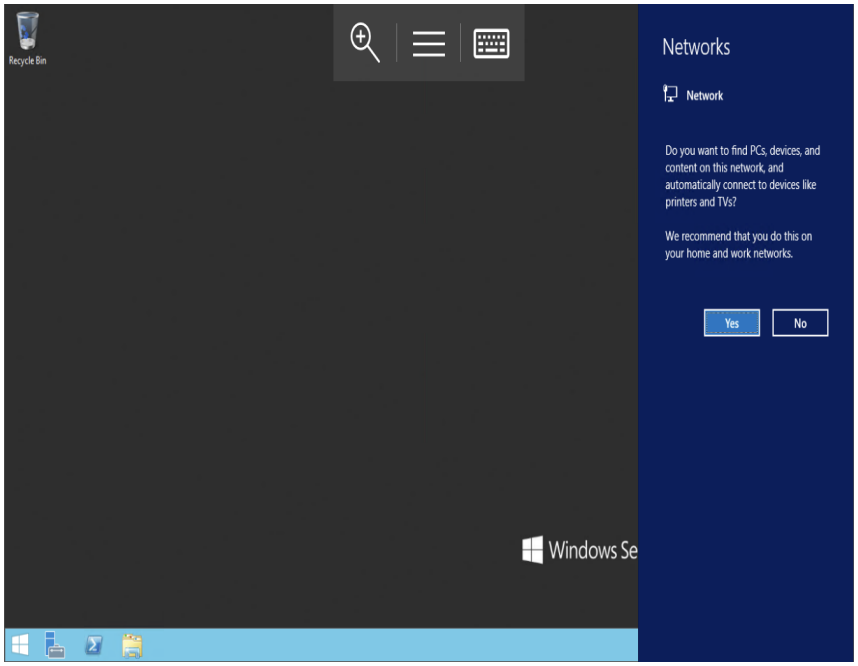

1.5 Logging In to a Linux ECS

1.5.1 Login Overview

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Linux ECS is **root**.
- If the login password is forgotten, reset the password on the ECS console.

  To reset a password, locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list. For details, see **Resetting the Password for Logging In to an ECS on the Management Console**.

Login Modes

You can choose from a variety of login modes based on your local OS type.

Table 1-14 Linux ECS login modes

ECS OS	Local OS	Connection Method	Requirement
Linux	Windows	(Recommended) Use CloudShell available on the management console to log in to the ECS. **Login Using CloudShell**	The target ECS has an EIP bound. (If you log in to an ECS through an intranet, for example, through VPN or Direct Connect, the ECS does not require an EIP.)
	Windows	Use a remote login tool, such as PuTTY or Xshell. • Password-authenticated: **Logging In to a Linux ECS from a Local Windows Server** • Key-pair-authenticated: **Logging In to a Linux ECS from a Local Windows Server**	
	Linux	Run commands. • Password-authenticated: **Logging In to a Linux ECS from a Local Linux Server** • Key-pair-authenticated: **Logging In to a Linux ECS from a Local Linux Server**	
	Mobile terminal	Use an SSH client tool, such as Termius or JuiceSSH, to log in to the ECS. **Logging In to a Linux ECS from a Mobile Terminal**	
	macOS	Use the terminal included in the macOS. **Logging In to a Linux ECS from a macOS Server**	
	Windows	Use the remote login function available on the management console. For details, see **Login Using VNC**.	No EIP is required.

Helpful Links

- **What Can I Do If I Forget My Password for Remote Login?**
- **Why Can't I Log In to My Linux ECS?**

1.5.2 Login Using CloudShell

Scenarios

This section describes how to use CloudShell provided on the management console to log in to an ECS.

For instructions about how to copy and paste data on CloudShell pages after the ECS login, see **Common CloudShell Operations**.

Constraints

For details about the supported regions, see **Function Overview**.

Prerequisites

- The ECS is running.
- The login port (port 22 by default) was opened. If a different port is required, you have used the default port to log in to the ECS and changed the port number.

  For details about how to change the remote login port, see **How Can I Change a Remote Login Port?** For details about how to configure security group rules, see **Configuring Security Group Rules**.
- The password for logging in to the target ECS has been set. If you did not set a password when creating the ECS, reset the password before logging in to the ECS.
- You can use CloudShell to connect to an ECS through a public or private network. When you choose to connect through a private network, service authorization is required.
  - If the **Service authorization** page is displayed, it means you have the Security Administrator permissions. Click **Agree**.

    The service authorization takes effect at the region level and is required only when you use CloudShell for the first time in a specific region.

    Figure 1-31 Service authorization

    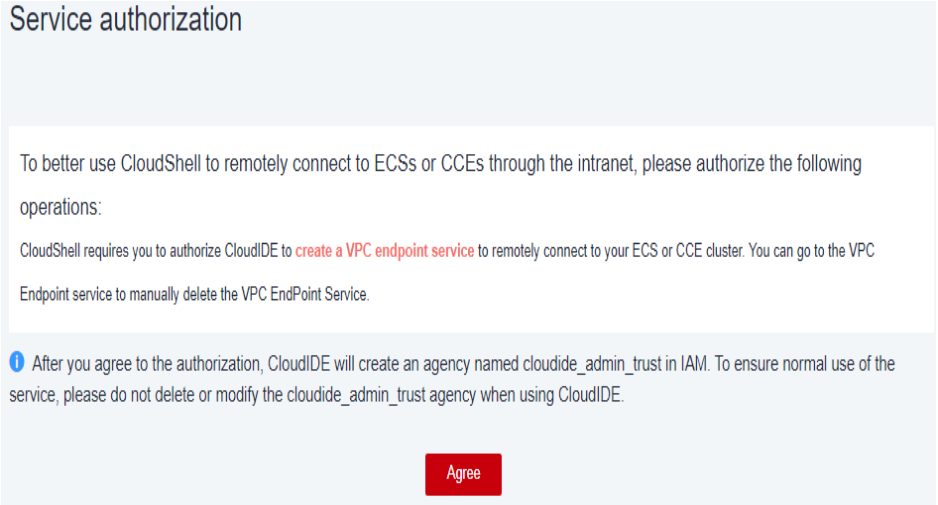

- If you do not have the Security Administrator permissions, a page will be displayed, requiring you to contact the administrator to assign permissions to you.

  Perform the following steps to assign permissions:

  i. Create a user group and assign the Security Administrator permissions to the user group. For details, see **Creating a User Group and Assigning Permissions**.

  ii. Add the user to the user group. For details, see **Adding Users to or Removing Users from a User Group**.

  NOTE

  When you use CloudShell to remotely connect to an ECS through a public network, service authorization is not required.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  and choose **Compute** > **Elastic Cloud Server**.
4. In the **Operation** column of the target ECS, click **Remote Login**.
5. In the **Logging In to a Linux ECS** dialog box, click **Log In** in **CloudShell-based Login**.
6. On the CloudShell page, configure the ECS information.

   Upon the first login, the CloudShell configuration wizard is displayed by default. Enter the ECS parameters to connect to the ECS.

   NOTE

   You can use the EIP or private IP address of the ECS to log in.

   - If you select the EIP bound to the ECS:

     i. In the CloudShell configuration wizard, set the port (22 by default), username, authentication mode, and password (or key) of the ECS.

     ii. Click **Connect** to log in to the ECS.

     If the system does not respond, the login password is incorrect or the password has not been set. In such a case, reset the password and attempt to log in to the ECS again.

Figure 1-32 CloudShell configuration wizard (EIP)

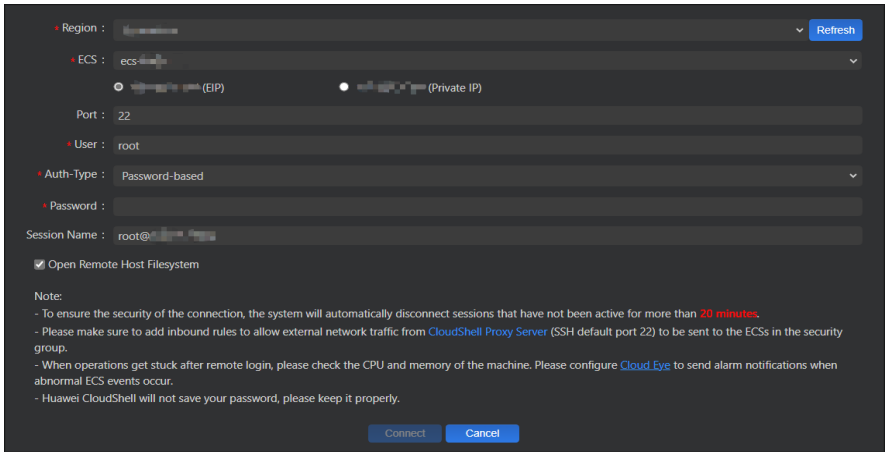

After the login is successful, the following figure is displayed.

Figure 1-33 Successful login

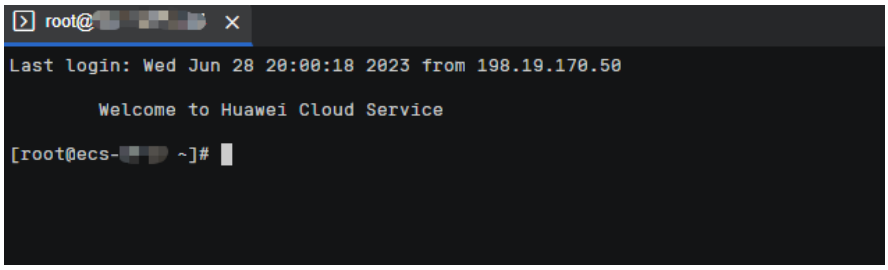

- If you select the private IP address of the ECS:

  i. Click **Go** to open the CloudShell configuration wizard.

  Figure 1-34 CloudShell configuration wizard (private IP) 1

  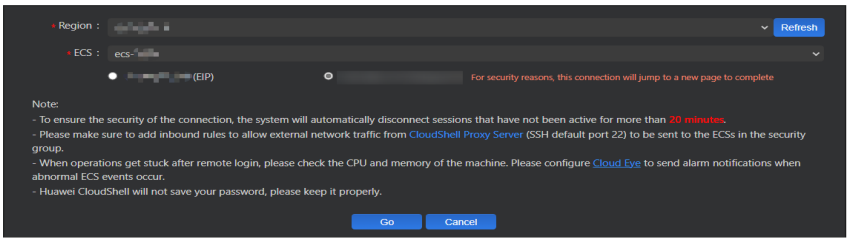

  ii. In the CloudShell configuration wizard, set the port (22 by default), username, authentication mode, and password (or key) of the ECS.

  iii. Click **Connect** to log in to the ECS.

  If the system does not respond, the login password is incorrect or the password has not been set. In such a case, reset the password and attempt to log in to the ECS again.

  After the login is successful, the following figure is displayed.

Figure 1-35 Successful login

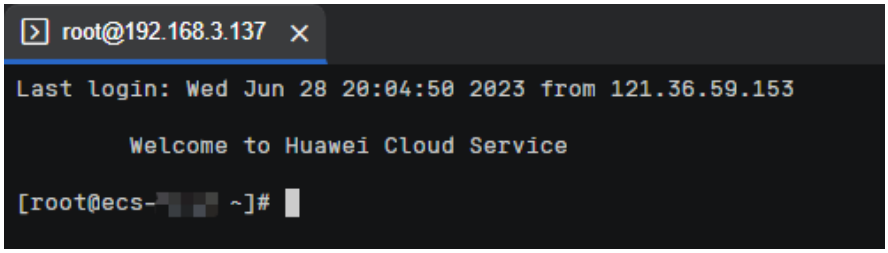

Common CloudShell Operations

- **New remote terminal**

  Choose **Terminal** > **New Terminal** to open a new terminal.
- **New session**

  Choose **Terminal** > **New Session** in the top navigation bar to open a new session.
- **Keyboard shortcuts**

  Use keyboard shortcuts to edit commands.

  Table 1-15 Keyboard shortcuts for CloudShell

Shortcut	Action
Ctrl+L	Moves the current line to the first line.
Ctrl+U	Clear the current line.
Ctrl+H	Delete one character forward.
Ctrl+A	Move the cursor to the beginning of the command line.
Ctrl+E	Move the cursor to the end of the command line.

- **Copy & Paste**

  Data can be copy-pasted across local and remote terminals by right-clicking the target file and choosing **Copy** and **Paste**, or using keyboard shortcuts **Ctrl +C** and **Ctrl+V**.
- **Historical records**

  Scroll up or down the terminal to view historical records. By default, only the latest 1000 lines of historical records are retained for terminals. This value is changeable.
- **Customized layout for multiple terminals**

  You can create multiple CloudShell terminals on the same page and drag panes to customize the layout.

1.5.3 Login Using VNC

Scenarios

This section describes how to use VNC provided on the management console to log in to an ECS.

For instructions about how to copy and paste data on VNC pages after the ECS login, see **Follow-up Procedure**.

NOTE

Before using remote login (VNC) provided on the management console to log in to a Linux ECS authenticated using a key pair, log in to the ECS **using an SSH key** and set a login password.

Constraints

- When you log in to an ECS using VNC, the system does not support copy and paste operations, reducing the efficiency of using the ECS. Unless otherwise specified, you are advised to log in to the ECS using SSH. For details, see **Login Using an SSH Key** and **Login Using an SSH Password**.

Prerequisites

You have used an SSH key to log in to the Linux ECS authenticated using a key pair and set a login password.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the **Operation** column of the target ECS, click **Remote Login**.

Figure 1-36 Remote Login

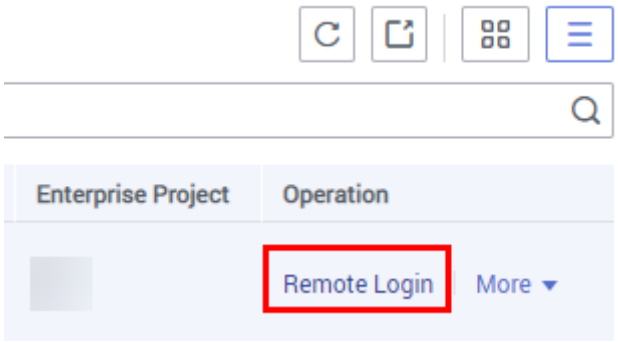

5. (Optional) When the system displays "Press CTRL+ALT+DELETE to log on", click **Ctrl+Alt+Del** in the upper part of the remote login page to log in to the ECS.

NOTE

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

6. Enter the ECS password as prompted.

Figure 1-37 Username (root as an example) and password

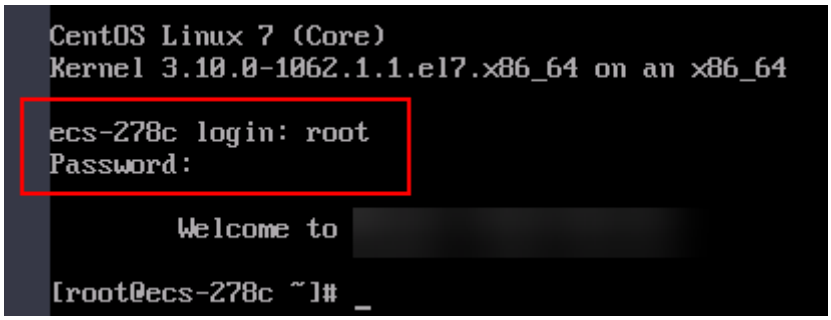

Follow-up Procedure

Local commands can be copied to an ECS. To do so, perform the following operations:

1. Log in to the ECS using VNC.
2. Click **Paste & Send** in the top area of the page.

Figure 1-38 Paste & Send

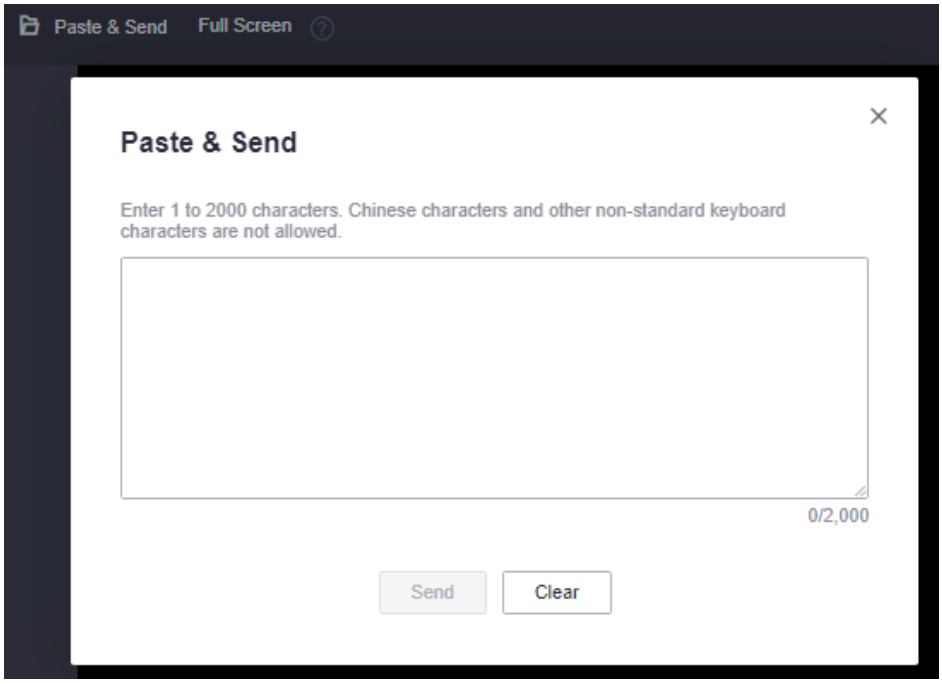

3. Press **Ctrl+C** to copy data from the local computer.
4. Press **Ctrl+V** to paste the local data to the **Paste & Send** window.
5. Click **Send**.

   Send the copied data to the CLI.

NOTE

There is a low probability that data is lost when you use Input Commands on the VNC page of a GUI-based Linux ECS. This is because the number of ECS vCPUs fails to meet GUI requirements. In such a case, it is a good practice to send a maximum of 5 characters at a time or switch from GUI to CLI (also called text interface), and then use the command input function.

Helpful Links

- **What Can I Do If I Forget My Password for Remote Login?**
- **Why Can't I Log In to My Linux ECS?**

1.5.4 Login Using an SSH Key

Scenarios

This section describes how to use an SSH key pair to remotely log in to a Linux ECS from a Windows and a Linux server, respectively.

Prerequisites

- You have obtained the private key file used for creating the ECS. For details about how to create a key pair, see **(Recommended) Creating a Key Pair on the Management Console**.
- You have bound an EIP to the ECS. For details, see **Viewing ECS Details (List View)**.
- You have configured the inbound rules of the security group. For details, see **Configuring Security Group Rules**.
- The network connection between the login tool (PuTTY) and the target ECS is normal. For example, the default port 22 is not blocked by the firewall.

Logging In to a Linux ECS from a Local Windows Server

You have two methods to log in to a Linux ECS from a local Windows server.

Method 1: Use PuTTY to log in to the ECS.

The following example shows how to convert the format of a private key file and use an SSH key to access a Linux ECS.

Figure 1-39 Accessing a Linux ECS using an SSH key

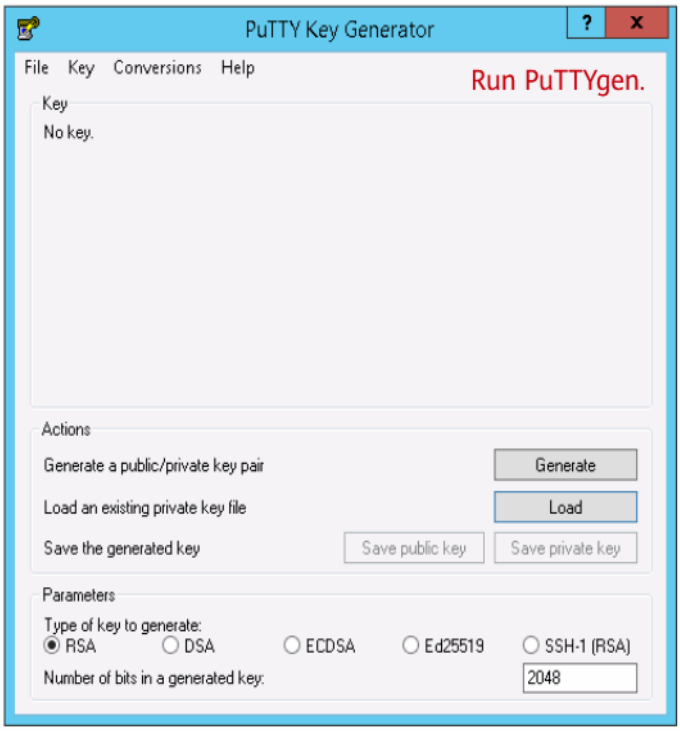

The following operations use PuTTY as an example. Before using PuTTY to log in, make sure that the private key file has been converted to .ppk format.

1. Check whether the private key file has been converted to .ppk format.
   - If yes, go to step 7.
   - If no, go to step 2.
2. Visit the following website and download PuTTY and PuTTYgen:

   https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html

   NOTE

   PuTTYgen is a key generator, which is used to create a key pair that consists of a public key and a private key for PuTTY.
3. Run PuTTYgen.
4. In the **Actions** pane, click **Load** and import the private key file that you stored during ECS creation.

   Ensure that the format of **All files (*.*)** is selected.
5. Click **Save private key**.
6. Save the converted private key, for example, **kp-123.ppk**, to the local computer.
7. Double-click **PUTTY.EXE**. The **PuTTY Configuration** page is displayed.
8. Choose **Session** and enter the EIP of the ECS under **Host Name (or IP address)**.

Figure 1-40 Configuring the EIP

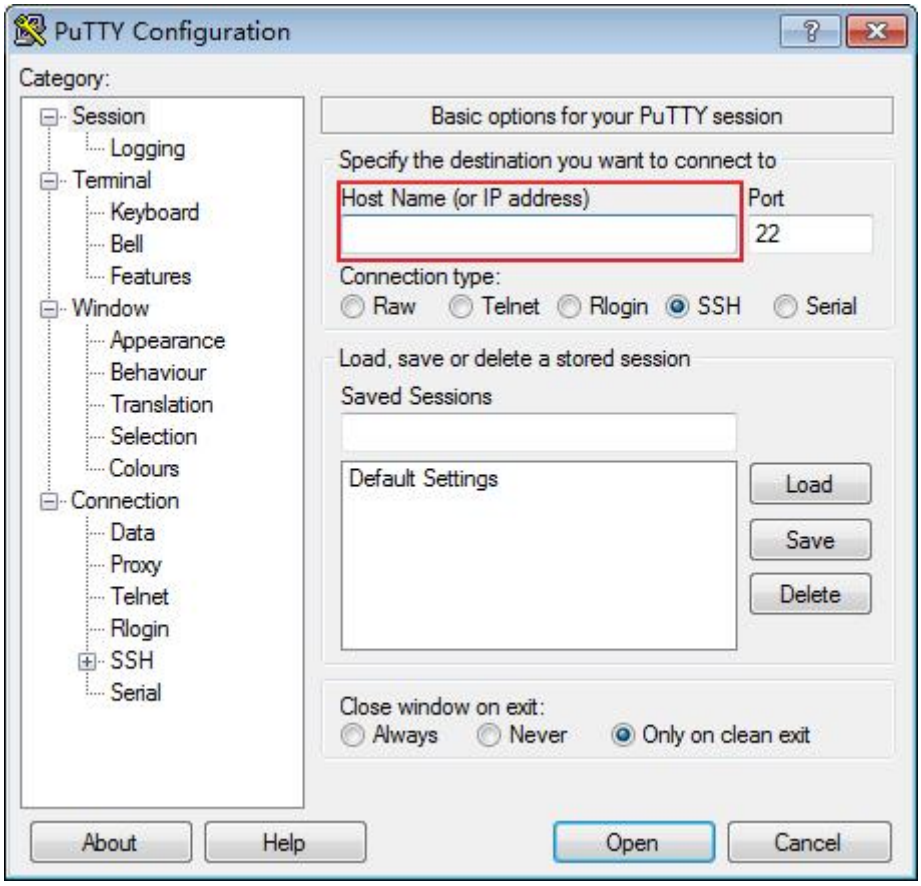

9. Choose **Connection** > **Data**. Enter the image username in **Auto-login username**.

   NOTE

   When you log in to an ECS using an SSH key:

   - The image username is **core** for a CoreOS public image.
   - The image username is **root** for a non-CoreOS public image.

10. Choose **Connection** > **SSH** > **Auth**. In the last configuration item **Private key file for authentication**, click **Browse** and select the private key converted in step 6.
11. Click **Open** to log in to the ECS.

Method 2: Use Xshell to log in to the ECS.

1. Start the Xshell tool.
2. Run the following command using the EIP to remotely log in to the ECS through SSH:

   ssh *Username@EIP*

   NOTE

   When you log in to an ECS using an SSH key:

   - The image username is **core** for a CoreOS public image.
   - The image username is **root** for a non-CoreOS public image.

3. (Optional) If the system displays the **SSH Security Warning** dialog box, click **Accept & Save**.

Figure 1-41 SSH Security Warning

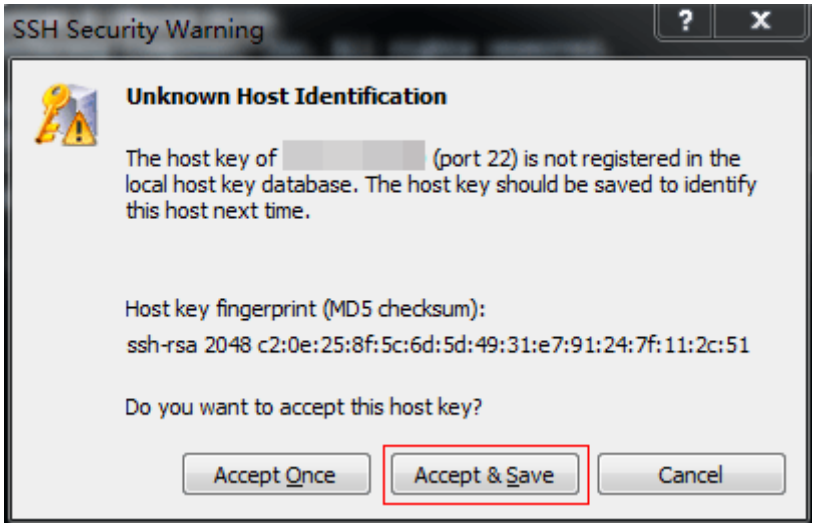

4. Select **Public Key** and click **Browse** beside the user key text box.
5. In the user key dialog box, click **Import**.
6. Select the locally stored key file and click **Open**.
7. Click **OK** to log in to the ECS.

Logging In to a Linux ECS from a Local Linux Server

To log in to the Linux ECS from local Linux, perform the operations described in this section. The following operations use private key file **kp-123.pem** as an example to log in to the ECS. The name of your private key file may differ.

1. On the Linux CLI, run the following command to change operation permissions:

   chmod 400 /*path*/kp-123.pem

   NOTE

   In the preceding command, replace *path* with the actual path where the key file is saved.

2. Run the following command to log in to the ECS:

   ssh -i /*path*/kp-123.pem *Default username*@*EIP*

   For example, if the default username is **root** and the EIP is **123.123.123.123**, run the following command:

   ssh -i /*path*/kp-123.pem root@123.123.123.123

   NOTE

   In the preceding command:

   - *path* refers to the path under which the key file is stored.
   - *EIP* is the EIP bound to the ECS.

Follow-up Procedure

- After logging in to the ECS using the SSH key, you can set a password (by using the **passwd** command) to log in to the ECS using VNC.

Helpful Links

- **What Can I Do If I Forget My Password for Remote Login?**
- **Why Can't I Log In to My Linux ECS?**

1.5.5 Login Using an SSH Password

Scenarios

This section describes how to remotely log in to a Linux ECS using an SSH password from a Windows and a Linux server, respectively.

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.
- 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**.
- The network connection between the login tool (PuTTY) and the target ECS is normal. For example, the default port 22 is not blocked by the firewall.

Logging In to a Linux ECS from a Local Windows Server

To log in to a Linux ECS from a local Windows server, perform the operations below.

The following example shows how to access a Linux ECS using an SSH password.

Figure 1-42 Accessing a Linux ECS using an SSH password

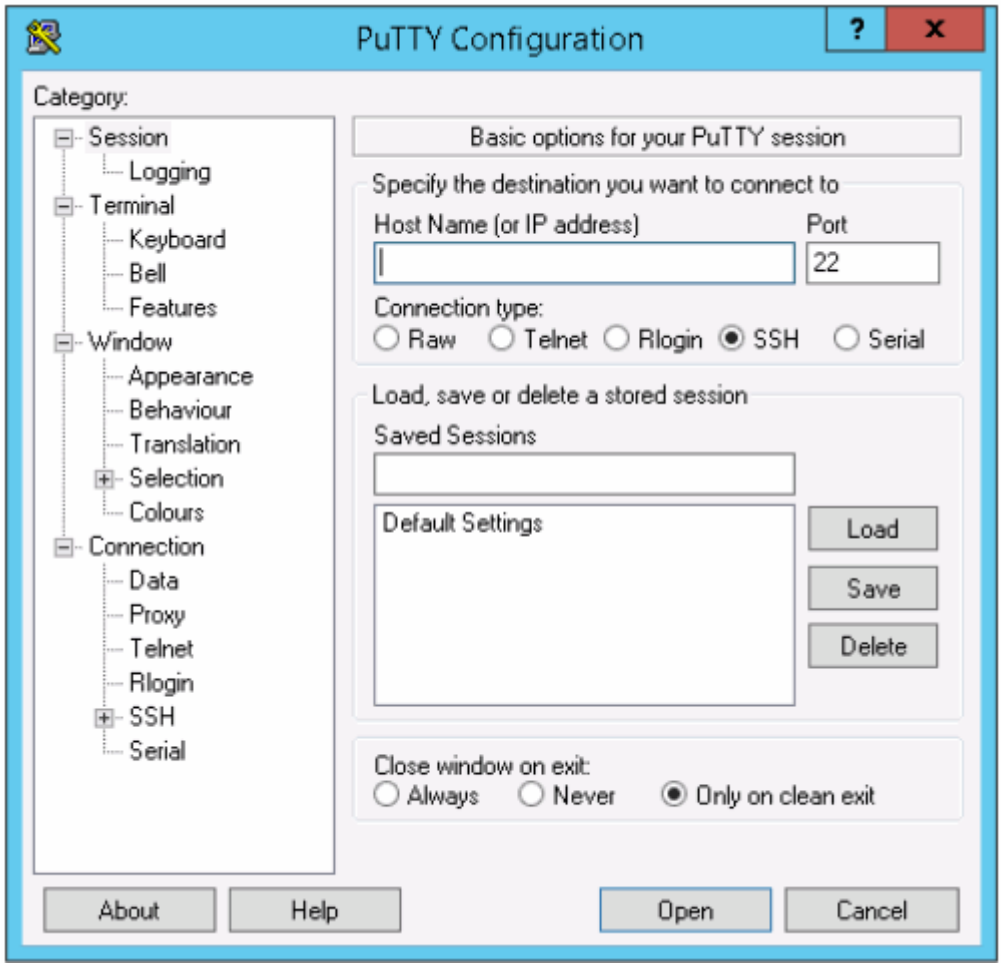

The following operations use PuTTY as an example to log in to the ECS.

1. Visit the following website and download PuTTY and PuTTYgen:
   https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html
2. Run PuTTY.
3. Choose **Session**.
   a. **Host Name (or IP address)**: Enter the EIP bound to the ECS.
   b. **Port**: Enter **22**.
   c. **Connection type**: Click **SSH**.
   d. **Saved Sessions**: Enter the task name, which can be clicked for remote connection when you use PuTTY next time.

Figure 1-43 Session

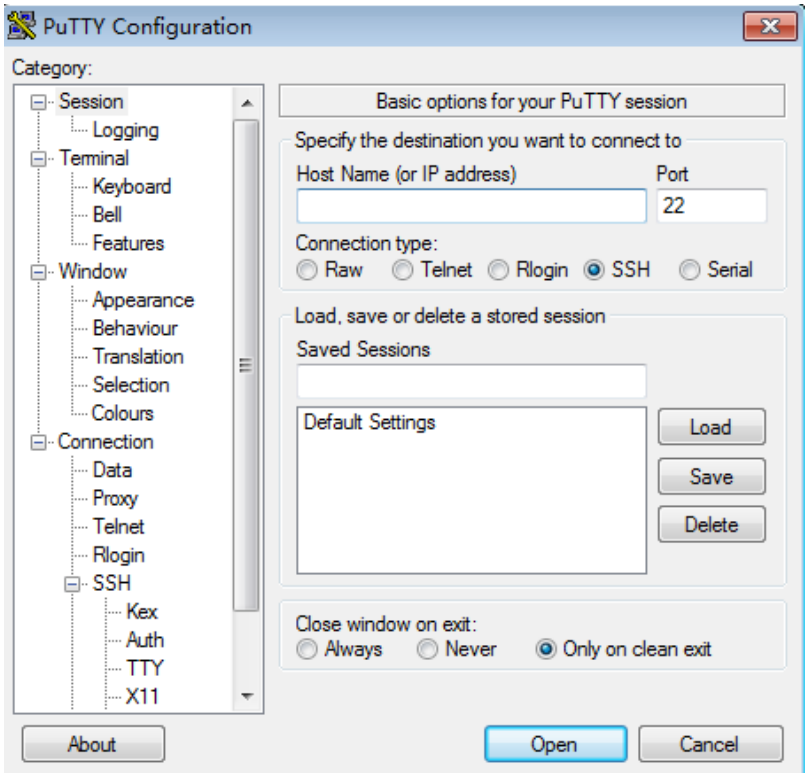

4. Choose **Window**. Then, select **UTF-8** for **Received data assumed to be in which character set:** in **Translation**.
5. Click **Open**.

   If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.
6. After the SSH connection to the ECS is set up, enter the username and password as prompted to log in to the ECS.

   NOTE

   The username and password for the first login to the ECS created using a public image (including CoreOS) are as follows:

   - Username: **root**
   - Password: the one you set when you purchased the ECS

     If you did not set a password when purchasing the ECS, see **Resetting the Password for Logging In to an ECS on the Management Console**.

Logging In to a Linux ECS from a Local Linux Server

To log in to a Linux ECS from a local Linux server, perform the operations below.

1. On the Linux CLI, run the following command to log in to the ECS:

   ssh *xx.xx.xx.xx*

   NOTE

   xx.xx.xx.xx indicates the EIP bound to the ECS.

2. Verify the SSH fingerprint of the ECS and enter **yes**.
```
The authenticity of host 'xx.xx.xx.xx (xx.xx.xx.xx)' can't be established.
ECDSA key fingerprint is SHA256:rnKuzrUSYS03MCoaxxxxxxxxxxxxxxxxxxxxxxxxxxx.
ECDSA key fingerprint is MD5:cf:64:5b:5e:74:30:xx:xx:xx:xx:xx:xx:xx:xx:xx:xx.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'xx.xx.xx.xx' (ECDSA) to the list of known hosts.
```
3. Enter the password for logging in to ECS.
```
root@xx.xx.xx.xx's password:

    Welcome to Huawei Cloud Service
```

Helpful Links

- **What Can I Do If I Forget My Password for Remote Login?**
- **Why Can't I Log In to My Linux ECS?**

1.5.6 Logging In to a Linux ECS from a Mobile Terminal

Scenarios

This section describes how to access a Linux ECS from a mobile terminal.

- For instructions about how to log in to a Linux ECS from an iOS terminal through iTerminal-SSH Telnet, see **Logging In to a Linux ECS from an iOS Terminal**.
- For instructions about how to log in to a Linux ECS from an Android terminal through JuiceSSH, see **Logging In to a Linux ECS from an Android Terminal**.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If the password is forgotten, reset the password by referring to **Resetting the Password for Logging In to an ECS on the Management Console**.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.
- 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**.

Logging In to a Linux ECS from an iOS Terminal

Before performing the operation, make sure that you have installed an SSH client tool, for example, Termius, on the iOS terminal. In this example, the Linux ECS runs CentOS 7.6, and it is authenticated using a username and password.

1. Start Termius and tap **New Host**.

Figure 1-44 New Host

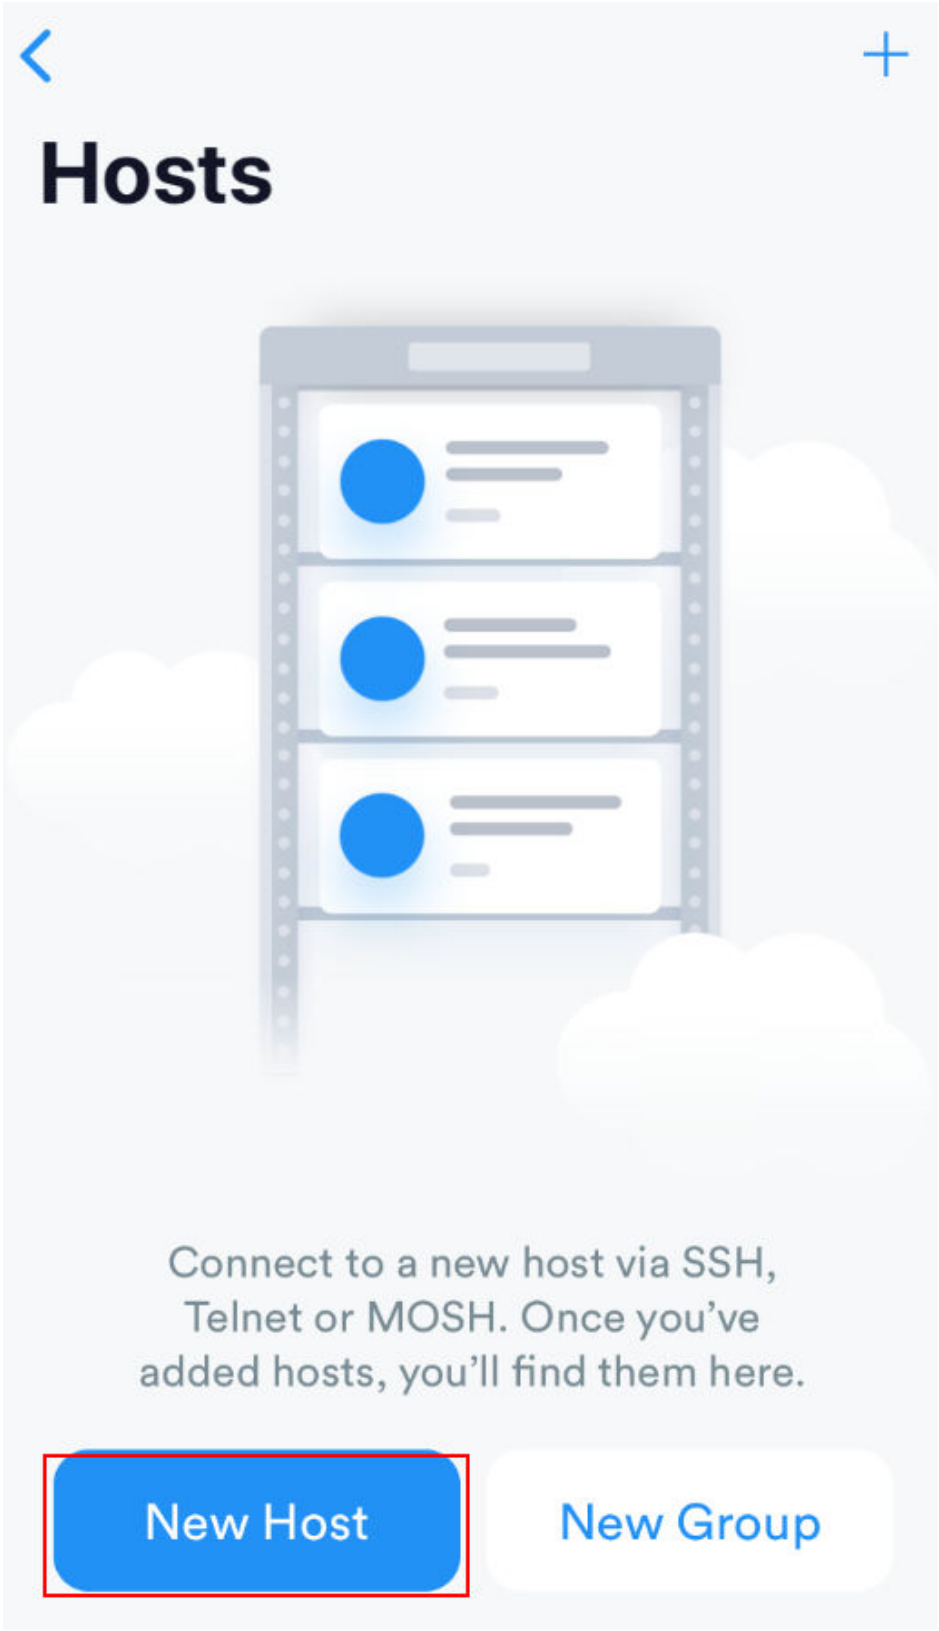

2. On the **New Host** page, set the following parameters:
   - **Alias**: Enter the hostname. In this example, set this parameter to **ecs01**.

- **Hostname**: Enter the EIP bound to the target ECS.
- **Use SSH**: Enable it.
- **Host**: Enter the EIP bound to the target ECS.
- **Port**: Enter port number **22**.
- **Username**: Enter **root**.
- **Password**: Enter the login password.

Figure 1-45 Setting parameters

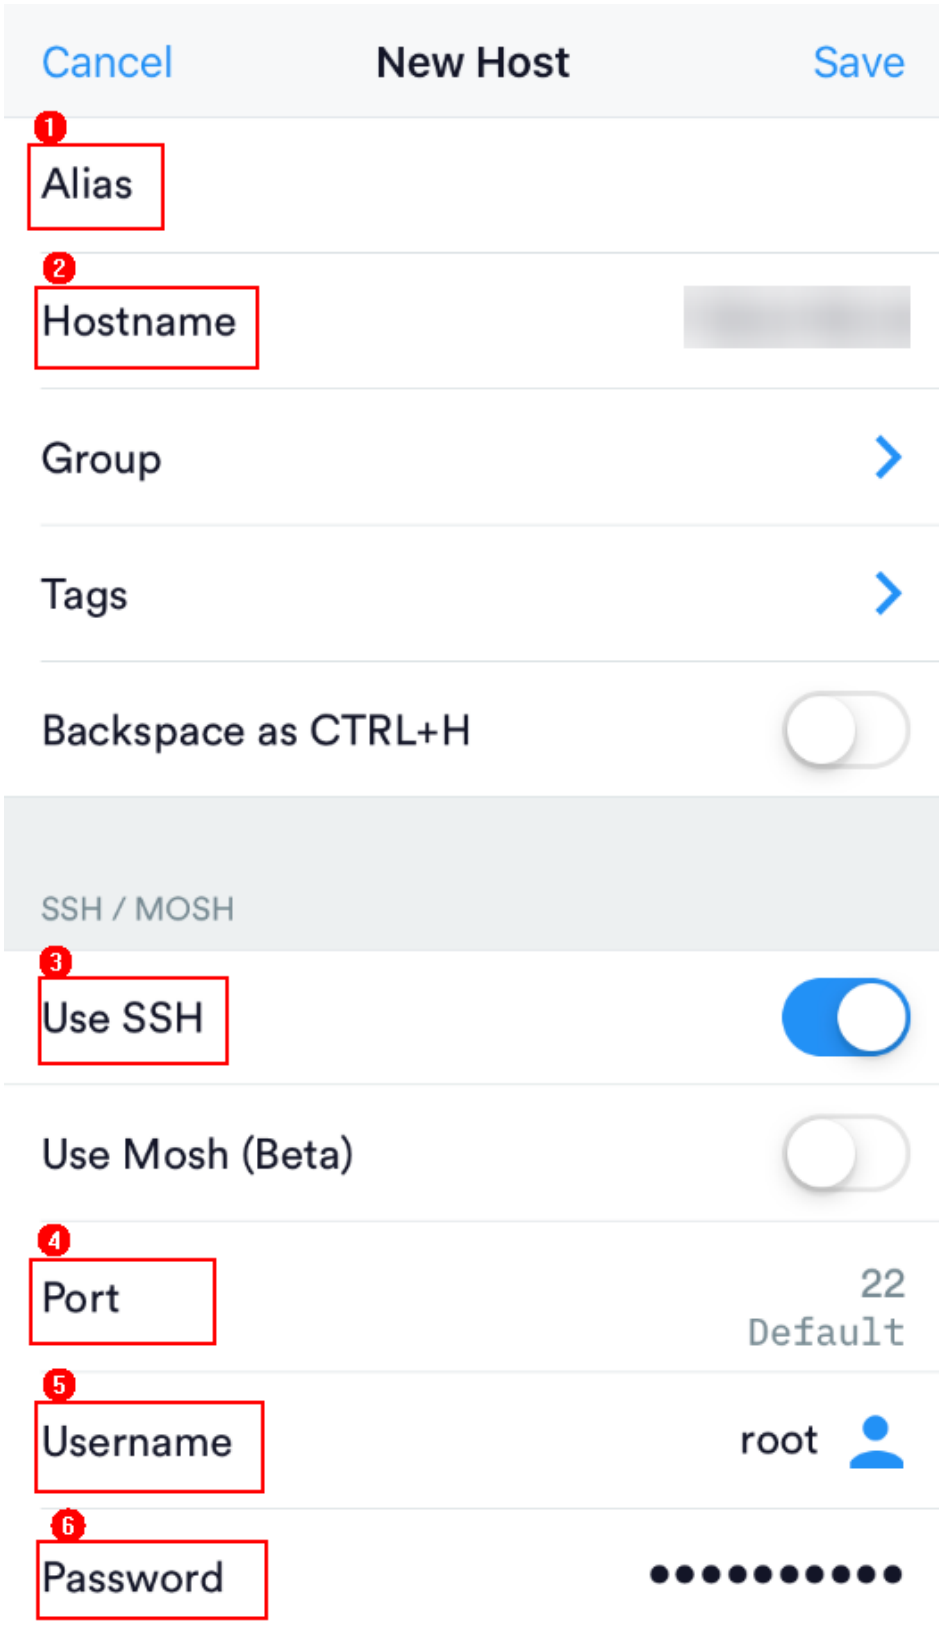

3. Tap **Save** in the upper right corner of the page to save the login settings. On the **Hosts** page, tap the name of the connection.

Figure 1-46 Login information

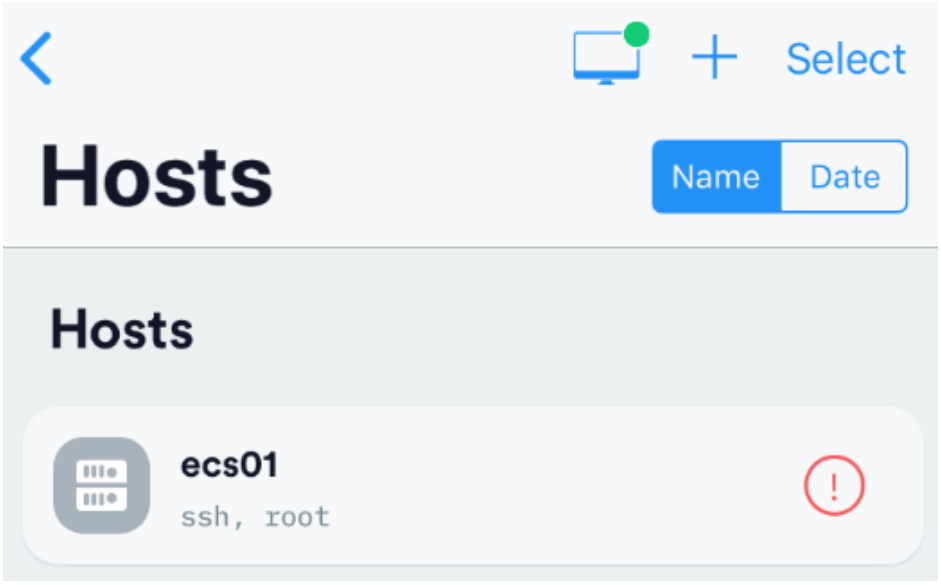

If the following page is displayed, you have connected to the Linux ECS.

Figure 1-47 Connected

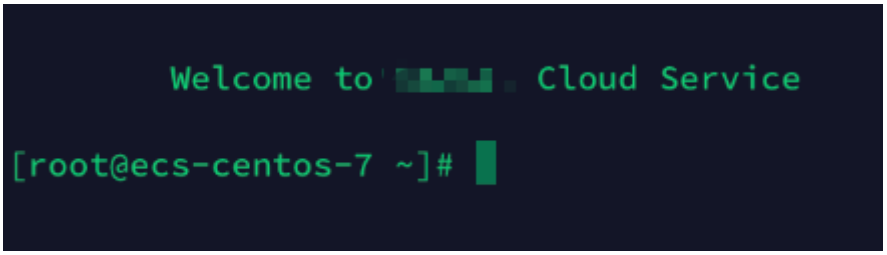

Logging In to a Linux ECS from an Android Terminal

Before performing the operation, make sure that you have installed JuiceSSH on the Android terminal. In this example, the Linux ECS runs CentOS 7.6, and it is authenticated using a username and password.

1. Start JuiceSSH and tap **Connections**.

Figure 1-48 Starting JuiceSSH

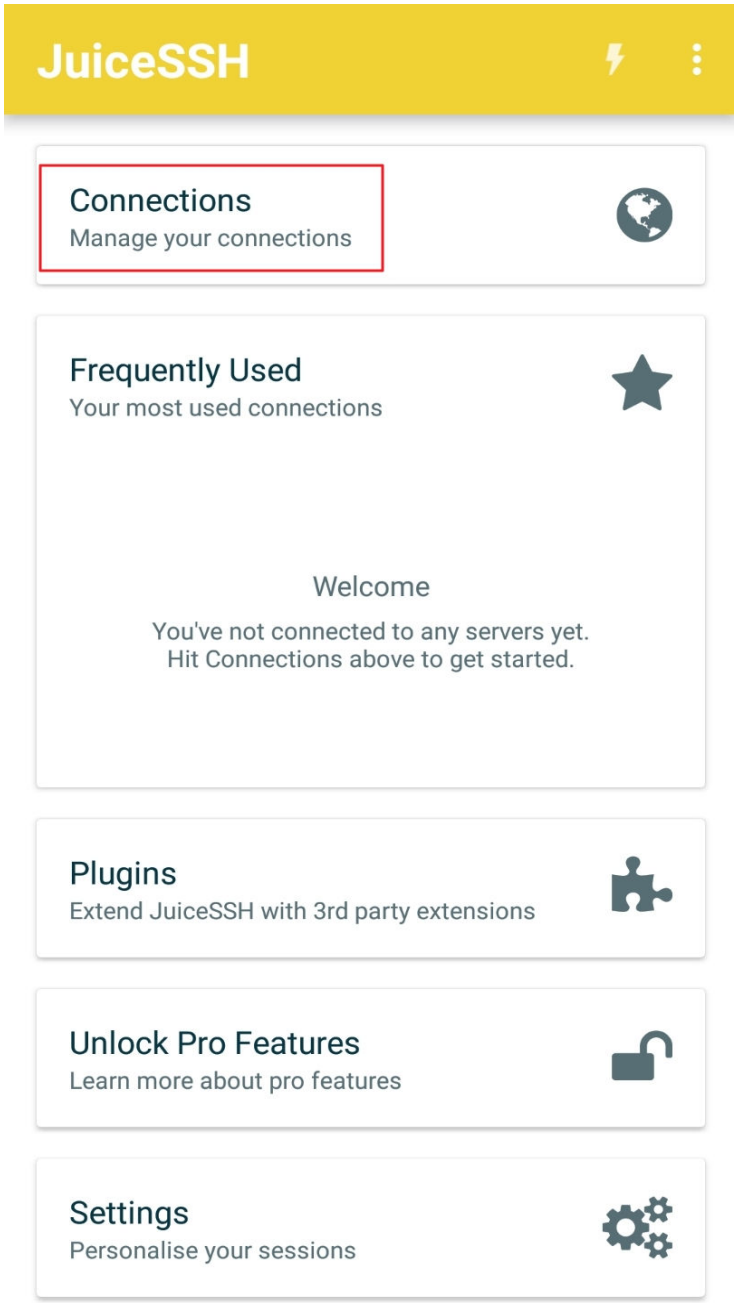

2. On the **Connections** page, tap .

Figure 1-49 Connections

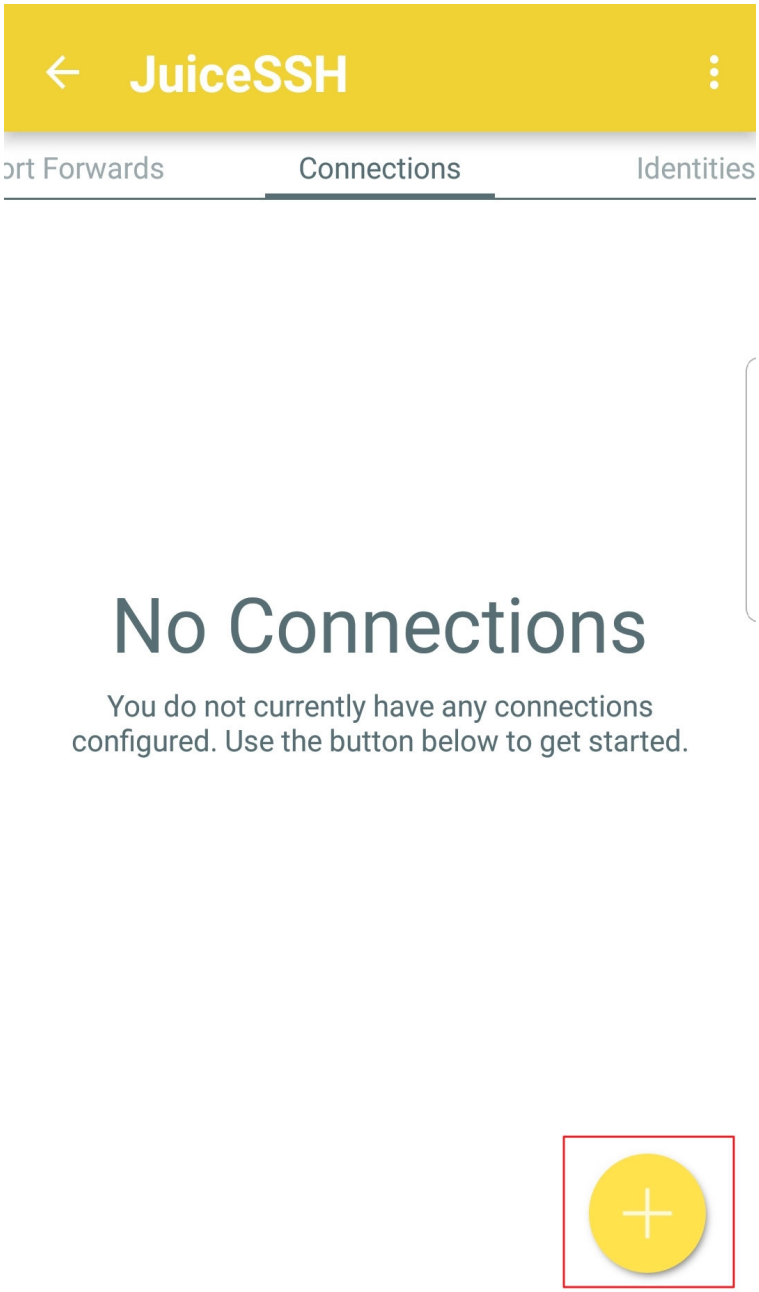

3. On the **New Connection** page, configure basic and advanced settings and save the settings. The parameters are as follows:
   - **Nickname**: Set the name of the login session. In this example, set this parameter to **linux_test**.
   - **Type**: Retain the default value **SSH**.
   - **Address**: Enter the EIP bound to the target Linux ECS.
   - Perform the following operations to set **Identity**:
     i. Tap **Identity** and choose **New** from the drop-down list.

ii. On the **New Identity** page, set the following parameters and tap  .

- **Nickname**: Set an identity name as required to facilitate subsequent management. This parameter is optional. In this example, set it to **linux_test**.
- **Username**: Enter **root**.
- **Password**: Tap **SET (OPTIONAL)**, enter the login password, and tap **OK**.

Figure 1-50 New Identity

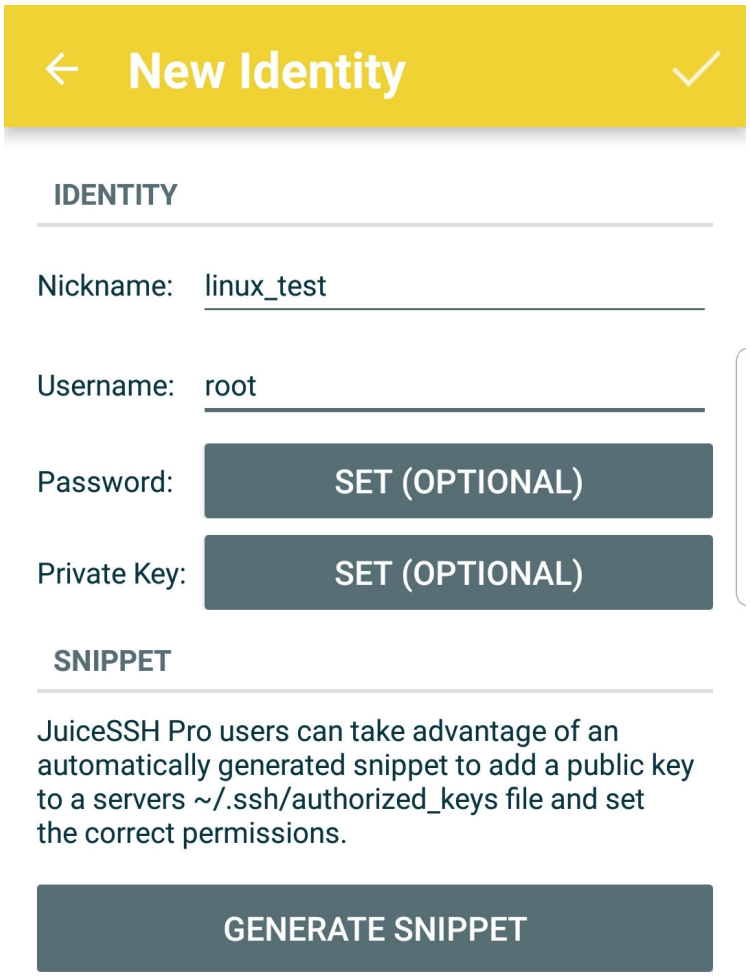

– **Port**: Enter port number **22**.

Figure 1-51 Port

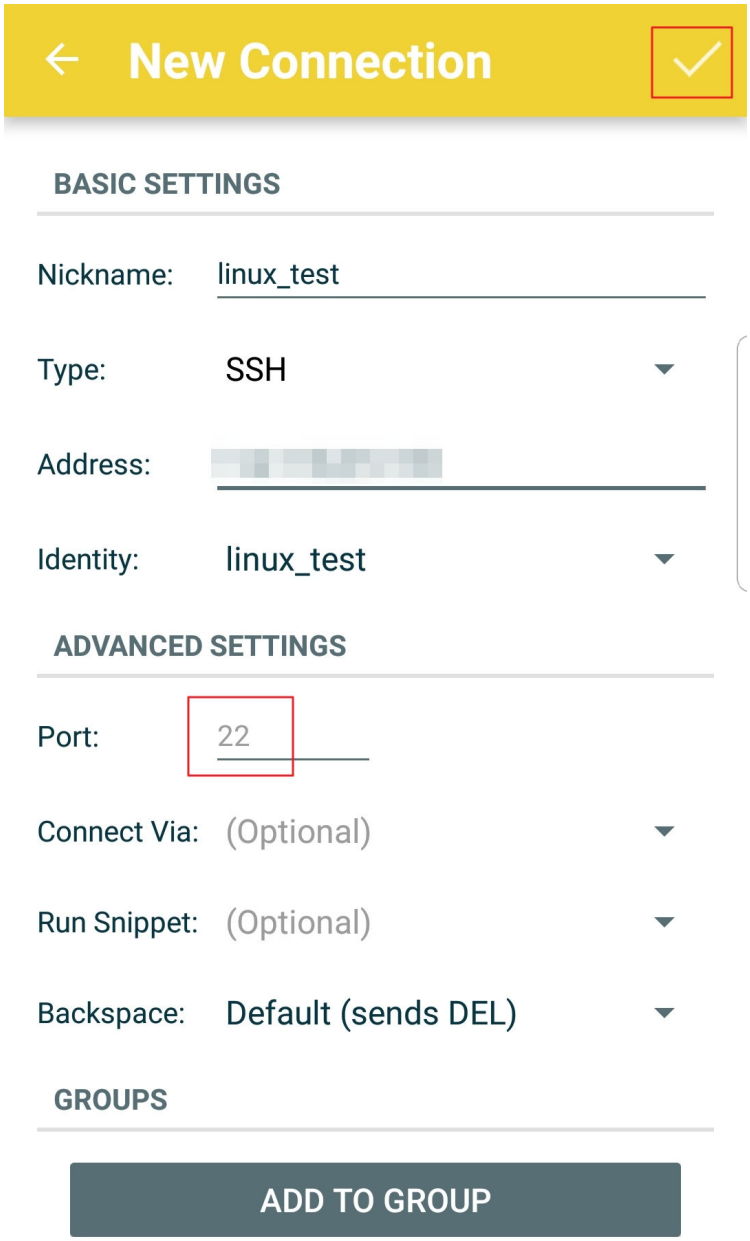

4. On the **Connections** page, tap the created connection.

Figure 1-52 Connections

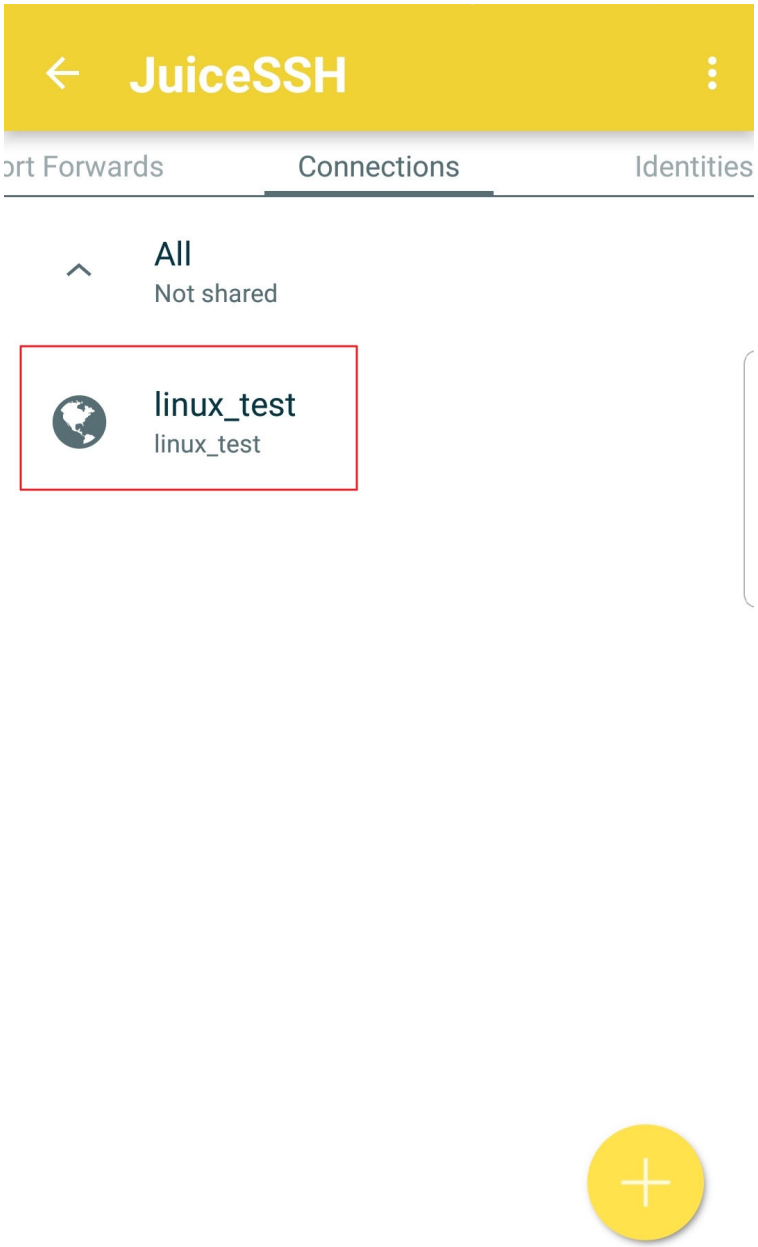

5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 1-53 Confirming the information

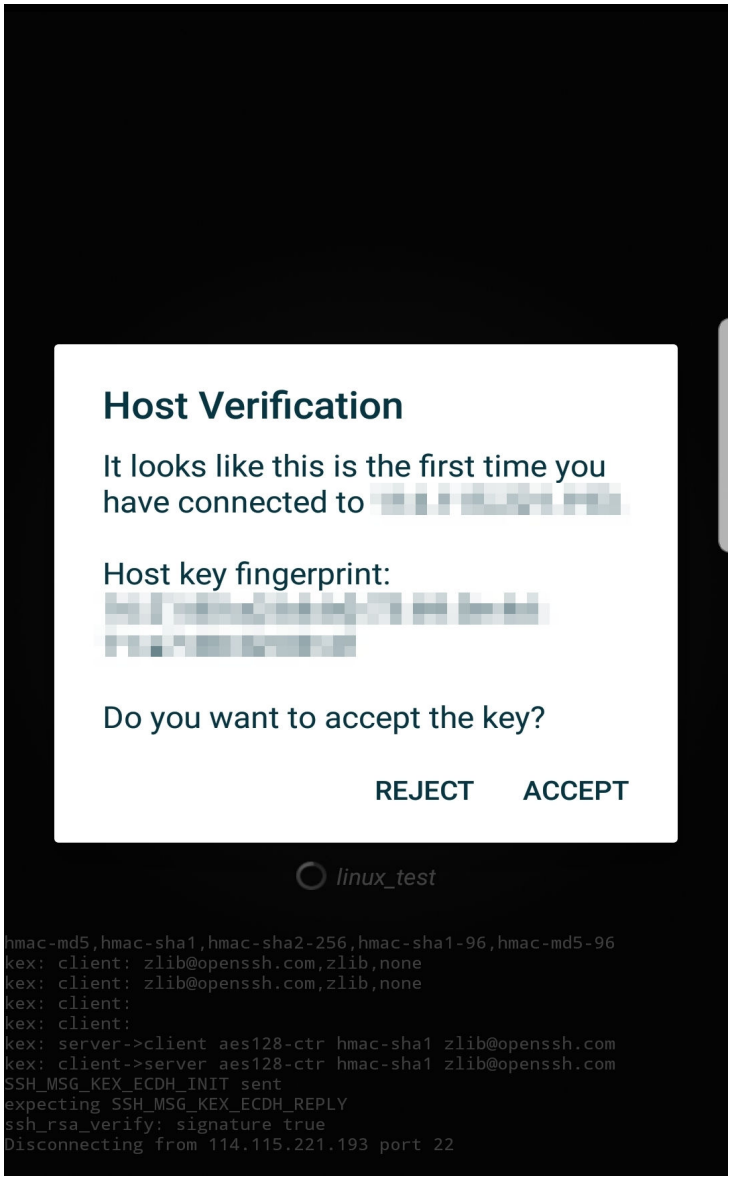

6. (Optional) When you log in to the ECS for the first time, JuiceSSH displays a tutorial for you, including setting the font size and popping up the keyboard. Confirm the information and click **OK - I'VE GOT IT**.

Figure 1-54 Tutorial

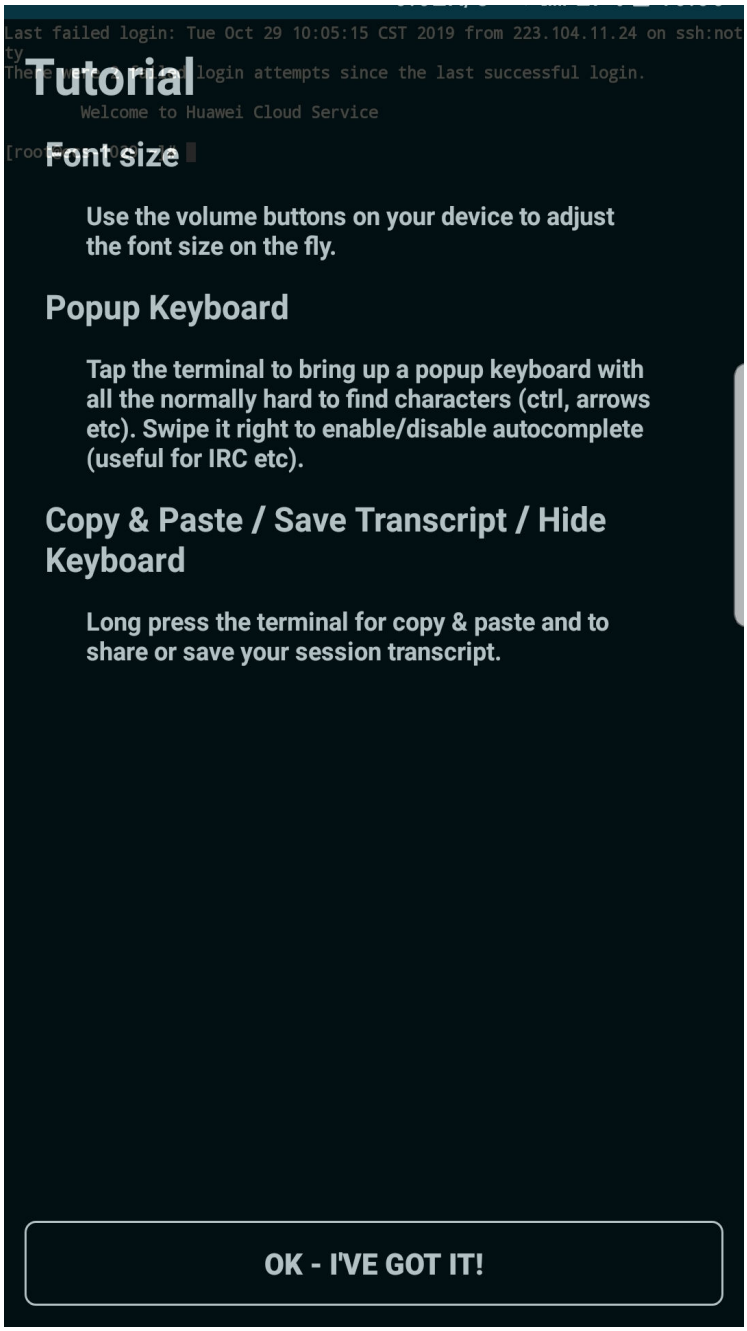

You have logged in to the Linux ECS.

Figure 1-55 Successful login

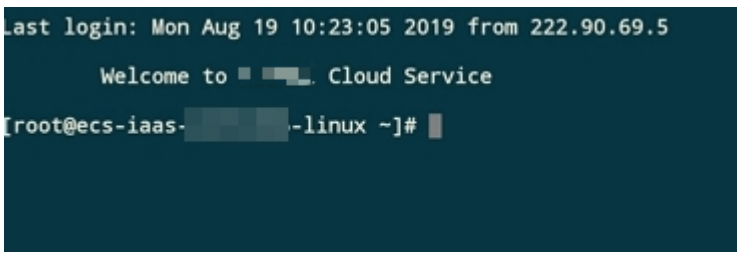

1.5.7 Logging In to a Linux ECS from a macOS Server

Scenario

This section describes how to log in to a Linux ECS from a macOS server.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If the password is forgotten, reset the password by referring to **Resetting the Password for Logging In to an ECS on the Management Console**.
- You have bound an EIP to the ECS. For details, see **Binding an EIP**.
- Port 22 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see **Configuring Security Group Rules**.

Procedure

You can log in to the Linux ECS through the terminal included in the macOS.

- Using an SSH password
  a. Open the terminal of the macOS and run the following command to log in to the ECS:

     ssh *Username*@*EIP*

     NOTE

     If a public image is used, the username is **root**.
- Using an SSH key
  a. Open the terminal of the macOS and run the following command to change permissions. The following operations use private key file **kp-123.pem** as an example. Replace it with your actual private key file.

     chmod 400 /*path*/kp-123.pem

     NOTE

     In the preceding command, *path* refers to the path where the key file is saved.

  b. Run the following command to log in to the ECS:

     ssh -i /*path*/kp-123.pem *Username*@*EIP*

     NOTE

     - The username is **core** for a CoreOS public image.
     - The username is **root** for a non-CoreOS public image.

Follow-up Procedure

- After logging in to the ECS using the SSH key, you can set a password (by using the **passwd** command) to log in to the ECS using VNC.

1.6 Managing ECSs

1.6.1 Changing ECS Names

Scenarios

The name of a created ECS can be changed to meet your service requirements.

Multiple ECS names can be changed in a batch. After the change, the ECS names are the same.

Changing the Name of a Single ECS

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. Click the name of the target ECS.
5. On the page providing details about the ECS, click next to the ECS name. Then, change the name as prompted.

   Allow duplicate name: allows ECS names to be duplicate. If **Allow duplicate name** is not selected and the new name you configure is the same as an existing ECS name, the system displays a message indicating that the name has been used and you need to change it to another name.
6. Click ECS next to the ECS name.
7. Click **OK**.

Changing the Names of Multiple ECSs in a Batch

1. Log in to the management console.
2. Log in to ManageOne as a VDC administrator or operator using a browser.

   URL: **https://***Address for accessing ManageOne Operation Portal*, for example, **https://console.demo.com**
3. Click in the upper left corner and select your region and project.
4. Click . Under **Compute**, click **Elastic Cloud Server**.
5. Select the target ECSs.
6. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.
7. Enter a new name.
8. Click **OK**.

   If you change ECS names in a batch, the new ECS names are the same, for example, all are **ecs-test**.

1.6.2 Reinstalling the OS

Scenarios

If the OS of an ECS fails to start or requires optimization, reinstall the OS.

Notes

- After the OS is reinstalled, the IP and MAC addresses of the ECS remain unchanged.
- Reinstalling the OS clears the data in all partitions of the EVS system disk, including the system partition. Therefore, back up data before reinstalling the OS.
- Reinstalling the OS does not affect data disks.
- Do not perform any operations on the ECS immediately after its OS is reinstalled. Wait for several minutes until the system successfully injects the password or key. Otherwise, the injection may fail, and the ECS cannot be logged in to.

Constraints

- The EVS disk quotas must be greater than 0.
- If the target ECS is created using a private image, ensure that the private image is available.
- If the target ECS is billed on a pay-per-use basis, ensure that your account has sufficient balance.
- If the target ECS is billed on a yearly/monthly basis, ensure that the subscribed resources are within the validity period.
- H2 ECSs do not support OS reinstallation.

Prerequisites

- The target ECS has a system disk attached.

Procedure

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. Locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Reinstall OS** in the **Operation** column.

   Before reinstalling the OS, stop the ECS first or select **Automatically stop the ECSs and reinstall their OSs** in the displayed dialog box.
5. Select the login mode.

   If the target ECS uses key pair authentication, you can replace the original key pair.

Figure 1-56 Reinstall OS

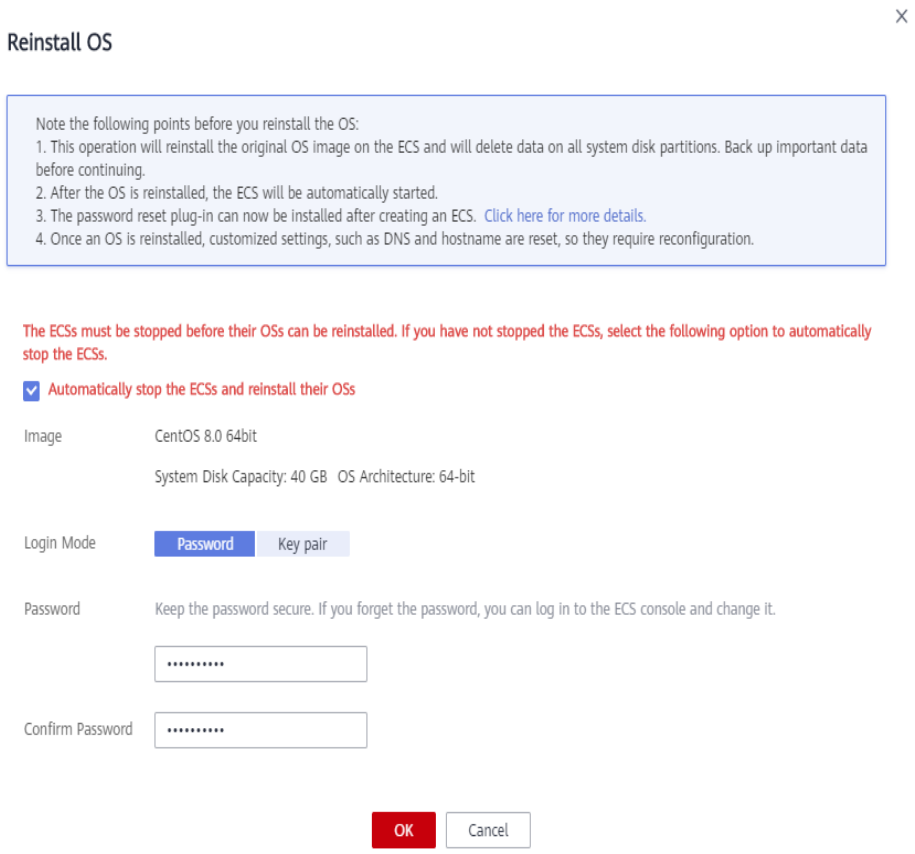

6. Click **OK**.
7. On the **Reinstall OS** page, confirm the settings, read and select the agreement or disclaimer, and click **OK**.

   After the request is submitted, the status **Reinstalling** is displayed. When this status disappears, the reinstallation is complete.

   NOTE

   A temporary ECS is created during the reinstallation process. After reinstallation, this ECS will be automatically deleted. Do not perform any operation on the temporary ECS during the reinstallation process.

Follow-up Procedure

If the reinstallation fails, perform steps 3 to 7 again to retry the OS installation.

If the second reinstallation attempt still fails, contact customer service for manual recovery at the backend.

1.6.3 Changing the OS

Scenarios

Changing an ECS OS will change the system disk attached to the ECS. After the change, the system disk ID of the ECS will be changed, and the original system disk will be deleted.

If the OS running on an ECS cannot meet service requirements, change the ECS OS.

The cloud platform supports changing between image types (public images, private images, and shared images) and between OSs. You can change your OS by changing your ECS image.

Constraints

- For a yearly/monthly ECS, the system disk capacity may be insufficient if you change the image type. You need to detach the system disk, expand the disk capacity, attach the expanded disk, and then change the OS.
- The OS of a yearly/monthly ECS can be changed:
  - Only changes between free OSs are supported.
  - If an ECS is created from a private image on Marketplace and is billed on a yearly/monthly basis, the OS cannot be changed.
  - OS change between Windows and Linux is supported only in the Chinese mainland regions.
- The EVS disk quota must be greater than 0.
- H2 ECSs do not support OS change.
- For details about the change between different OSs, see **Notes on Change Between Windows and Linux**.
- An ISO image created from an ISO file cannot be used to change the OS of an ECS. You need to install an OS and drivers on the ECS and use the ECS to create a system disk image first.
- The boot mode (BIOS or UEFI) cannot be changed.
- The OS cannot be changed between an x86 ECS and a Kunpeng ECS.

Notes

- After the OS is changed, the original OS is not retained, and the original system disk is deleted, including the data in all partitions of the system disk.
- Back up data before changing the OS. For details, see **Cloud Backup and Recovery**.
- Changing the OS does not affect data in data disks.
- After the OS is changed, your service running environment must be deployed in the new OS again.
- After the OS is changed, the ECS will be automatically started.
- After the OS is changed, the system disk type of the ECS cannot be changed.
- After the OS is changed, the IP and MAC addresses of the ECS remain unchanged.
- After the OS is changed, customized configurations, such as DNS and hostname of the original OS will be reset and require reconfiguration.
- It takes about 10 to 20 minutes to change the OS. During this process, the ECS is in **Changing OS** state.
- Do not perform any operations on the ECS immediately after its OS is changed. Wait for several minutes until the system successfully injects the password or key. Otherwise, the injection may fail, and the ECS cannot be logged in to.

Notes on Change Between Windows and Linux

When you change the OS from Windows to Linux or from Linux to Windows, note the following:

- To change Windows to Linux, install an NTFS partition tool, such as NTFS-3G for data reads and writes on the Windows ECS.
- To change Linux to Windows, install software, such as Ext2Read or Ext2Fsd to identify ext3 or ext4.

  NOTE

  If there are LVM partitions on the Linux ECS, these partitions may fail after the OS is changed to Windows. Therefore, a change from Linux to Windows is not recommended.

Billing Rules

- The new system disk may have a larger capacity after an OS change, so you may be billed more.

Prerequisites

- The target ECS has a system disk attached.
- Necessary data has been backed up. (Changing the OS clears the data in all partitions of the system disk, including the system partition.)
- If the original ECS uses password authentication while the new ECS uses key pair authentication, ensure that a key pair is available.
- If you plan to use a private image to change the OS, ensure that a private image is available. For details about how to create a private image, see **Image Management Service User Guide**.
  - If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
  - If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
  - If a private image from another region is required, make sure that the image has been copied.
  - If a private image from another user account is required, make sure that the image has been shared with you.

Procedure

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. Locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Change OS** in the **Operation** column.

   Before changing the OS, stop the ECS first or select **Automatically stop the ECSs and reinstall their OSs** in the displayed dialog box.

5. Modify related ECS parameters, such as **Image Type** and **Image**, based on service requirements.

   NOTE

   For a yearly/monthly ECS, if the system disk capacity is smaller than the size of your image, you must detach the system disk, expand its capacity, and attach it to the original ECS before changing the OS.

   For instructions about how to expand the system disk capacity, see **Disk Capacity Expansion**.

   Figure 1-57 Change OS

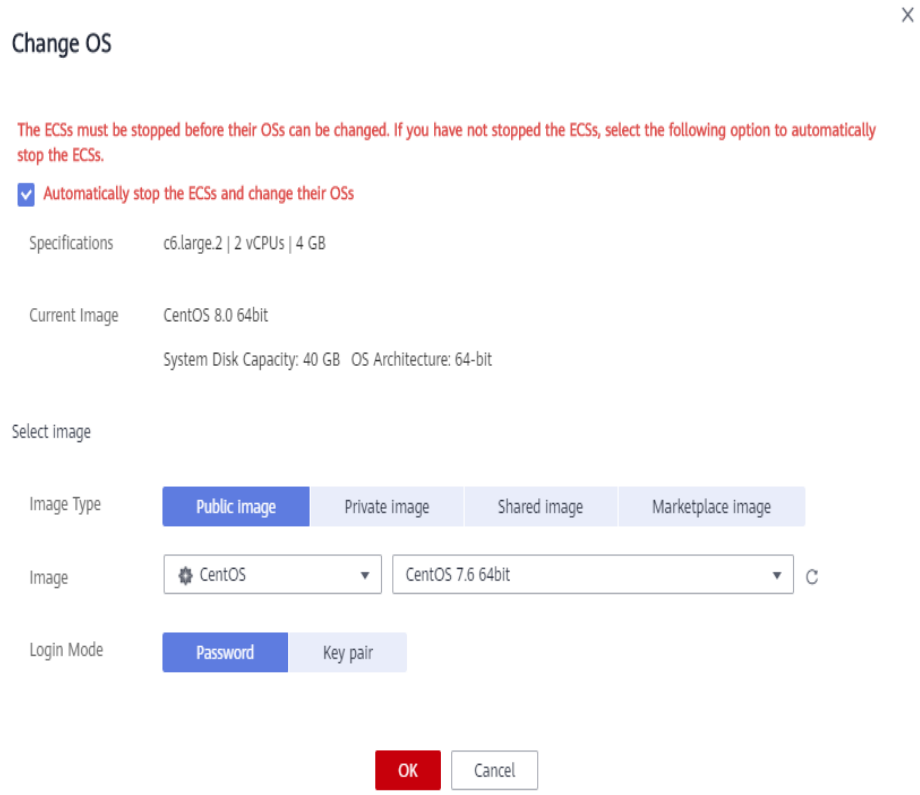

6. Configure the login mode.

   If the target ECS uses key pair authentication, you can replace the original key pair.

7. Click **OK**.

8. On the **Change OS** page, confirm the specifications, read and select the agreement or disclaimer, and click **OK**.

   After the application is submitted, the status **Changing OS** is displayed. When this status disappears, the OS change is complete.

   NOTE

   A temporary ECS is created during the OS change process. After the process is complete, this ECS will be automatically deleted.

Follow-up Procedure

- If the OSs before and after the OS change are both Linux, and automatic mounting upon system startup has been enabled for data disks, the data disk partition mounting information will be lost after the OS is changed. In such a case, you need to update the **/etc/fstab** configuration.
  a. Write the new partition information into **/etc/fstab**.
     It is a good practice to back up the **/etc/fstab** file before writing data into it.
     To enable automatic partition mounting upon system startup, see **Initializing a Linux Data Disk (fdisk)**.
  b. Mount the partition so that you can use the data disk.
     mount *Disk partition Device name*
  c. Check the mount result.
     df -TH
- If the OS change is unsuccessful, perform steps **3** to **8** again to retry the OS change.
- If the second OS change attempt is unsuccessful, contact customer service for manual recovery at the backend.

1.6.4 Managing ECS Groups

Scenarios

An ECS group logically groups ECSs. ECSs in an ECS group comply with the same policy associated with the ECS group.

Currently, only the anti-affinity policy is supported.

This policy enables ECSs in the same ECS group to run on different hosts for improved reliability, high availability, and disaster recovery.

You can perform the following operations on an ECS group:

- **Creating an ECS Group**
- **Adding an ECS to an ECS Group**
  - Add an ECS to an ECS group during ECS creation.
    For details, see **Step 3: Configure Advanced Settings**.
  - Add an existing ECS to an ECS group.
- **Removing an ECS from an ECS Group**
- **Deleting an ECS Group**

Constraints

- ECS groups support the anti-affinity policy only.
- ECSs are deployed on different physical hosts.
- If the maximum number of ECS groups is reached, you need to contact customer service to increase the quota.

- The maximum number of ECSs that can be added to an ECS group varies depending on the region. You can view the quota on the **ECS Group** page, as shown in **Figure 1-58**.

  Figure 1-58 Maximum number of ECSs that can be added to an ECS group

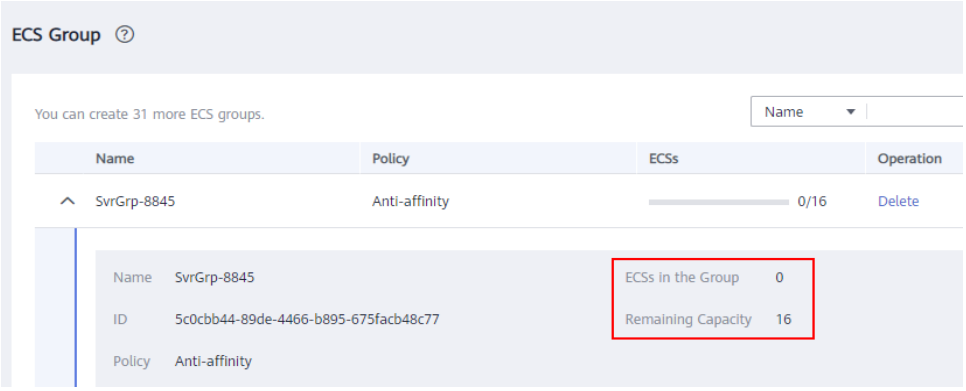

Creating an ECS Group

Create an ECS group and associate the same policy to all group members. ECS groups are independent from each other.

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. In the navigation pane on the left, choose **ECS Group**.
5. On the **ECS Group** page, click **Create ECS Group**.
6. Enter the name of an ECS group.
7. Select a policy for the ECS group.
8. Click **OK**.

Adding an ECS to an ECS Group

To improve service reliability, you can add ECSs to an ECS group so that these ECSs in this group can run on different hosts.

NOTE

- After an ECS is added to an ECS group, the system reallocates a host to run this ECS to ensure that ECSs in this group are running on different hosts. When the ECS is being restarted, the startup may fail due to insufficient resources. In such a case, remove the ECS from the ECS group and try to restart the ECS again.
- ECSs that have local disks attached can be added to an ECS group only during the creation process. Once created, they can no longer be added to any ECS groups.
- Existing ECSs cannot be added to any ECS groups if they have local disks attached (such as disk-intensive, H2, P1, or P2 ECSs), local NVMe SSD disks attached (such as ultra-high I/O ECSs), GPU cards attached (such as G3 ECSs), or FPGA cards attached (such as Fp1 or Fp1c ECSs). They can only be added to an ECS group during the creation process.

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. In the navigation pane on the left, choose **ECS Group**.
5. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
6. On the **Add ECS** page, select an ECS to be added.
7. (Optional) Stop the ECS.

   Certain types of ECSs need to be stopped before being added to an ECS group. If **Stop** is displayed in the row containing the selected ECS, you need to stop the ECS first.

   a. Click **Stop** in the **Operation** column.
   b. Select a stop option.
      - **Stop**: The ECS will be stopped normally.
      - **Forcibly stop the preceding ECSs**: This operation will cause loss of unsaved data. Exercise caution when performing this operation.
   c. Click **Yes**.
8. Click **OK**. The ECS is added to the ECS group.

Removing an ECS from an ECS Group

After an ECS is removed from an ECS group, the ECS does not comply with the anti-affinity policy anymore.

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. In the navigation pane on the left, choose **ECS Group**.
5. Expand the ECS group information and view the ECSs in the ECS group.
6. Locate the ECS to be removed and click **Remove** in the **Operation** column.
7. In the displayed dialog box, click **Yes**.

   The ECS is removed from the ECS group.

Deleting an ECS Group

After an ECS group is deleted, the policy does not apply to the ECSs in the ECS group anymore.

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. In the navigation pane on the left, choose **ECS Group**.

5. Locate the ECS group to be deleted and click **Delete** in the **Operation** column.
6. In the displayed dialog box, click **Yes**.

1.6.5 Changing the Time Zone for an ECS

Scenarios

The default time zone for an ECS is the one you selected when creating the image that was used to create the ECS. This section describes how to change the time zone for an ECS to the local one or to another time zone in your network.

After you log in to your ECS, if you find that the time on the ECS is different from the local time, you can change the time zone for the ECS so that the time on the ECS is the same as the local time.

For Linux ECSs

The process of changing the time zone for a Linux ECS depends on the OS. In this section, the CentOS 6.x 64bit OS is used to demonstrate how to change the time zone for a Linux ECS.

1. Log in to the ECS.
2. Run the following command to switch to user **root**:

   su - root
3. Run the following command to obtain the time zones supported by the ECS:

   ls /usr/share/zoneinfo/

   In the terminal display, the **/user/share/zoneinfo** directory contains a hierarchy of time zone data files. Use the directory structure to obtain your desired time zone file.

   The directory structure shown in **/user/share/zoneinfo** includes both time zones and directories. The directories contain time zone files for specific cities. Locate the time zone for the city in which the ECS is located.

   For example:

   - If you are to use the time zone for Shanghai, China, run the **ls /usr/share/zoneinfo/Asia** command to obtain the directory **/usr/share/zoneinfo/Asia/Shanghai**.
   - If you are to use the time zone for Paris, France, run the **ls /usr/share/zoneinfo/Europe** command to obtain the directory **/usr/share/zoneinfo/Europe/Paris**.
4. Set the target time zone.

   a. Run the following command to open the **/etc/sysconfig/clock** file:

      vim /etc/sysconfig/clock

   b. Locate the **ZONE** entry and change its value to the name of the desired time zone file.

      For example:

      - If the target time zone is for Shanghai, China, change the **ZONE** entry value as follows:

ZONE="Asia/Shanghai"

- If the target time zone is for Paris, France, change the **ZONE** entry value as follows:

  ZONE="Europe/Paris"

5. Press **Esc**. Then, run the following command to save and exit the **/etc/sysconfig/clock** file:

   :wq

6. Run the following command to check whether the **/etc/localtime** file is available on the ECS:

   ls /etc/localtime

   – If the file is available, go to step **7**.

   – If the file is not available, go to step **8**.

7. Run the following command to delete the existing **/etc/localtime** file:

   rm /etc/localtime

8. Run the following command to create a symbolic link between **/etc/localtime** and your time zone file so that the ECS can find this time zone file when it references the local time:

   ln -sf /usr/share/zoneinfo/*Asia/city1*/etc/localtime

9. Run the following command to restart the ECS so that all services and applications running on the ECS use the new time zone:

   reboot

10. Log in to the ECS again and run the following command as user **root** to check whether the time zone has been changed:

    ls -lh /etc/localtime

    The following information is displayed:

```
# ls -lh /etc/localtime
lrwxrwxrwx 1 root root 33 Nov 27 11:01 /etc/localtime -> /usr/share/zoneinfo/Asia/city1
```

For Windows ECSs

1. Log in to the ECS.
2. Click the time display on the far right side of the task bar located at the bottom of your screen. In the dialog box that is displayed, click **Change date and time settings**.

   The **Date and Time** page is displayed.

Figure 1-59 Date and Time

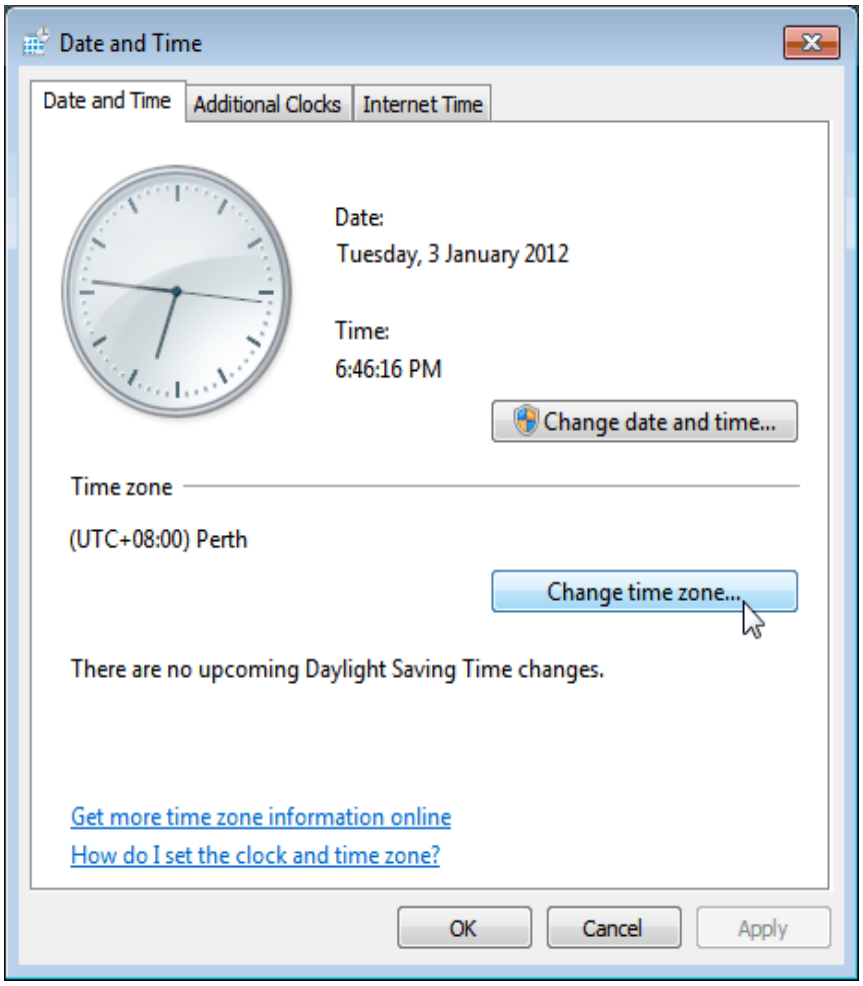

3. Click **Change time zone**.

   The **Time Zone Settings** page is displayed.
4. In the **Set the time zone** pane, choose the target time zone from the **Time zone** drop-down list.
5. Click **OK**.

1.6.6 Starting and Stopping ECSs

You can start, stop, restart, and delete ECSs.

- To prevent load increase when starting or stopping a large number of ECSs at the same time, you are advised to select a small number of ECSs for each batch.
- If an ECS remains in the **Restarting** or **Stopping** state for a long time, you can forcibly restart or stop it. In such a case, any unsaved data on the ECS will be lost. Therefore, exercise caution when forcibly restarting or stopping an ECS.

NOTE

To start or stop ECSs whose flavors contain physical, see this section.

If you use the OS commands such as shutdown, poweroff, or halt to stop ECSs, the commands may be invalid or the ECS may fail to be started after being stopped.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select the target ECSs.
5. Above the list, click **Start**, **Stop**, or choose **More** > **Restart** or **Delete**.
6. In the displayed window, click **Yes**.

   NOTE

   Contact the administrator if an ECS has been in any of the following intermediate states for more than 30 minutes:

   - Starting
   - Stopping
   - Restarting
   - Forcibly restarting
   - Deleting

1.7 Modifying ECS Specifications

1.7.1 General Operations

Scenarios

If ECS specifications do not meet service requirements, you can modify the ECS specifications, including vCPUs and memory. Certain ECSs allow you to change their types when you modify their specifications.

- Before changing a Xen ECS to a KVM ECS, you need to manually install the required drivers on the ECS first, or the ECS will be unavailable after the modification is complete. For example, starting the OS will fail. The following section describes how to change a Xen ECS to a KVM ECS. For Linux, you are recommended to use a script to automatically change a Xen ECS to a KVM ECS.
  - **Changing a Xen ECS to a KVM ECS (Windows)**
  - **Automatically Changing a Xen ECS to a KVM ECS (Linux)**
  - **Manually Changing a Xen ECS to a KVM ECS (Linux)**

NOTE

- ECSs can be classified as the following based on the virtualization types:
  - Xen ECSs: S1, C1, C2, and M1 ECSs.
  - KVM ECSs: See the **Virtualization** column in **ECS Specifications**.

Notes

- When modifying the specifications of an ECS, sold-out CPU and memory resources are unavailable for selection.

- Downgrading ECS specifications (vCPU or memory) will reduce performance.
- Certain ECS types do not support specifications modification currently. For details about available ECS types as well as their functions and usage, see "Notes" in **ECS Types**.
- When the disk status is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.
- Before modifying the specifications of a Windows ECS, modify the SAN policy by following the instructions provided in **Why Does a Disk Attached to a Windows ECS Go Offline?** to prevent disks from going offline after the specifications are modified.

Fees Description

Modifying specifications will lead to fee changes. For details, see **Pricing of a Changed Specification**.

Preparations

After ECS specifications are modified, NIC flapping may occur. Before modifying the specifications, perform the following operations:

NOTE

NIC flapping occurs because NIC retaining is enabled in the image from which the ECS is created.

For more information about NIC flapping, see **What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?**

- Linux

  Run the following commands on the ECS to delete the files with **persistent** and **net** included in their names in the network rule directory:

  rm -fr /etc/udev/rules.d/*net*persistent*.rules

  rm -fr /etc/udev/rules.d/*persistent*net*.rules

- Windows

  Delete the following directories in the registry on the ECS:

  HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\NetworkList\Profiles

  HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion\NetworkList\Signatures\Unmanaged

Figure 1-60 Registry

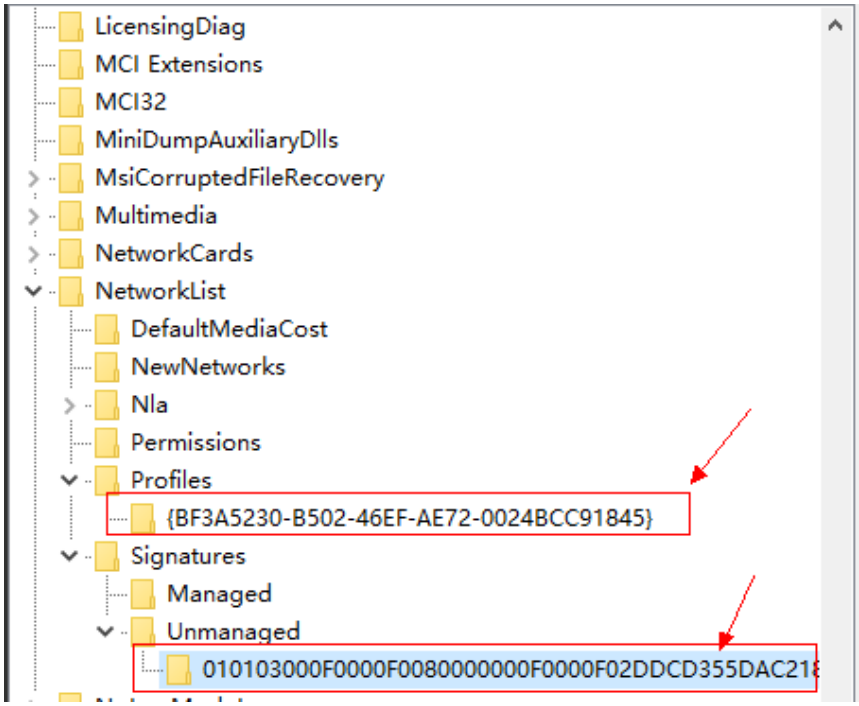

Step 1: Modify Specifications

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Click **More** in the **Operation** column and select **Modify Specifications**.
   The **Modify ECS Specifications** page is displayed.
5. Select the new ECS type, vCPUs, and memory as prompted.
   Before modifying the specifications, stop the ECS or select **Automatically stop the ECSs and then modify specifications**.

Figure 1-61 Modifying specifications

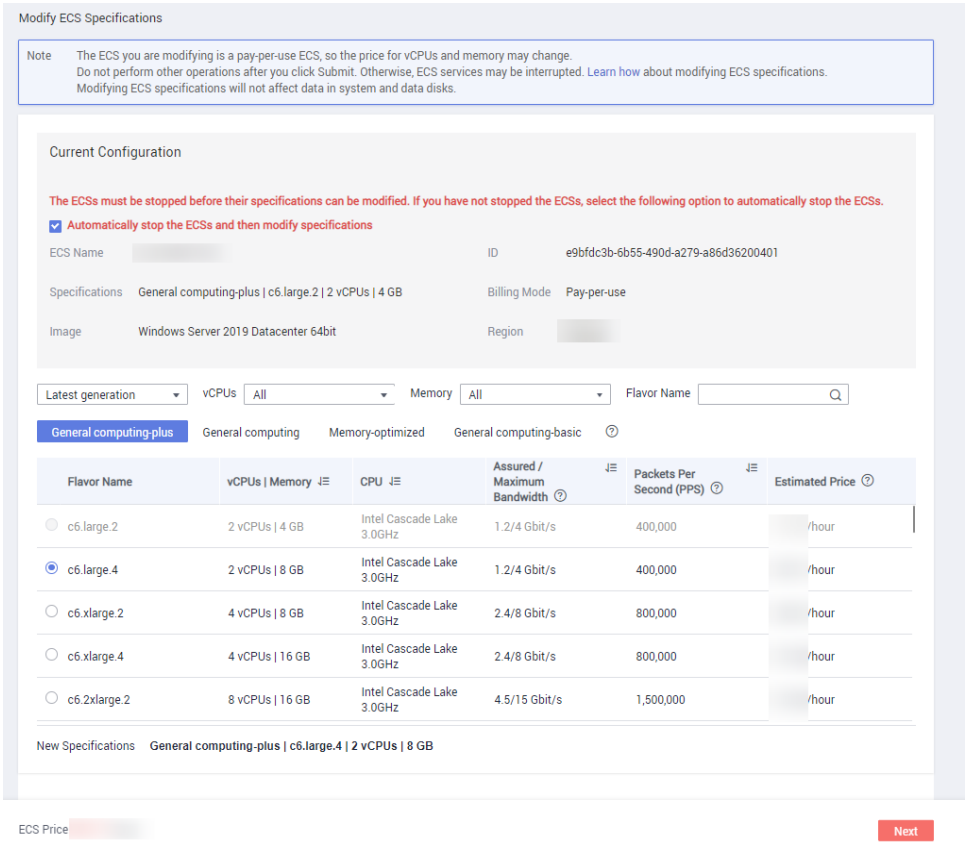

6. Click **Next**.
7. Confirm the settings, read and select the disclaimer, and then click **Submit Application**.
8. Check whether the specifications have been modified.

   After modifying the specifications, you can check whether the specifications have been modified in **Failures**.

   a. Check whether **Failures** is displayed on the management console. For details, see **Viewing Failed Tasks**.

      - If yes, go to step **8.b**.
      - If no, the specifications have been modified.

   b. Click **Failures**. Then, in the **Failures** dialog box, click **Operation Failures** and check whether the task is contained in the list by **Name/ID**, **Operated At**, or **Task**.

      - If yes, the specifications modification failed. See **Follow-up Procedure** for failure causes.
      - If no, the specifications have been modified.

Step 2: Check Disk Attachment

After specifications are modified, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Windows ECS

  For details, see **Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?**
- Linux ECS

  For details, see **Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?**

Follow-up Procedure

Perform the following operations in the event of a specifications modification failure:

1. Log in to the management console.
2. Under **Management & Governance**, choose **Cloud Trace Service**.
3. In the navigation pane on the left, choose **Trace List**.
4. In the **Trace Name** column, locate the **resizeServer** event by resource ID.

   The resource ID is the ID of the ECS on which the specifications modification failed.
5. Click **View Trace** in the **Operation** column to view the failure cause.

   If the fault cannot be rectified based on logs, contact customer service.

1.7.2 Changing a Xen ECS to a KVM ECS (Windows)

Scenarios

Before changing a Xen ECS that runs Windows to a KVM ECS, make sure that PV driver and UVP VMTools have been installed on the ECS.

This section describes how to install the PV driver and UVP VMTools and change Xen to KVM.

NOTE

- ECSs can be classified as the following based on the virtualization types:
  - Xen ECSs: S1, C1, C2, and M1 ECSs.
  - KVM ECSs: See the **Virtualization** column in **ECS Specifications**.

Constraints

- If a Windows ECS is attached with a cross-region disk, the ECS specifications cannot be modified. Otherwise, ECS data may be lost.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 1-62 shows the flowchart for changing a Xen ECS to a KVM ECS.

Figure 1-62 Flowchart for changing a Xen ECS to a KVM ECS

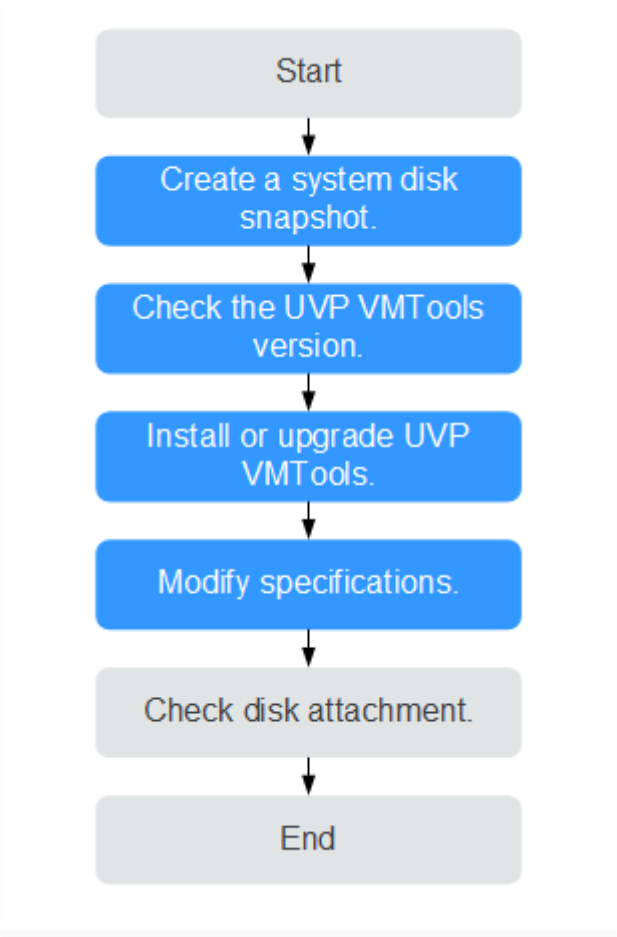

Table 1-16 describes the operations for changing a Xen ECS to a KVM ECS.

Table 1-16 Procedure for changing a Xen ECS to a KVM ECS

Step	Operation
1	**Step 1: Create a System Disk Snapshot**
2	**Step 2: Check the UVP VMTools Version**
3	**Step 3: Install or Upgrade UVP VMTools**
4	**Step 4: Modify Specifications**
5	**(Optional) Step 5: Check Disk Attachment**

Step 1: Create a System Disk Snapshot

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, create a snapshot for the system disk first.

1. Before you create a system disk snapshot, check the ECS.

   Stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. For instructions about how to create a system disk snapshot, see **Creating a Snapshot** in *Elastic Volume Service User Guide*.

NOTE

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

Step 2: Check the UVP VMTools Version

Before modifying specifications, check the UVP VMTools version.

1. Log in to the ECS.
2. Download the driver check script.

   Execute the script as the administrator and wait for the check result.

   URL for downloading the script: **https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/windows/server_resize/check_kvm_drivers.vbs**

   After checking that the required driver has been installed, the system automatically tags the ECS. The specifications of only the tagged ECSs can be modified.

   - If the check result is "Check version success!", the driver version meets service requirements and the ECS is tagged. Then, go to **Step 4: Modify Specifications**.
   - If the check result is "Check version success but set metadata failed! Please run this script again later.", the driver version meets service requirements but tagging the ECS failed. In such a case, try again later.
   - If the check result is "Check version failed! Please install drivers at first.", the driver version does not meet service requirements. In such a case, install or upgrade UVP VMTools by following the instructions provided in **Step 3: Install or Upgrade UVP VMTools**.

Step 3: Install or Upgrade UVP VMTools

When you install or upgrade UVP VMTools, if the PV driver has been installed on the ECS, the system will check the PV driver version. Ensure that the PV driver version meets service requirements. Otherwise, installing UVP VMTools will fail on the ECS. This section describes how to check the installation of the PV driver and UVP VMTools.

CAUTION

Before installing the PV driver or upgrading UVP VMTools, ensure that the ECS meets the following requirements:

- The available system disk size of the ECS is greater than 2 GB.
- Third-party virtualization platform tools, such as Citrix Xen Tools and VMware Tools, have been uninstalled to prevent driver installation failures. For instructions about how to uninstall the tools, see the official documents of the tools.
- Antivirus software or intrusion detection software has been disabled. You can enable them after the driver is installed.

1. Check whether the PV driver version meets the UVP VMTools dependency requirements.

   Switch to the **C:\Program Files (x86)\Xen PV Drivers\bin** directory, open the **version.ini** file, and view the PV driver version.

   ```
   pvdriverVersion=5.0.104.010
   ```

   - If the directory is available and the driver version is 5.0 or later, the PV driver meeting service requirements has been installed. In such a case, go to step **6** to install UVP VMTools.
   - If the directory is unavailable or the driver version is earlier than 5.0, the PV driver has not been properly installed or the version does not meet service requirements. Then, see the following steps to uninstall the PV driver and install a new one.

2. Record the User Account Control (UAC) configuration of the ECS.

   NOTE

   If the PV driver version is earlier than 5.0, DisableLUA is added to the registry during PV driver installation to prevent too many pop-up windows during driver upgrade, and EnableLUA is added to the registry during PV driver uninstallation (this has been resolved in PV driver 5.0 and later versions). To prevent adverse impact on your services, you need to record the UAC configuration before uninstalling the PV driver. Then check and restore the EnableLUA configuration in the registry after installing the new version. For details about UAC configurations, see **official Microsoft documents**.

   a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.

   b. Record the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\EnableLUA** value.

   Figure 1-63 EnableLUA

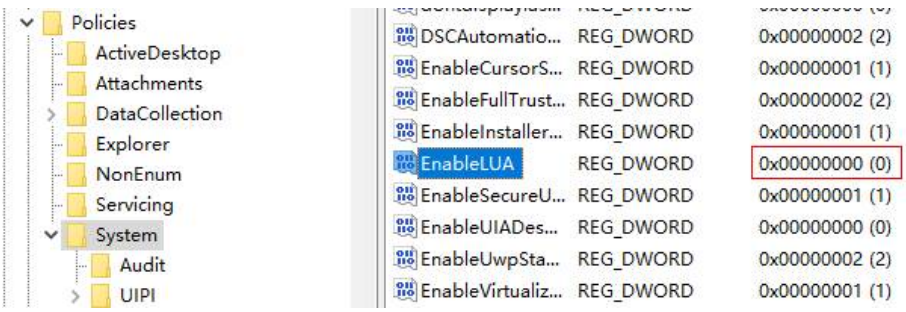

3. Uninstall the PV driver of the old version.
   a. On the ECS OS, choose **Start** > **Control Panel**.
   b. Click **Uninstall a program**.
   c. Uninstall **GPL PV Drivers for Windows** *x.x.x.xx* as prompted.
   d. Restart the ECS on the management console.
4. Install the PV driver of the new version.
   a. Download the PV driver installation package.

      Download the PV driver at **https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/pvdriver-windows.zip**.
   b. Decompress the PV driver software package.
   c. Double-click **pvdriver-windows.iso**.
   d. Run **Setup.exe** and install the PV driver as prompted.

      Wait until the driver installation is complete. Do not click **Setup.exe** during the installation.
   e. Restart the ECS as prompted for the PV driver to take effect.
5. Check and restore the UAC configuration.
   a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
   b. Check the **HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\EnableLUA** value and compare it with the value you recorded. If they are different, change the value to the one recorded in step **2**.
6. Install or upgrade UVP VMTools.
   a. Download the UVP VMTools installation package.

      Download UVP VMTools at **https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/vmtools-windows.zip**.
   b. Decompress the UVP VMTools installation package.
   c. Double-click **vmtools-windows.iso**.
   d. Run **Setup.exe** and install UVP VMTools as prompted.

      The installation program will automatically adapt to the OS version and identify whether UVP VMTools is newly installed or upgraded.

      Wait until the installation is complete. Do not click **Setup.exe** during the installation.
   e. Restart the ECS as prompted for UVP VMTools to take effect.
   f. Check whether UVP VMTools has been installed. For details, see **Step 2: Check the UVP VMTools Version**.

Step 4: Modify Specifications

1. Log in to management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.

4. On the **Elastic Cloud Server** page, view the status of the target ECS.

   If the ECS is not in **Stopped** state, click **More** in the **Operation** column and select **Stop**.

5. Click **More** in the **Operation** column and select **Modify Specifications**.

   The **Modify ECS Specifications** page is displayed.

6. Select the new ECS type, vCPUs, and memory as prompted.
7. (Optional) Set **DeH**.

   If the ECS is created on a DeH, you can change the DeH where the ECS resides.

   To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.

8. Select the check box to confirm that step ECS has been performed.
9. Click **OK**.

   NOTE

   - The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
   - If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is not affected.

(Optional) Step 5: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Windows ECS

  For details, see **Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?**

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, contact customer service or reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

NOTE

Reinstalling the OS will clear the system disk data, but the data on data disks is not affected.

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

1.7.3 Automatically Changing a Xen ECS to a KVM ECS (Linux)

Scenarios

Before changing a Xen ECS that runs Linux to a KVM ECS, make sure that the required drivers have been installed and configured on the ECS.

This section describes how to use a script to automatically install drivers on the ECS, configure the device name, and change Xen to KVM.

NOTE

- Xen ECSs include S1, C1, C2, and M1 ECSs.
- To obtain KVM ECSs, see the **Virtualization** column in **ECS Specifications**.
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

Constraints

- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 1-64 shows the flowchart for automatically changing a Xen ECS to a KVM ECS.

Figure 1-64 Flowchart for automatically changing a Xen ECS to a KVM ECS

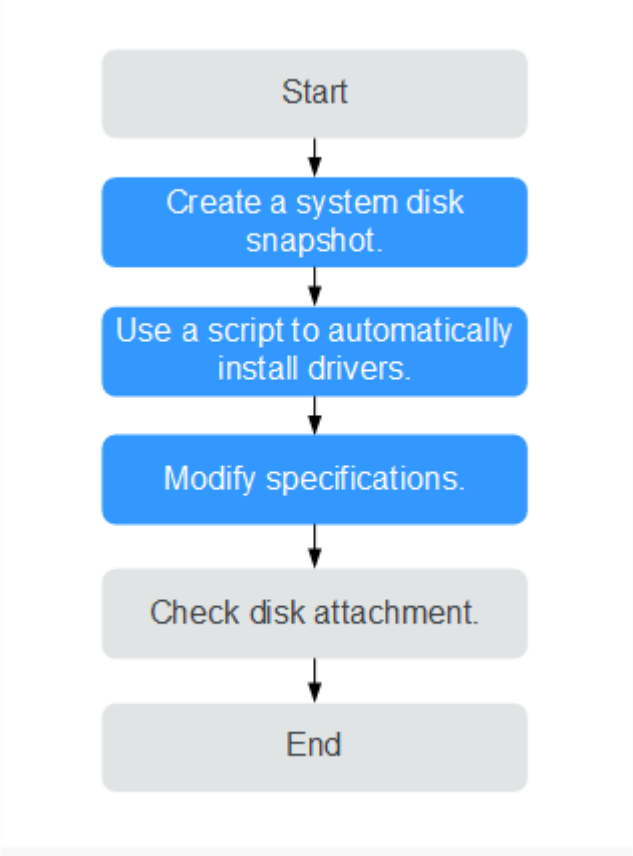

Table 1-17 describes the operations for automatically changing a Xen ECS to a KVM ECS using a script.

Table 1-17 Procedure for automatically changing a Xen ECS to a KVM ECS using a script

Step	Operation
1	**Step 1: Create a System Disk Snapshot**
2	**Step 2: Using a Script to Automatically Install Drivers**
3	**Step 3: Modify Specifications**
4	**(Optional) Step 4: Check Disk Attachment**

Step 1: Create a System Disk Snapshot

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, create a snapshot for the system disk first.

1. Before you create a system disk snapshot, check the ECS.

Stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. For instructions about how to create a system disk snapshot, see **Creating a Snapshot** in *Elastic Volume Service User Guide*.

NOTE

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

Step 2: Using a Script to Automatically Install Drivers

Use a script to install drivers on an ECS. If your ECS does not support configuration using a script, manually configure it by referring to **Manually Changing a Xen ECS to a KVM ECS (Linux)**.

1. Log in to the ECS.
2. Run the following command to download the driver installation script to the **root** directory:

   curl *URL* **> ~/resize_ecs_modify_linux.sh**

   In the preceding command, *URL* is the address for downloading the specifications modification script.

   Select an address for downloading the optimization script based on the region where the ECS is located:

   URL for downloading the script: **https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/linux/server_resize/resize_ecs_modify_linux.sh**
3. Run the following command to execute the script which automatically checks and installs the native Xen PV driver and virtio driver:

   bash resize_ecs_modify_linux.sh

   Figure 1-65 Executing the script

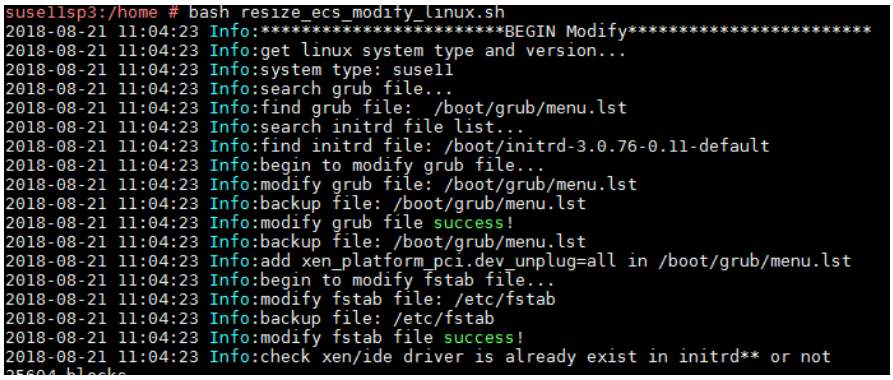

4. Wait until the script is executed.

   After checking that the required driver has been installed, the system automatically tags the ECS. The specifications of only the tagged ECSs can be modified.

   If the check result is "{*Image name*} already contain xen and virtio driver", the driver has been installed.

- If the check result is "Success to set kvm meta!" or "this server already has kvm meta.", the ECS has been tagged. Then, go to **Step 3: Modify Specifications**.
- If the check result is "Failed to set metadata, please try again.", tagging the ECS failed. In such a case, try again later.

If the installation failed, manually configure the ECS by following the instructions provided in **Manually Changing a Xen ECS to a KVM ECS (Linux)** or contact customer service.

Figure 1-66 Successful script execution

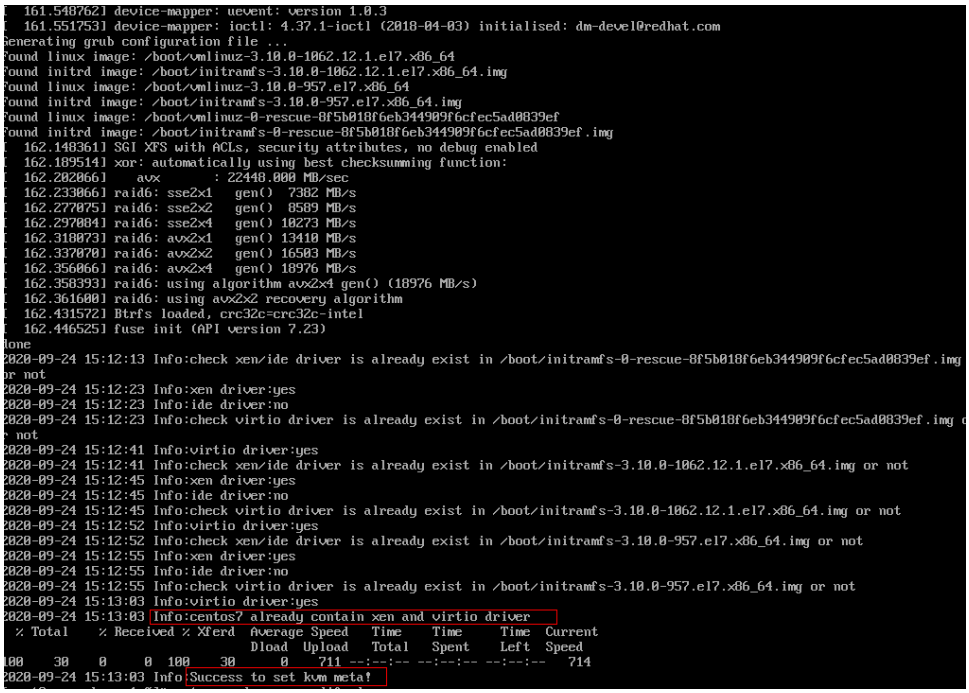

NOTE

- Make sure that the ECS has been configured successfully, or the ECS may become unavailable after the specifications are modified. If the operation failed, follow the instructions provided in **Manually Changing a Xen ECS to a KVM ECS (Linux)** for manual operations.
- FAQs related to a script installation failure:
  - **What Should I Do If Executing a Driver Installation Script Failed on an ECS Running CentOS 5?**
  - **What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify the Specifications of a Linux ECS?**

Step 3: Modify Specifications

1. Log in to management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, view the status of the target ECS.

If the ECS is not in **Stopped** state, click **More** in the **Operation** column and select **Stop**.

5. Click **More** in the **Operation** column and select **Modify Specifications**.

   The **Modify ECS Specifications** page is displayed.

6. Select the new ECS type, vCPUs, and memory as prompted.
7. (Optional) Set **DeH**.

   If the ECS is created on a DeH, you can change the DeH where the ECS resides.

   To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.

8. Select the check box to confirm that the configuration is complete.
9. Click **OK**.

   NOTE

   - The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
   - If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is not affected.

(Optional) Step 4: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Linux ECS

  For details, see **Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?**

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, contact customer service or reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

NOTE

Reinstalling the OS will clear the system disk data, but the data on data disks is not affected.

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

1.7.4 Manually Changing a Xen ECS to a KVM ECS (Linux)

Scenarios

Before changing a Xen ECS that runs Linux to a KVM ECS, install and configure required drivers.

This section describes how to manually install drivers on a Linux ECS, configure automatic disk attachment, and change Xen to KVM.

For instructions about how to use a script to automatically install drivers, see **Automatically Changing a Xen ECS to a KVM ECS (Linux)**.

NOTE

- Xen ECSs include S1, C1, C2, and M1 ECSs.
- To obtain KVM ECSs, see the **Virtualization** column in **ECS Specifications**.
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

Constraints

- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

Procedure

Figure 1-67 shows the flowchart for manually changing a Xen ECS to a KVM ECS.

Figure 1-67 Flowchart for manually changing a Xen ECS to a KVM ECS

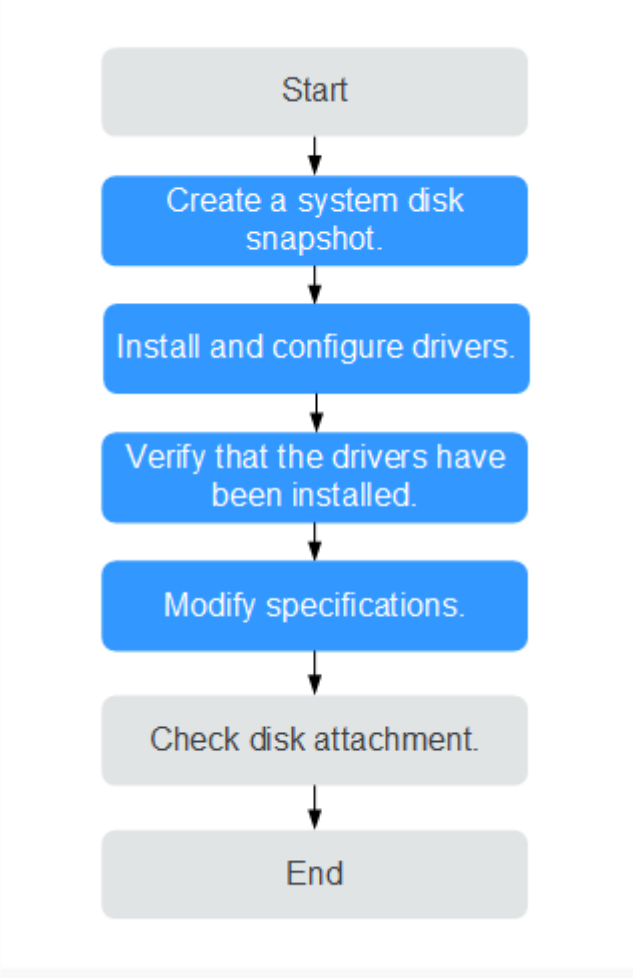

Table 1-18 Procedure for manually changing a Xen ECS to a KVM ECS

Step	Task
1	**Step 1: Create a System Disk Snapshot**
2	**Step 2: Install Drivers**
3	**Step 3: Check Whether the ECS Is Configured Correctly**
4	**Step 4: Modify Specifications**
5	**(Optional) Step 5: Check Disk Attachment**

Step 1: Create a System Disk Snapshot

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. Therefore, create a snapshot for the system disk first.

1. Before you create a system disk snapshot, check the ECS.

   Stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. For instructions about how to create a system disk snapshot, see **Creating a Snapshot** in *Elastic Volume Service User Guide*.

NOTE

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

Step 2: Install Drivers

Perform the following operations to manually install drivers on an ECS.

1. Log in to the ECS.
2. Uninstall tools from the ECS.

   For details, see **Uninstalling PV Drivers from a Linux ECS**.
3. Change the GRUB disk ID to UUID.

   For details, see **Changing the Disk Identifier in the GRUB Configuration File to UUID**.
4. Change the fstab disk ID to UUID.

   For details, see **Changing the Disk Identifier in the fstab File to UUID**.
5. Install native Xen and KVM drivers.

   For details, see **How Do I Install the Native Xen and KVM Drivers?**

Step 3: Check Whether the ECS Is Configured Correctly

Perform the following operations to check whether the drivers have been installed and the configuration files have been modified.

NOTE

Before manually modifying specifications, make sure that the ECS has been configured correctly.

1. Log in to the ECS.
2. Run the following command to check whether the root partition is in UUID format:

   cat /boot/grub/grub.cfg

   - If yes, the disk ID in the GRUB configuration file has been changed to UUID.
   - If no, the modification failed. In such a case, change the GRUB disk ID to UUID again by referring to **Step 2: Install Drivers**.

```
...menuentry 'Ubuntu Linux, with Linux 3.13.0-24-generic' --class ubuntu --class gnu-linux --class gnu
--class os --unrestricted $menuentry_id_option 'gnulinux-3.13.0-24-generic-advanced-
ec51d860-34bf-4374-ad46-a0c3e337fd34' {
recordfail
load_video
gfxmode $linux_gfx_mode
insmod gzio
insmod part_msdos
insmod ext2
```



```
if [ x$feature_platform_search_hint = xy ]; then
search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34
else
search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34
fi
echo 'Loading Linux 3.13.0-24-generic ...'
linux /boot/vmlinuz-3.13.0-24-generic root=UUID=ec51d860-34bf-4374-ad46-a0c3e337fd34 ro
echo 'Loading initial ramdisk ...'
initrd /boot/initrd.img-3.13.0-24-generic
}
```

NOTE

The path in which the GRUB configuration file is stored varies depending on the OS. For example, the path can be **/boot/grub/menu.lst**, **/boot/grub/grub.cfg**, **/boot/grub2/grub.cfg**, or **/boot/grub/grub.conf**.

3. Run the following command to check whether the disk ID in the fstab configuration file is UUID:

   cat /etc/fstab

   - If yes, the disk ID has been changed to UUID.
   - If no, the modification failed. In such a case, change the fstab disk ID to UUID again by referring to **Step 2: Install Drivers**.

```
[root@****** ~]# cat /etc/fstab
UUID=4eb40294-4c6f-4384-bbb6-b8795bbb1130 /      xfs     defaults    0 0
UUID=2de37c6b-2648-43b4-a4f5-40162154e135 swap    swap    defaults    0 0
```

4. Check whether the native Xen and KVM drivers have been installed.

   - If the boot virtual file system is initramfs, run the following commands:

     lsinitrd /boot/initramfs-`uname -r`.img | grep ` uname -r ` | grep xen

     lsinitrd /boot/initramfs-`uname -r`.img | grep ` uname -r ` |grep virtio

   - If the boot virtual file system is initrd, run the following commands:

     lsinitrd /boot/initrd-`uname -r` | grep ` uname -r ` | grep xen

     lsinitrd /boot/initrd-`uname -r` | grep ` uname -r ` | grep virtio

   If the names of the native Xen and KVM drivers are displayed in the command output, the drivers have been installed.

```
[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep ` uname -r`| grep xen
-rwxr--r--   1 root     root        54888 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
block/xen-blkfront.ko
-rwxr--r--   1 root     root        45664 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/
drivers/net/xen-netfront.ko

[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep ` uname -r`| grep virtio
-rwxr--r--   1 root     root        23448 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
block/virtio_blk.ko
-rwxr--r--   1 root     root        50704 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/
drivers/net/virtio_net.ko
-rwxr--r--   1 root     root        28424 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
scsi/virtio_scsi.ko
drwxr-xr-x   2 root     root            0 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio
-rwxr--r--   1 root     root        14544 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio.ko
-rwxr--r--   1 root     root        21040 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio_pci.ko
-rwxr--r--   1 root     root        18016 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio_ring.ko
```

NOTE

Make sure that the ECS has been configured successfully, or the ECS may become unavailable after the specifications are modified.

Step 4: Modify Specifications

1. Log in to management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, view the status of the target ECS.

   If the ECS is not in **Stopped** state, click **More** in the **Operation** column and select **Stop**.
5. Click **More** in the **Operation** column and select **Modify Specifications**.

   The **Modify ECS Specifications** page is displayed.
6. Select the new ECS type, vCPUs, and memory as prompted.
7. (Optional) Set **DeH**.

   If the ECS is created on a DeH, you can change the DeH where the ECS resides.

   To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with specifications modified.
8. Select the check box to confirm that the configuration is complete.
9. Click **OK**.

   NOTE

   - The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
   - If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is not affected.

(Optional) Step 5: Check Disk Attachment

After a Xen ECS is changed to a KVM ECS, disk attachment may fail. Therefore, check disk attachment after specifications modification. If disks are properly attached, the specifications modification is successful.

- Linux ECS

  For details, see **Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?**

Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, contact customer service or reinstall the ECS OS to resolve this issue. For details, see **Reinstalling the OS**.

NOTE

Reinstalling the OS will clear the system disk data, but the data on data disks is not affected.

After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.

1.8 Migrating an ECS

Scenarios

ECSs can be migrated:

- Between Dedicated Hosts (DeHs)
- From a DeH to a public resource pool
- From a public resource pool to a DeH

This section describes how to migrate ECSs from a public resource pool to a DeH.

NOTE

- Before migrating an ECS, ensure that there are available DeH resources.
- For details about migrating ECSs from a DeH to another DeH or to a public resource pool, see **Migrating ECSs**.

Constraints

- Only stopped ECSs can be migrated.
- To ensure that the migration is successful, there must be an available DeH.
- ECS IDs remain unchanged after a migration.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row that contains the target ECS and choose **More** > **Migrate ECS** in the **Operation** column.
5. In the displayed dialog box, select the target DeH.

   NOTE

   If no DeHs are available, create a DeH first. For details, see **Buying DeHs**.
6. Click **OK**.

1.9 Obtaining Metadata and Passing User Data

1.9.1 Obtaining Metadata

Scenarios

ECS metadata includes basic information of an ECS on the cloud platform, such as the ECS ID, hostname, and network information. ECS metadata can be obtained using either OpenStack or EC2 compatible APIs, as shown in **Table 1-19**. The following describes the URI and methods of using the supported ECS metadata.

Notes

If the metadata contains sensitive data, take appropriate measures to protect the sensitive data, for example, controlling access permissions and encrypting the data.

Perform the following configuration on the firewall:

- Windows

  If you need to assign permissions only to the administrator to access custom data, enable the firewall as an administrator and run the following commands in PowerShell:

  PS C:\>$RejectPrincipal = New-Object -TypeName System.Security.Principal.NTAccount ("Everyone")

  PS C:\>$RejectPrincipalSID = $RejectPrincipal.Translate([System.Security.Principal.SecurityIdentifier]).Value

  PS C:\>$ExceptPrincipal = New-Object -TypeName System.Security.Principal.NTAccount ("Administrator")

  PS C:\>$ExceptPrincipalSID = $ExceptPrincipal.Translate([System.Security.Principal.SecurityIdentifier]).Value

  PS C:\>$PrincipalSDDL = "O:LSD:(D;;CC;;;$ExceptPrincipalSID)(A;;CC;;;$RejectPrincipalSID)"

  PS C:\>New-NetFirewallRule -DisplayName "Reject metadata service for $($RejectPrincipal.Value), exception: $($ExceptPrincipal.Value)" -Action block -Direction out -Protocol TCP -RemoteAddress 169.254.169.254 -LocalUser $PrincipalSDDL

- Linux

  If you need to assign permissions only to user **root** to access custom data, run the following command as user **root**:

  iptables --append OUTPUT --proto tcp --destination 169.254.169.254 --match owner ! --uid-owner root --jump REJECT

ECS Metadata Types

Table 1-19 does not contain the following metadata items: ami-id, ami-launch-index, ami-manifest-path, block-device-mapping/, instance-action, instance-id, reservation-id, ramdisk-id, and kernel-id. These metadata items are meaningless and are not recommended.

Table 1-19 ECS metadata types

Metadata Type	Metadata Item	Description
OpenStack	/meta_data.json	Displays ECS metadata. For the key fields in the ECS metadata, see **Table 1-20**.
OpenStack	/password	Displays the password for logging in to an ECS. This metadata is used by Cloudbase-Init to store ciphertext passwords during initialization of key-pair-authenticated Windows ECSs.
OpenStack	/user_data	Displays ECS user data. This metadata allows you to specify scripts and configuration files for initializing ECSs. For details, see **Passing User Data to ECSs**. For password-authenticated Linux ECSs, this metadata is used to save password injection scripts.
OpenStack	/network_data.json	Displays ECS network information.
OpenStack	/securitykey	Obtains temporary AKs and SKs. Before enabling an ECS to obtain a temporary AK and SK, make sure that the **op_svc_ecs** account has been authorized on IAM and that the desired ECS resources have been authorized for management.
OpenStack	/spot/instance-action	Queries the prompt of stopping a spot ECS.
EC2	/meta-data/hostname	Displays the name of the host accommodating an ECS. To remove the suffix **.novalocal** from an ECS, see: **Is an ECS Hostname with Suffix .novalocal Normal?**
EC2	/meta-data/local-hostname	The meaning of this field is the same as that of hostname.
EC2	/meta-data/public-hostname	The meaning of this field is the same as that of hostname.
EC2	/meta-data/instance-type	Displays an ECS flavor.

Metadata Type	Metadata Item	Description
EC2	/meta-data/local-ipv4	Displays the fixed IP address of an ECS. If there are multiple NICs, only the IP address of the primary NIC is displayed.
EC2	/meta-data/placement/availability-zone	Displays the AZ accommodating an ECS.
EC2	/meta-data/public-ipv4	Displays the EIP bound to an ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.
EC2	/meta-data/public-keys/0/openssh-key	Displays the public key of an ECS.
EC2	/user-data	Displays ECS user data.
EC2	/meta-data/security-groups	Displays the security group of an ECS.

Table 1-20 Metadata key fields

Parameter	Type	Description
uuid	String	Specifies an ECS ID.
availability_zone	String	Specifies the AZ where an ECS locates.
meta	Dict	Specifies the metadata information, including the image name, image ID, and VPC ID.
hostname	String	Specifies the name of the host accommodating an ECS. To remove the suffix **.novalocal** from an ECS, see: **Is an ECS Hostname with Suffix .novalocal Normal?**
enterprise_project_id	String	Specifies the ID of the enterprise project accommodating an ECS.

Prerequisites

- The target ECS has been logged in.
- Security group rules in the outbound direction meet the following requirements:

- **Protocol**: **TCP**
- **Port**: **80**
- **Destination**: **169.254.0.0/16**

NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. Default security group rules for the outbound direction are as follows:

- **Protocol**: **All**
- **Port**: **All**
- **Destination**: **0.0.0.0/0**

Metadata (OpenStack Metadata API)

This API is used to query ECS metadata.

- URI

  /169.254.169.254/openstack/latest/meta_data.json
- Usage method

  Supports GET requests.
- Example

  To use cURL to view Linux ECS metadata, run the following command:

  curl http://169.254.169.254/openstack/latest/meta_data.json

  To use Invoke-RestMethod to view Windows ECS metadata, run the following command:

  Invoke-RestMethod http://169.254.169.254/openstack/latest/meta_data.json | ConvertTo-Json


```
{
    "random_seed": "rEocCViRS+dNwlYdGIxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRNy4kKGoNPEVBCc05Hg1TcDblAPfJwgJS1okqEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGgArucn/
WzDcy19DGioKPE7F8ILtSQ4Ww3VClK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbO3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmxth7IGNmRMQKV6+miI78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNiHA/NvlEsxJDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmtCNzw3Ra0hiKchGhqK3BIeToV/kVx5DdF081xrEA
+qyoM6CVyfJtEoz1zlRRyoo9bJ65Eg6JJd8dj1UCVsDqRY1pIjgzE/
Mzsw6AaaCVhaMJL7u7YMVdyKzA6z65Xtvujz0Vo=",
    "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
    "availability_zone": "lt-test-1c",
    "enterprise_project_id" : "0",
    "hostname": "ecs-ddd4.novalocal",
    "launch_index": 0,
    "instance_type": "s3.medium.2",
    "meta": {
        "metering.image_id": "3a64bd37-955e-40cd-ab9e-129db56bc05d",
        "metering.imagetype": "gold",
        "metering.resourcespeccode": "s3.medium.2.linux",
        "metering.cloudServiceType": "hws.service.type.ec2",
        "image_name": "CentOS 7.6 64bit",
        "metering.resourcetype": "1",
        "vpc_id": "3b6c201f-aeb3-4bce-b841-64756e66cb49",
        "os_bit": "64",
        "cascaded.instance_extrainfo": "pcibridge:1",
        "os_type": "Linux",
```



```
        "charging_mode": "0"
    },
    "region_id": "xxx",
    "project_id": "6e8b0c94265645f39c5abbe63c4113c6",
    "name": "ecs-ddd4"
}
```

User Data (OpenStack Metadata API)

This API is used to query ECS user data. The value is configured only when you create an ECS. It cannot be changed after the configuration.

- URI

  /169.254.169.254/openstack/latest/user_data

- Usage method

  Supports GET requests.

- Example

  Linux:

  curl http://169.254.169.254/openstack/latest/user_data

  Windows:

  Invoke-RestMethod http://169.254.169.254/openstack/latest/user_data

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIH
BsYWNlIHRvIGdvIG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJu
cyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3csIHRvbywgd2hlbiB5b3UgbGlmdCB5b3
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA==
```

> **NOTE**
>
> If user data was not passed to the ECS during ECS creation, the query result is 404.

Figure 1-68 404 Not Found

```
[root@pythonsdktempest--server-1519783681 ~]# curl http://169.254.169.254/openstack/latest/user_data
<html>
 <head>
  <title>404 Not Found</title>
 </head>
 <body>
  <h1>404 Not Found</h1>
  The resource could not be found.<br /><br />


 </body>
</html>
```

Network Data (OpenStack Metadata API)

This API is used to query information about all NICs attached to an ECS, including their DNS server addresses, network bandwidth, IDs, private IP addresses, EIPs, and MAC addresses.

- URI

  /openstack/latest/network_data.json

- Usage method

  Supports GET requests.

- Example

NOTE

instance_max_bandwidth and **instance_min_bandwidth** are in the unit of Mbit/s. If the value is **-1**, the bandwidth is not limited.

Linux:

curl http://169.254.169.254/openstack/latest/network_data.json

Windows:

Invoke-RestMethod http://169.254.169.254/openstack/latest/network_data.json | ConvertTo-Json

```
{
   "services": [{
      "type": "dns",
      "address": "xxx.xx.x.x"
   },
   {
      "type": "dns",
      "address": "100.125.21.250"
   }],
   "qos":{
      "instance_min_bandwidth": 100,
      "instance_max_bandwidth": 500
   },
   "networks": [{
      "network_id": "67dc10ce-441f-4592-9a80-cc709f6436e7",
      "type": "ipv4_dhcp",
      "link": "tap68a9272d-71",
      "id": "network0"
   }],
   "links": [{
      "vif_id": "68a9272d-7152-4ae7-a138-3ef53af669e7",
      "public_ipv4": "100.100.xx.xx",
      "ethernet_mac_address": "fa:16:3e:f7:c1:47",
      "mtu": null,
      "local_ipv4": "192.169.10.10",
      "type": "cascading",
      "id": "tap68a9272d-71"
   }]
}
```

Security Key (OpenStack Metadata API)

This API is used to obtain temporary AKs and SKs.

NOTE

- If an ECS needs to obtain a temporary AK and SK, go to the ECS details page, and configure **Agency** for the ECS in the **Management Information** area so that the ECS is authorized on IAM.

  For details, see **Cloud Service Delegation**.
- The validity period of a temporary AK and SK is one hour. The temporary AK and SK are updated 10 minutes ahead of the expiration time. During the 10 minutes, both the new and old temporary AKs and SKs can be used.
- When using temporary AKs and SKs, add **'X-Security-Token':{securitytoken}** in the message header. **securitytoken** is the value returned when a call is made to the API.

- URI

  /openstack/latest/securitykey
- Usage method

  Supports GET requests.

- Examples

  Linux:

  curl http://169.254.169.254/openstack/latest/securitykey

  Windows:

  Invoke-RestMethod http://169.254.169.254/openstack/latest/securitykey

Instance Action (OpenStack Metadata API)

This API is used to query the prompt of stopping a spot ECS.

NOTE

If your spot ECS is about to be interrupted, this API returns the estimated time of stopping that spot ECS.

- URI

  /openstack/latest/spot/instance-action

- Usage method

  Supports GET requests.

- Example

  Linux:

  curl http://169.254.169.254/openstack/latest/spot/instance-action

  Windows:

  Invoke-RestMethod http://169.254.169.254/openstack/latest/spot/instance-action

```
{"action":"terminate","timestamp":"2023-06-01 09:15:00"}
```

User Data (EC2 Compatible API)

This API is used to query ECS user data. The value is configured only when you create an ECS. It cannot be changed after the configuration.

- URI

  /169.254.169.254/latest/user-data

- Usage method

  Supports GET requests.

- Example

  Linux:

  curl http://169.254.169.254/latest/user-data

  Windows:

  Invoke-RestMethod http://169.254.169.254/latest/user-data

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbiBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhlIH
BsYWNlIHRvIGdvIG5vdy4gQnV0IHRoZSBza3kga25vd3MgdGhlIHJlYXNvbnMgYW5kIHRoZSBwYXR0ZXJu
cyBiZWhpbmQgYWxsIGNsb3VkcywgYW5kIHlvdSB3aWxsIGtub3csIHRvbywgd2hlbiB5b3UgbGlmdCB5b3
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA==
```

Hostname (EC2 Compatible API)

This API is used to query the name of the host accommodating an ECS. The **.novalocal** suffix will be added later.

- URI

  /169.254.169.254/latest/meta-data/hostname
- Usage method

  Supports GET requests.
- Example

  Linux:

  curl http://169.254.169.254/latest/meta-data/hostname

  Windows:

  Invoke-RestMethod http://169.254.169.254/latest/meta-data/hostname

  ```
  vm-test.novalocal
  ```

Instance Type (EC2 Compatible API)

This API is used to query an ECS flavor.

- URI

  /169.254.169.254/latest/meta-data/instance-type
- Usage method

  Supports GET requests.
- Example

  Linux:

  curl http://169.254.169.254/latest/meta-data/instance-type

  Windows:

  Invoke-RestMethod http://169.254.169.254/latest/meta-data/instance-type

  ```
  s3.medium.2
  ```

Local IPv4 (EC2 Compatible API)

This API is used to query the fixed IP address of an ECS. If there are multiple NICs, only the IP address of the primary NIC is displayed.

- URI

  /169.254.169.254/latest/meta-data/local-ipv4
- Usage method

  Supports GET requests.
- Example

  Linux:

  curl http://169.254.169.254/latest/meta-data/local-ipv4

  Windows:

  Invoke-RestMethod http://169.254.169.254/latest/meta-data/local-ipv4

  ```
  192.1.1.2
  ```


Availability Zone (EC2 Compatible API)

This API is used to query the AZ accommodating an ECS.

- URI
  /169.254.169.254/latest/meta-data/placement/availability-zone
- Usage method
  Supports GET requests.
- Example
  Linux:
  curl http://169.254.169.254/latest/meta-data/placement/availability-zone
  Windows:
  Invoke-RestMethod http://169.254.169.254/latest/meta-data/placement/availability-zone

```
az1.dc1
```

Public IPv4 (EC2 Compatible API)

This API is used to query the EIP bound to an ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.

- URI
  /169.254.169.254/latest/meta-data/public-ipv4
- Usage method
  Supports GET requests.
- Example
  Linux:
  curl http://169.254.169.254/latest/meta-data/public-ipv4
  Windows:
  Invoke-RestMethod http://169.254.169.254/latest/meta-data/public-ipv4

```
46.1.1.2
```

Public Keys (EC2 Compatible API)

This API is used to query the public key of an ECS.

- URI
  /169.254.169.254/latest/meta-data/public-keys/0/openssh-key
- Usage method
  Supports GET requests.
- Example
  Linux:
  curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key
  Windows:
  Invoke-RestMethod http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key


```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAABAQDI5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSiIc/
hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/
WRenxIwR00KkczHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXAjH4eKoKTVNtMXAvPP9aMy2SLgsJNt
Mb9ArfziAiblQynq7UIfLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwlL6K4i
+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBIklOBfuUENIJUhAB
Generated-by-Nova
```

Helpful Links

Why My Linux ECS Cannot Obtain Metadata?

1.9.2 Passing User Data to ECSs

Scenarios

Specify **User Data** to pass user data to ECSs to:

- Simplify ECS configuration.
- Initialize the ECS OS configuration.
- Upload your scripts to ECSs during ECS creation.
- Perform other tasks using scripts.

Use Restrictions

- Linux
  - The image that is used to create ECSs must have Cloud-Init installed.
  - The user data to be specified must be less than or equal to 32 KB.
  - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
  - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloud-Init installed.
  - The format of the customized scripts must be supported by Linux ECSs.
  - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.
  - When the password login mode is selected, user data cannot be passed.
- Windows
  - The image that is used to create ECSs must have Cloudbase-Init installed.
  - The user data to be specified must be less than or equal to 32 KB.
  - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
  - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloudbase-Init installed.
  - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.

Passing User Data

1. Create a user data script, the format of which complies with user data script specifications. For details, see **Helpful Links**.
2. When creating an ECS, set **Advanced Options** to **Configure now**, and paste the content of the user data script to the **User Data** text box or upload the user data file.

   NOTE

   You can pass user data to an ECS as text or as a file.

   Text: Copy the content of the user data script to the text box.

   File: Save the user data script to a text file and then upload the file.

   Figure 1-69 User data injection

   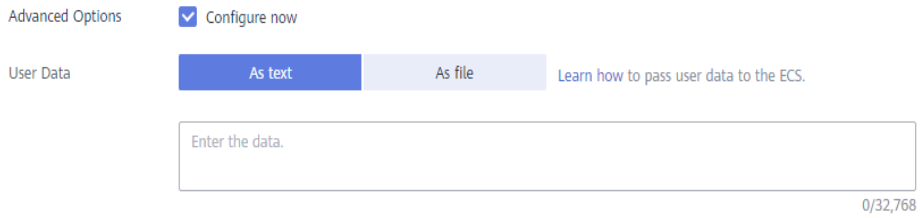

3. The created ECS automatically runs Cloud-Init/Cloudbase-Init and reads the user data script upon startup.

User Data Scripts of Linux ECSs

Customized user data scripts of Linux ECSs are based on the open-source Cloud-Init architecture. This architecture uses ECS metadata as the data source for automatically configuring the ECSs. The customized script types are compatible with open-source Cloud-Init. For details about Cloud-Init, see **http://cloudinit.readthedocs.io/en/latest/topics/format.html**.

- Script execution time: A customized user data script is executed after the status of the target ECS changes to **Running** and before **/etc/init** is executed.

  NOTE

  By default, the scripts are executed as user **root**.

- Script type: Both user-data scripts and Cloud-Config data scripts are supported.

  Table 1-21 Linux ECS script types

-	User-Data Script	Cloud-Config Data Script
Description	Scripts, such as Shell and Python scripts, are used for custom configurations.	Methods pre-defined in Cloud-Init, such as the yum repository and SSH key, are used for configuring certain ECS applications.

-	User-Data Script	Cloud-Config Data Script
Format	The first line must start with **#!** (for example, **#!/bin/bash** or **#!/usr/bin/env python**) and no spaces are allowed at the beginning. When a script is started for the first time, it will be executed at the rc.local-like level, indicating a low priority in the boot sequence.	The first line must be **#cloud-config**, and no space is allowed in front of it.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.
Frequency	The script is executed only once when the ECS is started for the first time.	The execution frequency varies according to the applications configured on the ECS.

- How can I view the customized user data passed to a Linux ECS?
  a. Log in to the ECS.
  b. Run the following command to view the customized user data as user **root**:

     curl http://169.254.169.254/openstack/latest/user_data
- Script usage examples

  This section describes how to inject scripts in different formats into Linux ECSs and view script execution results.

  Example 1: Inject a user-data script.

  When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

  ```
  #!/bin/bash
  echo "Hello, the time is now $(date -R)" | tee /root/output.txt
  ```

  After the ECS is created, start it and run the **cat** *[file]* command to check the script execution result.

  ```
  [root@XXXXXXXX ~]# cat /root/output.txt
  Hello, the time is now Mon, 16 Jul 2016 16:03:18+0800
  ```

  Example 2: Inject a Cloud-Config data script.

  When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

  ```
  #cloud-config
  bootcmd:
  - echo 192.168.1.130 us.archive.ubuntu.com >> /etc/hosts
  ```

  After the ECS is created, start it and run the **cat /etc/hosts** command to check the script execution result.

Figure 1-70 Viewing operating results

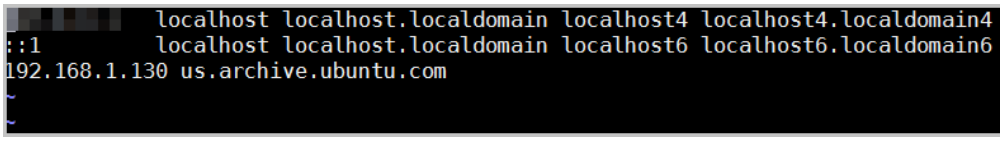

User Data Scripts of Windows ECSs

Customized user data scripts of Windows ECSs are based on the open-source Cloudbase-Init architecture. This architecture uses ECS metadata as the data source for initializing and automatically configuring the ECSs. The customized script types are compatible with open-source Cloudbase-Init. For details about Cloudbase-Init, see **https://cloudbase-init.readthedocs.io/en/latest/userdata.html**.

- Script type: Both batch-processing program scripts and PowerShell scripts are supported.

  Table 1-22 Windows ECS script types

-	Batch-Processing Program Script	PowerShell Script
Format	The script must be started with **rem cmd**, which is the first line of the script. No space is allowed at the beginning of the first line.	The script must be started with **#ps1**, which is the first line of the script. No space is allowed at the beginning of the first line.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.

- How can I view the customized user data passed into a Windows ECS?

  a. Log in to the ECS.

  b. Access the following URL in the address box of the browser and view the user data:

  http://169.254.169.254/openstack/latest/user_data

- Script usage examples

  This section describes how to inject scripts in different formats into Windows ECSs and view script execution results.

  Example 1: Inject a batch-processing program script.

  When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
rem cmd
echo "Hello, BAT Test" > C:\1111.txt
```

  After the ECS is created, start it and check the script execution result. In this example, a text file named **1111** is added to disk C:\.

Figure 1-71 Creating text file (Batch)

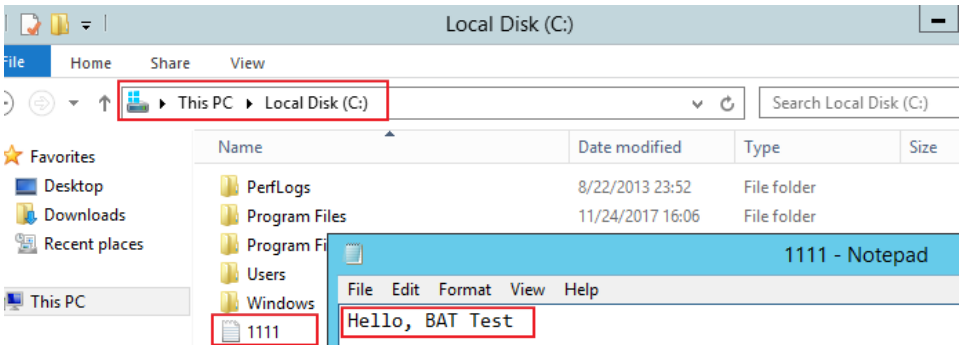

To view the user data passed to the Windows ECS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 1-72 Viewing user data (Batch)

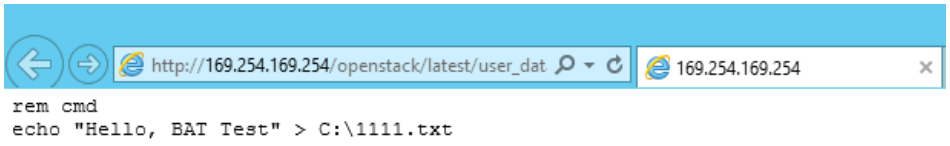

Example 2: Inject a PowerShell script.

When creating an ECS, set **User Data** to **As text** and enter the customized user data script.

```
#ps1
echo "Hello, Powershell Test" > C:\aaaa.txt
```

After the ECS is created, start it and check the script execution result. In this example, a text file named **aaaa** is added to disk C:\.

Figure 1-73 Creating text file (PowerShell)

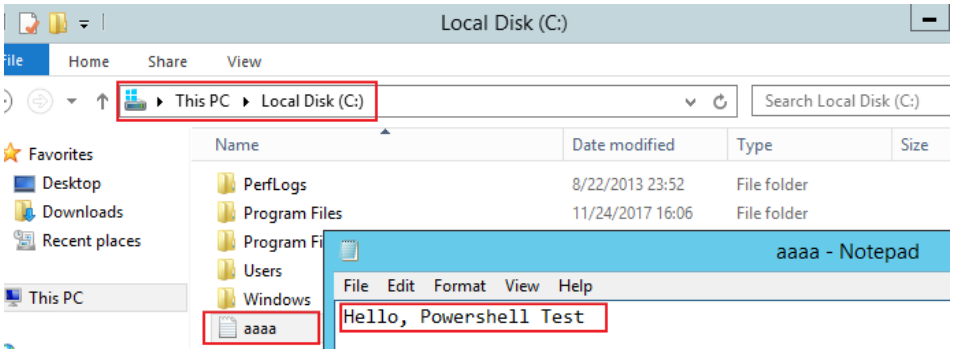

To view the user data passed to the Windows ECS, log in at http://169.254.169.254/openstack/latest/user_data.

Figure 1-74 Viewing user data (PowerShell)

Case 1

This case illustrates how to pass user data to simplify Linux ECS configuration.

In this example, vim is configured to enable syntax highlighting, display line numbers, and set the tab stop to **4**. The .vimrc configuration file is created and injected into the **/root/.vimrc** directory during ECS creation. After the ECS is created, vim is automatically configured based on your requirements. This improves ECS configuration efficiency, especially in batch ECS creation scenarios.

User data example:

```
#cloud-config
write_files:
  - path: /root/.vimrc
    content: |
      syntax on
      set tabstop=4
      set number
```

Case 2

This case illustrates how to use the user data passing function to set the password for logging in to a Linux ECS.

NOTE

The new password must meet the password complexity requirements listed in **Table 1-23**.

Table 1-23 Password complexity requirements

Parameter	Requirement	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Windows: $!@%-_=+[]:./,? – Special characters for Linux: !@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not use it.

User data example:

Using a ciphertext password (recommended)

```
#!/bin/bash
echo 'root:$6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig' | chpasswd -e;
```

In the preceding command output, **6V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig** is the ciphertext password, which can be generated as follows:

1. Run the following command to generate an encrypted ciphertext value:

   python -c "import crypt, getpass, pwd;print crypt.mksalt()"

   The following information is displayed:

   ```
   $6$V6azyeLwcD3CHlpY
   ```

2. Run the following command to generate a ciphertext password based on the salt value:

   python -c "import crypt, getpass, pwd;print crypt.crypt('Cloud.1234','\$6\$V6azyeLwcD3CHlpY')"

   The following information is displayed:

   ```
   $6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig
   ```

After the ECS is created, you can use the password to log in to it.

Case 3

This case illustrates how to use the user data passing function to reset the password for logging in to a Linux ECS.

In this example, the password of user **root** is reset to ******.

NOTE

The new password must meet the password complexity requirements listed in **Table 1-24**.

Table 1-24 Password complexity requirements

Parameter	Requirement	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Windows: $!@%-_=+[]:./,? – Special characters for Linux: !@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not use it.

User data example (Retain the indentation in the following script):

```
#cloud-config
chpasswd:
  list: |
    root:******
  expire: False
```

After the ECS is created, you can use the reset password to log in to it. To ensure system security, change the password of user **root** after logging in to the ECS for the first time.

Case 4

This case illustrates how to use the user data passing function to create a user on a Windows ECS and configure the password for the user.

In this example, the user's username is **abc**, its password is **********, and the user is added to the **administrators** user group.

NOTE

The new password must meet the password complexity requirements listed in **Table 1-24**.

User data example:

```
rem cmd
net user abc ****** /add
net localgroup administrators abc /add
```

After the ECS is created, you can use the created username and password to log in to it.

Case 5

This case illustrates how to use the user data passing function to update system software packages for a Linux ECS and enable the HTTPd service. After the user data is passed to an ECS, you can use the HTTPd service.

User data example:

```
#!/bin/bash
yum update -y
service httpd start
chkconfig httpd on
```

Case 6

This case illustrates how to use the user data passing function to assign user **root** permission for remotely logging in to a Linux ECS. After passing the file to an ECS, you can log in to the ECS as user **root** using SSH key pair authentication.

User data example:

```
#cloud-config
disable_root: false
runcmd:
- sed -i 's/^PermitRootLogin.*$/PermitRootLogin without-password/' /etc/ssh/sshd_config
```



```
- sed -i '/^KexAlgorithms.*$/d' /etc/ssh/sshd_config
- service sshd restart
```

Helpful Links

For more information about user data passing cases, visit the official Cloud-init/Cloudbase-init website:

- **https://cloudinit.readthedocs.io/en/latest/**
- **https://cloudbase-init.readthedocs.io/en/latest/**

1.10 (Optional) Configuring Mapping Between Hostnames and IP Addresses

ECSs in the same VPC can communicate with each other using hostnames. In such a case, you are required to configure the mapping between hostnames and IP addresses. The communication using hostnames is more convenient than that using IP addresses.

Constraints

This method applies only to Linux ECSs.

Procedure

For example, there are two ECSs in a VPC, ecs-01 and ecs-02. Perform the following operations to enable communication using hostnames between ecs-01 and ecs-02:

Step 1 Log in to ecs-01 and ecs-02 and obtain their private IP addresses.

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, obtain the private IP address in the **IP Address** column.

   For example, the obtained private IP addresses are as follows:

   ecs-01: 192.168.0.1

   ecs-02: 192.168.0.2

Step 2 Obtain the hostnames for the two ECSs.

1. Log in to an ECS.
2. Run the following command to view the ECS hostname:

   sudo hostname

   For example, the obtained hostnames are as follows:

   ecs-01: hostname01

   ecs-02: hostname02

Step 3 Create a mapping between the hostnames and IP addresses and add information about other ECSs in the same VPC.

1. Log in to ecs-01.
2. Run the following command to switch to user **root**:
   sudo su -
3. Run the following command to edit the hosts configuration file:
   vi /etc/hosts
4. Press **i** to enter editing mode.
5. Add the statement in the following format to set up the mapping:
   Private IP address hostname
   For example, add the following statement:
   192.168.0.1 hostname01
   192.168.0.2 hostname02
6. Press **Esc** to exit editing mode.
7. Run the following command to save the configuration and exit:
   :wq
8. Log in to ecs-02.
9. Repeat **Step 3.2** to **Step 3.7**.

Step 4 Check whether the ECSs can communicate with each other using hostnames.

Log in to an ECS in the same VPC, run the following command to ping the added host, and check whether the operation is successful:

ping *Hostname*

----End

1.11 (Optional) Installing a Driver and Toolkit

1.11.1 GPU Driver

Overview

Before using a GPU-accelerated ECS, make sure that a GPU driver has been installed on the ECS for GPU acceleration.

GPU-accelerated ECSs support GRID and Tesla drivers.

- To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.
  - A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be purchased and configured separately. Before using such an ECS, check whether the desired driver has been installed on it and whether the version of the installed driver meets service requirements.
  - To install a GRID driver on a GPU-accelerated ECS created using a private image, see **Installing a GRID Driver on a GPU-accelerated ECS**.

- To use computing acceleration, install a Tesla driver.
  - A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
  - To install a Tesla driver on a GPU-accelerated ECS created using a private image, see **Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS**.

Table 1-25 Acceleration supported by GPU drivers

Driver	License	CUDA	OpenGL	DirectX	Vulkan	Application Scenario	Description
GRID	Required	Supported	Supported	Supported	Supported	3D rendering, graphics workstation, and game acceleration	The GRID driver must be paid and requires a license to accelerate graphics and image applications.
Tesla	Not required	Supported	Not supported	Not supported	Not supported	Scientific computing, deep learning training, and inference	The Tesla driver is downloaded free of charge and usually used with NVIDIA CUDA SDKs to accelerate general computing applications.

1.11.2 Installing a GRID Driver on a GPU-accelerated ECS

Scenarios

To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.

- A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be purchased and configured separately.
- If a GPU-accelerated ECS is created using a private image, install a GRID driver and separately purchase and configure a GRID license.

This section describes how to install a GRID driver, purchase or apply for a GRID license, and configure the license server.

Process of installing a GRID driver:

1. **Purchasing a GRID License**
2. **Downloading GRID Driver and Software License Packages**
3. **Deploying and Configuring the License Server**
4. **Installing the GRID Driver and Configuring the License**

NOTE

- NVIDIA allows you to apply for a 90-day trial license.
- For details about GPU-accelerated ECSs with different specifications and application scenarios, see **GPU-accelerated ECSs**.

Purchasing a GRID License

- Purchase a license.

  To obtain an official license, contact NVIDIA or their NVIDIA agent in your local country or region.

- Apply for a trial license.

  Log in at the **official NVIDIA website** and enter desired information.

  For details about how to register for an account and apply for a trial license, see **official NVIDIA help page**.

  NOTE

  The method of using a trial license is the same as that of using an official license. You can use an official license to activate an account with a trial license to prevent repetitive registration. The trial license has a validity period of 90 days. After the trial license expires, it cannot be used anymore. Purchase an official license then.

Figure 1-75 Applying for a trial license

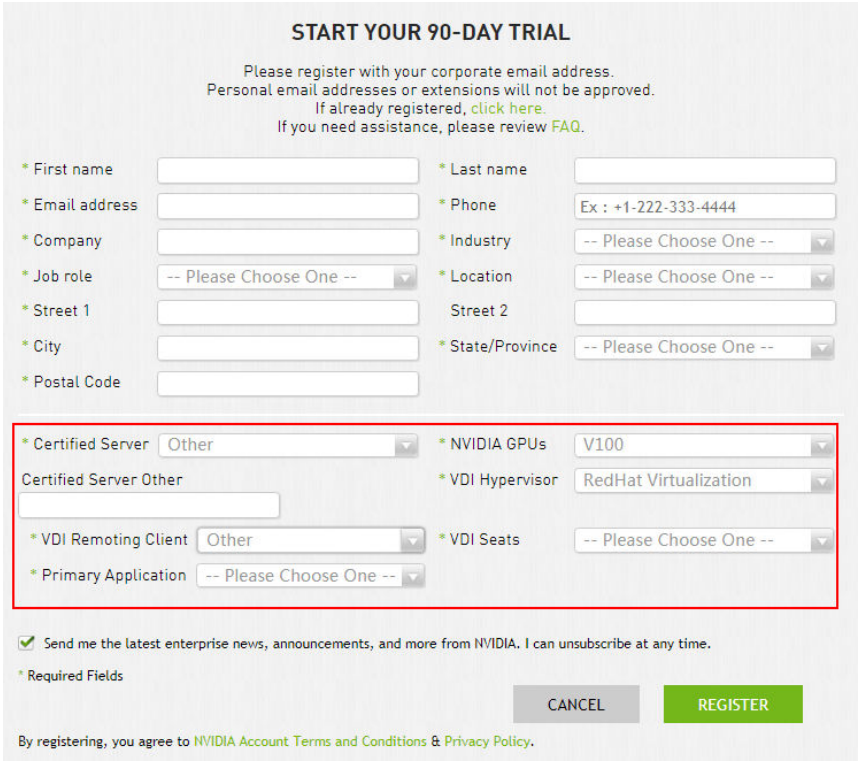

Downloading GRID Driver and Software License Packages

1. Obtain the driver installation package required for an OS. For details, see **Table 1-26**.

   For more information about the GRID driver, see **NVIDIA vGPU Software Documentation**.

   NOTE

   For a GPU passthrough ECS, select a GRID driver version as required.

   For a GPU virtualization ECS, select a driver version based on the following table.

   Table 1-26 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
G5.8xlarge.4	GPU passthrough	• Windows Server 2019 Standard 64bit • Windows Server 2016 Standard 64bit • CentOS 8.2 64bit • CentOS 7.6 64bit • CentOS 7.5 64bit • Ubuntu Server 20.04 64bit • Ubuntu Server 18.04 64bit	Select a version as required.	x86_64
P2s	GPU passthrough	• Windows Server 2019 Standard 64bit • Windows Server 2016 Standard 64bit • CentOS 8.2 64bit • CentOS 7.6 64bit • Ubuntu Server 20.04 64bit • Ubuntu Server 18.04 64bit	Select a version as required.	x86_64

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
P2v	GPU passthrough	• Windows Server 2019 Standard 64bit • Windows Server 2016 Standard 64bit • Ubuntu Server 16.04 64bit • CentOS 7.4 64bit • EulerOS 2.2 64bit	Select a version as required.	x86_64
PI2	GPU passthrough	• Windows Server 2019 Standard 64bit • Windows Server 2016 Standard 64bit • CentOS 7.5 64bit	Select a version as required.	x86_64
PI1	GPU passthrough	• Windows Server 2019 Standard 64bit • Windows Server 2016 Standard 64bit • CentOS 7.3 64bit • Ubuntu Server 16.04 64bit • Ubuntu Server 14.04 64bit	Select a version as required.	x86_64

2. After the registration, log in at the **official NVIDIA website** and enter the account.
3. Check whether NVIDIA is used for the first time.
   a. If yes, go to step **4**.
   b. If no, go to step **6**.
4. Obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 1-76 PAK

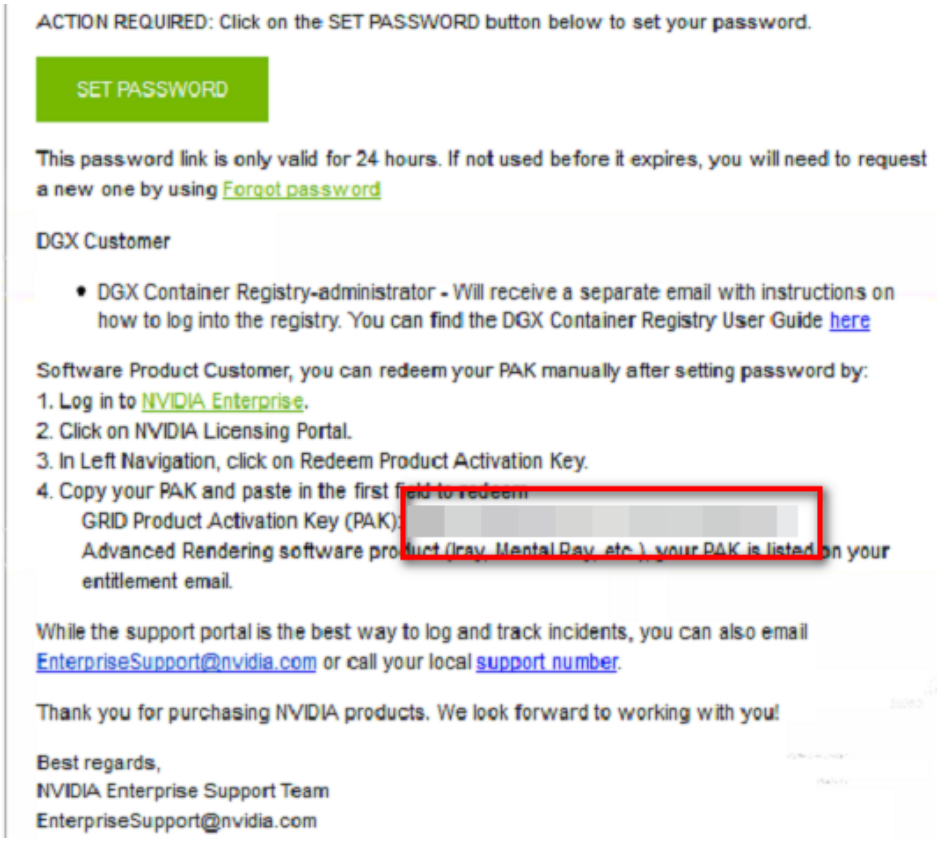

5. Enter the PAK obtained in step **4** on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 1-77 Redeem Product Activation Keys

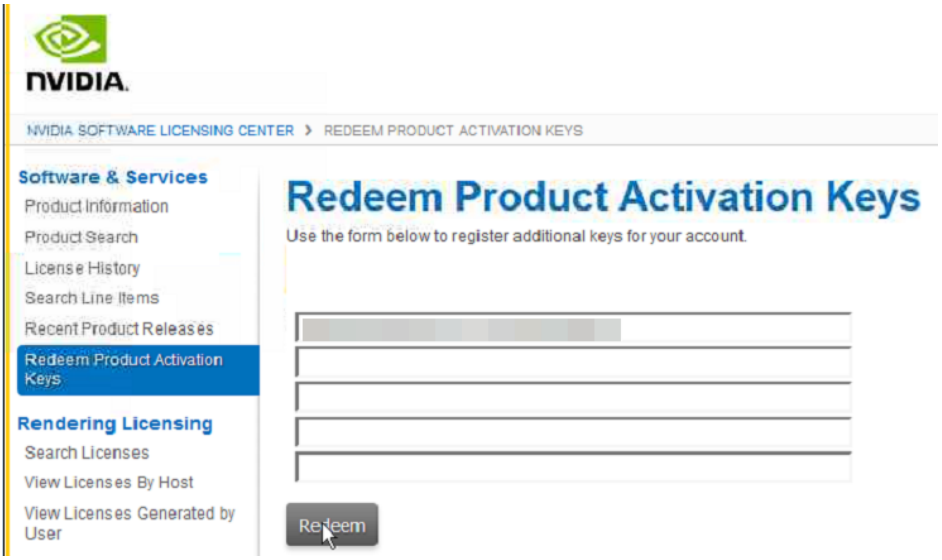

6. Specify **Username** and **Password** and click **LOGIN**.

Figure 1-78 Logging in to the official NVIDIA website

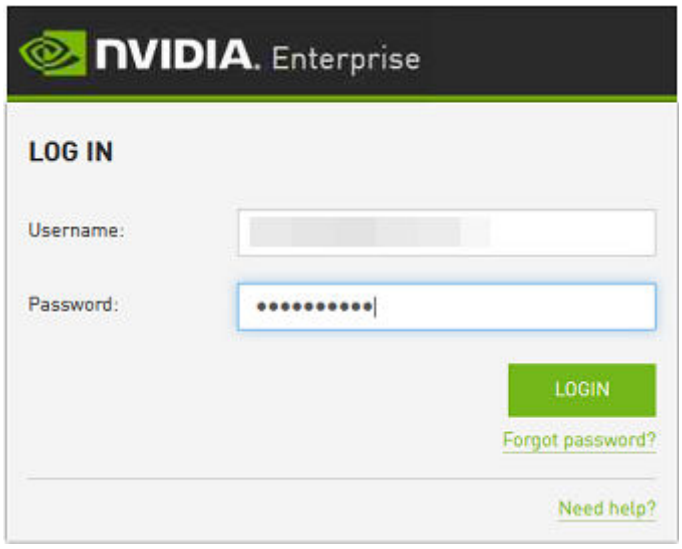

7. Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.

Figure 1-79 SOFTWARE DOWNLOADS page

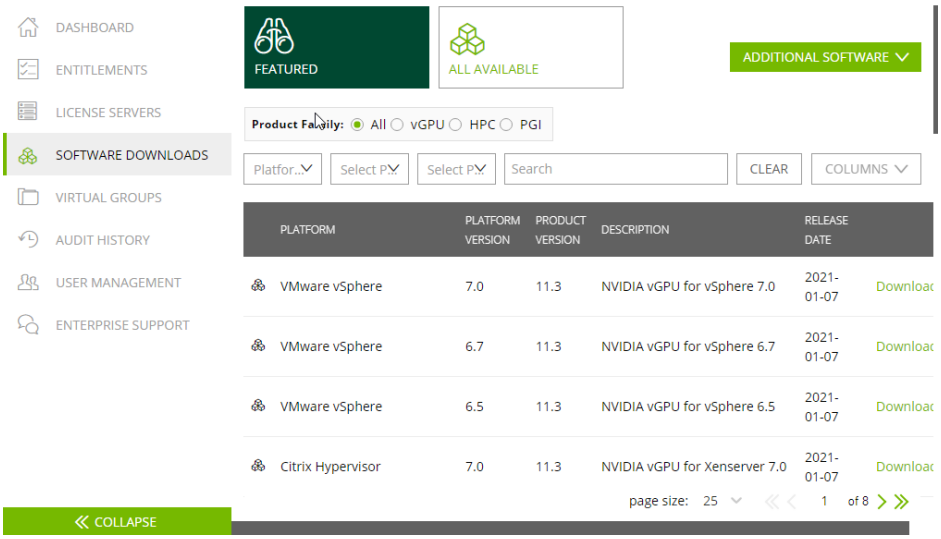

8. Download the GRID driver of the required version. For details, see **Table 1-26**.
9. Decompress the GRID driver installation package and install the driver that matches your ECS OS.
10. On the **SOFTWARE DOWNLOADS** page, click **ADDITIONAL SOFTWARE** to download the license software package.

Figure 1-80 ADDITIONAL SOFTWARE

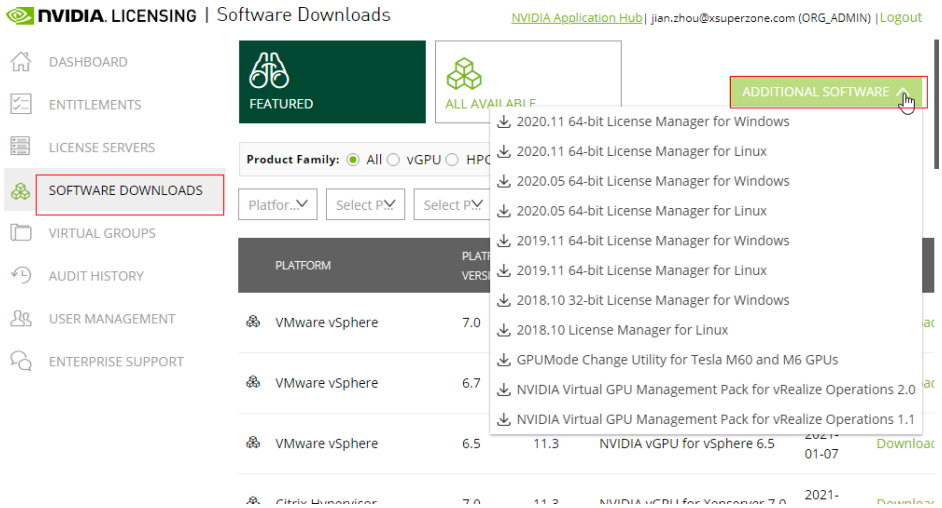

Deploying and Configuring the License Server

The following uses an ECS running CentOS 7.5 as an example to describe how to deploy and configure the license server on the ECS.

NOTE

- The target ECS must have at least 2 vCPUs and 4 GiB of memory.
- Ensure that the MAC address of the target ECS has been recorded.
- If the license server is used in the production environment, deploy it in high availability mode. For details, see **official NVIDIA documentation for license server high availability**.

1. Configure the network.
   - If the license server is to be accessed using the VPC, ensure that the license server and the GPU-accelerated ECS with the GRID driver installed are in the same VPC subnet.
   - If the license server is to be accessed using a public IP address, configure the security group to which license server belongs and add inbound rules for TCP 7070 and TCP 8080.
2. Install the license server.
   a. Run the following command to decompress the installation package. The **Installer.zip** in the command indicates the name of the software package obtained in **10**.

      unzip Installer.zip

   b. Run the following command to assign execution permissions to the installer:

      chmod +x setup.bin

   c. Run the installer as user **root**:

      sudo ./setup.bin -i console

   d. In the Introduction section, press **Enter** to continue.


```
===============================================================================
Introduction
------------

InstallAnywhere will guide you through the installation of License Server.

It is strongly recommended that you quit all programs before continuing with
this installation.

Respond to each prompt to proceed to the next step in the installation.  If
you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE:
```

e. In the License Agreement section, press **Enter** to turn to last pages and accept the license agreement.

   Enter **Y** and press **Enter**.

```
DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N): Y
```

f. In the Choose Install Folder section, press **Enter** to retain the default path for installing the License Server software.

g. In the Choose Local Tomcat Server Path section, enter the Tomcat's local path in the "/var/lib/*Tomcat version*" format, for example, /var/lib/tomcat8.

h. In the Choose Firewall Options section, confirm the port to be enabled in the firewall and press **Enter**.

```
Choose Firewall Options
----------------------

The license server listens on port 7070. This port must be opened in the
firewall for other machines to obtain licenses from this server.

The license server's management interface listens on port 8080. Leave this
port closed to prevent unauthorized access to the management interface.

  ->1- License server (port 7070)
    2- Management interface (port 8080)

ENTER A COMMA-SEPARATED LIST OF NUMBERS REPRESENTING THE DESIRED CHOICES, OR
   PRESS <ENTER> TO ACCEPT THE DEFAULT:
```

i. In the Pre-Installation Summary section, confirm the information and press **Enter** to start the installation.

```
Pre-Installation Summary
------------------------

Please Review the Following Before Continuing:

Product Name:
    License Server

Install Folder:
    /opt/flexnetls/nvidia

Link Folder:
    /root/NVIDIA Corporation/License Server

Disk Space Information (for Installation Target):
    Required:     105,216,774 Bytes
    Available: 35,501,248,512 Bytes

PRESS <ENTER> TO CONTINUE:
```

j. In the Install Complete section, press **Enter** to end the installation.

```
Install Complete
----------------

License Server has been successfully installed to:

   /opt/flexnetls/nvidia

PRESS <ENTER> TO EXIT THE INSTALLER:
```

3. Obtain the license file.
   a. Log in to the **NVIDIA website** on a new tab and select **LICENSE SERVERS**.

Figure 1-81 LICENSE SERVERS

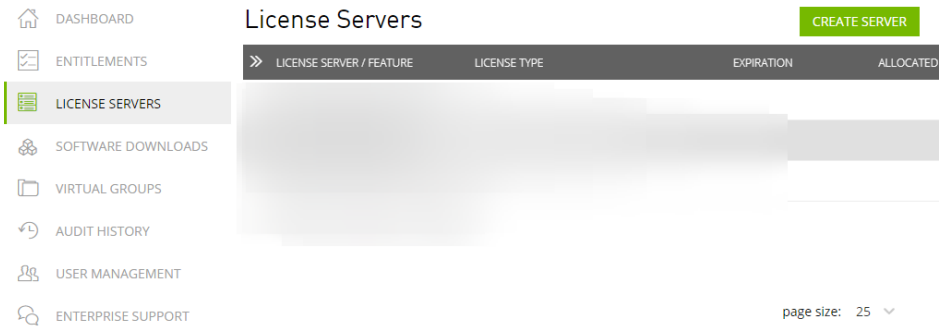

   b. Click **CREATE SERVER**.
   c. On the displayed **Create License Server** page, configure parameters.

Figure 1-82 Create License Server

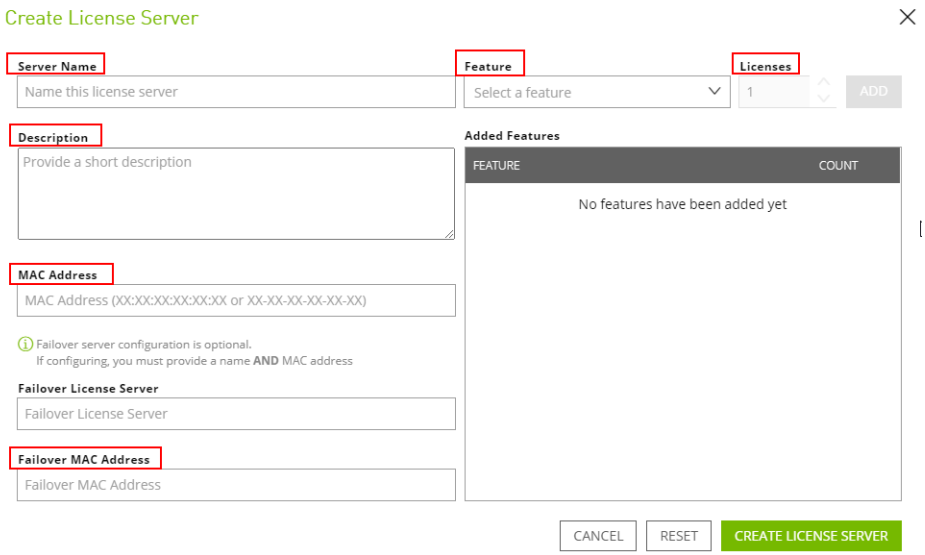

Table 1-27 Parameters for creating a license server

Parameter	Description
Server Name	License server name, which can be customized.
Description	License description information.
MAC Address	MAC address of the ECS where the license server is deployed. You can log in to the ECS and run **ipconfig -a** to query the MAC address.
Feature	Select a feature, enter the number of required licenses in the **Licenses** text box, and click **ADD**. In active/standby deployment, enter the name of the standby server in **Failover License Server** and enter the MAC address in **Failover MAC Address**.

d. Click **CREATE LICENSE SERVER**.

e. Download the license file.

Figure 1-83 Downloading the license file

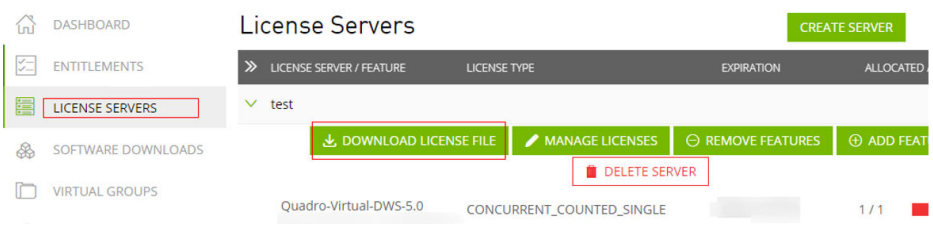

4. In the web browser, access the homepage of the license server management page using the link configured during the installation.

   Default URL: http://*IP address of the EIP*:8080/licserver

5. In the navigation pane on the left, click **License Server** > **License Management**.
6. Select the .bin license file to be uploaded and click **Upload**.

Figure 1-84 Uploading a license file

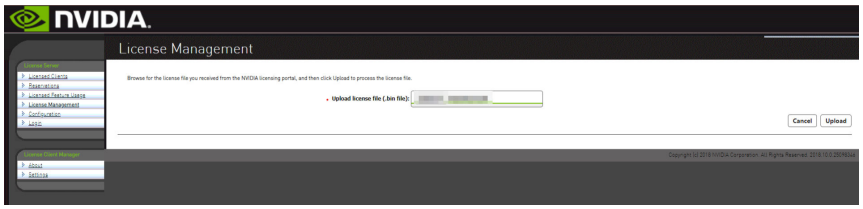

Installing the GRID Driver and Configuring the License

1. Install the GRID driver of a desired version, for example, on a GPU-accelerated Windows ECS.

   NOTE

   Microsoft remote login protocols do not support GPU 3D hardware acceleration. To use this function, install third-party desktop protocol-compliant software, such as VNC, PCoIP, or NICE DCV, and access the ECS through the client.

2. Open the NVIDIA control panel on the Windows control panel.
3. Enter the IP address and port number of the deployed license server in the level-1 license server, and then click **Apply**. If the message indicating that you have obtained a GRID license is displayed, the installation is successful. Additionally, the MAC address of the GPU-accelerated ECS with the GRID driver installed is displayed on the **Licensed Clients** page of the license server management console.

Figure 1-85 License server management console

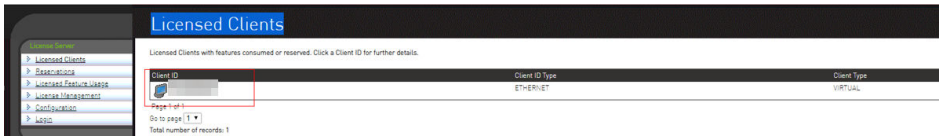

1.11.3 Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS

Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS for computing acceleration.

- A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
- After a GPU-accelerated ECS is created using a private image, it must have a Tesla driver installed. Otherwise, computing acceleration will not take effect.

This section describes how to install a Tesla driver and CUDA toolkit on a GPU-accelerated ECS.

Notes

- The target ECS has an EIP bound.
- The Tesla driver and CUDA toolkit have not been installed on the ECS.

NOTE

- Download the CUDA toolkit from the official NVIDIA website and install it. A Tesla driver matching the CUDA version will be automatically installed then. However, if there are specific requirements or dependencies on the Tesla driver version, download the matching Tesla driver from the official NVIDIA website first and then install the driver before installing the CUDA toolkit.
- If a Tesla driver has been installed on the ECS, check the driver version. Before installing a new driver version, uninstall the original Tesla driver to prevent an installation failure due to driver conflicts.

Installation process:

- **Obtaining a Tesla Driver and CUDA Toolkit**
- Installing a Tesla Driver
  - **Installing a Tesla Driver on a Linux ECS**
  - **Installing a Tesla Driver on a Windows ECS**
- Installing a CUDA Toolkit
  - **Installing the CUDA Toolkit on a Linux ECS**
  - **Installing the CUDA Toolkit on a Windows ECS**

Installing a Tesla Driver on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the Tesla driver matching CUDA 10.1 on a GPU-accelerated ECS.

NOTE

The Linux kernel version is compatible with the driver version. If installing the driver failed, check the driver installation log, which is generally stored in **/var/log/nvidia-installer.log**. If the log shows that the failure was caused by a driver compilation error, for example, the **get_user_pages** parameter setting is incorrect, the kernel version is incompatible with the driver version. In such a case, select the desired kernel version and driver version and reinstall them. It is recommended that the release time of the kernel version and driver version be the same.

1. Log in to the ECS.
2. Update the system software based on the OS.
   - Ubuntu

     Update the software installation source: **apt-get -y update**

     Install necessary programs: **apt-get install gcc g++ make**
   - CentOS

     Update the software installation source: **yum -y update --exclude=kernel* --exclude=centos-release* --exclude=initscripts***

     Install the desired program: **yum install -y kernel-devel-`uname -r` gcc gcc-c++**

3. Download the NVIDIA driver package.

   Select a driver version at **NVIDIA Driver Downloads** based on the ECS type. Click **SEARCH**.

   Figure 1-86 Selecting a NVIDIA driver version

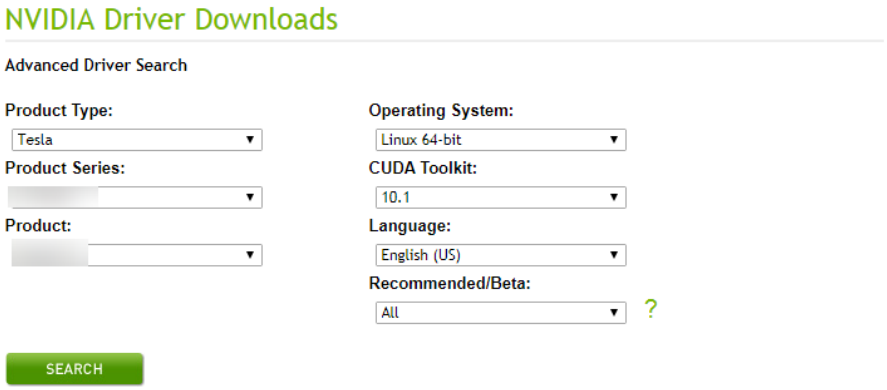

4. Select a driver version as required. The following uses Tesla 418.67 as an example.

   Figure 1-87 Selecting a driver version

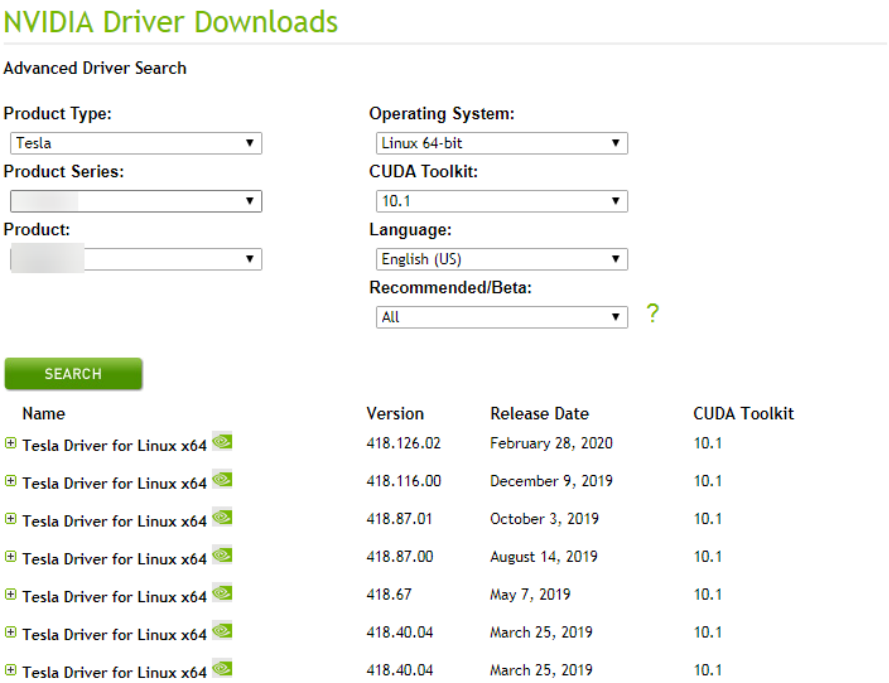

5. Click the driver to be downloaded. On the **TESLA DRIVER FOR LINUX X64** page that is displayed, click **DOWNLOAD**.
6. Copy the download link.

Figure 1-88 Copying the download link

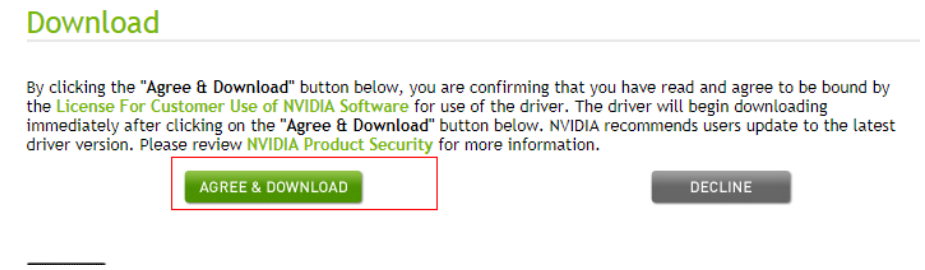

7. Run the following command on the ECS to download the driver:

   wget *Copied link*

   For example, **wget http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run**

Figure 1-89 Obtaining the installation package

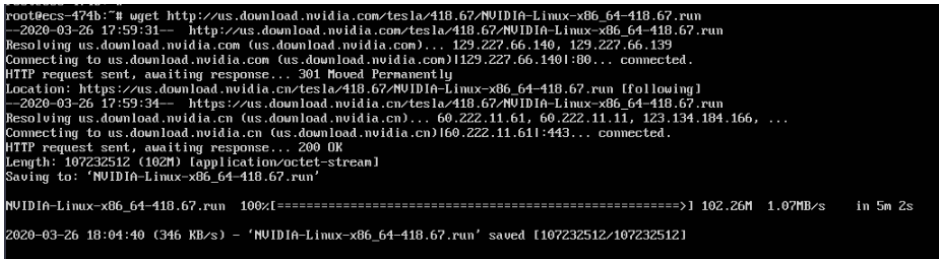

8. Run the following command to install the driver:

   sh NVIDIA-Linux-x86_64-418.67.run

9. (Optional) If the following information is displayed after the command for installing the driver is executed, disable the Nouveau driver.

Figure 1-90 Disabling the Nouveau driver

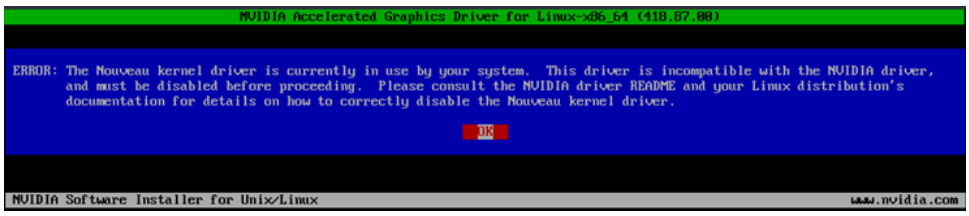

   a. Run the following command to check whether the Nouveau driver has been installed:

      lsmod | grep nouveau

      - If the command output contains information about the Nouveau driver, the Nouveau driver has been installed and must be disabled. Then, go to step **9.b**.
      - If the command output does not contain information about the Nouveau driver, the Nouveau driver has been disabled. Then, go to step **10**.

   b. Edit the **blacklist.conf** file.

      If the **/etc/modprobe.d/blacklist.conf** file is unavailable, create it.

vi /etc/modprobe.d/blacklist.conf

Add the following statement to the end of the file:

```
blacklist nouveau
options nouveau modeset=0
```

c. Run the following command to back up and create an initramfs application:

- Ubuntu

  sudo update-initramfs -u

- CentOS:

  mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).img.bak

  dracut -v /boot/initramfs-$(uname -r).img $(uname -r)

d. Restart the ECS:

reboot

10. Select **OK** for three consecutive times as prompted to complete the driver installation.

Figure 1-91 Completing the NVIDIA driver installation

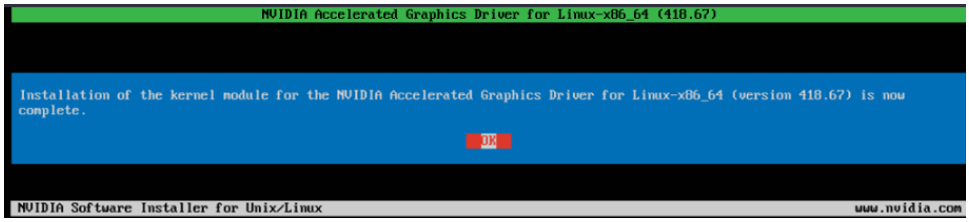

11. Run the following command to set systemd:

    systemctl set-default multi-user.target

12. Run the **reboot** command to restart the ECS.
13. Log in to the ECS and run the **nvidia-smi** command. If the command output contains the installed driver version, the driver has been installed.

Figure 1-92 Viewing the NVIDIA driver version

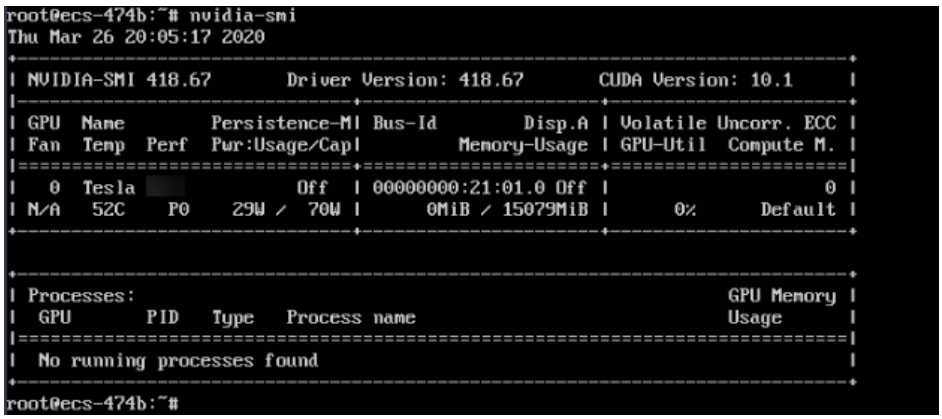

Installing a Tesla Driver on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install a Tesla driver on a GPU-accelerated ECS.

1. Log in to the ECS.
2. Download the NVIDIA driver package.

   Select a driver version at **NVIDIA Driver Downloads** based on the ECS type.

   Figure 1-93 Selecting a driver type (Windows)

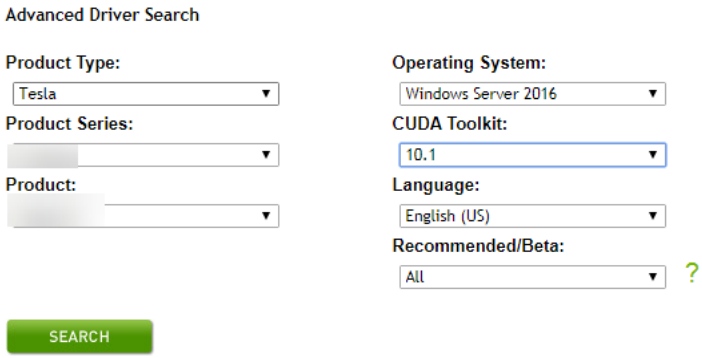

3. Select a driver version as required. The following uses Tesla 425.25 as an example.

   Figure 1-94 Selecting a driver version (Windows)

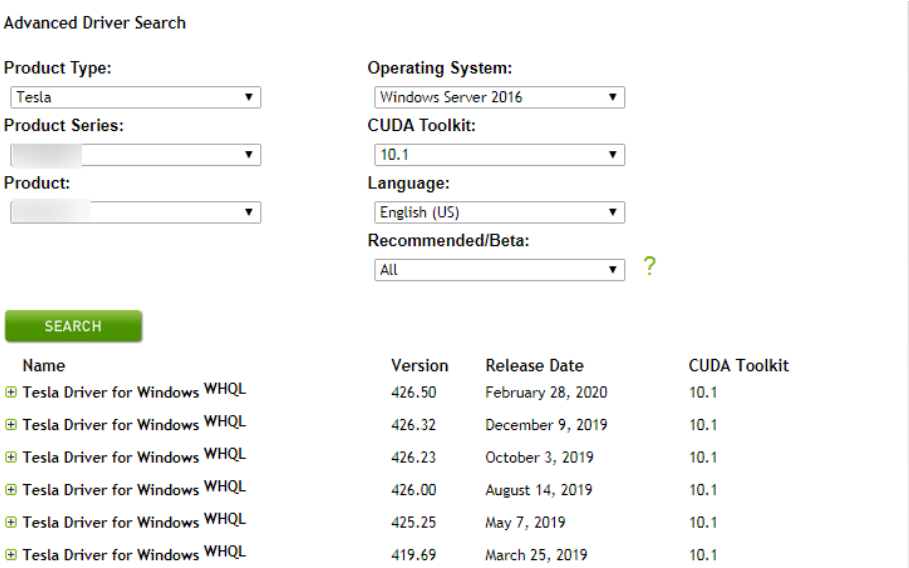

4. Click the driver to be downloaded. On the **TESLA DRIVER FOR WINDOWS** page that is displayed, click **DOWNLOAD**.
5. Click **Agree & Download** to download the installation package.

Figure 1-95 Downloading the driver installation package

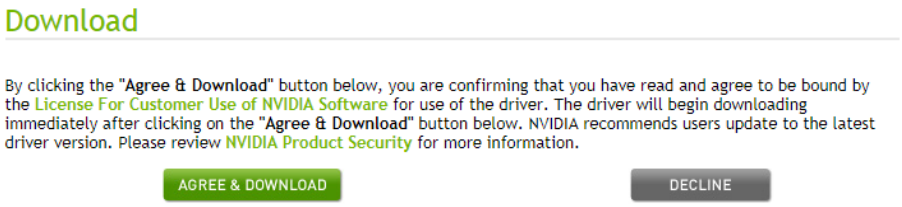

6. Double-click the driver and click **Run**.

Figure 1-96 Running the NVIDIA driver installation program

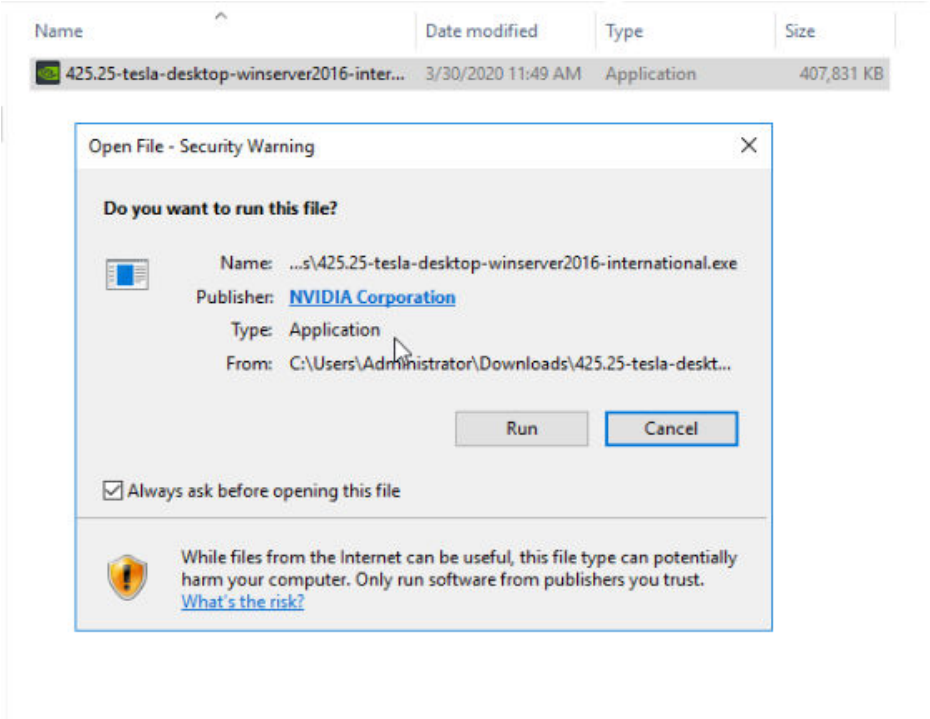

7. Select an installation path and click **OK**.

Figure 1-97 Selecting an installation path

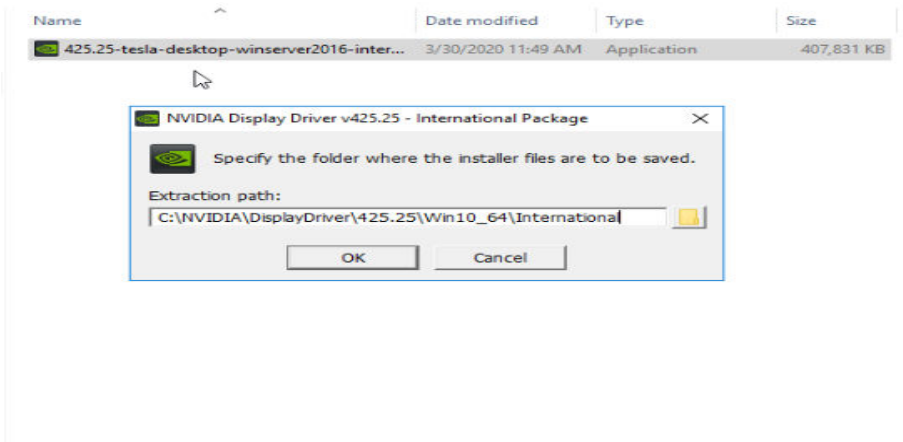

8. Install the NVIDIA program as prompted.

Figure 1-98 Completing the driver installation

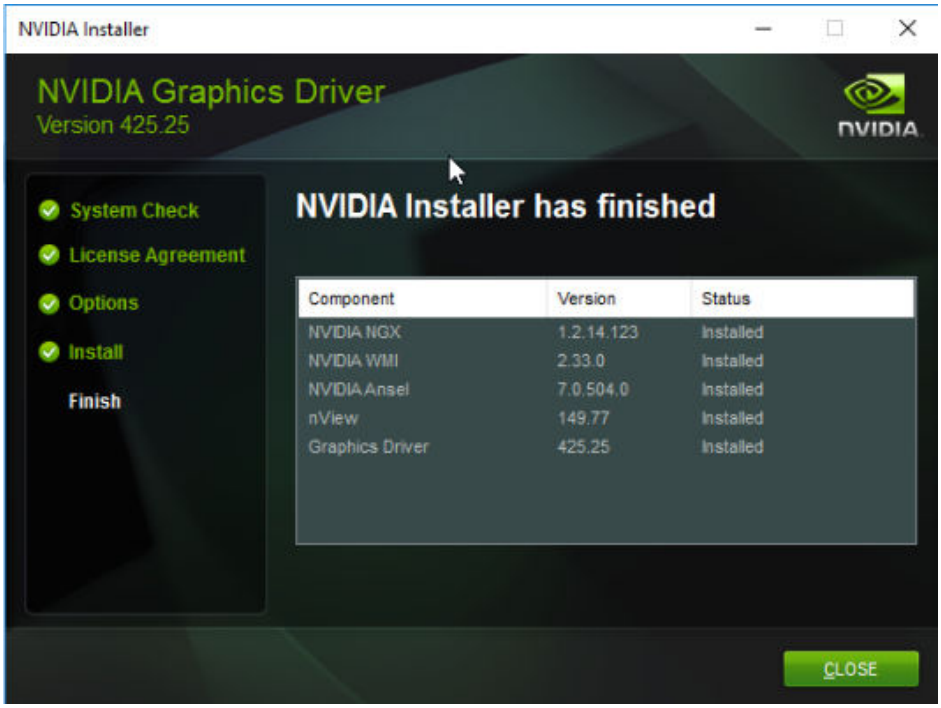

9. Restart the ECS.
10. Check whether the NVIDIA driver has been installed.
    a. Switch to **Device Manager** and click **Display adapters**.

Figure 1-99 Display adapters

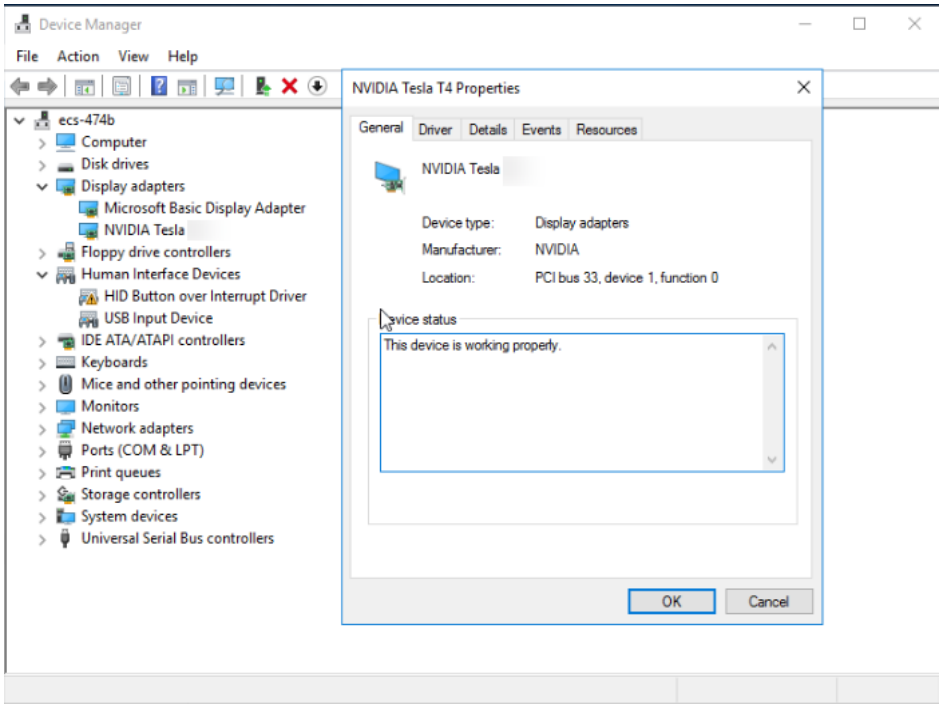

    b. Open the **cmd** window on the ECS and run the following commands:
    cd C:\Program Files\NVIDIA Corporation\NVSMI

nvidia-smi

If the command output contains the installed driver version, the driver has been installed.

Figure 1-100 Viewing the NVIDIA driver version

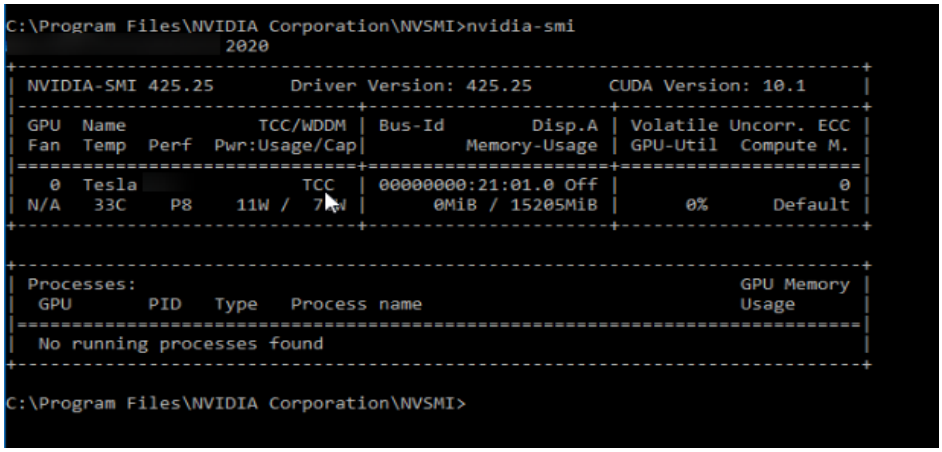

Installing the CUDA Toolkit on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.

1. Log in to the ECS.
2. Update the system software based on the OS.
   - Ubuntu

     Update the software installation source: **apt-get -y update**

     Install necessary programs: **apt-get install gcc g++ make**
   - CentOS

     Update the software installation source: **yum -y update --exclude=kernel* --exclude=centos-release* --exclude=initscripts***

     Install the desired program: **yum install -y kernel-devel-`uname -r` gcc gcc-c++**
3. On the CUDA download page, set parameters according to the information shown in **Obtaining a Tesla Driver and CUDA Toolkit**.

Figure 1-101 Selecting a CUDA version

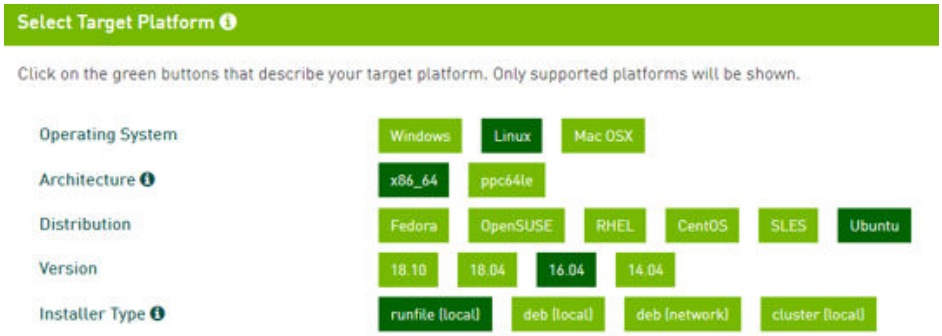

4. Find the link for downloading CUDA 10.1 and copy the link.

Figure 1-102 Copying the link for downloading CUDA

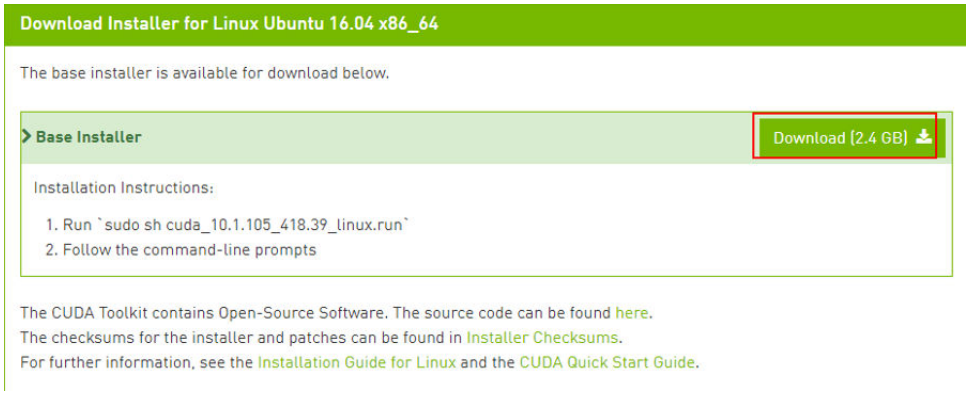

5. Run the following command on the ECS to download CUDA:

   wget *Copied link*

   For example, **wget https://developer.nvidia.com/compute/cuda/10.1/Prod/local_installers/cuda_10.1.105_418.39_linux.run**

Figure 1-103 Downloading CUDA

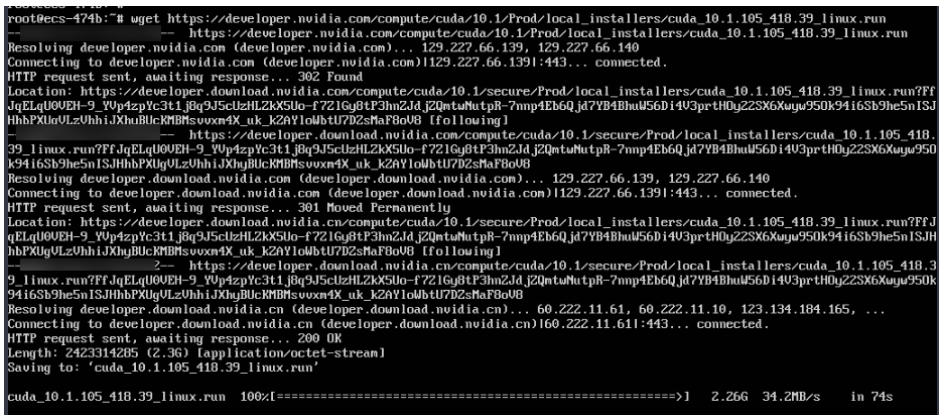

6. Install CUDA.

   Follow the instructions provided on the official NVIDIA website.

Figure 1-104 Installing CUDA

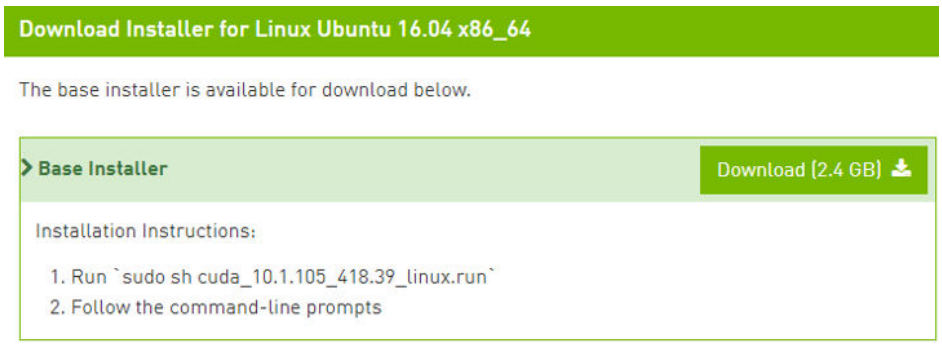

7. Run the following command to install CUDA:

   sh cuda_10.1.243_418.87.00_linux.run

8. Select **accept** on the installation page and press **Enter**.

Figure 1-105 Installing CUDA_1

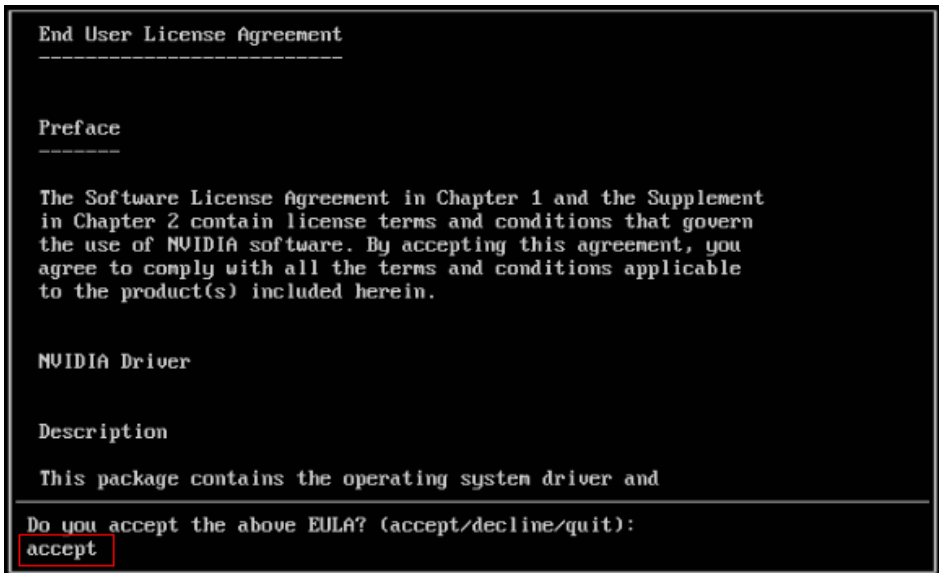

9. Select **Install** and press **Enter** to start the installation.

Figure 1-106 Installing CUDA_2

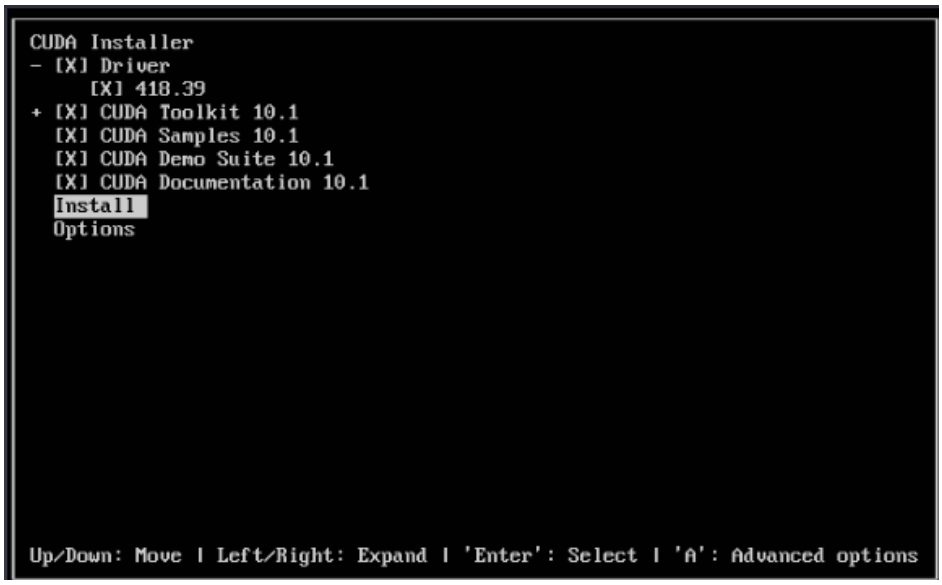

Figure 1-107 Completing the installation

10. Run the following command to switch to **/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery**:

    cd /usr/local/cuda-10.1/samples/1_Utilities/deviceQuery

11. Run the **make** command to automatically compile the deviceQuery program.
12. Run the following command to check whether CUDA has been installed:

    ./deviceQuery

    If the command output contains the CUDA version, CUDA has been installed.

    Figure 1-108 deviceQuery common output

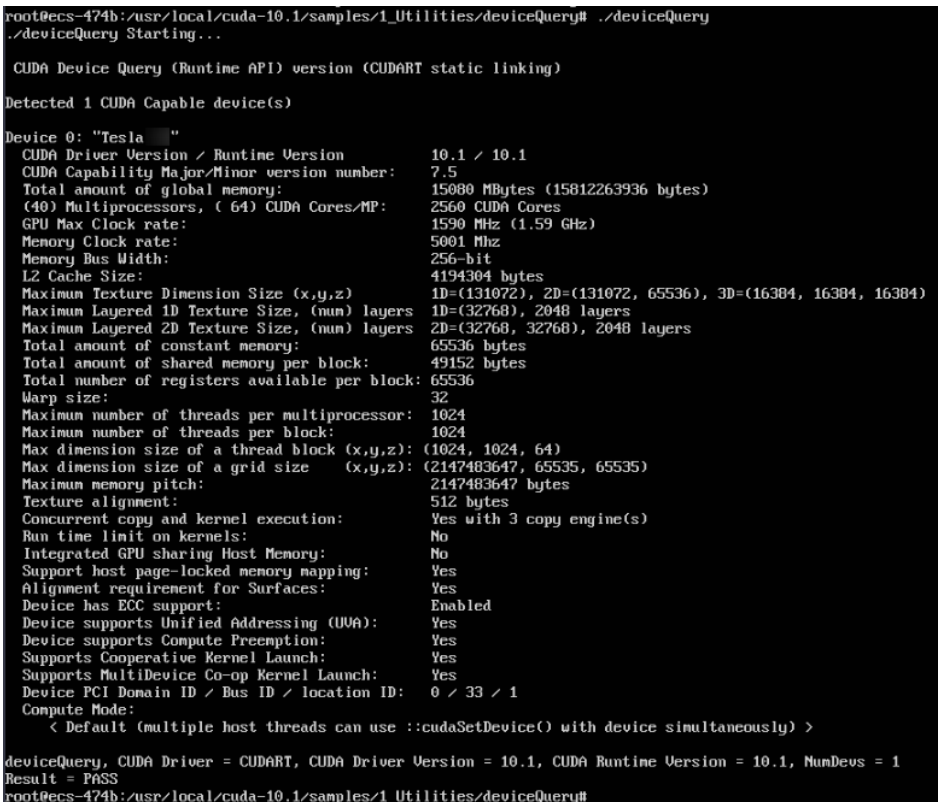

13. Check the CUDA version.

    /usr/local/cuda/bin/nvcc -V

    Figure 1-109 Checking the CUDA version

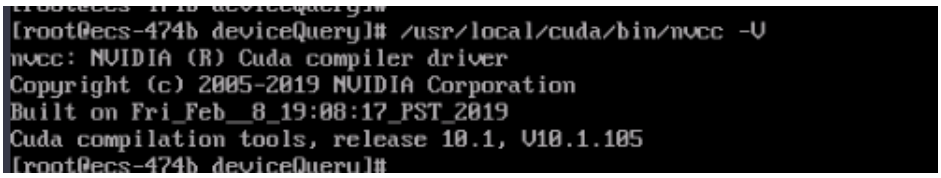

14. Run the following command to enable the persistent mode:

    sudo nvidia-smi -pm 1

    Enabling the persistent mode optimizes the GPU performance on Linux ECSs.

Installing the CUDA Toolkit on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.

1. Log in to the ECS.
2. On the CUDA download page, set parameters according to the information shown in **Downloading a CUDA Toolkit**.

Figure 1-110 Selecting a CUDA version

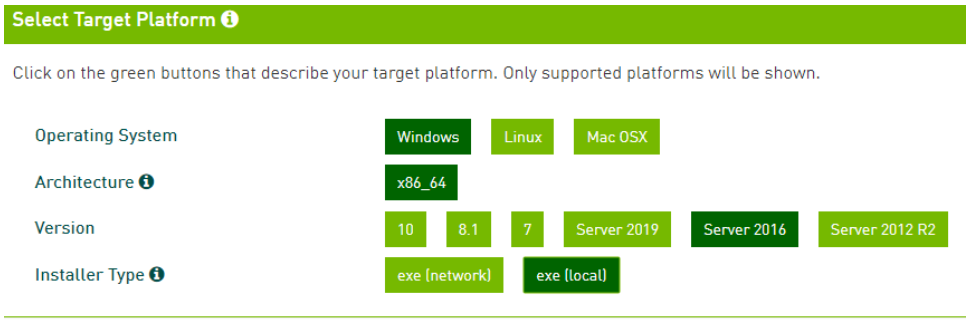

3. Find the link for downloading CUDA 10.1.

Figure 1-111 Finding the link for downloading CUDA

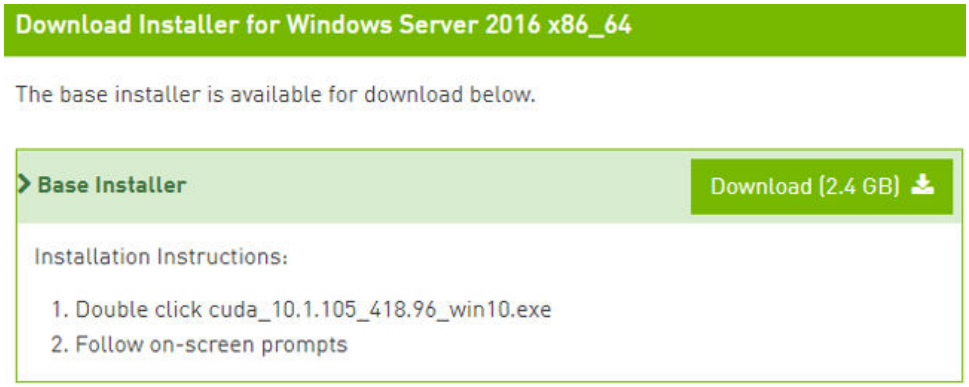

4. Click **Download** to download the CUDA toolkit.
5. Double-click the installation file and click **Run** to install the CUDA toolkit.

Figure 1-112 Installing CUDA

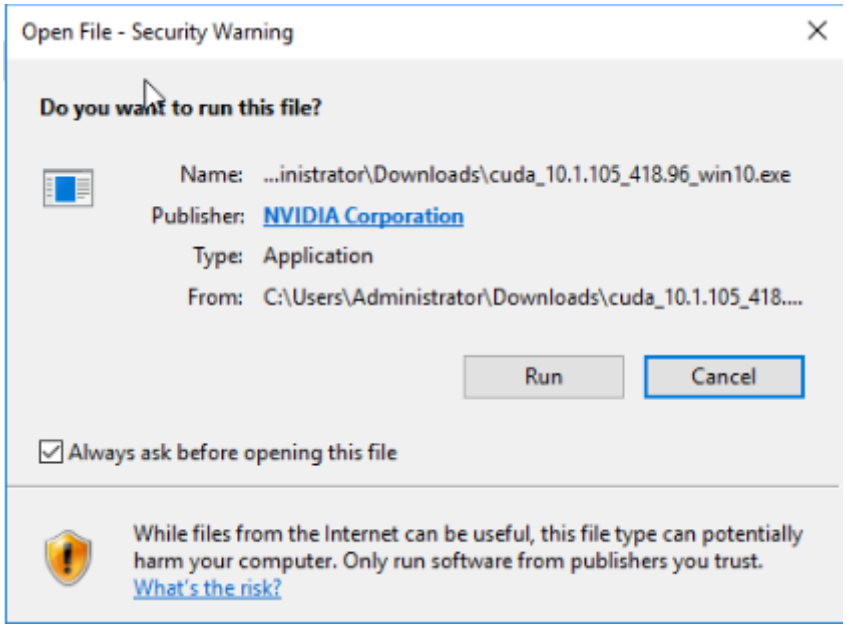

6. On the **CUDA Setup Package** page, select an installation path and click **OK**.

Figure 1-113 Selecting an installation path

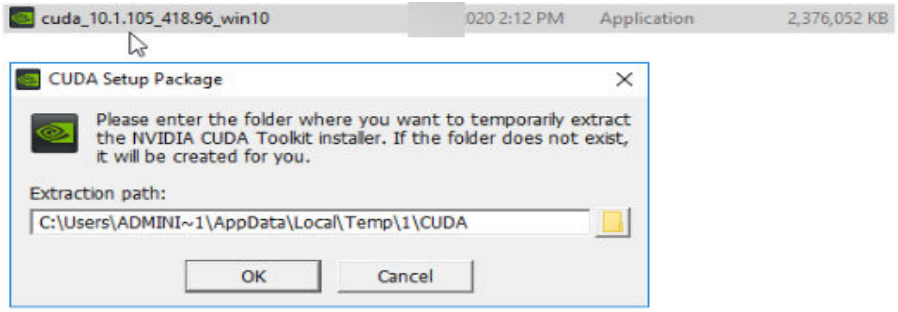

7. Install the CUDA toolkit as prompted.

Figure 1-114 Completing the installation

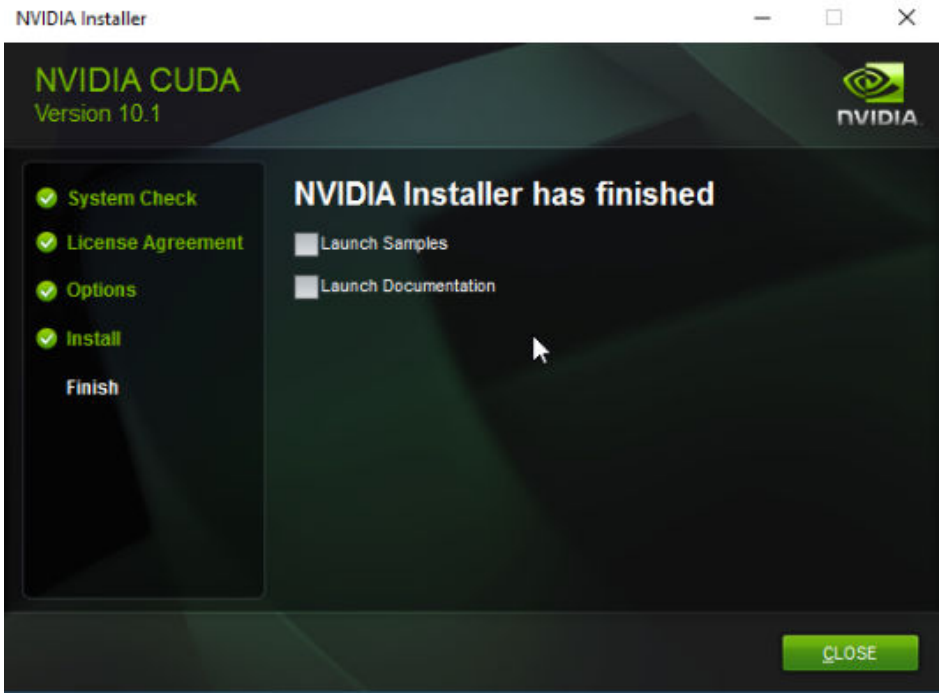

8. Check whether CUDA has been installed

   Open the **cmd** window and run the following command:

   nvcc -V

   If the command output contains the CUDA version, CUDA has been installed.

Figure 1-115 Successful installation

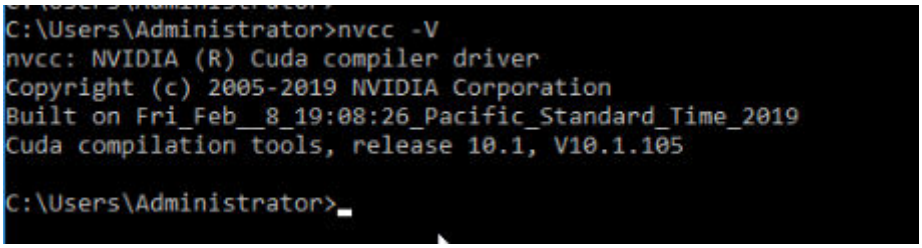

1.11.4 Obtaining a Tesla Driver and CUDA Toolkit

Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS. Otherwise, computing acceleration will not take effect. This section describes how to obtain a Tesla driver and CUDA toolkit. Select a driver version based on your ECS type.

For instructions about how to install the Tesla driver and CUDA toolkit, see **Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS**.

Downloading a Tesla Driver

Download a driver based on your ECS type.

Table 1-28 Mapping between Tesla drivers and ECS types

ECS Type	Driver	Product Series	Product
P2s	Tesla	V	V100
P2v	Tesla	V	V100
Pi2	Tesla	T	T4
PI1	Tesla	P	P4

Downloading a CUDA Toolkit

Download the **CUDA software package** and select the corresponding CUDA Toolkit software package based on the instance type and driver version.

NOTE

NVIDIA Driver Downloads provides the mapping between the driver version and CUDA Toolkit. If the versions do not match, the driver may be unavailable.

The following uses Tesla T4 as an example to describe how to download the driver package and CUDA Toolkit.

1. Select the Linux operating system and the CUDA Toolkit 11.6 version.

   Figure 1-116 Selecting the CUDA Toolkit version

   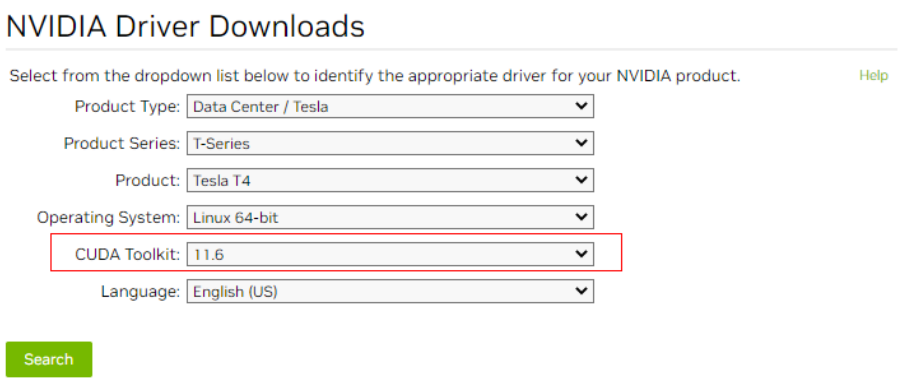

2. Select a CUDA Toolkit 11.6 package to download.

Figure 1-117 Downloading a CUDA Toolkit 11.6 package

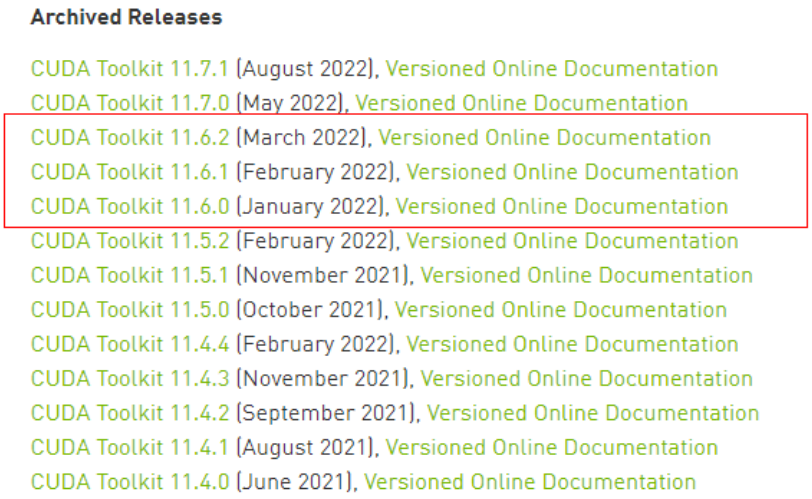

1.11.5 Uninstalling a GPU Driver from a GPU-accelerated ECS

Scenarios

You can manually uninstall the GPU driver from a GPU-accelerated ECS.

This section describes how to uninstall a GPU driver from a Windows ECS and a Linux ECS.

- **Uninstalling a GPU Driver from a Windows ECS**
- **Uninstalling a GPU Driver from a Linux ECS**

Uninstalling a GPU Driver from a Windows ECS

This section uses Windows Server 2016 Datacenter Edition 64-bit as an example to describe how to uninstall the NVIDIA driver (driver version: 462.31) from a GPU-accelerated ECS.

1. Log in to the ECS.
2. Click **Start** in the task bar and choose **Control Panel**.
3. In Control Panel, click **Uninstall a program** under **Programs**.

Figure 1-118 Uninstalling a program.

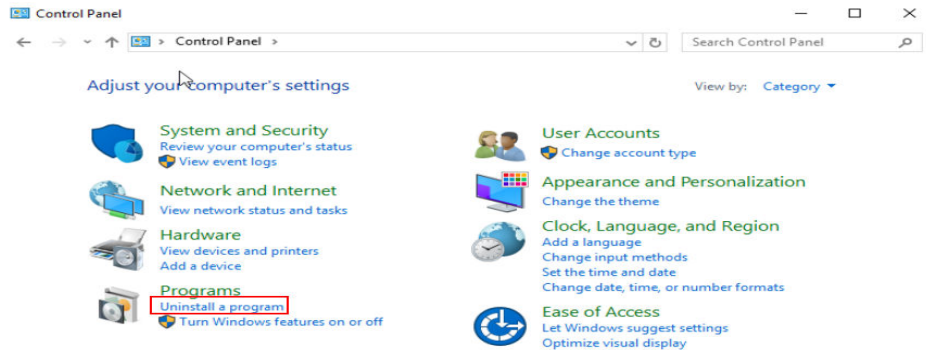

4. Right-click the NVIDIA driver to be uninstalled and choose **Uninstall/Change** from the shortcut menu.

Figure 1-119 Uninstalling a NVIDIA driver

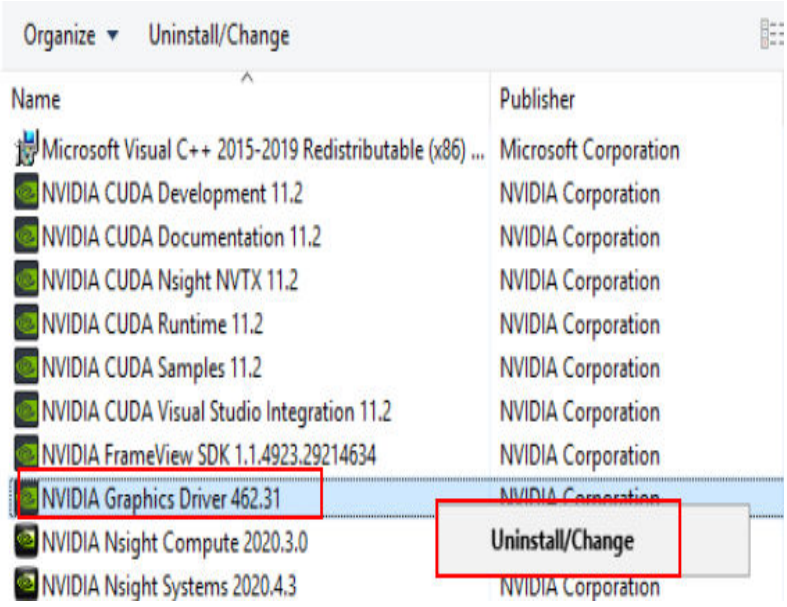

5. In the displayed **NVIDIA Uninstaller** window, click **UNINSTALL**.

Figure 1-120 Confirming the uninstallation

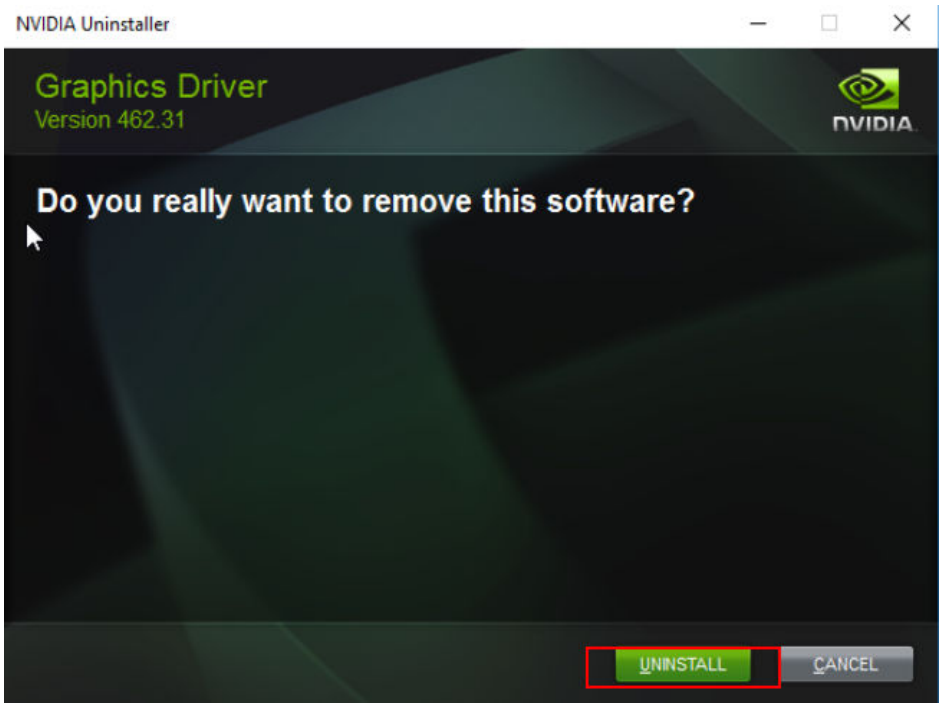

6. After the uninstallation is complete, click **RESTART LATER**.
7. Check whether the NVIDIA driver has been uninstalled.
   a. In Control Panel, click **Device Manager**.

      If no NVIDIA graphics cards are not displayed under **Display adapters**, the driver is uninstalled successfully.

      Figure 1-121 Viewing Display adapters

      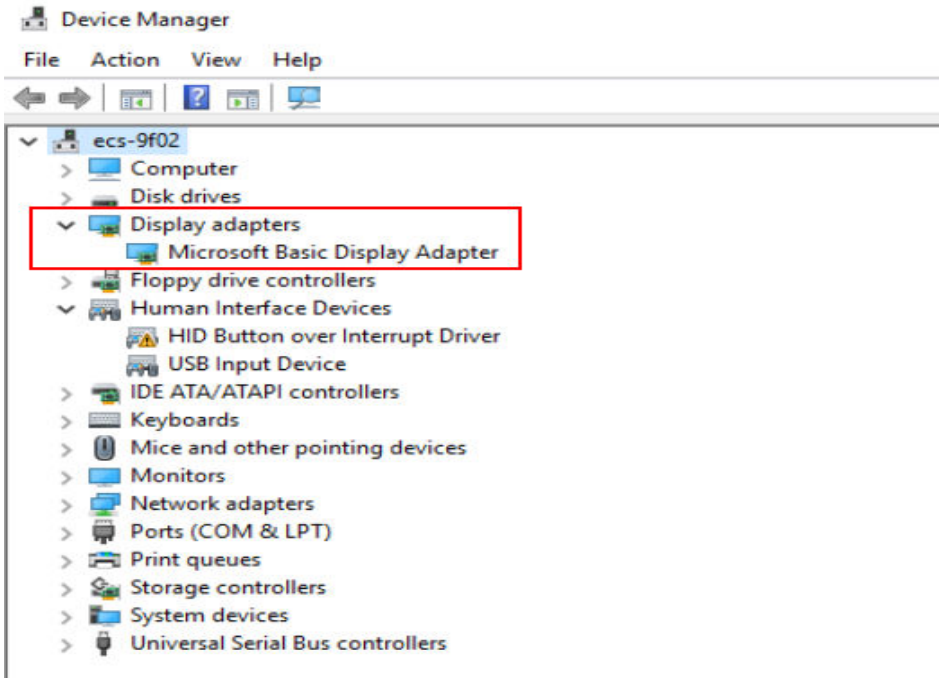

   b. Open the cmd window of the ECS and run the following commands:

cd C:\Program Files\NVIDIA Corporation\NVSMI

nvidia-smi.exe

Figure 1-122 Command output

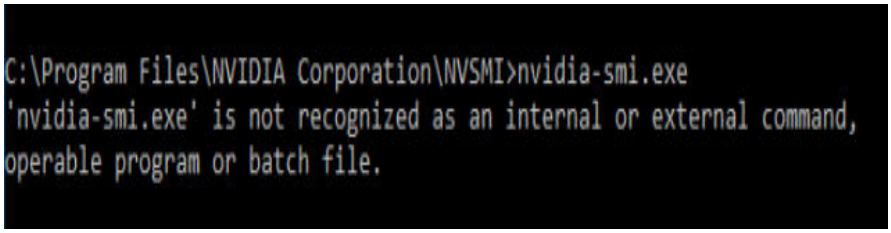

If the command output indicates that the file does not exist, the driver is uninstalled successfully.

After the NVDIA driver is uninstalled, you can install a new NVIDIA driver without restarting the ECS.

Uninstalling a GPU Driver from a Linux ECS

For NVIDIA Tesla drivers installed using .run Packages, you are advised to perform the following steps to uninstall it.

NOTE

If you use .run Packages to install the NVIDIA Grid driver, you only need to perform **step 1** to uninstall the NVIDIA driver.

The following uses 64-bit Ubuntu Server 20.04 as an example to describe how to uninstall Tesla 460.73.01 and CUDA 11.2.

1. Uninstall the NVIDIA driver.
   a. Query the path where **nvidia-uninstall** is stored.

      whereis nvidia-uninstall

      Generally, **nvidia-uninstall** is stored in the **/usr/bin/** directory.

      Figure 1-123 Querying the nvidia-uninstall path

      

   b. Uninstall the driver from the path where **nvidia-uninstall** is stored.

      /usr/bin/nvidia-uninstall

   c. Select **Yes** and press **Enter**.

Figure 1-124 NVIDIA driver uninstallation (1)

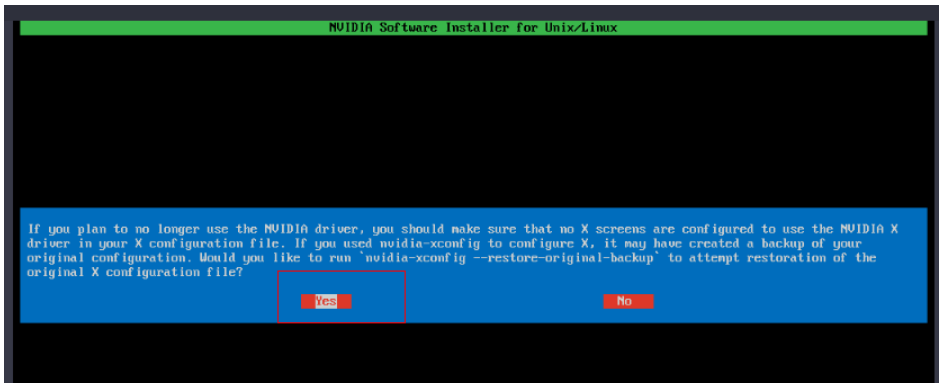

d. Select **OK** and press **Enter**.

Figure 1-125 NVIDIA driver uninstallation (2)

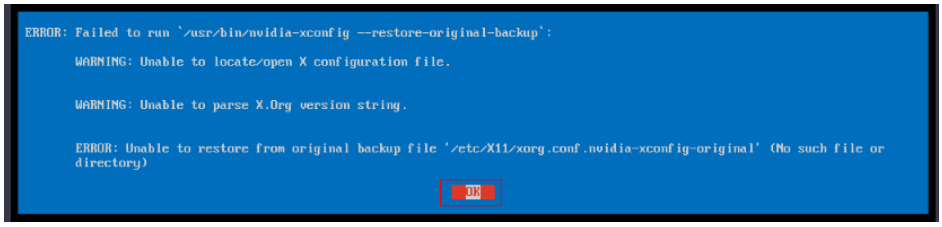

e. After the driver is uninstalled, press **Enter**.

Figure 1-126 NVIDIA driver uninstallation (3)

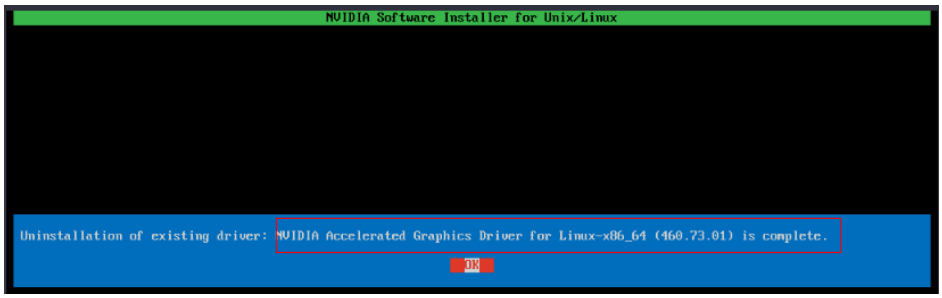

2. Uninstall the CUDA and CUDA Deep Neural Network (cuDNN) libraries.

   To upgrade the CUDA driver version, uninstall the corresponding CUDA library and then install a new one with the target version.

   a. Uninstall the CUDA library.

      /usr/local/cuda/bin/cuda-uninstaller

      Generally, **cuda-uninstaller** is stored in the **/usr/local/cuda/bin** directory.

      NOTE

      The uninstallation command varies depending on CUDA versions. If the **cuda-uninstaller** file is not found, check whether a file starting with **uninstall_cuda** exists in the **/usr/local/cuda/bin/** directory.

      If such a file exists, replace **cuda-uninstaller** in the preceding command with the file name.

   b. On the uninstallation page, select all options, move the cursor to **Done**, and press **Enter**.

Figure 1-127 Uninstalling a CUDA driver

If the CUDA library is uninstalled, the message "Successfully uninstalled" is displayed.

c. Remove the CUDA and cuDNN libraries.

rm –rf /usr/local/cuda-11.2

2 Images

2.1 Overview

Image

An image is an ECS or BMS template that contains an OS or service data. It may also contain proprietary software and application software, such as database software. Images are classified into public, private, and shared images.

Image Management Service (IMS) allows you to easily create and manage images. You can create an ECS using a public image, private image, or shared image. You can also use an existing ECS or external image file to create a private image.

Public Image

A public image is a standard, widely used image that contains a common OS, such as Ubuntu, CentOS, or Debian, and preinstalled public applications. This image is available to all users. Select your desired public image. Alternatively, create a private image based on a public image to copy an existing ECS or rapidly create ECSs in a batch. You can customize a public image by configuring the application environment or software.

For more information about public images, see **Overview**.

Private Image

A private image contains an OS or service data, preinstalled public applications, and private applications. It is available only to the user who created it.

Table 2-1 Private image types

Image Type	Description
System disk image	Contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud.
Data disk image	Contains only service data. You can use a data disk image to create EVS disks and migrate your service data to the cloud.
Full-ECS image	Contains an OS, application software, and data for running services. A full-ECS image contains the system disk and all data disks attached to it.
ISO image	Created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.

If you plan to use a private image to change the OS, ensure that the private image is available. For instructions about how to create a private image, see **Image Management Service User Guide**.

- If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
- If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
- If a private image from another region is required, make sure that the image has been copied.
- If a private image from another user account is required, make sure that the image has been shared with you.

Shared Image

A shared image is a private image shared by another user and can be used as your own private image. For details, see **Sharing Images**.

- Only the private images that have not been published in Marketplace can be shared.
- Images can be shared within a region only.
- Each image can be shared to a maximum of 128 tenants.
- You can stop sharing images anytime without notifying the recipient.
- You can delete shared image anytime without notifying the recipient.
- Encrypted images cannot be shared.
- Only the full-ECS images created using CBR can be shared.

Marketplace Image

A Marketplace image is a third-party image that has an OS, application environment, and software preinstalled. You can use the images to deploy websites and application development environments with a few clicks. No additional configuration is required.

A Marketplace image can be free of charge or paid, based on image service providers. When you use a paid image to create an ECS, you need to pay for the Marketplace image and ECS.

Helpful Links

- **Creating a Private Image**
- **Image Source Management**

2.2 Creating an Image

Scenarios

You can use an existing ECS to create a system disk image, data disk image, and full-ECS image.

- System disk image: contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud.
- Data disk image: contains only service data. You can create a data disk image from an ECS data disk. You can also use a data disk image to create EVS disks and migrate your service data to the cloud.
- Full-ECS image: contains all the data of an ECS, including the data on the data disks attached to the ECS. A full-ECS image can be used to rapidly create ECSs with service data.
- ISO image: is created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.

You can use a private image to change the OS. For instructions about how to create a private image, see **Image Management Service User Guide**.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Create Image** in the **Operation** column.
5. Configure the following information:

   Table 2-2 and **Table 2-3** list the parameters in the **Image Type and Source** and **Image Information** areas, respectively.

Table 2-2 Image type and source

Parameter	Description
Type	Select **Create Image**.

Parameter	Description
Region	Select a region close to your business. If you select an incorrect region here, you can replicate the created image to your desired region. For details, see **Replicating Images Across Regions**.
Image Type	Select **System disk image**.
Source	Click the **ECS** tab and select an ECS with required configurations.

Table 2-3 Image information

Parameter	Description
Encryption	This parameter specifies whether the image will be encrypted. The value is provided by the system and cannot be changed. • Only an unencrypted private image can be created from an unencrypted ECS. • Only an encrypted private image can be created from an encrypted ECS.
Name	Set a name for the image.
Enterprise Project	Select an enterprise project from the drop-down list. This parameter is available only if you have enabled enterprise projects or your account is an enterprise account. To enable this function, contact your customer manager. An enterprise project provides central management of cloud resources on a project.
Tag	(Optional) Set a tag key and a tag value for the image to make identification and management of your images easier.
Description	(Optional) Enter a description of the image.

6. Click **Next** and and submit the request.

3 EVS Disks

3.1 Overview

What Is Elastic Volume Service?

Elastic Volume Service (EVS) offers scalable block storage for ECSs. With high reliability, high performance, and rich specifications, EVS disks can be used for distributed file systems, development and test environments, data warehouses, and high-performance computing (HPC) scenarios to meet diverse service requirements.

Disk Types

EVS disk types differ in performance. Choose a disk type based on your requirements.

For more information about EVS disk specifications and performance, see **Elastic Volume Service Service Overview**.

Helpful Links

- **Attaching an EVS Disk to an ECS**
- **Introduction to Data Disk Initialization Scenarios and Partition Styles**
- **Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?**
- **How Can I Adjust System Disk Partitions?**
- **Can I Attach Multiple Disks to an ECS?**
- **What Are the Requirements for Attaching an EVS Disk to an ECS?**

3.2 Adding a Disk to an ECS

Scenarios

The disks attached to an ECS include one system disk and one or more data disks. The system disk is automatically created and attached when the ECS is created.

You do not need to purchase it again. The data disks can be added in either of the following ways:

- If you add data disks when purchasing an ECS, the data disks will be automatically attached to the ECS.
- If you purchase data disks after an ECS is created, the data disks need to be manually attached to the ECS.

This section describes how to add a data disk after purchasing an ECS.

Procedure

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. Locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Add Disk** in the **Operation** column.

   The page for adding a disk is displayed.
5. Set parameters for the new EVS disk as prompted.

   For instructions about how to set EVS disk parameters, see **Purchasing an EVS Disk**.

   NOTE

   - By default, the billing mode of the new disk is the same as that of the ECS.
   - By default, the new disk is in the same region as the ECS.
   - By default, the new disk is in the same AZ as the ECS, and the AZ of the disk cannot be changed.
   - After the new disk is created, it is attached to the ECS by default.
   - The expiration time of a new disk billed on a yearly/monthly basis is the same as that of the ECS.
6. Click **Next** to confirm the order and click **Submit** to complete the payment.

   The system automatically switches back to the **Disks** tab on the ECS management console. Then, you can view the information of the new disk.

Follow-up Procedure

The system automatically attaches the new disk to the ECS, but the disk can be used only after it is initialized. To do so, log in to the ECS and initialize the disk.

For details about how to initialize a data disk, see **Initializing an EVS Data Disk**.

3.3 Attaching an EVS Disk to an ECS

Scenarios

If the existing disks of an ECS fail to meet service requirements, for example, due to insufficient disk space or poor disk performance, you can attach more available EVS disks to the ECS, or purchase more disks (in **Storage** > **Elastic Volume Service**) and attach them to the ECS.

Prerequisites

- EVS disks are available.

  For instructions about how to purchase an EVS disk, see **Purchasing an EVS Disk**.

Procedure

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. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS.

   The page providing details about the ECS is displayed.
6. Click the **Disks** tab. Then, click **Attach Disk**.

   The **Attach Disk** dialog box is displayed.

   Figure 3-1 Attach Disk (KVM)

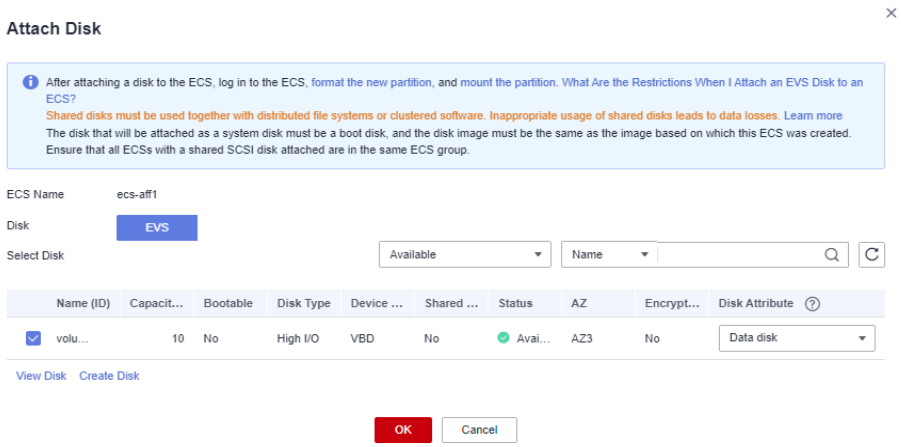

7. Select the target disk and specify the disk as the system disk or data disk
   - For KVM ECSs, you can specify a disk as a system disk or data disk but cannot specify a device name for the disk.
   - For Xen ECSs, you can specify the device name of a disk, such as **/dev/vdb**.

   NOTE

   - If no EVS disks are available, click **Create Disk** in the lower part of the list.
   - For the restrictions on attaching disks, see **What Are the Requirements for Attaching an EVS Disk to an ECS?**
8. Click **OK**.

   After the disk is attached, you can view the information about it on the **Disks** tab.

Follow-up Procedure

If the attached disk is newly created, the disk can be used only after it is initialized.

For details about how to initialize a data disk, see **Initializing an EVS Data Disk**.

3.4 Adding a Yearly/Monthly EVS Disk

Scenarios

You are allowed to add yearly/monthly EVS disks to a yearly/monthly ECS. The expiration time of the newly added EVS disks is the same as that of the ECS.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS.

   The page providing details about the ECS is displayed.
6. Click the **Disks** tab. Then, click **Add Disk**.

   The system switches to the EVS disk purchase page.
7. Configure parameters for the new EVS disk as prompted.
8. Click **Next**.
9. Verify that the disk is correctly configured, select the agreement, and click **Submit**.

   The new EVS disk is automatically attached to the target ECS.

   NOTE

   After the new disk is detached, it can only be attached to the original ECS.

3.5 Detaching an EVS Disk from a Running ECS

Scenarios

You can detach EVS disks from an ECS.

- System disks (mounted to **/dev/sda** or **/dev/vda**) can only be detached offline. They must be stopped before being detached.
- Data disks (mounted to points other than **dev/sda**) can be detached online if the attached ECS is running certain OSs. You can detach these data disks without stopping the ECS.

This section describes how to detach a disk from a running ECS.

Constraints

- The EVS disk to be detached must be mounted to a point other than **/dev/sda** or **/dev/vda**.

  EVS disks mounted to **/dev/sda** or **/dev/vda** are system disks and cannot be detached from running ECSs.
- Before detaching an EVS disk from a running Windows ECS, make sure that UVP VMTools have been installed on the ECS and that the tools are running properly.
- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- SCSI EVS disks cannot be detached from running Windows ECSs.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Notes

- On a Windows ECS, if the disk is in non-offline state, the system forcibly detaches the EVS disk. If this occurs, the system may generate a xenvbd alarm. You can ignore this alarm.

  NOTE

  To view the status of an EVS disk, perform the following operations:

  1. Click **Start** in the task bar. In the displayed **Start** menu, right-click **Computer** and choose **Manage** from the shortcut menu.

     The **Server Manager** page is displayed.
  2. In the navigation pane on the left, choose **Storage** > **Disk Management**.

     The EVS disk list is displayed in the right pane.
  3. View the status of each EVS disk.
- Do not detach an EVS disk from an ECS that is being started, stopped, or restarted.
- Do not detach an EVS disk from a running ECS whose OS does not support this feature. OSs supporting EVS disk detachment from a running ECS are listed in **OSs Supporting EVS Disk Detachment from a Running ECS**.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and then attached to it again. This is a normal case due to the drive letter allocation mechanism of the Linux system.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and the ECS is restarted. This is a normal case due to the drive letter allocation mechanism of the Linux system.

OSs Supporting EVS Disk Detachment from a Running ECS

OSs supporting EVS disk detachment from a running ECS include two parts:

- For the first part, see **External Image File Formats and Supported OSs**.
- **Table 3-1** lists the second part of supported OSs.

Table 3-1 OSs supporting EVS disk detachment from a running ECS

OS	Version
CentOS	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Debian	8.6.0 64bit
	8.5.0 64bit
Fedora	25 64bit
	24 64bit
SUSE	SUSE Linux Enterprise Server 12 SP2 64bit
	SUSE Linux Enterprise Server 12 SP1 64bit
	SUSE Linux Enterprise Server 11 SP4 64bit
	SUSE Linux Enterprise Server 12 64bit
OpenSUSE	42.2 64bit
	42.1 64bit
Oracle Linux Server release	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Ubuntu Server	16.04 64bit
	14.04 64bit
	14.04.4 64bit
Windows	Windows Server 2008 R2 Enterprise 64bit
	Windows Server 2012 R2 Standard 64bit
	Windows Server 2016 R2 Standard 64bit
Red Hat Linux Enterprise	7.3 64bit
	6.8 64bit

NOTE

Online detachment is not supported by the ECSs running OSs not listed in the preceding table. For such ECSs, stop the ECSs before detaching disks from them to prevent any possible problems from occurring.

Procedure

1. On the **Elastic Cloud Server** page, click the name of the ECS from which the EVS disk is to be detached. The page providing details about the ECS is displayed.
2. Click the **Disks** tab. Locate the row containing the EVS disk to be detached and click **Detach**.

3.6 Expanding the Capacity of an EVS Disk

Scenarios

You can expand the disk capacity if the disk space is insufficient.

Procedure

The capacity of an EVS disk can be expanded in either of the following ways:

- Apply for an EVS disk and attach it to an ECS.
- Expand the capacity of an existing EVS disk. The capacities of both system disks and data disks can be expanded.

For details, see **Expansion Overview**.

NOTE

After the capacity is expanded through the management console, only the storage capacity of the EVS disk is expanded. To use the expanded capacity, you also need to log in to the ECS and expand the partition and file system.

Related Operations

For a Windows ECS, if you want to expand the disk capacity by clearing disk files, you can reduce the size of the WinSxS folder using tools built into Windows. For details, see **Clean Up the WinSxS Folder**.

3.7 Expanding the Local Disks of a Disk-intensive ECS

Scenarios

Disk-intensive ECSs can use both local disks and EVS disks to store data. Local disks are generally used to store service data and feature higher throughput than EVS disks.

Disk-intensive ECSs do not support specifications modification. When the capacity of local disks is insufficient, you can create a new disk-intensive ECS with higher specifications for capacity expansion. The data stored in the original ECS can be migrated to the new ECS through EVS.

Procedure

1. Create an EVS disk according to the volume of data to be migrated.

2. Attach the EVS disk to the disk-intensive ECS for which you want to expand the capacity.
3. Back up the data stored in the local disks to the EVS disk that is newly attached to the disk-intensive ECS.
4. Detach the EVS disk from the ECS.
   a. On the **Elastic Cloud Server** page, select this disk-intensive ECS and ensure that it has been stopped.

      If the ECS is running, choose **More** > **Stop** to stop the ECS.
   b. Click the name of the disk-intensive ECS. The page providing details about the ECS is displayed.
   c. Click the **Disks** tab. Locate the row containing the EVS data disk and click **Detach** to detach the disk from the ECS.
5. Ensure that a new disk-intensive ECS with higher specifications than the original one is available.

   The local disk capacity is sufficient enough to meet your requirements.
6. Attach the EVS disk to the new disk-intensive ECS.

   On the **Elastic Cloud Server** page, click the name of the ECS described in step **5** to view details.
7. Click the **Disks** tab. Then, click **Attach Disk**.

   In the displayed dialog box, select the EVS disk detached in step **4** and the device name.
8. Migrate the data from the EVS disk to the local disks of the new disk-intensive ECS.

3.8 Enabling Advanced Disk

Scenarios

- Disk functions have been upgraded on the platform. Newly created ECSs can have up to 60 attached disks. However, an existing ECS can still have a maximum of 24 attached disks (40 for certain ECSs). To allow such ECSs to have up to 60 attached disks, enable advanced disk.
- After advanced disk is enabled, you can view the mapping between device names and disks. For details, see "What Is the Mapping Between Device Names and Disks?"

This section describes how to enable advanced disk on an ECS.

Procedure

1. Log in to management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . 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 **Disks** tab.
6. View the current number of disks that can be attached to the ECS and enable advanced disk as prompted.

   The **Enable Advanced Disk** dialog box is displayed.
7. Click **OK**.
8. Stop and then start the target ECS.

   This operation allows advanced disk to take effect.
9. Switch to the page providing details about the ECS again, click the **Disks** tab, and check whether the number of disks that can be attached to the ECS has been changed.
   - If yes, advanced disk has been enabled.
   - If no, enabling advanced disk failed. In such a case, try again later or contact customer service.

4 Backup Using CBR

4.1 Overview

What Is CBR?

Cloud Backup and Recovery (CBR) enables you to back up cloud servers and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

CBR protects your services by ensuring the security and consistency of your data.

What Are the Differences Between Backup, Snapshot, and Image?

You can use the cloud server backup function to create ECSs and the cloud disk backup function to create EVS disks.

An image can be a system disk image, data disk image, or full-ECS image.

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Cloud server backup	All disks (system and data disks) on an ECS	• **Hacker attacks and viruses** You can use cloud server backup to restore data to the latest backup point at which the ECS has not been affected by hacker attacks and viruses. • **Accidental data deletion** You can use cloud server backup to restore data to the backup point prior to the accidental deletion. • **Application update errors** You can use cloud server backup to restore data to the backup point prior to the application update. • **System breakdown** You can use cloud server backup to restore an ECS to the backup point in time prior to system breakdown.	All disks on an ECS are backed up at the same time, ensuring data consistency. In addition, you can configure backup policies for automatic backup.	**Creating a Cloud Server Backup**	• **Restoring Data Using a Cloud Server Backup** • **How Do I Restore Data on the Original Server to a New Server?**

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Cloud disk backup	One or more specified disks (system or data disks)	• **Only data disks need to be backed up, because the system disk does not contain users' application data.** You can use cloud disk backup to back up and restore data if an EVS disk is faulty or encounters a logical error, for example, accidental deletion, hacker attacks, and virus infection. • **Use backups as baseline data.** After a backup policy has been set, the EVS disk data can be automatically backed up based on the policy. You can use the backups created on a timely basis as the baseline data to create new EVS disks or to restore the backup data to EVS disks.	Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption. Backup cost is reduced without compromising data security.	**Creating a Cloud Disk Backup**	• **Restoring Data Using a Cloud Disk Backup** • **Using a Backup to Create a Disk**

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Snapshot	One or more specified disks (system or data disks)	• **Routine data backup** You can create snapshots for disks on a timely basis and use snapshots to recover your data in case that data loss or data inconsistency occurred due to unintended actions, viruses, or attacks. • **Rapid data restoration** You can create a snapshot or multiple snapshots before an application software upgrade or a service data migration. If an exception occurs during the upgrade or migration, service data can be rapidly restored to the time point when the snapshot was created. For example, if ECS A cannot be started due to a fault occurred in system disk A, you can create disk B using an existing snapshot of system disk A and attach disk B to a properly running ECS, for example ECS B. In this case, ECS B can read the data of system disk A from the disk B. • **Rapid deployment of multiple services** You can use a snapshot to create multiple EVS disks containing the same	• The snapshot data is stored with the disk data to facilitate rapid data backup and restoration. • You can create snapshots to rapidly save disk data as it was at specified points in time. You can also use snapshots to create new disks so that the created disks will contain the snapshot data in the beginning.	**Creating a Snapshot**	**Rolling Back Data from a Snapshot**

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
		initial data, and these disks can be used as data resources for various services, for example data mining, report query, and development and testing. This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements. **NOTE** ● A snapshot can be rolled back only to its source disk. Rollback to another disk is not possible. ● If you have reinstalled or changed the ECS OS, snapshots of the system disk are automatically deleted. Snapshots of the data disks can be used as usual.			

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
System disk image	System disk	• **Rapid system recovery** You can create a system disk image for the system disk of an ECS before OS change, application software upgrade, or service data migration. If an exception occurs during the migration, you can use the system disk image to change ECS OS or create a new ECS. • **Rapid deployment of multiple services** You can use a system disk image to quickly create multiple ECSs with the same OS, thereby quickly deploying services these ECSs.	A system disk image can help an ECS with OS damaged to quickly change its OS.	**Creating a System Disk Image**	• **Changing the OS of a Faulty ECS Using a System Disk Image** • **Creating an ECS from a System Disk Image'**
Data disk image	Specific data disk	**Rapid data replication** You can use a data disk image to create multiple EVS disks containing the same initial data, and then attach these disks to ECSs to provide data resources for multiple services.	A data disk image can replicate all data on a disk and create new EVS disks. The EVS disks can be attached to other ECSs for data replication and sharing.	**Creating a Data Disk Image**	**Creating a Data Disk Using a Data Disk Image**

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Full-ECS image	All disks (system and data disks) on an ECS	• **Rapid system recovery** You can create a full-ECS image for the system disk and data disks of an ECS before OS change, application software upgrade, or service data migration. If an exception occurs during the migration, you can use the full-ECS image to change ECS OS or create a new ECS. • **Rapid deployment of multiple services** You can use a full-ECS image to quickly create multiple ECSs with the same OS and data, thereby quickly deploying services these ECSs.	A full-ECS image facilitates service migration.	Creating a Full-ECS Image	Creating an ECS from a Full-ECS Image

CBR Architecture

CBR consists of backups, vaults, and policies.

- **Backup**

  A backup is a copy of a particular chunk of data and is usually stored elsewhere so that it may be used to restore the original data in the event of data loss. CBR supports the following backup types:

  - Cloud server backup: This type of backup uses the consistency snapshot technology for disks to protect data of ECSs and BMSs. The backups of servers without deployed databases are common server backups, and those of servers with deployed databases are application-consistent backups.
  - Cloud disk backup: This type of backup provides snapshot-based data protection for EVS disks.

- **Vault**

  CBR uses vaults to store backups. Before creating a backup, you need to create at least one vault and associate the resource you want to back up with the vault. Then the backup of the resource is stored in the associated vault.

Vaults can be classified into two types: backup vaults and replication vaults. Backup vaults store backups, whereas replication vaults store replicas of backups.

The backups of different types of resources must be stored in different types of vaults.

- **Policy**

  Policies are divided into backup policies and replication policies.

  - Backup policies: To perform automatic backups, configure a backup policy by setting the execution times of backup tasks, the backup cycle, and retention rules, and then apply the policy to a vault.
  - Replication policies: To automatically replicate backups or vaults, configure a replication policy by setting the execution times of replication tasks, the replication cycle, and retention rules, and then apply the policy to a vault. Replicas of backups must be stored in replication vaults.

Backup Mechanism

A full backup is performed only for the first backup and backs up all used data blocks.

For example, if the size of a disk is 100 GB and the used space is 40 GB, the 40 GB of data is backed up.

An incremental backup backs up only the data changed since the last backup, which is storage- and time-efficient.

When a backup is deleted, only the data blocks that are not depended on by other backups are deleted, so that other backups can still be used for restoration. Both a full backup and an incremental backup can restore data to the state at a given backup point in time.

When creating a backup of a disk, CBR also creates a snapshot for it. Every time a new disk backup is created, CBR deletes the old snapshot and keeps only the latest snapshot.

CBR stores backup data in OBS, enhancing backup data security.

Backup Options

CBR supports one-off backup and periodic backup. A one-off backup task is manually created by users and is executed only once. Periodic backup tasks are automatically executed based on a user-defined backup policy.

Table 4-1 One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks driven by a backup policy

Item	One-Off Backup	Periodic Backup
Backup name	User-defined backup name, which is **manualbk**_*xxxx* by default	System-assigned backup name, which is **autobk**_*xxxx* by default
Backup mode	Full backup for the first time and incremental backup subsequently, by default	Full backup for the first time and incremental backup subsequently, by default
Application scenario	Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails.	Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

4.2 Backing Up an ECS Data

Scenarios

CBR enhances data integrity and service continuity. For example, if an ECS or EVS disk is faulty or a misoperation causes data loss, you can use data backups to quickly restore data. This section describes how to back up ECSs and EVS disks.

For more information, **CBR Architecture**, **Backup Mechanism**, and **Backup Options**.

You can back up ECS data using Cloud Server Backup or Cloud Disk Backup.

- Cloud Server Backup (recommended): Use this backup function if you want to back up the data of all EVS disks (system and data disks) on an ECS. This prevents data inconsistency caused by time difference in creating a backup.
- Cloud Disk Backup: Use this backup function if you want to back up the data of one or more EVS disks (system or data disk) on an ECS. This minimizes backup costs on the basis of data security.

ECS Backup Procedure

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. In the ECS list, locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Create Server Backup** in the **Operation** column.
   - If the ECS has been associated with a vault, configure the backup information as prompted.

- **Server List**: The ECS to be backed up is selected by default.
- **Name**: Customize your backup name.
- **Description**: Supplementary information about the backup.
- **Full Backup**: If this option is selected, the system will perform full backup for the ECS to be associated. The storage capacity used by the backup increases accordingly.

– If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.

For details, see **Purchasing a Server Backup Vault**.

5. Click **OK**. The system automatically creates a backup for the ECS.

On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

The ECS can be restarted if the backup progress of an ECS exceeds 10%. However, to ensure data integrity, restart it after the backup is complete.

After the backup is complete, you can restore server data or create images on the **Backups** tab page. For details, see **Restoring Data Using a Cloud Server Backup** and **Using a Backup to Create an Image**.

EVS Disk Backup Procedure

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. In the ECS list, locate the row containing the target ECS and choose **More** > **Manage Image/Disk/Backup** > **Create Disk Backup** in the **Operation** column.
   – If the ECS has been associated with a vault, configure the backup information as prompted.
     - **Server List**: The ECS to be backed up is selected by default. Click to view the disks attached to the ECSs. Select the disks to be backed up.
     - **Name**: Customize your backup name.
     - **Description**: Supplementary information about the backup.
     - **Full Backup**: If this option is selected, the system will perform full backup for the disks to be associated. The storage capacity used by the backup increases accordingly.
   – If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.

     For details, see **Purchasing a Disk Backup Vault**.
5. Click **OK**. The system automatically creates a backup for the disk.

   On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.

If some files are deleted from the disk during the backup, the deleted files may fail to be backed up. Therefore, to ensure data integrity, delete the target data after the backup is complete.

After the backup is complete, you can restore disk data on the **Backups** tab page. For details, see **Restoring Data Using a Cloud Disk Backup**.

5 NICs

5.1 Overview

VPC

Virtual Private Cloud (VPC) allows you to create customized virtual networks in your logically isolated AZ. Such networks are dedicated zones that are logically isolated, providing secure network environments for your ECSs. You can define security groups, virtual private networks (VPNs), IP address segments, and bandwidth for a VPC. This facilitates internal network configuration and management and allows you to change your network in a secure and convenient network manner. You can also customize the ECS access rules within a security group and between security groups to improve ECS security.

For more information about VPC, see **Virtual Private Cloud User Guide**.

Network Interface Types

- A primary network interface is created together with an instance by default, and cannot be detached from the instance.
- A supplementary network interface is created on the **Network Interfaces** console, and can be attached to or detached from an instance.

Notes and Constraints

- The number of supplementary network interfaces that can be attached to an ECS is determined by the ECS specifications. For details, see **ECS Specifications**.
- Supplementary network interfaces cannot be used to directly access Huawei Cloud services, such as DNS. You can use VPCEP to access these services. For details, see **Buying a VPC Endpoint**.

NIC

A NIC is a virtual network adapter that can be bound to an ECS in a VPC. Through the NIC, you can manage the ECS network. A NIC can be a primary NIC or an extension NIC.

- Primary NIC

  When you create an ECS, the NIC automatically created with the ECS is the primary NIC. The primary NIC cannot be unbound. It is preferentially used for the default route generally.

- Extension NIC

  A NIC that can be separately added is an extension NIC, which can be bound to or unbound from an ECS.

5.2 Attaching a Network Interface

Scenarios

If your ECS requires multiple network interfaces, you can attach them to your ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.

   The page providing details about the ECS is displayed.
5. On the **Network Interfaces** tab, click **Attach Network Interface**.
6. Select either of the following methods to attach the network interface.
   - Use an existing network interface.
     i. (Optional) Search for the network interface by name, ID, or private IP address.
     ii. In the network interface list, select the target one.
   - Create a new network interface.

     Set the subnet and security group for the network interface to be attached.

Figure 5-1 Configuring the subnet and security group

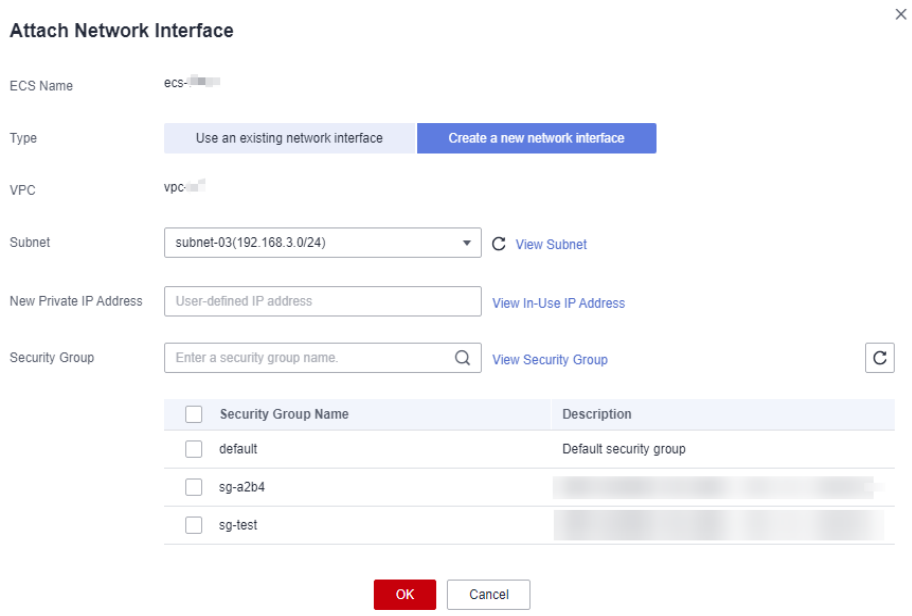

- **Subnet**: specifies the subnet which the network interface belongs to.
- **Private IP Address**: If you want to add a network interface with a specified IP address, enter an IP address into the **Private IP Address** field.
- **Security Group**: You can select multiple security groups. In such a case, the rules of these security groups are effectively aggregated to create one set of rules to apply to the ECS.

7. Click **OK**.

Follow-up Procedure

Some OSs cannot identify newly added network interfaces. In this case, you must manually activate the network interfaces. Ubuntu is used as an example in the following network interface activation procedure. Required operations may vary among systems. For additional information, see the documentation for your OS.

1. Locate the row containing the target ECS and click **Remote Login** in the **Operation** column.

   Log in to the ECS.
2. Run the following command to view the network interface name:

   ifconfig -a

   In this example, the network interface name is **eth2**.
3. Run the following command to switch to the target directory:

   cd /etc/network
4. Run the following command to open the **interfaces** file:

   vi interfaces
5. Add the following information to the **interfaces** file:

auto *eth2*

iface *eth2* **inet dhcp**

6. Run the following command to save and exit the **interfaces** file:

   :wq

7. Run either the **ifup eth2** command or the **/etc/init.d/networking restart** command to make the newly added network interface take effect.

   X in the preceding command indicates the network interface name and SN, for example, **ifup eth2**.

8. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:

   ifconfig

   For example, check whether **eth2** is displayed in the command output.

   - If yes, the newly added network interface has been activated, and no further action is required.
   - If no, the newly added network interface failed to be activated. Go to step 9.

9. Log in to the management console. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.
10. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:
    - If yes, no further action is required.
    - If no, contact customer service.

5.3 Detaching a Network Interface

Scenarios

An ECS can have up to 12 network interfaces, including one primary network interface that cannot be deleted and extension network interfaces. This section describes how to detach a network interface.

Procedure

1. Log in to the management console.
2. Click  . 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. On the **Network Interfaces** tab, locate the target network interface and click **Detach**.

   NOTE

   You are not allowed to delete the primary ECS network interface. By default, the primary ECS network interface is the first network interface displayed in the network interface list.

5. In the displayed dialog box, click Yes.

NOTE

Certain ECSs do not support network interface detachment when they are running. For details, see the GUI display. To detach a network interface from such an ECS, stop the ECS first.

5.4 Changing a VPC

Scenarios

This section describes how to change a VPC.

Constraints

- A VPC can be changed for an ECS only if the ECS has one NIC.
- If you have reinstalled or changed the OS of an ECS before changing the VPC, log in to the ECS and check whether the password or key pair configured during the reinstallation or change is successfully injected.
  - If the login is successful, the password or key pair is injected. Perform operations as required.
  - Otherwise, the system is injecting the password or key pair. During this period, do not perform any operations on the ECS.
- During the change process, do not perform operations on the ECS, including its EIP.
- If an ECS NIC has an IPv6 address, the VPC of the ECS cannot be changed.

Notes

- A VPC can be changed only on a running ECS. However, ECS network connections will be interrupted during the change process.
- After the VPC is changed, the subnet, private IP address, MAC address, and OS NIC name of the ECS will change.
- After the VPC is changed, the source/destination check and virtual IP address must be configured again.
- After the VPC is changed, you are required to reconfigure network-related application software and services, such as ELB, VPN, NAT, and DNS.

Procedure

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Change VPC**.

   The **Change VPC** dialog box is displayed.

Figure 5-2 Change VPC

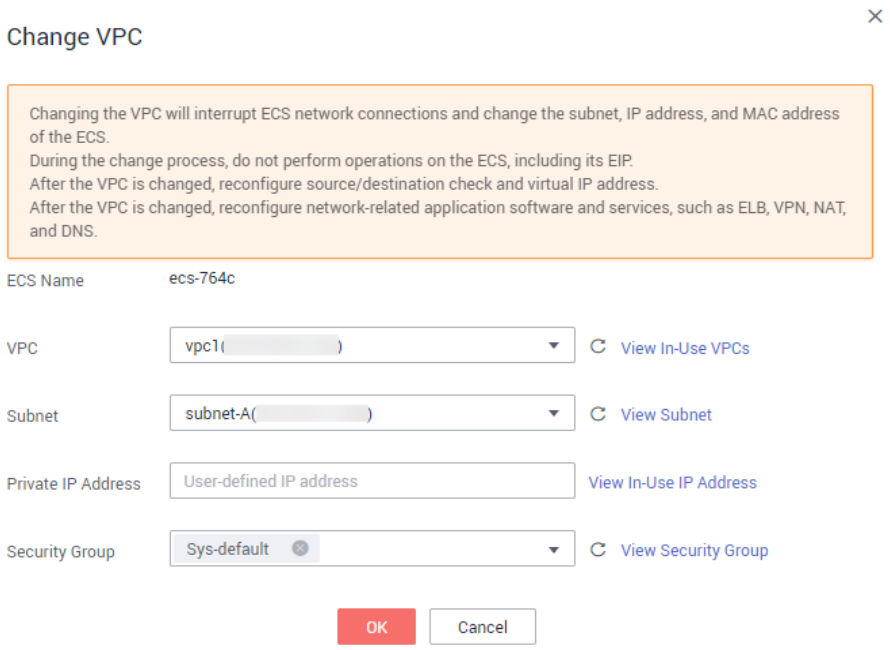

4. Select an available VPC and subnet from the drop-down lists, and set the private IP address and security group as prompted.

   You can select multiple security groups. In such a case, the access rules of all the selected security groups apply on the ECS.

   NOTE

   Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click **OK**.

5.5 Modifying a Private IP Address

Scenarios

You can modify the private IP address of the primary NIC. If you want to modify the private IP address of an extension NIC, delete the NIC and attach a new NIC.

Constraints

- The ECS must be stopped.
- If a virtual IP address or DNAT rule has been configured for the NIC, cancel the configuration before modifying the private IP address.
- If the NIC has an IPv6 address, its private IP address (IPv4 or IPv6 address) cannot be modified.
- Before changing the private IP address of an ELB backend server, delete the backend server group.

Procedure

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. Click the name of the target ECS.

   The ECS details page is displayed.
4. Click the **Network Interfaces** tab. Locate the row containing the primary network interface and click **Modify Private IP**.

   The **Modify Private IP** dialog box is displayed.
5. Change the subnet and private IP address of the primary NIC as required.

   NOTE

   Subnets can be changed only within the same VPC.

   If the target private IP address is not specified, the system will automatically assign one to the primary NIC.

5.6 Managing Virtual IP Addresses

Scenarios

A virtual IP address provides the second IP address for one or more ECS NICs, improving high availability between the ECSs.

Procedure

1. Log in to the management console.
2. Click . 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. On the **Network Interfaces** tab, locate the target virtual IP address and click **Manage Virtual IP Address**.
5. On the **IP Addresses** tab of the displayed page, locate the row containing the target virtual IP address and select **Bind to EIP** or **Bind to Server** in the **Operation** column.

   Multiple ECSs deployed to work in active/standby mode can be bound with a virtual IP address to improve DR performance.
6. Click **OK**.
7. Manually configure the virtual IP address bound to an ECS.

   After a virtual IP address is bound to an ECS NIC, you need to manually configure the virtual IP address on the ECS.

   Linux OS (CentOS 7.2 64bit is used as an example.)

   a. Run the following command to obtain the NIC to which the virtual IP address is to be bound and the connection of the NIC:

      nmcli connection

Information similar to the following is displayed:

The command output in this example is described as follows:

- **eth0** in the **DEVICE** column indicates the NIC to which the virtual IP address is to be bound.
- **Wired connection 1** in the **NAME** column indicates the connection of the NIC.

b. Run the following command to add the virtual IP address for the target connection:

nmcli connection modify "*CONNECTION***" ipv4.addresses** *VIP*

Configure the parameters as follows:

- CONNECTION: connection of the NIC obtained in **7.a**.
- VIP: virtual IP address to be added.
  - If you add multiple virtual IP addresses at a time, separate them with commas (,).
  - If a virtual IP address already exists and you need to add a new one, the command must contain both the new and original virtual IP addresses.

Example commands:

- Adding a single virtual IP address: **nmcli connection modify "Wired connection 1" ipv4.addresses 172.16.0.125**
- Adding multiple virtual IP addresses: **nmcli connection modify "Wired connection 1" ipv4.addresses 172.16.0.125,172.16.0.126**

c. Run the following command to make the configuration take effect:

nmcli connection up "*CONNECTION***"**

In this example, run the following command:

nmcli connection up "Wired connection 1"

Information similar to the following is displayed:

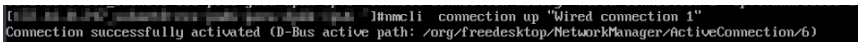

d. Run the following command to check whether the virtual IP address has been bound:

ip a

Information similar to the following is displayed. In the command output, the virtual IP address 172.16.0.125 is bound to NIC eth0.

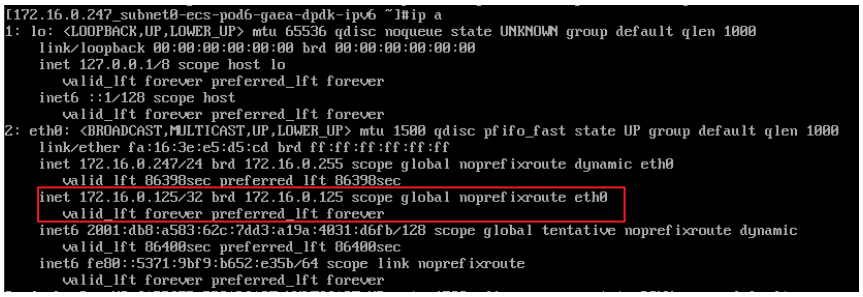

Windows OS (Windows Server is used as an example here.)

a. In **Control Panel**, click **Network and Sharing Center**, and click the corresponding local connection.

b. On the displayed page, click **Properties**.

c. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.

d. Click **Properties**.

e. Select **Use the following IP address** and set **IP address** to the private IP address of the ECS, for example, 10.0.0.101.

Figure 5-3 Configuring private IP address

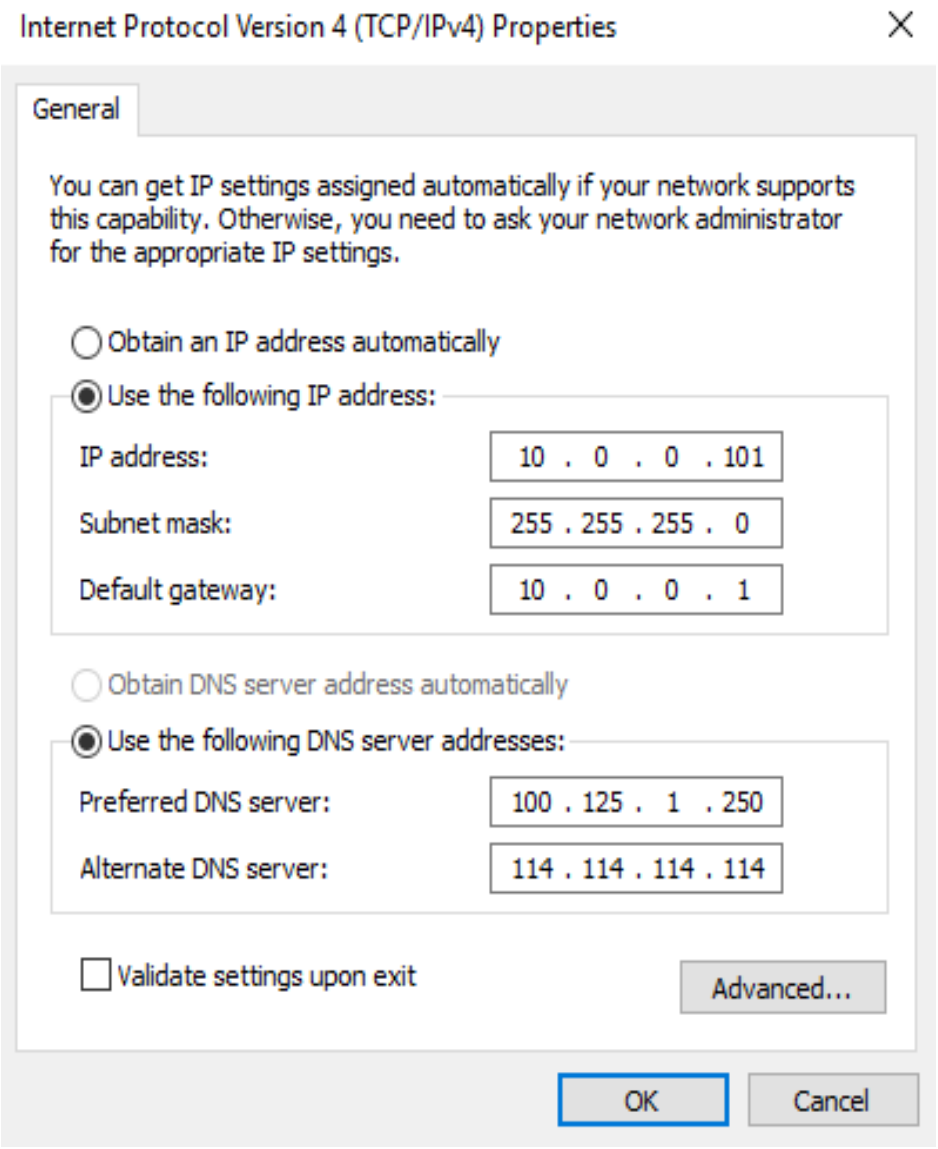

f. Click **Advanced**.

g. On the **IP Settings** tab, click **Add** in the **IP addresses** area.
Add the virtual IP address. For example, 10.0.0.154.

Figure 5-4 Configuring virtual IP address

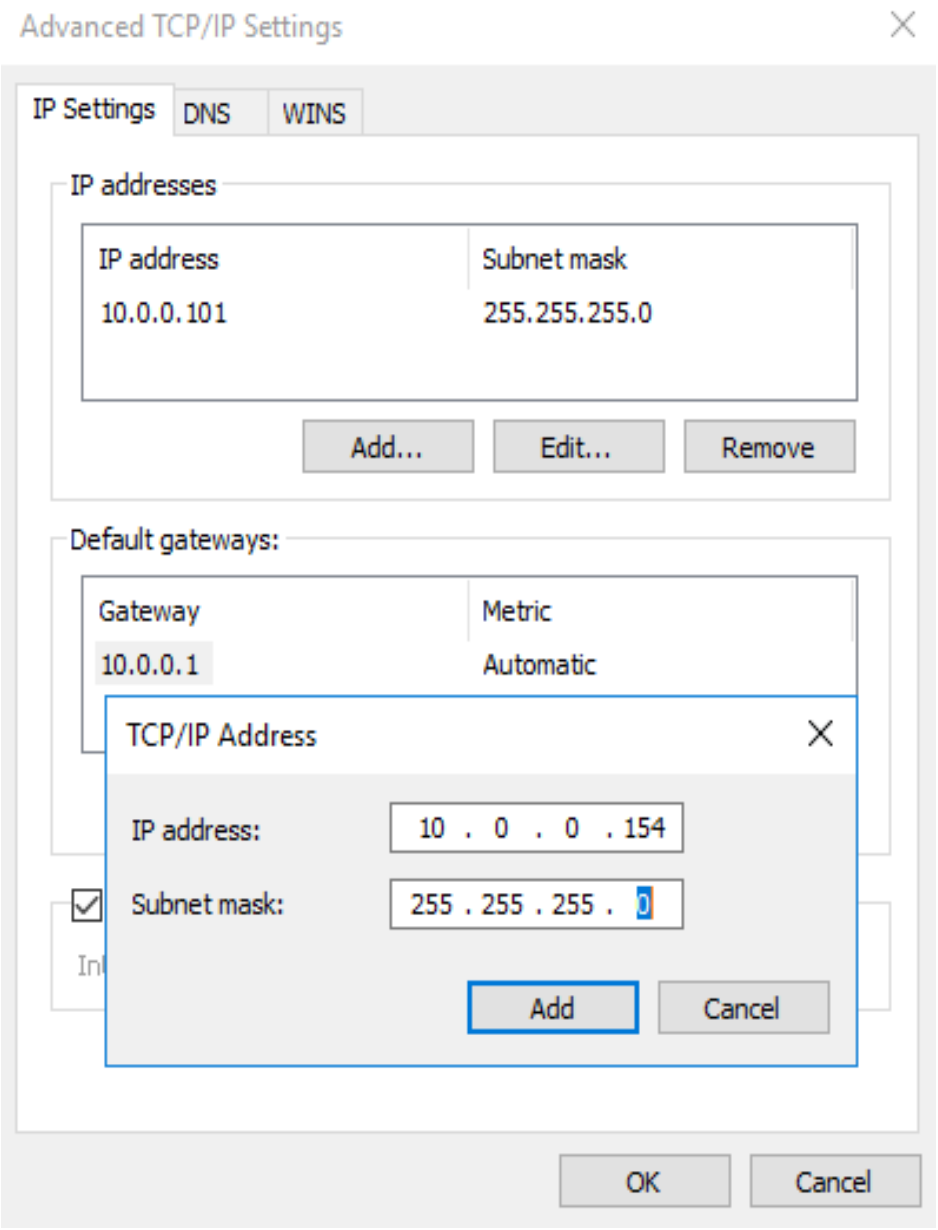

h. Click **OK**.

i. In the **Start** menu, open the Windows command line window and run the following command to check whether the virtual IP address has been configured:

ipconfig /all

In the command output, **IPv4 Address** is the virtual IP address 10.0.0.154, indicating that the virtual IP address of the ECS NIC has been correctly configured.

5.7 Enabling NIC Multi-Queue

Scenarios

Single-core CPU performance cannot meet the requirement of processing NIC interruptions incurred with the increase of network I/O bandwidth. NIC multi-queue enables multiple CPUs to process ECS NIC interruptions, thereby improving PPS and I/O performance.

The ECS described in this section is assumed to comply with the requirements on specifications and virtualization type.

- If the ECS was created using a public image listed in **Support of NIC Multi-Queue**, NIC multi-queue has been enabled on the ECS by default. Therefore, you do not need to perform the operations described in this section.
- If the ECS was created using a private image and the OS of the external image file is listed in **Support of NIC Multi-Queue**, perform the following operations to enable NIC multi-queue:
  a. **Importing the External Image File to the IMS Console**
  b. **Setting NIC Multi-Queue for the Image**
  c. **Creating an ECS Using a Private Image**
  d. **Running the Script for Configuring NIC Multi-Queue**

NOTE

After NIC multi-queue is enabled on an ECS, you need to enable this function on the ECS again after you add or delete a NIC or change the VPC for the ECS. For details, see **Running the Script for Configuring NIC Multi-Queue**.

Support of NIC Multi-Queue

NIC multi-queue can be enabled on an ECS only when the ECS specifications, virtualization type, and image OS meet the requirements described in this section.

- For details about the ECS specifications that support NIC multi-queue, see **ECS Types**.

  NOTE

  If the number of NIC queues is greater than 1, NIC multi-queue is supported.
- The virtualization type must be KVM.
- The Linux public images listed in **Table 5-2** support NIC multi-queue.

NOTE

- The PV driver of a Windows ECS dynamically adjusts the number of NIC queues based on the number of vCPUs of the ECS, and you do not need to set the number of Windows NIC multi-queues.
- Public images that contain Windows Server 2008 are no longer available. However, you can still use private images that contain Windows Server 2008.
- It is a good practice to upgrade the kernel version of the Linux ECS to 2.6.35 or later. Otherwise, NIC multi-queue is not supported.

  Run the **uname -r** command to obtain the kernel version. If the kernel version is earlier than 2.6.35, contact customer service to upgrade the kernel.

Table 5-1 Support of NIC multi-queue for Windows ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Windows Server 2008 R2 Standard/Enterprise/ DataCenter 64bit	Yes	Yes
Windows Server 2008 Enterprise SP2 64bit	Yes	Yes
Windows Server 2008 Web R2 64-bit	Yes	Yes
Windows Server 2008 R2 Enterprise 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard/ DataCenter 64 bit	Yes	Yes
Windows Server 2016 Standard/DataCenter 64 bit	Yes	Yes
Windows Server 2019 DataCenter 64 bit	Yes	Yes

Table 5-2 Support of NIC multi-queue for Linux ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Ubuntu 14.04/16.04/18.04/20.04 server 64bit	Yes	Yes
OpenSUSE 42.2/15.* 64bit	Yes	Yes
SUSE Enterprise 12 SP1/SP2 64bit	Yes	Yes
CentOS 6.8/6.9/7.*/8.* 64bit	Yes	Yes
Debian 8.0.0/8.8.0/8.9.0/9.0.0/10.0.0/10.2.0 64bit	Yes	Yes

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Fedora 24/25/30 64bit	Yes	Yes
EulerOS 2.2/2.3/2.5 64bit	Yes	Yes

Importing the External Image File to the IMS Console

For details, see "Registering an Image File as a Private Image" in *Image Management Service User Guide*. After the image file is imported, view the value of **NIC Multi-Queue** on the page providing details about the image.

- If the value is **Supported**, go to **Creating an ECS Using a Private Image**.
- If the value is **Not supported**, go to **Setting NIC Multi-Queue for the Image**.

Setting NIC Multi-Queue for the Image

Windows OSs have not commercially supported NIC multi-queue. If you enable NIC multi-queue in a Windows image, starting an ECS created using such an image may be slow.

Use one of the following methods to set the NIC multi-queue attribute:

Method 1:

1. Log in to the management console.
2. Click  . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab, locate the row containing the target image, click **Modify** in the **Operation** column.
4. Set the NIC multi-queue attribute of the image.

Method 2:

1. Log in to the management console.
2. Click  . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab. In the image list, click the name of the target image to switch to the page providing details about the image.
4. Click **Modify** in the upper right corner. In the displayed **Modify Image** dialog box, set the NIC multi-queue attribute.

Method 3: Add **hw_vif_multiqueue_enabled** to an image through the API.

1. For instructions about how to obtain the token, see **Authentication**.
2. For instructions about how to call an API to update image information, see **Updating Image Information (Native OpenStack API)**.
3. Add **X-Auth-Token** to the request header.

   The value of **X-Auth-Token** is the token obtained in step **1**.

4. Add **Content-Type** to the request header.

   The value of **Content-Type** is **application/openstack-images-v2.1-json-patch**.

   The request URI is in the following format:

   PATCH /v2/images/{image_id}

   The request body is as follows:

```
[
       {
        "op":"add",
        "path":"/hw_vif_multiqueue_enabled",
        "value": "true"
       }
]
```

   Figure 5-5 shows an example request body for modifying the NIC multi-queue attribute.

   Figure 5-5 Example request body

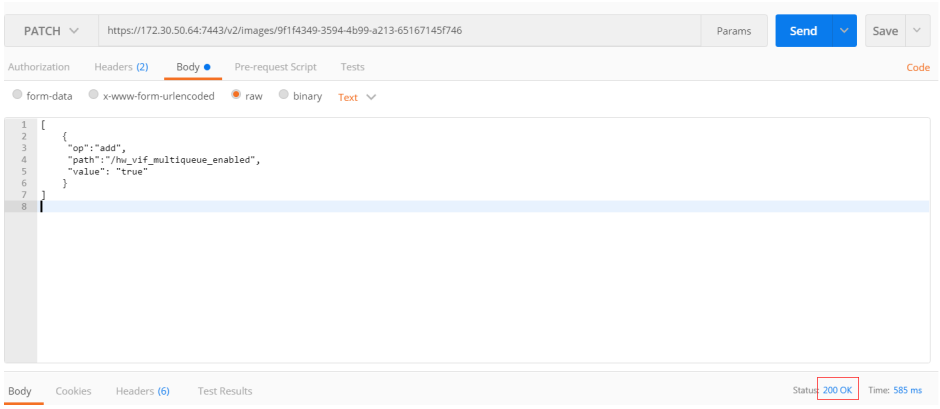

Creating an ECS Using a Private Image

Create an ECS using a registered private image. Note the following when setting the parameters:

- **Region**: Select the region where the private image is located.
- **Image**: Select **Private image** and then the desired image from the drop-down list.

Running the Script for Configuring NIC Multi-Queue

The PV driver of a Windows ECS dynamically adjusts the number of NIC queues based on the number of vCPUs of the ECS, and you do not need to set the number of Windows NIC multi-queues.

A script for automatically enabling NIC multi-queue on a Linux ECS is available. After the script is configured, the ECS supports NIC multi-queue.

1. Log in to the ECS and run the following command to check the number of queues supported by and enabled for a NIC:

   ethtool -l *NIC*

   Example:


```
[root@localhost ~]# ethtool -l eth0   #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
RX:             0
TX:             0
Other:          0
Combined:       4   #The NIC supports a maximum of four queues.
Current hardware settings:
RX:             0
TX:             0
Other:          0
Combined:       1   #One queue has been enabled for the NIC.
```

If the values of the two **Combined** fields are the same, NIC multi-queue has been enabled. No further action is required.

2. Run the following command to download the configuration script "multi-queue-hw".

   wget *URL to download the script*

   URL: **https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/multi-queue-hw**

3. Run the following command to assign execution permissions to the script:

   chmod +x multi-queue-hw

4. Run the following command to move the **multi-queue-hw** script to the **/etc/init.d** directory:

   mv multi-queue-hw /etc/init.d

5. Run the following command to run the script:

   /etc/init.d/multi-queue-hw start

   The script takes effect immediately after being executed. However, after the ECS is stopped, NIC multi-queue disables automatically.

6. Add startup configuration for each OS so that NIC multi-queue automatically enables upon the ECS startup.
   - For CentOS, Red Hat, Fedora, EulerOS, SUSE, and OpenSUSE, run the following command:

     chkconfig multi-queue-hw on
   - For Ubuntu, run the following command:

     update-rc.d multi-queue-hw defaults 90 10
   - For Debian, run the following command:

     systemctl enable multi-queue-hw

Viewing the Number of Queues of the NIC

NIC multi-queue has been enabled.

1. Log in to the ECS.
2. Run the following command to obtain the number of queues supported by the NIC and the number of queues with NIC multi-queue enabled:

   ethtool -l *NIC*

Example:

```
[root@localhost ~]# ethtool -l eth0  #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
```



```
RX:             0
TX:             0
Other:          0
Combined: 4  #Indicates that a maximum of four queues can be enabled for the NIC.
Current hardware settings:
RX:             0
TX:             0
Other:          0
Combined: 1 #Indicates that four queues have been enabled.
```

5.8 Dynamically Assigning IPv6 Addresses

Scenarios

IPv6 addresses are used to deal with IPv4 address exhaustion. If an ECS uses an IPv4 address, the ECS can run in dual-stack mode after IPv6 is enabled for it. Then, the ECS will have two IP addresses to access the intranet and Internet: an IPv4 address and an IPv6 address.

In some cases, an ECS cannot dynamically acquire an IPv6 address even if it meets all the requirements in **Constraints**. You need to configure the ECS to dynamically acquire IPv6 addresses. For public images:

- By default, dynamic IPv6 address assignment is enabled for Windows public images. You do not need to configure it. The operations in **Windows Server 2012** and **Windows Server 2008** are for your reference only.
- Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 has been enabled and then whether dynamic IPv6 address assignment has been enabled. Currently, IPv6 is enabled for all Linux public images but dynamic IPv6 address assignment is only enabled for Ubuntu 16 public images by default.

Constraints

- Ensure that IPv6 has been enabled on the subnet where the ECS works.

  For details about how to enable IPv6 on a subnet, see **Enabling IPv6 on the Subnet Where the ECS Works**.
- Ensure that the ECS flavor supports IPv6.

  The ECS flavors that support IPv6 vary depending on regions and AZs. Check whether an ECS flavor supports IPv6 after you select a region and AZ on the management console.

Figure 5-6 Checking whether an ECS flavor supports IPv6

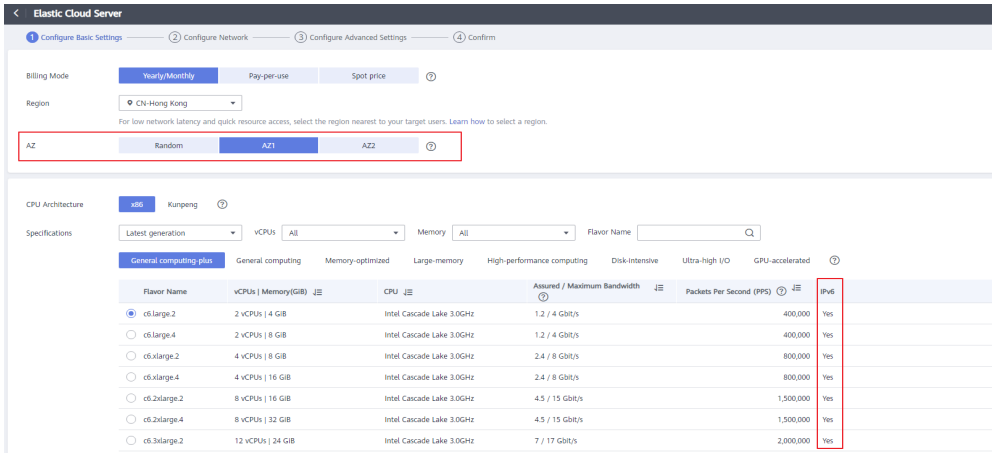

If the value of **IPv6** is **Yes** for an ECS flavor, the flavor supports IPv6.

NOTE

AZ and **Flavor** determine whether IPv6 is supported.

After you select an AZ, if **IPv6** is not displayed or the value of **IPv6** is **No**, IPv6 is not supported by any or certain flavors in the AZ.

- Ensure that **Self-assigned IPv6 address** is selected during ECS creation.

Figure 5-7 Self-assigned IPv6 address

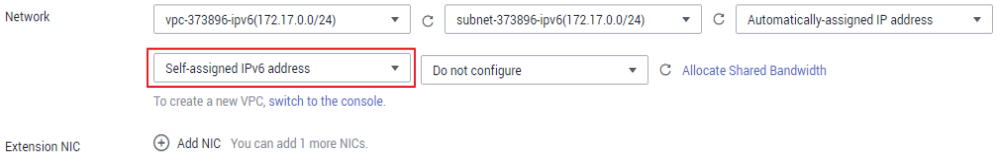

- After an ECS is started, its hot-swappable NICs cannot automatically acquire IPv6 addresses.
- Only ECSs can work in dual-stack mode and BMSs cannot.
- Only one IPv6 address can be bound to a NIC.

Procedure

- Windows: Windows Server 2012/2008 is used as an example to describe how to enable dynamic assignment of IPv6 addresses in Windows.
- Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually.

  For CentOS 6.x and Debian, after dynamic IPv6 address assignment is enabled for an ECS and the ECS is used to create an image, the new ECSs created from this image will start up slowly due to IPv6 address assignment timeout. You can rectify this issue by referring to **Setting the Timeout Duration for IPv6 Address Assignment**.

Table 5-3 Enabling dynamic assignment of IPv6 addresses for different OSs

OS	Automatically/ Manually Enabling	Reference
Windows Server 2012	Automatically	**Windows Server 2012**
Windows Server 2008	Automatically	**Windows Server 2008**
Linux	Automatically (recommended)	**Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)**
Linux	Manually	**Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)**

Enabling IPv6 on the Subnet Where the ECS Works

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. Click the target ECS to go to the detail page.
5. In the **ECS Information** area, click the VPC name.
6. Click the number in the **Subnets** column.
   The **Subnets** page is displayed.
7. In the subnet list, locate the target subnet and click its name.
   The subnet details page is displayed.
8. In the **Subnet Information** area, click **Enable** for **IPv6 CIDR Block**.
9. Click **Yes**.

Windows Server 2012

Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window to check it:

ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

Figure 5-8 Querying the IPv6 address

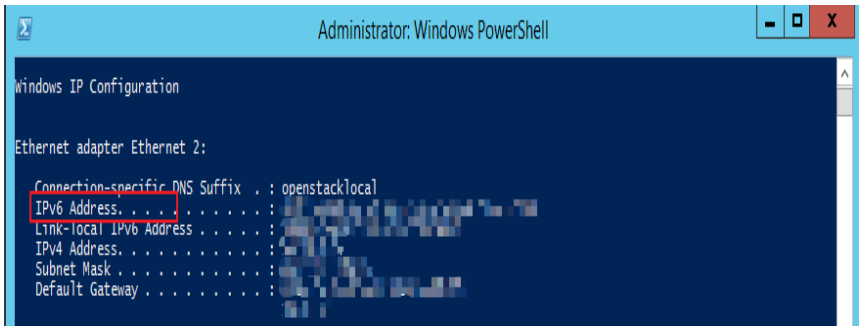

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 5-9 Link-local IPv6 address

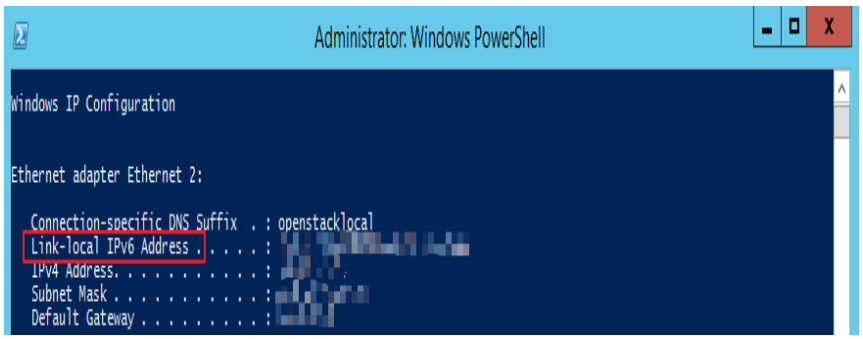

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to **Step 3**.

Figure 5-10 IPv6 disabled

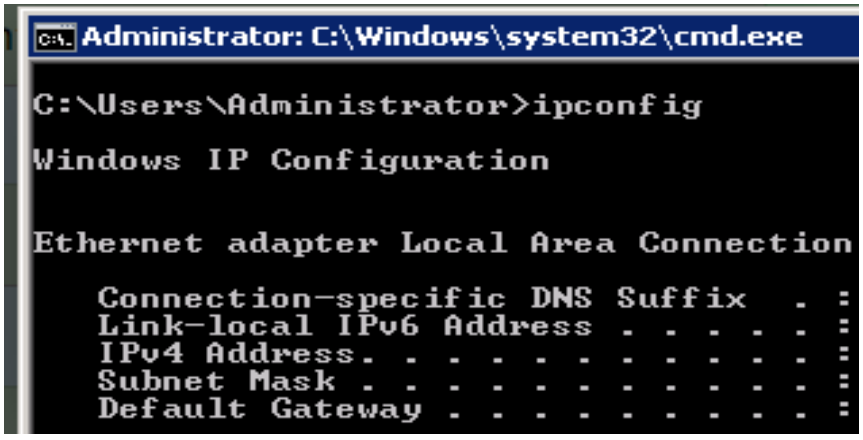

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 5-8**. No additional configuration is required.

Step 2 Enable dynamic IPv6 address assignment.

1. Choose **Start** > **Control Panel**.

2. Click **Network and Sharing Center**.
3. Click the Ethernet connection.

Figure 5-11 Ethernet connection

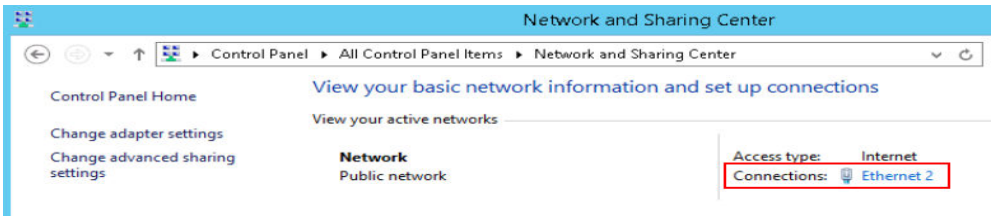

4. In the **Ethernet Status** dialog box, click **Properties** in the lower left corner.
5. Select **Internet Protocol Version 6 (TCP/IPv6)** and click **OK**.

Figure 5-12 Configuring dynamic IPv6 address assignment

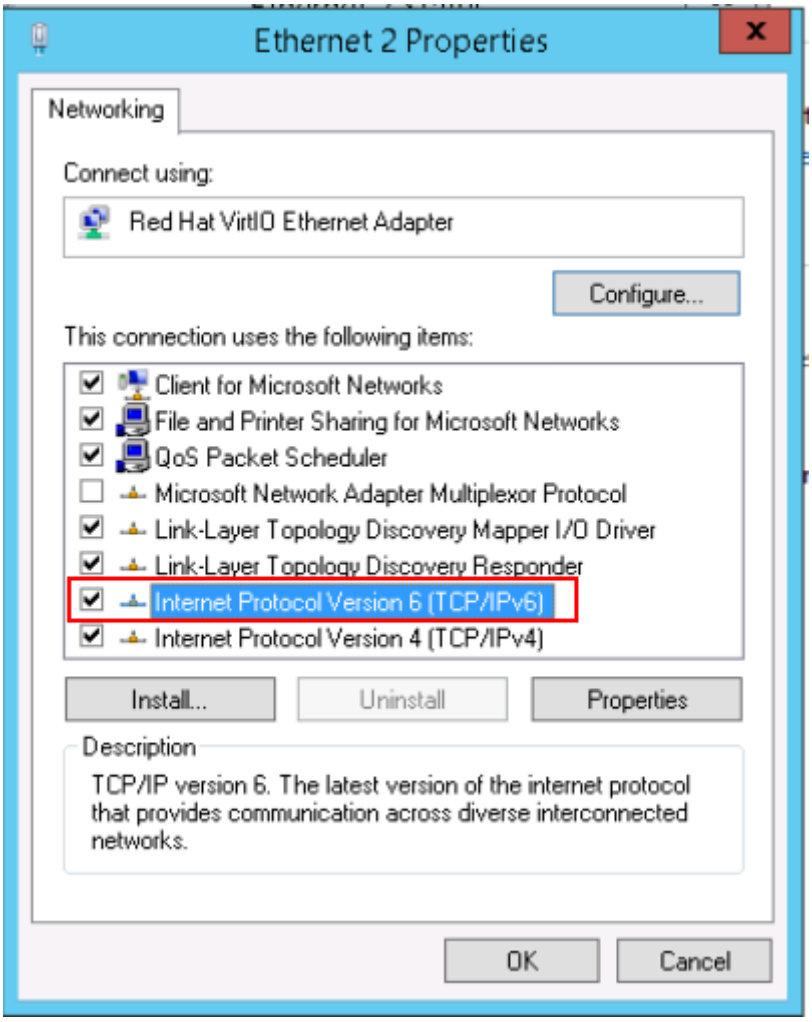

6. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

Step 3 Enable and configure IPv6.

1. In the **Internet Protocol Version 6 (TCP/IPv6) Properties** dialog box, configure an IPv6 address and a DNS server address.

- **IPv6 address**: IPv6 address allocated during ECS creation. Obtain the value from the ECS list on the console.
- **Subnet prefix length**: **64**
- **Preferred DNS server**: **240c::6666** (recommended)

Figure 5-13 Configuring an IPv6 address and a DNS server address

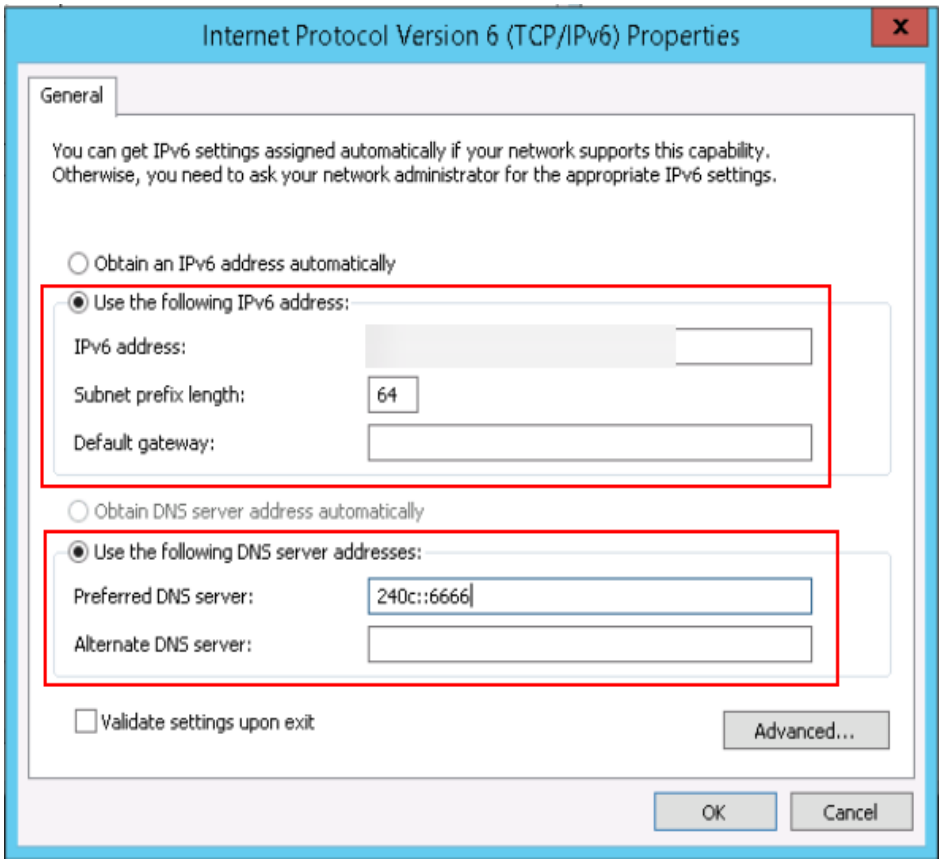

2. (Optional) Run the following command depending on your ECS OS.

   For Windows Server 2012, run the following command in PowerShell or CMD:

   Set-NetIPv6Protocol -RandomizeIdentifiers disabled

3. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

----End

Windows Server 2008

Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window to check it:

ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

Figure 5-14 Querying the IPv6 address

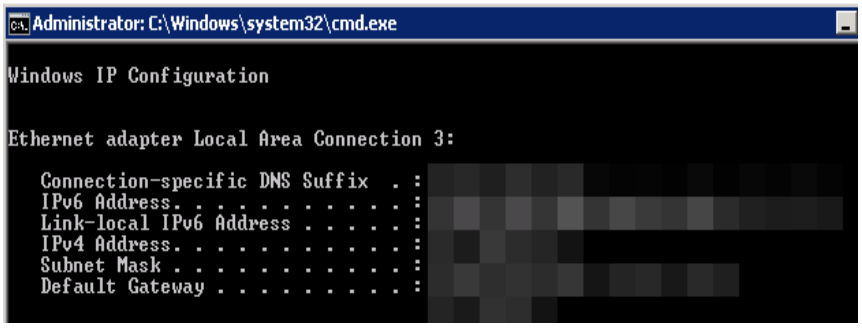

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 5-15 Link-local IPv6 address

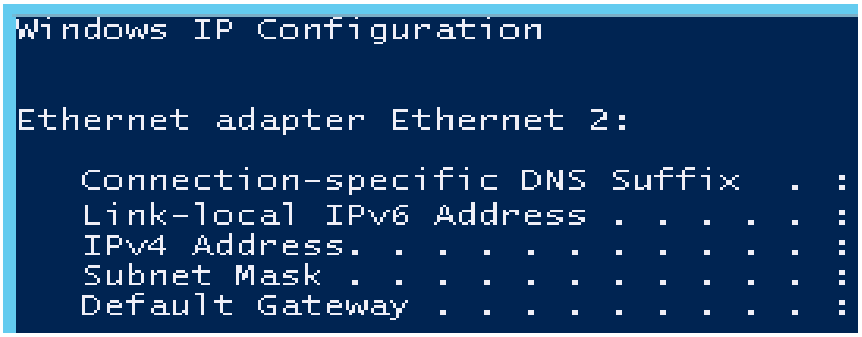

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to **Step 3**.

Figure 5-16 IPv6 disabled

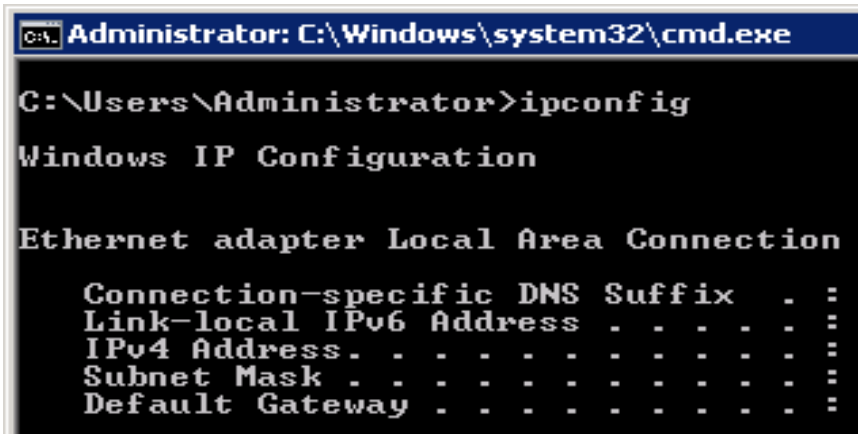

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 5-14**. No additional configuration is required.

Step 2 Enable dynamic IPv6 address assignment.

1. Choose **Start** > **Control Panel**.

2. Click **Network and Sharing Center**.
3. Click **Change adapter settings**.
4. Right-click the local network connection and choose **Properties**.
5. Select **Internet Protocol Version 6 (TCP/IPv6)** and click **OK**.

Figure 5-17 Configuring dynamic IPv6 address assignment

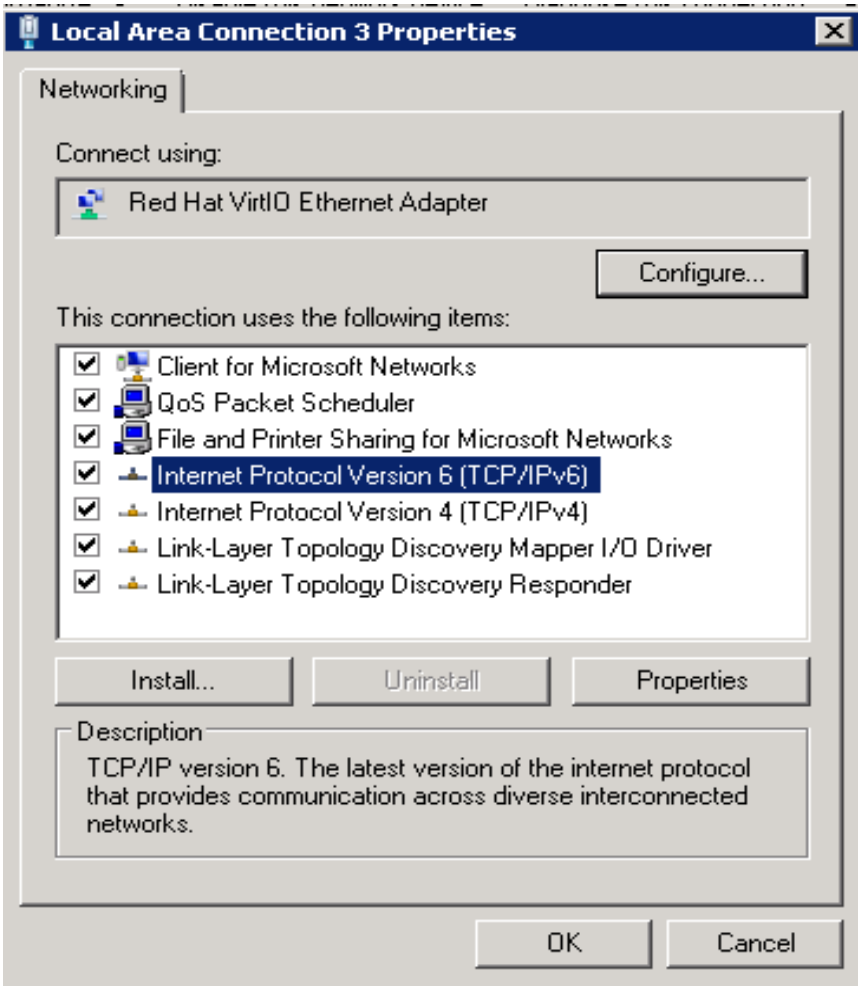

6. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

Step 3 Enable and configure IPv6.

1. Choose **Start** > **Control Panel** > **Network Connection** > **Local Connection**.
2. Select **Properties**, select the following options, and click **Install**.

Figure 5-18 Enabling and configuring IPv6

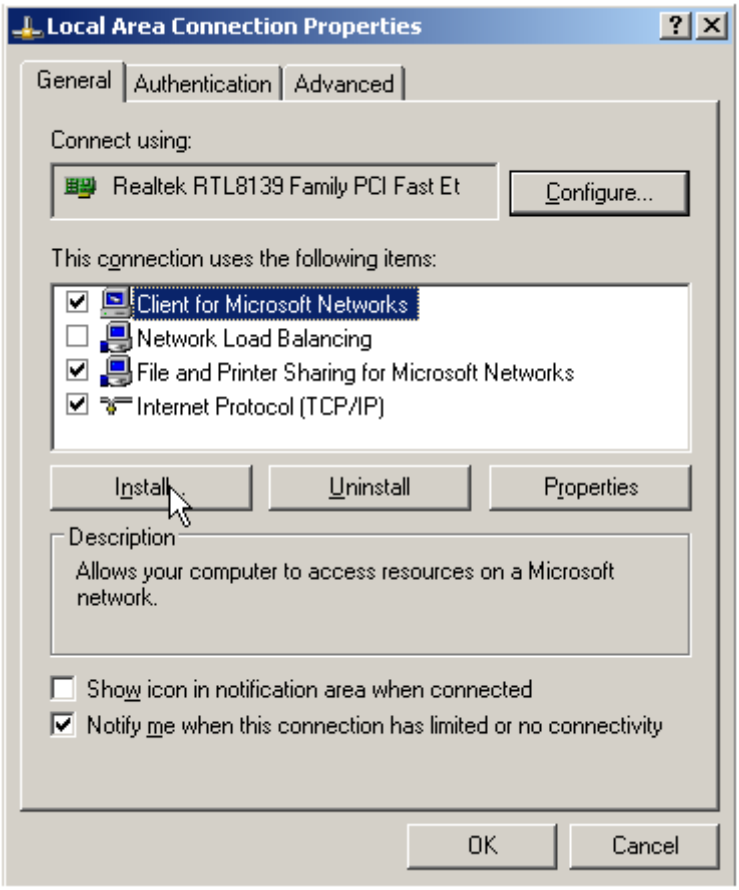

3. Select **Protocol** and click **Add**.

Figure 5-19 Adding the protocol

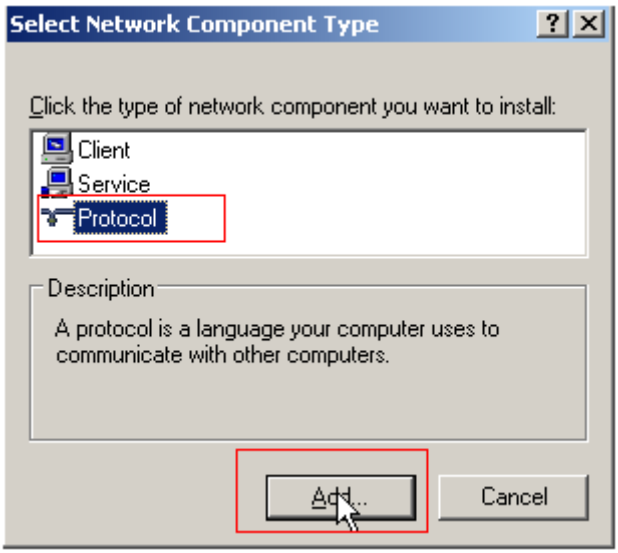

4. Select **Microsoft TCP/IP Version 6** and click **OK**.

Figure 5-20 Network protocols

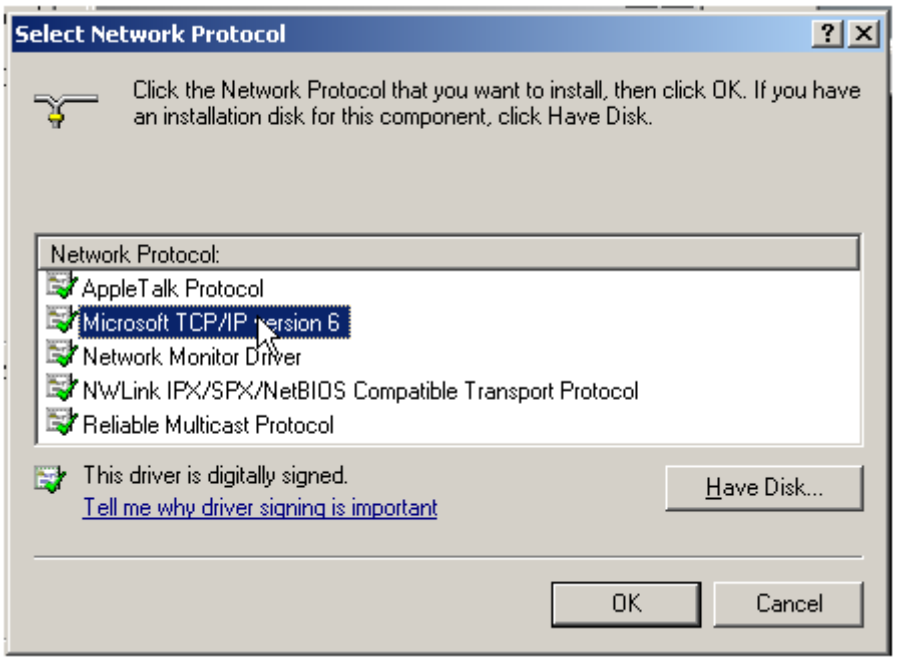

5. (Optional) Run the following commands depending on your ECS OS.

   For Windows Server 2008, run the following command in PowerShell or CMD:

   netsh interface ipv6 set global randomizeidentifiers=disable

   Disable the local connection and then enable it again.

   To disable the local connection, choose **Start** > **Control Panel** > **Network and Internet** > **Network and Sharing Center** > **Change Adapter Options**. Right-click the local connection and choose **Disable** from the shortcut menu.

   To enable the local connection, choose **Start** > **Control Panel** > **Network and Internet** > **Network and Sharing Center** > **Change Adapter Options**. Right-click the local connection and choose **Enable** from the shortcut menu.

6. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

----End

Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)

The **ipv6-setup-***xxx* tool can be used to enable Linux OSs to automatically acquire IPv6 addresses. *xxx* indicates a tool, which can be rhel or debian.

You can also enable dynamic IPv6 address assignment by following the instructions in **Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)**.

CAUTION

- When you run **ipv6-setup-***xxx*, the network service will be automatically restarted. As a result, the network is temporarily disconnected.
- If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to **Setting the Timeout Duration for IPv6 Address Assignment** and try to create a new private image again.

Step 1 Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

- If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to **Step 2**.

  Figure 5-21 IPv6 disabled

  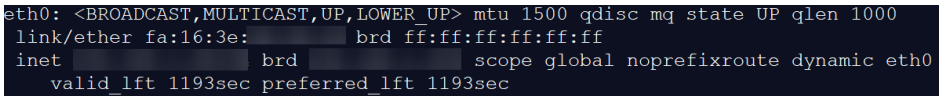

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

  Figure 5-22 IPv6 enabled

  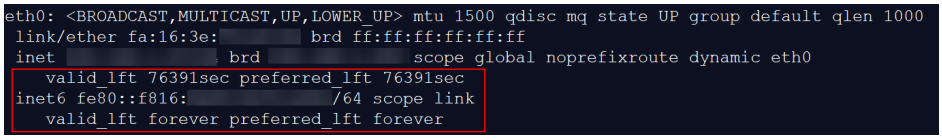

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

  Figure 5-23 IPv6 enabled and an IPv6 address assigned

  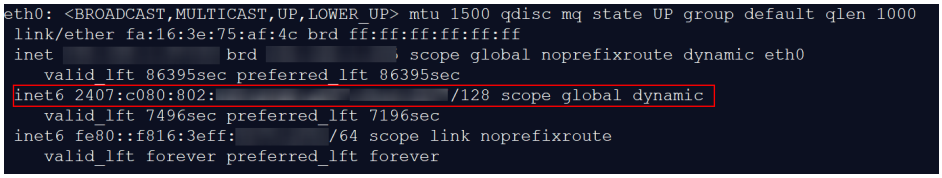

NOTE

IPv6 is enabled for Linux public images by default, as shown in **Figure 5-22**. IPv6 and dynamic IPv6 address assignment are enabled for Ubuntu 16 public images by default, as shown in **Figure 5-23**.

Step 2 Enable IPv6 for the ECS.

1. Run the following command to check whether IPv6 is enabled for the kernel:

   sysctl -a | grep ipv6

   – If a command output is displayed, IPv6 is enabled.

- If no information is displayed, IPv6 is disabled. Go to **Step 2.2** to load the IPv6 module.

2. Run the following command to load the IPv6 module:

   modprobe ipv6

3. Add the following content to the **/etc/sysctl.conf** file:

   net.ipv6.conf.all.disable_ipv6=0

4. Save the configuration and exit. Then, run the following command to load the configuration:

   sysctl -p

Step 3 Enable dynamic IPv6 address assignment for the ECS.

1. Download **ipv6-setup-rhel** or **ipv6-setup-debian** with a required version and upload it to the target ECS.

   ipv6-setup-*xxx* modifies the configuration file of a NIC to enable dynamic IPv6 address assignment or adds such a configuration file for a NIC, and then restarts the NIC or network service. **Table 5-4** lists the download paths of **ipv6-setup-rhel** and **ipv6-setup-debian**.

   Table 5-4 Download paths of ipv6-setup-rhel and ipv6-setup-debian

Series	Release Version	How to Obtain
RHEL	- CentOS 6/7 - EulerOS 2.2/2.3 - Fedora 25	**https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-rhel**
Debian	- Ubuntu 16/18/20 - Debian 8/9/10	**https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-debian**

2. Run the following command to make **ipv6-setup-***xxx* executable:

   chmod +x ipv6-setup-*xxx*

3. Run the following command to enable dynamic IPv6 address assignment for a NIC:

   ./ipv6-setup-*xxx* **--dev** [*dev*]

   Example:

   ./ipv6-setup-*xxx* **--dev eth0**

   NOTE

   - To enable dynamic IPv6 address assignment for all NICs, run the **./ipv6-setup-***xxx* command.
   - To learn how to use **ipv6-setup-***xxx*, run the **./ipv6-setup-***xxx* **--help** command.

----End

Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)

CAUTION

If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to **Setting the Timeout Duration for IPv6 Address Assignment** and try to create a new private image again.

Step 1 Run the following command to check whether IPv6 is enabled for the ECS:

ip addr

- If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to **Step 2**.

  Figure 5-24 IPv6 disabled

  ```
  eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
   link/ether fa:16:3e:        brd ff:ff:ff:ff:ff:ff
   inet                  brd                  scope global noprefixroute dynamic eth0
      valid_lft 1193sec preferred_lft 1193sec
  ```

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

  Figure 5-25 IPv6 enabled

  ```
  eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether fa:16:3e:        brd ff:ff:ff:ff:ff:ff
   inet                  brd                  scope global noprefixroute dynamic eth0
      valid_lft 76391sec preferred_lft 76391sec
   inet6 fe80::f816:                /64 scope link
      valid_lft forever preferred_lft forever
  ```

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

  Figure 5-26 IPv6 enabled and an IPv6 address assigned

  ```
  eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
   link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
   inet                  brd                  scope global noprefixroute dynamic eth0
      valid_lft 86395sec preferred_lft 86395sec
   inet6 2407:c080:802:                                /128 scope global dynamic
      valid_lft 7496sec preferred_lft 7196sec
   inet6 fe80::f816:3eff:          /64 scope link noprefixroute
      valid_lft forever preferred_lft forever
  ```

NOTE

IPv6 is enabled for Linux public images by default, as shown in **Figure 5-25**. IPv6 and dynamic IPv6 address assignment are enabled for Ubuntu 16 public images by default, as shown in **Figure 5-26**.

Step 2 Enable IPv6 for the ECS.

1. Run the following command to check whether IPv6 is enabled for the kernel:

   sysctl -a | grep ipv6

- If a command output is displayed, IPv6 is enabled.
- If no information is displayed, IPv6 is disabled. Go to **Step 2.2** to load the IPv6 module.

2. Run the following command to load the IPv6 module:

   modprobe ipv6

3. Add the following content to the **/etc/sysctl.conf** file:

   net.ipv6.conf.all.disable_ipv6=0

4. Save the configuration and exit. Then, run the following command to load the configuration:

   sysctl -p

Step 3 Enable dynamic IPv6 address assignment for the ECS.

- Ubuntu

  NOTE

  For Ubuntu 18.04 and 20.04, perform the following steps. For Ubuntu 16.04, skip these steps because dynamic IPv6 address assignment is enabled by default.

  a. Run the following command to access **/etc/netpaln/**:

     cd /etc/netplan

  b. Run the following command to list the configuration file:

     ls

     Figure 5-27 Configuration file name

     

  c. Run the following command to edit the configuration file:

     vi 01-network-manager-all.yaml

  d. Append the following content to the configuration file (pay attention to the yaml syntax and text indentation):

     ```
     ethernets:
      eth0:
       dhcp6: true
     ```

     Figure 5-28 Edited configuration file

     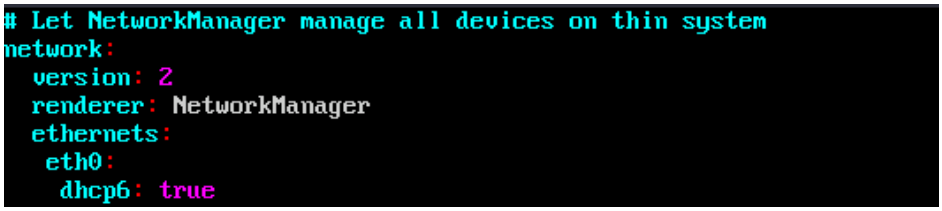

     Save the changes and exit.

  e. Run the following command to make the changes take effect:

     sudo netplan apply

- Debian

  a. Add the following content to the **/etc/network/interfaces** file:


```
auto lo
iface lo inet loopback
auto eth0
iface eth0 inet dhcp
iface eth0 inet6 dhcp
   pre-up sleep 3
```

b. Add configurations for each NIC to the **/etc/network/interfaces** file. The following uses eth1 as an example:

```
auto eth1
iface eth1 inet dhcp
iface eth1 inet6 dhcp
   pre-up sleep 3
```

c. Run the following command to restart the network service:

service networking restart

NOTE

If no IPv6 address is assigned after the NICs are brought down and up, you can run this command to restart the network.

d. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

- CentOS, EulerOS, or Fedora

a. Open the configuration file **/etc/sysconfig/network-scripts/ifcfg-eth0** of the primary NIC.

Add the following configuration items to the file:

```
IPV6INIT=yes
DHCPV6C=yes
```

b. Edit the **/etc/sysconfig/network** file to add or modify the following line:

```
NETWORKING_IPV6=yes
```

c. For an ECS running CentOS 6, you need to edit the configuration files of its extension NICs. For example, if the extension NIC is eth1, you need to edit **/etc/sysconfig/network-scripts/ifcfg-eth1**.

Add the following configuration items to the file:

```
IPV6INIT=yes
DHCPV6C=yes
```

In CentOS 6.3, dhcpv6-client requests are filtered by **ip6tables** by default. So, you also need to add a rule allowing the dhcpv6-client request to the **ip6tables** file.

i. Run the following command to add the rule to **ip6tables**:

ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT

ii. Run the following command to save the rule in **ip6tables**:

service ip6tables save

Figure 5-29 Example command

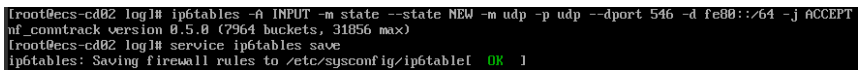

d. (Optional) For CentOS 7/CentOS 8, change the IPv6 link-local address mode of extension NICs to EUI64.

i. Run the following command to query the NIC information:

nmcli con

Figure 5-30 Querying NIC information

```
[root@ecs-166b ~]# nmcli con
NAME                UUID                                  TYPE      DEVICE
System eth0         5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03  ethernet  eth0
Wired connection 1  9c92fad9-6ecb-3e6c-eb4d-8a47c6f50c04  ethernet  eth1
Wired connection 1  3a73717e-65ab-93e8-b518-24f5af32dc0d  ethernet  eth2
```

ii. Run the following command to change the IPv6 link-local address mode of eth1 to EUI64:

nmcli con modify "*Wired connection 1***" ipv6.addr-gen-mode eui64**

NOTE

The NIC information varies depending on the CentOS series. In the command, *Wired connection 1* needs to be replaced with the value in the **NAME** column of the queried NIC information.

iii. Run the following commands to bring eth1 down and up:

ifdown eth1

ifup eth1

e. Restart the network service.

i. For CentOS 6, run the following command to restart the network service:

service network restart

ii. For CentOS 7/EulerOS/Fedora, run the following command to restart the network service:

systemctl restart NetworkManager

f. Perform **Step 1** to check whether dynamic IPv6 address assignment is enabled.

- SUSE, openSUSE, or CoreOS

  SUSE 11 SP4 does not support dynamic IPv6 address assignment.

  No additional configuration is required for SUSE 12 SP1 or SUSE 12 SP2.

  No additional configuration is required for openSUSE 13.2 or openSUSE 42.2.

  No additional configuration is required for CoreOS 10.10.5.

----End

Setting the Timeout Duration for IPv6 Address Assignment

Set the timeout duration depending on the OS type.

- CentOS 6.*x*:
  a. Run the following command to edit the **dhclient.conf** file:

     vi /etc/dhcp/dhclient.conf

  b. Press **i** to enter editing mode and add the timeout attribute to the file.

     ```
     timeout 30;
     ```

  c. Enter **:wq** to save the settings and exit.
- Debian 7.5:
  a. Run the following command to edit the **networking** file:

     vi /etc/init.d/networking

b. Press **i** to enter editing mode and add the timeout attribute.

Figure 5-31 Modification 1

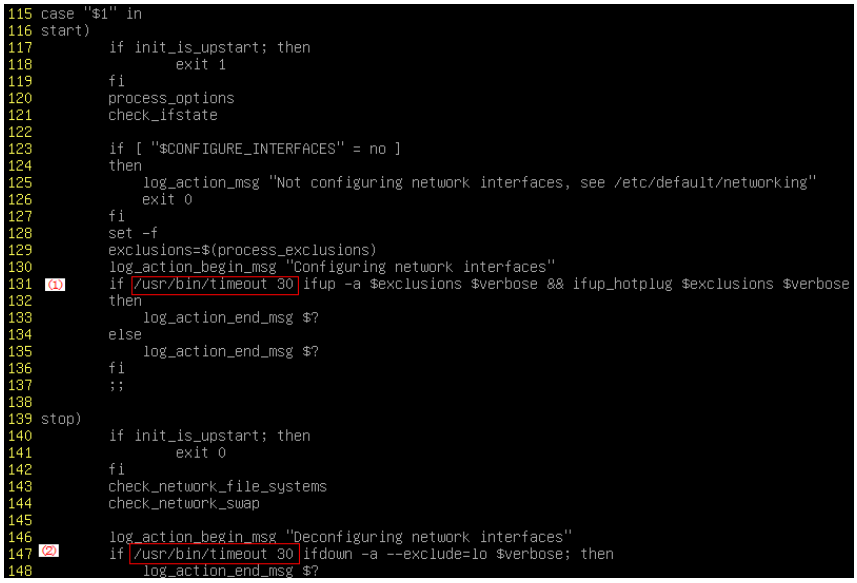

Figure 5-32 Modification 2

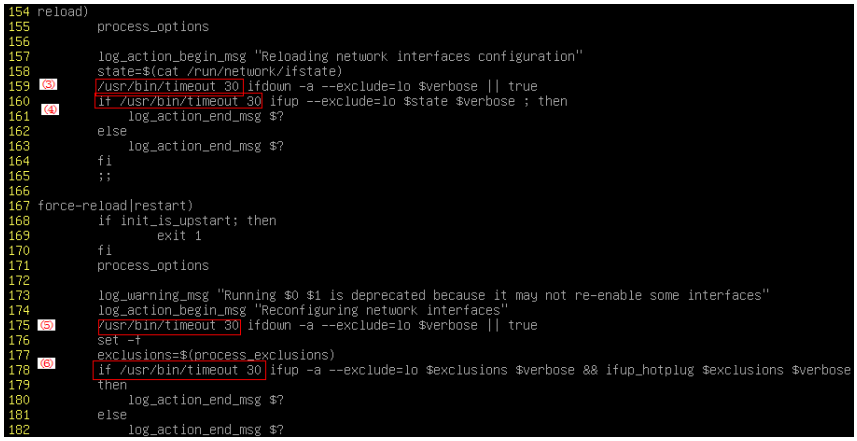

- Debian 8.2.0/8.8.0

  a. Run the following command to edit the **network-pre.conf** file:

  vi /lib/systemd/system/networking.service.d/network-pre.conf

  b. Press *i* to enter editing mode and add the timeout attribute to the file.

  ```
  [Service]
  TimeoutStartSec=30
  ```

- Debian 9.0

  a. Run the following command to edit the **networking.service** file:

  vi /etc/system/system/network-online.target.wants/networking.service

  b. Press **i** to enter editing mode and change **TimeoutStartSec=5min** to **TimeoutStartSec=30**.

6 EIPs

6.1 Overview

EIP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers. Various billing modes are provided to meet different service requirements.

Each EIP can be used by only one cloud resource at a time.

Figure 6-1 Accessing the Internet using an EIP

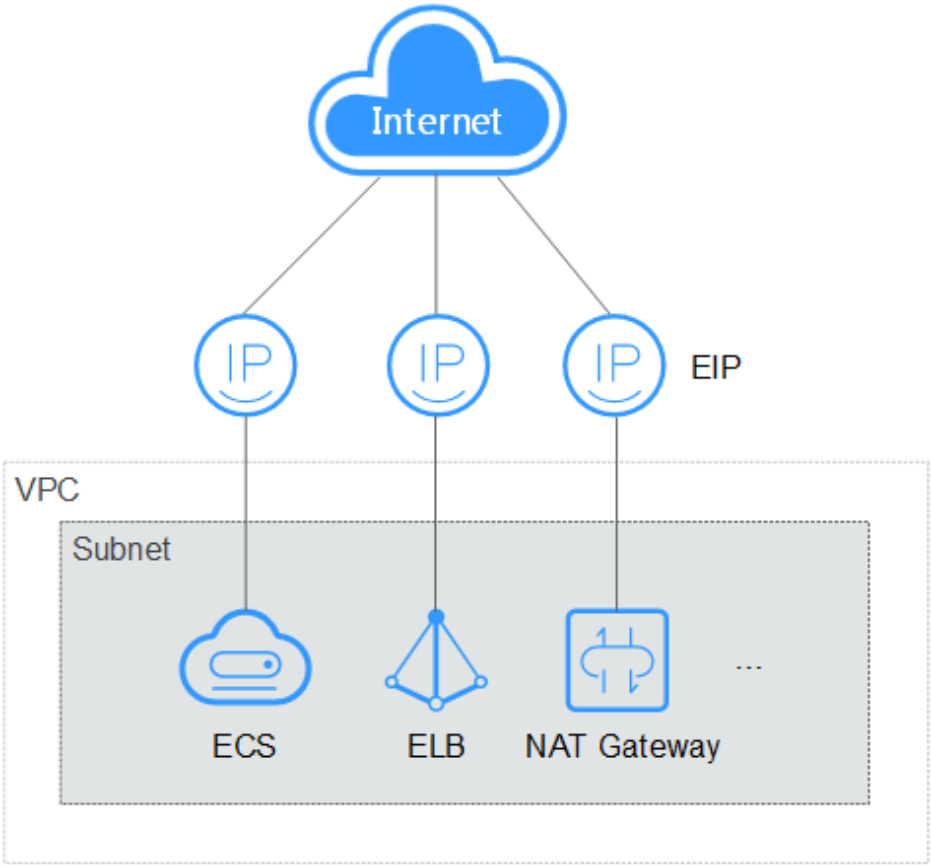

6.2 Binding an EIP

Scenarios

You can assign an EIP and bind it to an ECS to enable the ECS to access the Internet.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, select the ECS to which an EIP is to be bound, and choose **More** > **Manage Network** > **Bind EIP** in the **Operation** column.
5. In the displayed dialog box, select an EIP

   NOTE

   If no EIP is available in the current region, the EIP list is empty. In such a case, purchase an EIP and then bind it.

Figure 6-2 Binding an EIP

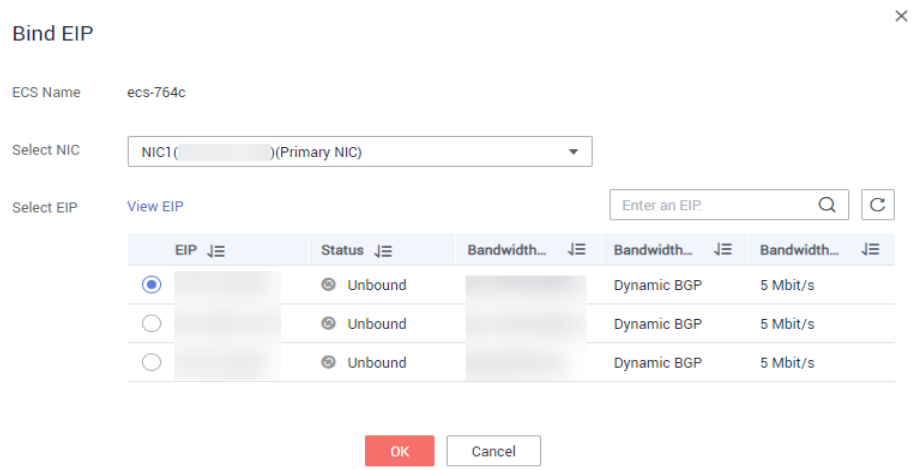

6. Click **OK**.

   After the EIP is bound, view it in the ECS list on the **Elastic Cloud Server** page.

6.3 Unbinding an EIP

Scenarios

This section describes how to unbind an EIP from an ECS.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Unbind EIP**.
5. Verify the EIP to be unbound and click **Yes**.

   NOTE

   Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

6.4 Changing an EIP

Scenarios

You can change the EIP bound to your ECS as needed.

The management console does not allow you to directly change the EIP bound to an ECS. Therefore, to change an EIP, unbind it from the ECS and bind the desired one to the ECS.

Restrictions

To avoid unintended actions, the system caches the EIP that you released for 24 hours. If you change the EIP within this period, the system preferentially assigns this EIP.

If you do not want to use the EIP that you released earlier, purchase another EIP first and then release the current one.

For details, see **What Is the EIP Assignment Policy?**

Prerequisites

An EIP has been assigned.

For details, see **Assigning an EIP**.

Unbinding an EIP

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Unbind EIP**.
4. Confirm the displayed information and click **Yes**.

   NOTE

   Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

Binding a New EIP

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Bind EIP**.
4. Select the desired EIP and click **OK**.

   NOTE

   If no EIP is available in the current region, the EIP list is empty. In such a case, purchase an EIP and then bind it.

Figure 6-3 Binding a new EIP

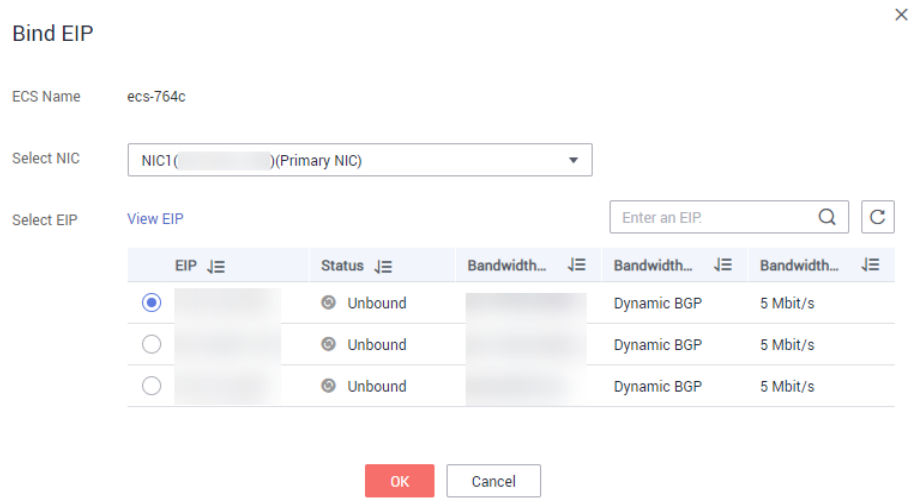

6.5 Changing an EIP Bandwidth

Scenarios

If an EIP has been bound to the ECS, the ECS can access the Internet using the bandwidth associated with the EIP. This section describes how to adjust the bandwidth of an ECS.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Modify Bandwidth**.
5. Change the bandwidth name, billing mode, or bandwidth size as prompted.

6.6 Enabling Internet Connectivity for an ECS Without an EIP

Scenarios

To ensure platform security and conserve EIPs, EIPs are assigned only to specified ECSs. ECSs without EIPs cannot access the Internet directly. If these ECSs need to access the Internet (for example, to perform a software upgrade or install a patch), you can select an ECS with an EIP bound to function as a proxy ECS, providing an access channel for these ECSs.

NOTE

NAT Gateway is recommended, which provides both the SNAT and DNAT functions for your ECSs in a VPC and allows the ECSs to access or provide services accessible from the Internet. For details, see **NAT Gateway**.

Prerequisites

- A proxy ECS with an EIP bound is available.
- The IP address of the proxy ECS is in the same network and same security group as the ECSs that need to access the Internet.

Linux Proxy ECS

In this example, the proxy ECS runs CentOS 6.5.

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
5. Click the name of the proxy ECS. The page providing details about the ECS is displayed.
6. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

Figure 6-4 Disabling source/destination check

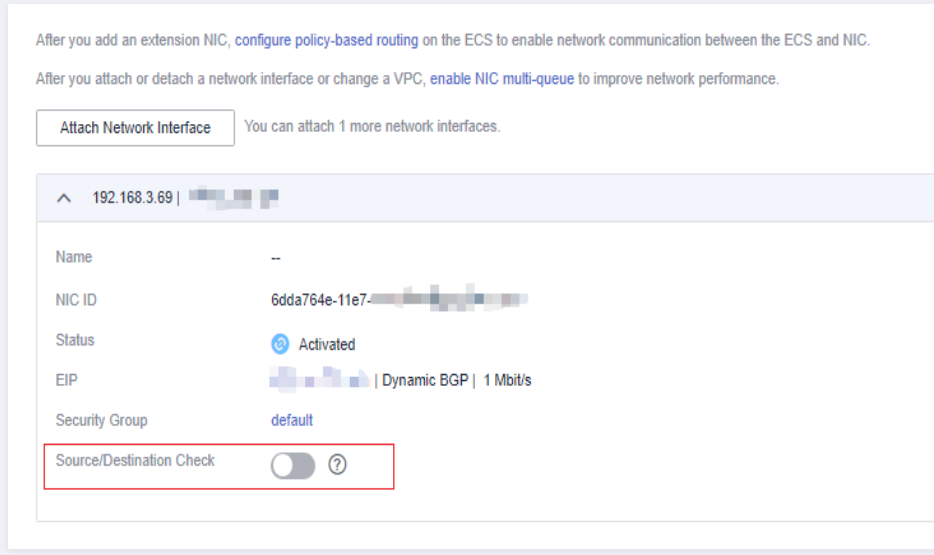

By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents

packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. Therefore, disable the source/destination check.

7. Log in to the proxy ECS.

   For more details, see **Login Overview**.

8. Run the following command to check whether the proxy ECS can access the Internet:

   ping www.huaweicloud.com

   The proxy ECS can access the Internet if information similar to the following is displayed:

   Figure 6-5 Checking connectivity

   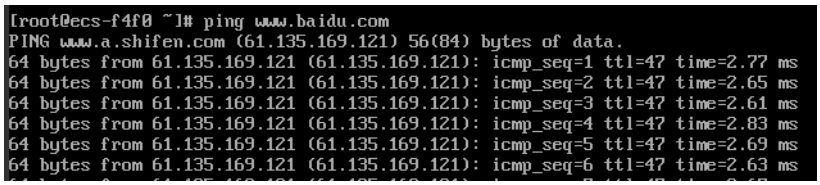

9. Run the following command to check whether IP forwarding is enabled on the proxy ECS:

   cat /proc/sys/net/ipv4/ip_forward

   - If **0** (disabled) is displayed, go to **10**.
   - If **1** (enabled) is displayed, go to **15**.

10. Run the following command to open the IP forwarding configuration file in the vi editor:

    vi /etc/sysctl.conf

11. Press **i** to enter editing mode.

12. Set the **net.ipv4.ip_forward** value to **1**.

    Set the **net.ipv4.ip_forward** value to **1**.

    NOTE

    If the **sysctl.conf** file does not contain the **net.ipv4.ip_forward** parameter, run the following command to add it:

    echo net.ipv4.ip_forward=1 >> /etc/sysctl.conf

13. Press **Esc**, type **:wq**, and press **Enter**.

    The system saves the configurations and exits the vi editor.

14. Run the following command to make the modification take effect:

    sysctl -p /etc/sysctl.conf

15. Run the following commands to configure default **iptables** rules:

    iptables -P INPUT ACCEPT

    iptables -P OUTPUT ACCEPT

    iptables -P FORWARD ACCEPT

CAUTION

Running **iptables -P INPUT ACCEPT** will set default INPUT policy to ACCEPT, which poses security risks. You are advised to set security group rules to restrict inbound access.

16. Run the following command to configure source network address translation (SNAT) to enable ECSs in the same network segment to access the Internet through the proxy ECS:

    iptables -t nat -A POSTROUTING -o eth0 -s *subnet/netmask-bits* **-j SNAT --to** *nat-instance-ip*

    For example, if the proxy ECS is in network 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

    iptables -t nat -A POSTROUTING -o eth0 -s *192.168.125.0/24* **-j SNAT --to 192.168.125.4**

    NOTE

    To retain the preceding configuration even after the ECS is restarted, run the **vi /etc/rc.local** command to edit the **rc.local** file. Specifically, copy the rule described in step 16 into **rc.local**, press **Esc** to exit Insert mode, and enter **:wq** to save the settings and exit.

17. Run the following commands to save the iptables configuration and make it start up automatically upon ECS startup:

    service iptables save

    chkconfig iptables on

18. Run the following command to check whether SNAT has been configured:

    iptables -t nat --list

    SNAT has been configured if information similar to **Figure 6-6** is displayed.

    Figure 6-6 Successful SNAT configuration

    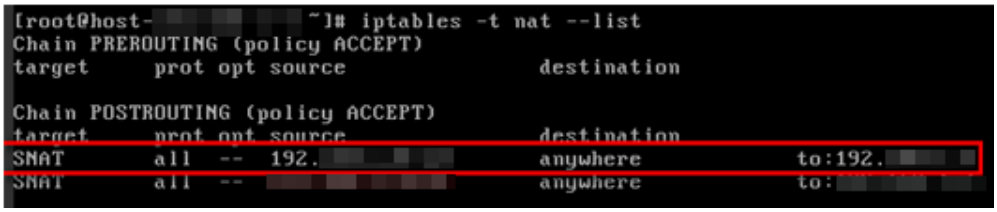

19. Add a route.
    a. Log in to the management console.
    b. Click in the upper left corner and select your region and project.
    c. Under **Networking**, click **Virtual Private Cloud**.
    d. Choose **Route Tables** in the left navigation pane. In the route table list, click a target route table. On the displayed page, click **Add Route**.
    e. Set route information on the displayed page.
        - **Destination**: indicates the destination network segment. The default value is **0.0.0.0/0**.

- **Next Hop**: indicates the private IP address of the proxy ECS.

  You can obtain the private IP address of the ECS on the **Elastic Cloud Server** page.

20. To delete the added iptables rules, run the following command:

    iptables -t nat -D POSTROUTING -o eth0 -s *subnet/netmask-bits* **-j SNAT --to** *nat-instance-ip*

    For example, if the proxy ECS is in network segment 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

    iptables -t nat -D POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4

7 Security

7.1 Methods for Improving ECS Security

Scenarios

If ECSs are not protected, they may be attacked by viruses, resulting in data leakage or data loss.

You can use the methods introduced below to protect your ECSs from viruses or attacks.

Protection Types

ECS can be protected externally and internally.

Table 7-1 Methods for improving ECS security

Type	Description	Protection Method
External security	DDoS attacks and Trojan horses or other viruses are common external security issues. To address these issues, you can choose services such as Host Security Service (HSS) based on your service requirements:	• **Enabling HSS** • **Monitoring ECSs** • **Enabling Anti-DDoS** • **Backing Up Data Periodically**
Internal security	Weak passwords and incorrect ports opening may cause internal security issues. Improving the internal security is the key to improving the ECS security. If the internal security is not improved, external security solutions cannot effectively intercept and block various external attacks.	• **Enhancing the Login Password Strength** • **Improving the Port Security** • **Periodically Upgrading the Operating System**

Enabling HSS

HSS is designed to improve the overall security for ECSs. It helps you identify and manage the information on your ECSs, eliminate risks, and defend against intrusions and web page tampering.

Before using the HSS service, install the HSS agent on your ECSs first so that your ECSs are protected by the HSS cloud protection center. You will be able to check the security statuses and risks (if any) of all ECSs in a region on the HSS console.

We provide different methods for you to install the HSS agent depending on whether your ECSs are to be created or already exist.

- **Scenario 1: An ECS is to be created.**

  When you use certain public images to create ECSs, you are advised to use HSS to protect your ECSs.

  Select one of the following options:

  - **HSS basic edition (free)**: provides HSS basic edition (1-month free trial), account cracking protection, weak password detection, and malicious program detection.

    NOTE

    After the free trial period expires, the HSS basic edition quotas will be automatically released, and HSS will not protect your servers.

    If you want to retain or upgrade HSS security capabilities, you are advised to purchase HSS. For details, see **Editions and Features**.

    This option is selected by default.
  - **Advanced HSS edition (paid)**: provides HSS enterprise edition, vulnerability patches, virus scan and removal, and graded protection.
  - **None**: Do not use security protection.

  After you enable HSS, the system automatically installs the HSS agent, enables account cracking prevention, and offers host security functions.

  HSS provides basic, enterprise, premium, and WTP editions. For details, see **Edition Details**.

  If the basic or enterprise edition does not meet service requirements, you can **Purchasing an HSS Quota** and switch the edition on the HSS console to obtain advanced protection without reinstalling the agent.

  Figure 7-1 Enabling HSS

  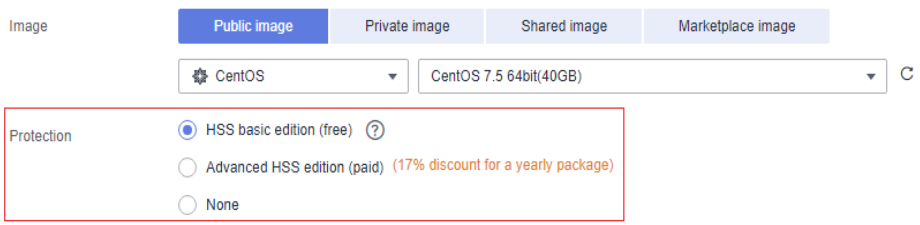

- **Scenario 2: An ECS is already created and HSS is not configured for it.**

  For an existing ECS without HSS configured, you can manually install an Agent on it.

For details, see **Installing an Agent on the Linux OS** and **Enabling Protection**.

Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

Server monitoring includes basic monitoring, OS monitoring, and process monitoring for servers.

- Basic monitoring

  Basic monitoring does not require the agent to be installed and automatically reports ECS metrics to Cloud Eye. Basic monitoring for KVM ECSs is performed every 5 minutes.

- OS monitoring

  By installing the Agent on an ECS, OS monitoring provides system-wide, active, and fine-grained monitoring. OS monitoring for KVM ECSs is performed every minute.

  To enable OS monitoring when purchasing an ECS:

  Select **Enable Detailed Monitoring** when purchasing an ECS. After this option is selected, the cloud platform automatically installs the agent required for OS monitoring.

  NOTE

  Currently, you can enable OS monitoring only when you purchase ECSs running specific OSs in specific regions.

  Figure 7-2 Enabling OS monitoring when purchasing an ECS

  

  To enable OS monitoring for a created ECS:

  You need to manually install the agent if **Enable Detailed Monitoring** is not selected during the creation.

  For instructions about how to install and configure the Agent, see **Agent Installation and Configuration**.

- Process monitoring

  Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. Processes are monitored at an interval of 1 minute (for KVM ECSs).

After server monitoring is enabled, you can set ECS alarm rules to customize the monitored objects and notification policies and learn about the ECS running status at any time.

On the ECS console, click  to view monitoring metrics.

Figure 7-3 Viewing ECS metrics

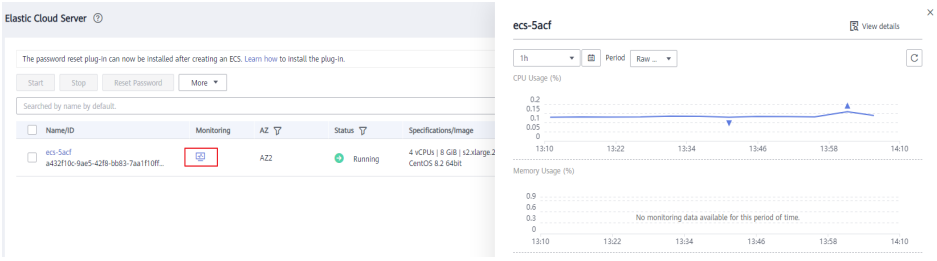

Enabling Anti-DDoS

To defend against DDoS attacks, HUAWEI CLOUD provides multiple security solutions. You can select an appropriate one based on your service requirements. Anti-DDoS Service on HUAWEI CLOUD provides three sub-services: Cloud Native Anti-DDoS (CNAD) Basic (also known as Anti-DDoS), CNAD Pro, and Advanced Anti-DDoS (AAD).

Anti-DDoS is free while CNAD Pro and AAD are paid services.

For details about CNAD Pro and AAD, see **What Is Anti-DDoS?**

If you choose to purchase an EIP when purchasing an ECS, the console will display a message indicating that you have enabled free-of-charge Anti-DDoS protection.

Figure 7-4 Enabling anti-DDoS protection

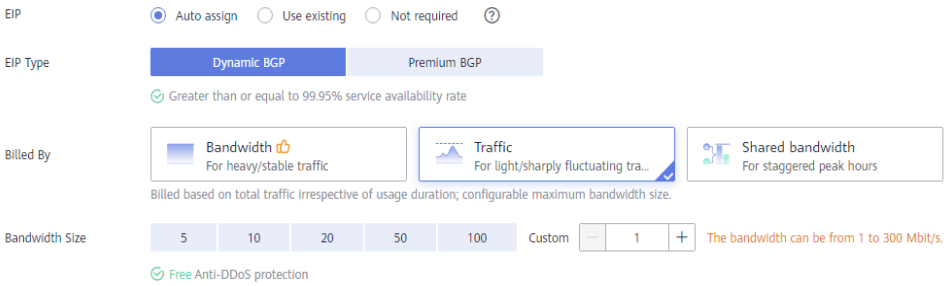

Anti-DDoS defends ECSs against DDoS attacks and sends alarms immediately when detecting an attack. In addition, Anti-DDoS improves the bandwidth utilization to further safeguard user services.

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.

Backing Up Data Periodically

Data backup is a process of storing all or part of data in different ways to prevent data loss. The following uses Cloud Backup and Recovery (CBR) as an example. For more backup methods, see **Overview**.

CBR enables you to back up ECSs and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up. CBR protects your services by ensuring the security and consistency of your data.

To enable CBR when purchasing an ECS:

Set CBR when purchasing an ECS. The system will associate the ECS with a cloud backup vault and the selected backup policy to periodically back up the ECS.

- Auto assign
  a. Set the name of the cloud backup vault, which is a character string consisting of 1 to 64 characters, including letters, digits, underscores (_), and hyphens (-). For example, **vault-f61e**. The default naming rule is **vault**_*xxxx*.
  b. Enter the vault capacity, which is required for backing up the ECS. The vault capacity cannot be smaller than that of the ECS to be backed up. Its value ranges from the total capacity of the ECS to 10,485,760 in the unit of GB.
  c. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Use existing
  a. Select an existing cloud backup vault from the drop-down list.
  b. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Not required: Skip the CBR setting. If you require this function after purchasing the ECS, log in to the CBR console and bind the desired cloud backup vault to your ECS.

Figure 7-5 Setting CBR

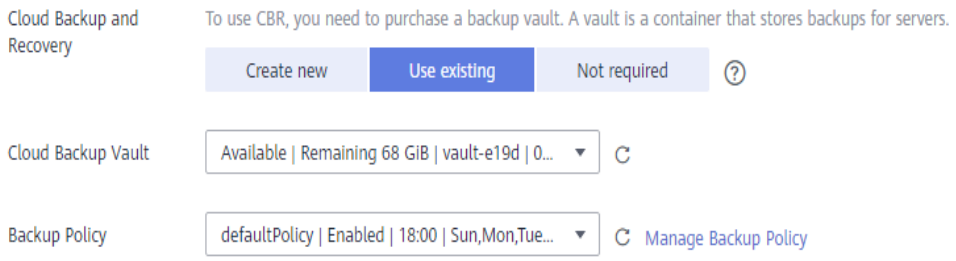

To back up data for a created ECS:

You can use the cloud server backup and cloud disk backup to **back up your ECS data**.

- Cloud server backup (recommended): Use this backup method if you want to back up the data of all EVS disks (system and data disks) attached to an ECS. This prevents data inconsistency caused by the time difference in creating a backup.
- Cloud disk backup: Use this backup method if you want to back up the data of one or more EVS disks (system or data disk) attached to an ECS. This minimizes backup costs on the basis of data security.

Enhancing the Login Password Strength

Key pair authentication is recommended because it is more secure than password-based authentication. If you select the password-based authentication, ensure that the password meets the strength requirements listed in **Table 7-2** to prevent malicious attacks.

The system does not periodically change the ECS password. It is recommended that you change your password regularly for security.

The password must conform to the following rules:

- The password must consist of at least 10 characters.
- Do not use easily guessed passwords (for example, passwords in common rainbow tables or passwords with adjacent keyboard characters). The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
- Do not include accounts in passwords, such as administrator, test, root, oracle, and mysql.
- Change the password at least every 90 days.
- Do not reuse the latest five passwords.
- Set different passwords for different applications. Do not use the same password for multiple applications.

Table 7-2 Password strength requirements

Parameter	Requirement	Example Value
Password	• Consists of 8 characters to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters: $!@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not use this example password.

Improving the Port Security

You can use security groups to protect the network security of your ECSs. A security group controls inbound and outbound traffic for your ECSs. Inbound traffic originates from the outside to the ECS, while outbound traffic originates from the ECS to the outside.

You can configure security group rules to grant access to or from specific ports. You are advised to disable high-risk ports and only enable necessary ports.

Table 7-3 lists common high-risk ports. You are advised to change these ports to non-high-risk ports. For details, see **Common Ports Used by ECSs**.

Table 7-3 Common high-risk ports

Protocol	Port
TCP	42, 135, 137, 138, 139, 444, 445, 593, 1025, 1068, 1434, 3127, 3128, 3129, 3130, 4444, 4789, 5554, 5800, 5900, and 9996
UDP	135 to 139, 1026, 1027, 1028, 1068, 1433, 1434, 4789, 5554, and 9996

Periodically Upgrading the Operating System

After ECSs are created, you need to maintain and periodically upgrade the operating system. The officially released vulnerabilities will be released in **Security Notices**.

7.2 Security Groups

7.2.1 Overview

Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and that are mutually trusted. After a security group is created, you can create various access rules for the security group, these rules will apply to all ECSs added to this security group.

You can also customize a security group or use the default one. The system provides a default security group for you, which permits all outbound traffic and denies inbound traffic. ECSs in a security group are accessible to each other. For details about the default security group, see **Default Security Group and Rules**.

NOTE

If two ECSs are in the same security group but in different VPCs, the security group does not take effect. You can use a VPC peering connection to connect the two VPCs first. For details, see **VPC Connectivity**.

Security Group Rules

After a security group is created, you can add rules to the security group. A rule applies either to inbound traffic (ingress) or outbound traffic (egress). After ECSs are added to the security group, they are protected by the rules of that group.

Each security group has default rules. For details, see **Default Security Group and Rules**. You can also customize security group rules. For details, see **Configuring Security Group Rules**.

Security Group Constraints

- For better network performance, you are advised to associate no more than five security groups to an instance.
- A security group can have no more than 6,000 instances associated, or performance will deteriorate.
- A security group can have up to 124 rules to be associated with IP address groups in one direction.
- If you specify an IP address group or inconsecutive ports for a security group rule, the rule takes effect only for certain ECSs. For details, see **Table 7-4**.

Table 7-4 Scenarios that security group rules do not take effect

Rule Configuration	ECS Type
Source or **Destination** is set to **IP address group**.	The following x86 ECS types are not supported: • General computing (S1, C1, and C2 ECSs) • Memory-optimized (M1 ECSs) • High-performance computing (H1 ECSs) • Disk-intensive (D1 ECSs) • GPU-accelerated (G1 and G2 ECSs) • Large-memory (E1, E2, and ET2 ECSs)
Port is set to non-consecutive ports.	The following x86 ECS types are not supported: • General computing (S1, C1, and C2 ECSs) • Memory-optimized (M1 ECSs) • High-performance computing (H1 ECSs) • Disk-intensive (D1 ECSs) • GPU-accelerated (G1 and G2 ECSs) • Large-memory (E1, E2, and ET2 ECSs)
	All Kunpeng ECS flavors do not support inconsecutive ports. If you use inconsecutive port numbers in a security group rule of a Kunpeng ECS, this rule and rules configured after this one do not take effect. If you configure security group rule A with inconsecutive ports **22,24** and then configure security group rule B with port 9096, both rule A and rule B do not take effect.

NOTE

- For details about x86 ECSs, see **ECS Specifications (x86)**.
- For details about Kunpeng ECSs, see **ECS Specifications (Kunpeng)**.

7.2.2 Default Security Group and Rules

If you have not created any security group, the system automatically creates a default security group for you and associates it with the instance (such as an ECS) when you create it. A default security group has the following rules:

- Inbound rules control incoming traffic to instances in a security group. Only instances in the same security group can communicate with each other, and all inbound requests are denied.
- Outbound rules allow all outbound traffic and response traffic to the outbound requests.

Figure 7-6 shows the default security group.

Figure 7-6 Default security group

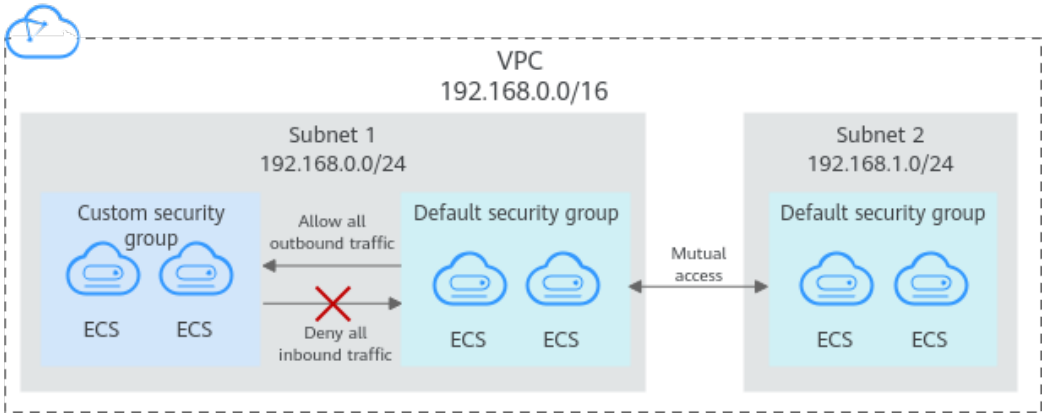

NOTE

- Both default and custom security groups are free of charge. The name of a default security group is **default**.
- You cannot delete the default security group, but you can modify existing rules or add rules to the group.
- The default security group is automatically created to simplify the process of creating an instance for the first time. The default security group denies all external requests. To log in to an instance, add a security group rule by referring to **Remotely Logging In to an ECS from a Local Server**.

Table 7-5 describes the rules in the default security group.

Table 7-5 Rules in the default security group

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	All	Source: Default security group (default)	This rule allows instances in the security group to communicate with each other.
Inbound	Allow	IPv6	All	Source: Default security group (default)	This rule allows instances in the security group to communicate with each other.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	This rule allows access from instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	This rule allows access from instances in the security group to any IPv6 address over any port.

When you create an ECS for the first time, the system automatically creates two security groups, **Sys-WebServer** and **Sys-FullAccess** with the newly created VPC **vpc-default**. The following table lists the default rules configured for the two security groups.

Table 7-6 Sys-WebServer security group rules

Transfer Direction	Protocol	Port Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-*xxxxx*)	Allows communication among ECSs within the security group and denies all inbound traffic (incoming data packets).
Inbound	TCP	22	Source: 0.0.0.0/0	Allows all IP addresses to access Linux ECSs through SSH.
Inbound	TCP	3389	Source: 0.0.0.0/0	Allows all IP addresses to access Windows ECSs through RDP.

Transfer Direction	Protocol	Port Range	Source/Destination	Description
Inbound	ICMP	All	Source: 0.0.0.0/0	Allows ping operations.
Inbound	TCP	443	Source: 0.0.0.0/0	Allows web page access through HTTPS.

Table 7-7 Sys-FullAccess security group rules

Transfer Direction	Protocol	Port Range	Source/Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-*xxxxx*)	Allows communication among ECSs within the security group and denies all inbound traffic (incoming data packets).
Inbound	All	All	Source: 0.0.0.0/0	Allows all inbound traffic.

7.2.3 Security Group Configuration Examples

Here are some common security group configuration examples for different scenarios, including remote login to ECSs, website access, and internal communication between instances in different security groups.

Generally, a security group denies all external requests by default. You need to add inbound rules to a security group based on the whitelist principle to allow specific external requests to access instances in the security group.

- **Remotely Logging In to an ECS from a Local Server**
- **Remotely Connecting to an ECS from a Local Server to Upload or Download Files**
- **Setting Up a Website on an ECS to Provide Services Externally**
- **Using ping Command to Verify Network Connectivity**
- **Enabling ECSs In Different Security Groups to Communicate Through an Internal Network**
- **ECS Providing Database Access Service**

- **Allowing ECSs to Access Only Specific External Websites**

By default, all outbound rules of a security group allow all requests from instances in the security group to access external networks. **Table 7-8** lists the rules.

Table 7-8 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	100	Allow	IPv4	All	0.0.0.0/0	This rule allows access from instances in the security group to any IPv4 address over any port.
Outbound	100	Allow	IPv6	All	::/0	This rule allows access from instances in the security group to any IPv6 address over any port.

Remotely Logging In to an ECS from a Local Server

A security group denies all external requests by default. To remotely log in to an ECS from a local server, add an inbound security group rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable the SSH (22) port. For details, see **Table 7-9**.
- To remotely log in to a Windows ECS using RDP, enable the RDP (3389) port. For details, see **Table 7-10**.

Table 7-9 Remotely logging in to a Linux ECS using SSH

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 22	IP address: 0.0.0.0/0

Table 7-10 Remotely logging in to a Windows ECS using RDP

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3389	IP address: 0.0.0.0/0

NOTICE

If the source is set to 0.0.0.0/0, remotely logging in to the ECS through any IP address is allowed. To ensure security, set the source to a specific IP address based on service requirements. For details about the configuration example, see **Table 7-11**.

Table 7-11 Remotely logging in to an ECS using a specified IP address

ECS Type	Direction	Priority	Action	Type	Protocol & Port	Source
Linux ECS	Inbound	1	Allow	IPv4	TCP: 22	IP address: 192.168.0.0/24
Windows ECS	Inbound	1	Allow	IPv4	TCP: 3389	IP address: 10.10.0.0/24

Remotely Connecting to an ECS from a Local Server to Upload or Download Files

By default, a security group denies all external requests. If you need to remotely connect to an ECS from a local server to upload or download files, you need to enable FTP ports 20 and 21.

Table 7-12 Remotely connecting to an ECS from a local server to upload or download files

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 0.0.0.0/0

NOTICE

You must first install the FTP server program on the ECSs and check whether ports 20 and 21 are working properly.

Setting Up a Website on an ECS to Provide Services Externally

A security group denies all external requests by default. If you have set up a website on an ECS that can be accessed externally, you need to add an inbound rule to the ECS security group to allow access over specific ports, such as HTTP (80) and HTTPS (443).

Table 7-13 Setting up a website on an ECS to provide services externally

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 80	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv4	TCP: 443	IP address: 0.0.0.0/0

Using ping Command to Verify Network Connectivity

By default, a security group denies all external requests. If you need to run the **ping** command on an ECS to verify network connectivity, add an inbound rule to the ECS security group to allow access over the ICMP port.

Table 7-14 Using **ping** command to verify network connectivity

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	ICMP: All	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv6	ICMP: All	IP address: ::/0

Enabling ECSs In Different Security Groups to Communicate Through an Internal Network

ECSs in the same VPC but associated with different security groups cannot communicate with each other. If you want to share data between ECSs in a VPC, for example, ECSs in security group sg-A need to access MySQL databases in security group sg-B, you need to add an inbound rule to security group sg-B to allow access from ECSs in security group sg-A over MySQL port 3306.

Table 7-15 Enabling instances in different security groups to communicate through an internal network

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A

ECS Providing Database Access Service

A security group denies all external requests by default. If you have deployed the database service on an ECS and need to allow other ECSs to access the database service through an internal network, you need to add an inbound rule to the security group of the ECS with the database service deployed to allow access over ports, for example, MySQL (3306), Oracle (1521), MS SQL (1433), PostgreSQL (5432) and Redis (6379).

Table 7-16 ECS providing database access service

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A	This rule allows ECSs in security group sg-A to access the MySQL database service.
Inbound	1	Allow	IPv4	TCP: 1521	Security group: sg-B	This rule allows ECSs in security group sg-B to access the Oracle database service.
Inbound	1	Allow	IPv4	TCP: 1433	IP address: 172.16.3.21/32	This rule allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service.
Inbound	1	Allow	IPv4	TCP: 5432	IP address: 192.168.0.0/24	This rule allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database service.
Inbound	1	Allow	IPv4	TCP: 6379	IP address group: ipGroup-A	This rule allows ECSs whose private IP addresses belong to IP address group ipGroup-A to access the PostgreSQL database service.

NOTICE

In this example, the source is for reference only. Set the source address based on actual requirements.

Allowing ECSs to Access Only Specific External Websites

By default, a security group allows all outbound traffic. **Table 7-18** lists the default rules. If you want to allow ECSs to access only specific websites, configure the security groups of the ECSs as follows:

1. First, add outbound rules to allow traffic over specific ports and to specific IP addresses.

Table 7-17 Enabling instances in different security groups to communicate through an internal network

Direction	Priority	Action	Type	Protocol & Port	Source
Outbound	1	Allow	IPv4	TCP: 80	IP address: 132.15.XX.XX
Outbound	1	Allow	IPv4	TCP: 443	IP address: 145.117.XX.XX

2. Then, delete the original outbound rules that allow all traffic shown in **Table 7-18**.

Table 7-18 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	100	Allow	IPv4	All	0.0.0.0/0	This rule allows access from instances in the security group to any IPv4 address over any port.
Outbound	100	Allow	IPv6	All	::/0	This rule allows access from instances in the security group to any IPv6 address over any port.

7.2.4 Configuring Security Group Rules

Scenarios

Similar to firewall, a security group is a logical group used to control network access. You can define access rules for a security group to protect the ECSs that are added to this security group.

- Inbound: Inbound rules allow external network traffic to be sent to the ECSs in the security group.
- Outbound: Outbound rules allow network traffic from the ECSs in the security group to be sent out of the security group.

For details about the default security group rules, see **Default Security Groups and Security Group Rules**. For details about configuration examples for security group rules, see **Security Group Configuration Examples**.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **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 details page.
7. Configure required parameters.

   You can click **+** to add more inbound rules.

Figure 7-7 Add Inbound Rule

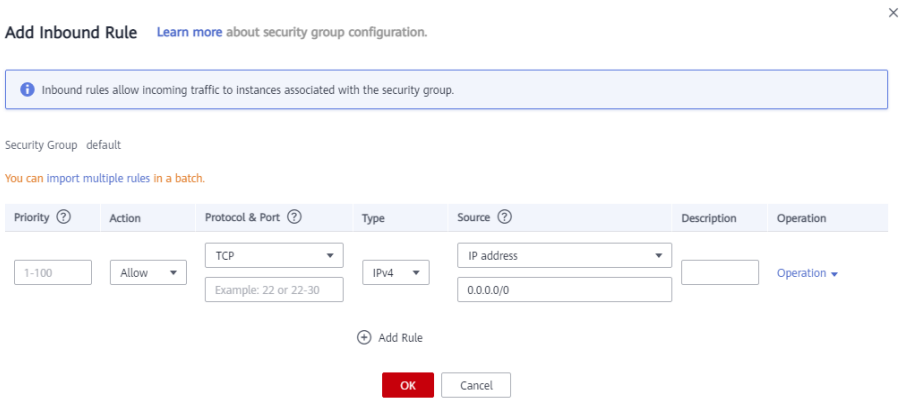

Table 7-19 Inbound rule parameter description

Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1

Parameter	Description	Example Value
Action	Allow or Deny • If the **Action** is set to **Allow**, access from the source is allowed to ECSs in the security group over specified ports. • If the **Action** is set to **Deny**, access from the source is denied to ECSs in the security group over specified ports. Deny rules take precedence over allow rules of the same priority.	Allow
Type	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be **All**, **TCP**, **UDP**, or **ICMP**, or others.	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Inbound rules control incoming traffic over specific ports to instances in the security group. Specify one of the following: • Individual port: Enter a port, such as 22. • Consecutive ports: Enter a port range, such as 22-30. • Non-consecutive ports: Enter ports and port ranges, such as **22,23-30**. You can enter a maximum of 20 ports and port ranges. Each port range must be unique. • All ports: Leave it empty or enter 1-65535.	22, or 22-30

Parameter	Description	Example Value
Source	The source in an inbound rule is used to match the IP address or address range of an external request. The source can be: • **IP address**: – Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128 – IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64 – All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses. • **Security group**: The source is from another security group. You can select a security group in the same region under the current account from the drop-down list. Instance A is in security group A and instance B is in security group B. If security group A has an inbound rule with **Action** set to **Allow** and **Source** set to security group B, access from instance B is allowed to instance A. • **IP address group**: An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.	IP address: 0.0.0.0/0
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

8. Configure required parameters.

   You can click **+** to add more outbound rules.

Figure 7-8 Add Outbound Rule

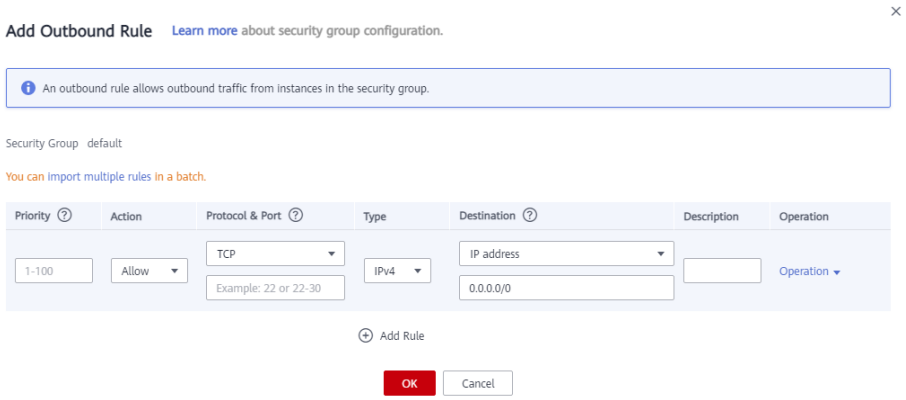

Table 7-20 Outbound rule parameter description

Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	Allow or Deny • If the **Action** is set to **Allow**, access from ECSs in the security group is allowed to the destination over specified ports. • If the **Action** is set to **Deny**, access from ECSs in the security group is denied to the destination over specified ports. Deny rules take precedence over allow rules of the same priority.	Allow
Type	Destination IP address version. You can select: • IPv4 • IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be **All**, **TCP**, **UDP**, or **ICMP**, or others.	TCP

Parameter	Description	Example Value
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Outbound rules control outgoing traffic over specific ports from instances in the security group. Specify one of the following: ● Individual port: Enter a port, such as 22. ● Consecutive ports: Enter a port range, such as 22-30. ● Non-consecutive ports: Enter ports and port ranges, such as **22,23-30**. You can enter a maximum of 20 ports and port ranges. Each port range must be unique. ● All ports: Leave it empty or enter 1-65535.	22, or 22-30
Destination	The destination in an outbound rule is used to match the IP address or address range of an internal request. The destination can be: ● **IP address** – Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128 – IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64 – All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses. ● **Security group**: The destination is from another security group. You can select a security group in the same region under the current account from the drop-down list. Instance A is in security group A and instance B is in security group B. If security group A has an outbound rule with **Action** set to **Allow** and **Destination** set to security group B, access from instance A is allowed to instance B. ● **IP address group**: An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.	IP address: 0.0.0.0/0

Parameter	Description	Example Value
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

9. Click **OK** to complete the security rule configuration.

Verifying Security Group Rules

After required security group rules are added, you can verify whether the rules take effect. For example, if you have deployed a website on an ECS and want users to access your website through HTTP (80), you need to add an inbound rule to the ECS security group to allow access over the port. **Table 7-21** shows the rule.

Table 7-21 Security group rule

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 80	0.0.0.0/0

Linux ECS

Check whether the security group rule takes effect on a Linux ECS:

1. Log in to the ECS.
2. Run the following command to check whether TCP port 80 is being listened on:

   netstat -an | grep 80

   If information similar to **Figure 7-9** is displayed, TCP port 80 is enabled.

   Figure 7-9 Command output for the Linux ECS

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

3. Enter **http://***ECS EIP* in the address box of the browser and press **Enter**.

   If the requested page can be accessed, the security group rule has taken effect.

Windows ECS

To verify the security group rule on a Windows ECS:

1. Log in to the ECS.
2. Choose **Start** > **Run**. Type cmd to open the Command Prompt.
3. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | findstr 80

If information similar to **Figure 7-10** is displayed, TCP port 80 is enabled.

Figure 7-10 Command output for the Windows ECS

```
TCP    0.0.0.0:80             0.0.0.0:0              LISTENING
```

4. Enter **http://***ECS EIP* in the address box of the browser and press **Enter**.

   If the requested page can be accessed, the security group rule has taken effect.

7.2.5 Changing a Security Group

Scenarios

To change the security group associated with an ECS network interface, perform the operations described in this section.

Constraints

- Changing the security group will overwrite the original security group settings.
- Using multiple security groups may deteriorate ECS network performance. You are advised to select no more than five security groups.

Procedure

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network** > **Change Security Group**.

   The **Change Security Group** dialog box is displayed.

Figure 7-11 Change Security Group

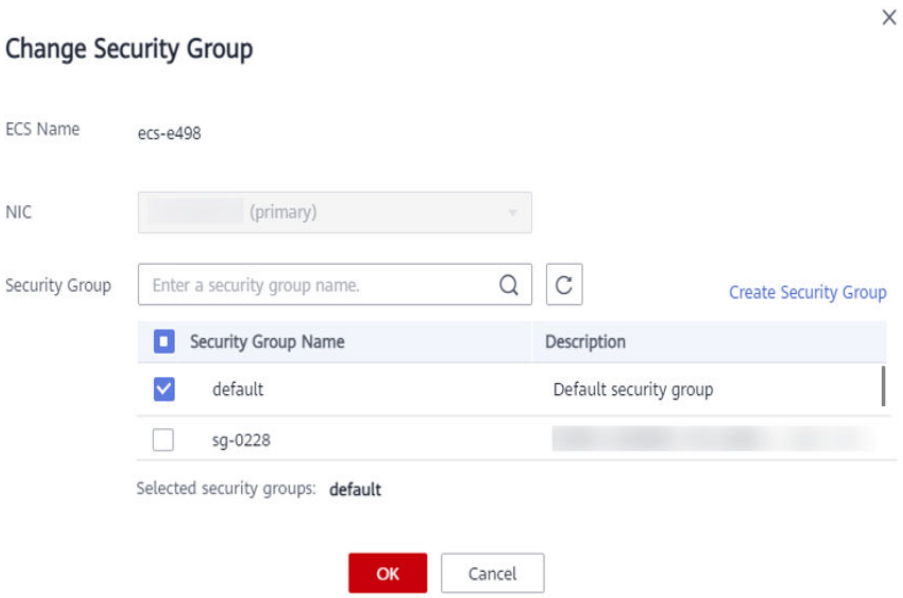

4. Select the target NIC and security groups.

   You can select multiple security groups. In such a case, the rules of all the selected security groups will be aggregated to apply on the ECS.

   To create a security group, click **Create Security Group**.

   NOTE

   Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click **OK**.

7.3 HSS

What Is HSS?

Host Security Service (HSS) is designed to improve the overall security for ECSs. It helps you identify and manage the information on your ECSs, eliminate risks, and defend against intrusions and web page tampering.

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.

How Do I Use HSS?

Before using the HSS service, install the HSS agent on your ECS. The installation method varies depending on whether your ECS is to be created or already exists.

- **Scenario 1: An ECS is to be created.**

  When you use certain public images to create ECSs, you are advised to use HSS to protect your ECSs.

  Select one of the following options:

  - **HSS basic edition (free)**: provides HSS basic edition (1-month free trial), account cracking protection, weak password detection, and malicious program detection.

    NOTE

    After the free trial period expires, the HSS basic edition quotas will be automatically released, and HSS will not protect your servers.

    If you want to retain or upgrade HSS security capabilities, you are advised to purchase HSS. For details, see **Editions and Features**.

    This option is selected by default.

  - **Advanced HSS edition (paid)**: provides HSS enterprise edition, vulnerability patches, virus scan and removal, and graded protection.
  - **None**: Do not use security protection.

  After you enable HSS, the system automatically installs the HSS agent, enables account cracking prevention, and offers host security functions.

  HSS provides basic, enterprise, premium, and WTP editions. For details, see **Edition Details**.

If the basic or enterprise edition does not meet service requirements, you can **Purchasing an HSS Quota** and switch the edition on the HSS console to obtain advanced protection without reinstalling the agent.

Figure 7-12 Enabling HSS

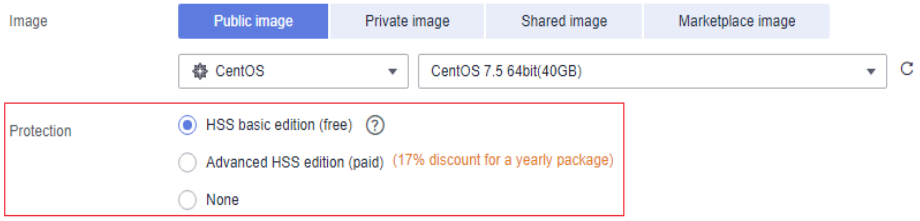

- **Scenario 2: An ECS is already created and HSS is not configured for it.**

  For an existing ECS without HSS configured, you can manually install an Agent on it.

  For details, see **Installing an Agent on the Linux OS** and **Enabling Protection**.

How Do I Check Host Security Statuses?

On the **Server** tab, you can view the ECS security statuses in the current region.

1. Log in to the management console.
2. Click  and choose **Security & Compliance** > **Host Security Service**.
3. On the **Server** tab, check the ECS security statuses.

Figure 7-13 ECS security statuses

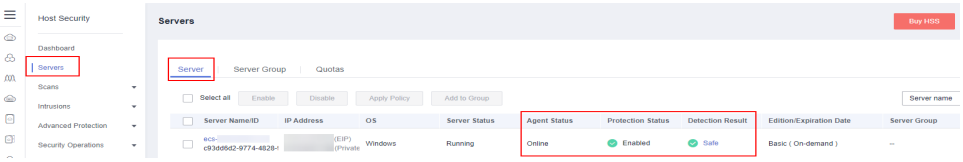

Table 7-22 Statuses

Parameter	Description
Agent Status	• **Not installed**: The agent has not been started or even has not been installed. • **Online**: The agent is running properly. • **Offline**: The agent fails to communicate with the HSS server. Therefore, HSS cannot protect your ECS. Click **Offline**. Then, the ECSs with agent being offline and the offline reasons are displayed.
Protection Status	• **Enabled**: The ECS is properly protected using HSS. • **Disabled**: HSS has been disabled on the ECS. If an ECS does not need protection, disable HSS on it to reduce its resource consumption.

Parameter	Description
Detection Result	• **Risky**: The ECS is risky. • **Safe**: No risks are detected. • **Pending risk detection**: HSS is not enabled for the ECS.

For more details, see **What Is HSS?**

7.4 Project and Enterprise Project

Creating a Project and Assigning Permissions

- **Creating a project**

  Log in to the management console, click the username in the upper right corner, and select **Identity and Access Management** from the drop-down list box. In the navigation pane on the left, choose **Projects**. In the right pane, click **Create Project**. On the displayed **Create Project** page, select a region and enter a project name.

- **Assigning permissions**

  You can assign permissions (of resources and operations) to user groups to associate projects with user groups. You can add users to a user group to control projects that users can access and the resources on which users can perform operations. To do so, perform the following operations:

  a. On the **User Groups** page of the IAM console, locate the target user group and click **Authorize** in the **Operation** column.

  b. Select policies or roles from the list.

  c. Click **Next** and select **Region-specific projects**.

  d. In the displayed regional project list, select one or more projects and click **OK**.

  e. On the **Users** page, locate the target user and click **Authorize** in the **Operation** column.

  f. Select **Inherit permissions from user groups** and select the user group authorized in **a**.

  g. Click **OK**.

Creating an Enterprise Project and Assigning Permissions

- **Creating an enterprise project**

  On the management console, choose **Enterprise** > **Project Management** in the upper right corner. On the **Enterprise Project Management** console, click **Create Enterprise Project**.

  NOTE

  Enterprise is available on the management console only if you have enabled the enterprise project, or your account is the primary account. To enable this function, contact customer service.

- **Assigning permissions**

  You can add a user group to an enterprise project and configure a policy to associate the enterprise project with the user group. You can add users to a user group to control projects users can access and the resources on which users can perform operations. To do so, perform the following operations:

  a. On the **Enterprise Management** console, click the name of an enterprise project to go to the enterprise project details page.

  b. On the **Permissions** tab, click **Authorize User Group** to go to the **User Groups** page on the IAM console. Associate the enterprise project with a user group and assign permissions to the group.

     For details, see **Creating a User Group and Assigning Permissions**.

- **Associating ECSs with enterprise projects**

  You can use enterprise projects to manage cloud resources.

  – Select enterprise projects when purchasing ECSs.

    On the page for buying an ECS, select an enterprise project from the **Enterprise Project** drop-down list.

  – Add ECSs to an enterprise project.

    On the **Enterprise Project Management** page, you can add existing ECSs to an enterprise project.

    Value **default** indicates the default enterprise project. Resources that are not allocated to any enterprise projects under your account are displayed in the default enterprise project.

For details, see **Enterprise Management User Guide**.

7.5 Protection for Mission-Critical Operations

Scenarios

ECS protects against mission-critical operations. If you want to perform a mission-critical operation on the management console, you must enter a credential for identity verification. You can perform the operation only after your identity is verified. For account security, it is a good practice to enable operation protection. The setting will take effect for both the account and users under the account.

The following operations can be protected: Stop, restart, or delete an ECS; reset the password for logging in to an ECS; detach a disk from an ECS; unbind an EIP from an ECS.

Enabling Operation Protection

Operation protection is disabled by default. Perform the following operations to enable it:

1. Log in to the management console.
2. Move the cursor to the username in the upper right corner of the page and select **Security Settings** from the drop-down list.

Figure 7-14 Security Settings

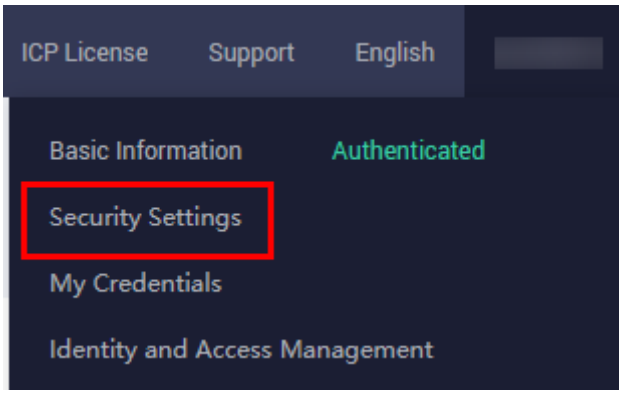

3. On the **Security Settings** page, choose **Critical Operations** > **Operation Protection** > **Enable**.

Figure 7-15 Critical Operations

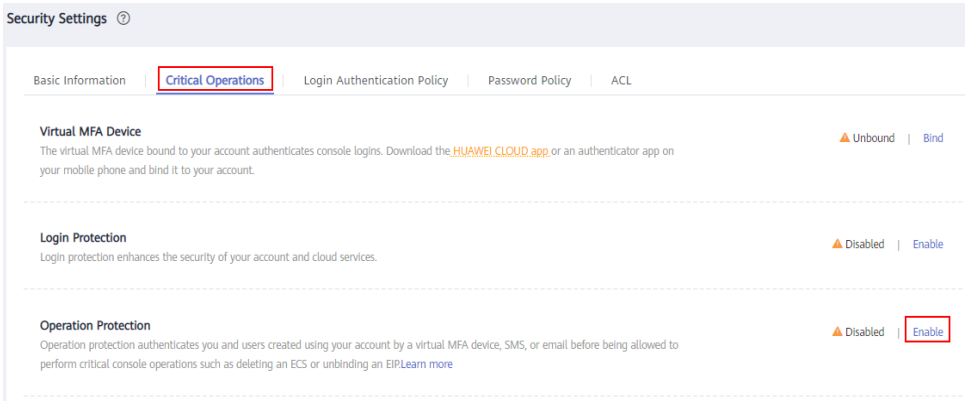

4. On the **Operation Protection** page, select **Enable** to enable operation protection.

   When you or the IAM users under your account perform critical operations, for example, deleting ECS resources, you are required to enter a verification code based on the selected verification method.

   NOTE

   - When performing a critical operation, you will be asked to choose a verification method from email, SMS, and virtual MFA device.
     - If you have bound only a mobile number, only SMS verification is available.
     - If you have bound only an email address, only email verification is available.
     - If you have not bound an email address, mobile number, or virtual MFA device, you are required to bind one to continue with the critical operation.
   - You can change the mobile number, email address, and virtual MFA device on the **Basic Information** page.

Verifying an Identity

After operation protection is enabled, when you perform a mission-critical operation, the system will verify your identity.

- If you have bound an email address, enter the email verification code.

- If you have bound a mobile number, enter the SMS verification code.
- If you have bound a virtual MFA device, enter a 6-digit dynamic verification code of the MFA device.

When you attempt to stop an ECS, select a verification method.

Figure 7-16 Identity Verification

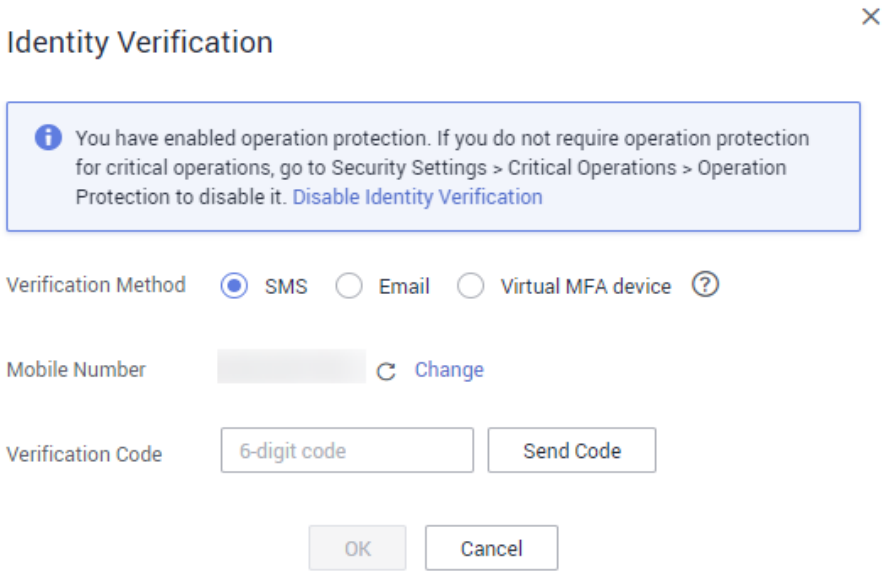

Disabling Operation Protection

Perform the following operations to disable operation protection.

1. Log in to the management console.
2. Move the cursor to the username in the upper right corner of the page and select **Security Settings** from the drop-down list.

   Figure 7-17 Security Settings

   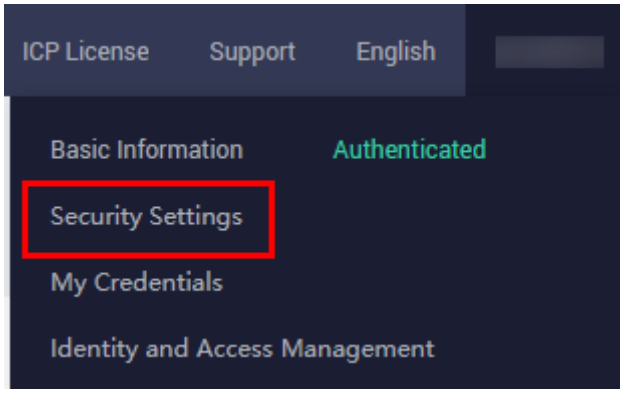

3. On the **Security Settings** page, choose **Critical Operations** > **Operation Protection** > **Change**.

Figure 7-18 Modifying operation protection settings

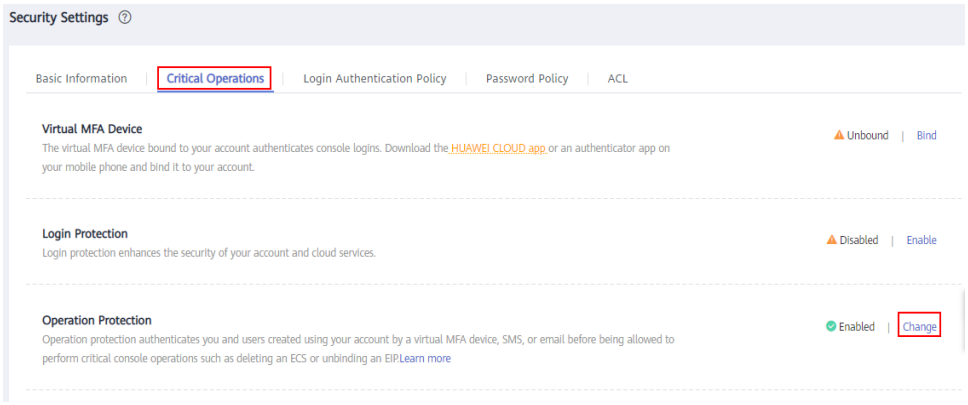

4. On the **Operation Protection** page, select **Disable** and click **OK**.

Helpful Links

- **How Do I Bind a Virtual MFA Device?**
- **How Do I Obtain a Virtual MFA Verification Code?**

8 Passwords and Key Pairs

8.1 Passwords

8.1.1 Application Scenarios for Using Passwords

The password for logging in to your ECS is important and please keep it secure. You can reset the password if it is forgotten or expires.

Table 8-1 provides guidance on how to reset your password in different scenarios.

Table 8-1 Resetting a password

Reference	Prerequisites
Resetting the Password for Logging In to an ECS on the Management Console	The password reset plug-in has been installed. **NOTE** • If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default. • The reference is for Windows or Linux ECSs.
Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed	The password reset plug-in has not been installed.
Resetting the Password for Logging In to a Linux ECS Without the Password Reset Plug-in Installed	The password reset plug-in has not been installed.

Background

Table 8-2 shows the ECS password complexity requirements.

Table 8-2 Password complexity requirements

Parameter	Requirement	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Windows: $!@%-_=+[]:./,? – Special characters for Linux: !@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not use it.

8.1.2 Resetting the Password for Logging In to an ECS on the Management Console

Scenarios

If you did not set a password when creating an ECS, or the password expires or is forgotten, reset the password by following the instructions provided in this section.

Prerequisites

- You have installed the password reset plug-in before your ECS password expires or is forgotten.
  - If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default.
  - If your ECS was created using a private image and has no password reset plug-in installed, see **Resetting the Password for Logging In to a Windows ECS** or **Resetting the Password for Logging In to a Linux ECS**.
- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will not be available.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP is enabled in the VPC to which the ECS belongs.
- The ECS network connectivity is normal.
- Ensure that the one-click password reset plug-in is not blocked by security software. Otherwise, the one-click password reset function is unavailable.

- After the password is reset, you must restart the ECS for the new password to take effect.

Procedure

Perform the following operations to change the login password of one or multiple ECSs in a batch on the management console.

NOTE

If you reset the password when the ECS is running, the new password takes effect only after the ECS is restarted. You can manually restart the ECS after resetting the password, or select **Auto Restart** when resetting the password.

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.

Figure 8-1 Reset Password

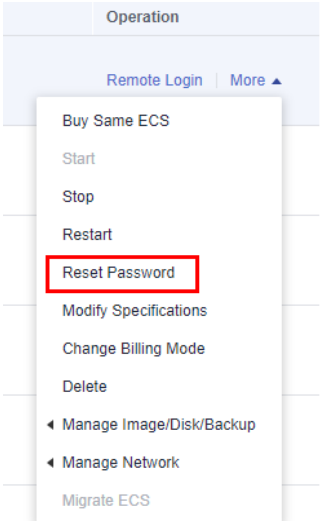

4. Set and confirm a new password as prompted.

NOTE

If the system displays a message indicating that the password cannot be reset, see **Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed** and **Resetting the Password for Logging In to a Linux ECS Without the Password Reset Plug-in Installed**.

The new password must meet the complexity requirements listed in **Table 8-3**.

Table 8-3 Password complexity requirements

Parameter	Requirement	Example Value
Password	• Consists of 8 to 26 characters. • Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Windows: $!@%-_=+[]:./,? – Special characters for Linux: !@%-_=+[]:./^,{}? • Cannot contain the username or the username spelled backwards. • Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)	YNbUwp!dUc9MClnv **NOTE** The example password is generated randomly. Do not use it.

5. Click **OK**.

   It takes about 10 minutes for the system to reset the password. Do not repeatedly perform this operation.

   – If the ECS is running when you reset the password, you need to manually restart the ECS for the new password to take effect.
   – If the ECS is stopped, the new password will take effect after you start the ECS.

8.2 One-Click Password Reset Plug-in

8.2.1 Obtaining the One-Click Password Reset Plug-in

Scenarios

If the password failed to be reset, this may be because the one-click password reset plug-in has not been installed. You can install the plug-in and verify its integrity following the instructions provided in this section.

Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Linux)

1. Log in to the ECS as user **root**.
2. Download the one-click password reset plug-in and SHA256 checksum.

   Obtain the download address from **Table 8-4** based on the region and OS (32- or 64-bit).

Example command to download the plug-in for a 32-bit x86 ECS in the CN North-Beijing1 region:

wget https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip

wget https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256

3. Obtain the hash value of your local one-click password reset plug-in.

   sha256sum *{Software package directory}***/CloudResetPwdAgent.zip**

   Replace ***Software package directory*** with the actual download directory.

4. Check whether the SHA256 hash value obtained in step 2 is consistent with that obtained in step 3.
   - If they are consistent, the verification is successful.
   - If they are inconsistent, download the one-click password reset plug-in of the corresponding version and repeat steps 2 to 4 to verify it.

Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Windows)

1. Log in to the ECS.
2. Download the one-click password reset plug-in and SHA256 checksum.

   Obtain the download address from Table 8-4 based on the region where the ECS resides.

3. Open Command Prompt as an administrator and run the following command to obtain the hash value of the local one-click password reset plug-in:

   certutil –hashfile {*Software package directory*}\CloudResetPwdAgent.zip SHA256

   Replace ***{Software package directory}*** with the actual download directory.

4. Check whether the SHA256 hash value obtained in step 2 is consistent with that obtained in step 3.
   - If they are consistent, the verification is successful.
   - If they are inconsistent, download the one-click password reset plug-in of the corresponding version and repeat steps 2 to 4 to verify it.

How to Obtain the One-Click Password Reset Plug-in and SHA256 Checksum

Table 8-4 Addresses for downloading the one-click password reset plug-in

Region	OS	Name	How to Obtain
CN North-Beijing1	Linux(x86_32)	CloudResetPwdAgent.zip	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip

Region	OS	Name	How to Obtain
		SHA256 checksum	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(x86_64)	CloudResetPwdAgent.zip	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(aarch64)	CloudResetPwdAgent.zip	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Windows	CloudResetPwdAgent.zip	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
CN North-Beijing4	Linux(x86_64)	CloudResetPwdAgent.zip	https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Windows	CloudResetPwdAgent.zip	https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256

Region	OS	Name	How to Obtain
CN East-Shanghai2	Linux(x86_32)	CloudResetPwdAgent.zip	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(x86_64)	CloudResetPwdAgent.zip	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(aarch64)	CloudResetPwdAgent.zip	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Windows	CloudResetPwdAgent.zip	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
CN South-Guangzhou	Linux(x86_32)	CloudResetPwdAgent.zip	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(x86_64)	CloudResetPwdAgent.zip	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256

Region	OS	Name	How to Obtain
	Linux(aarch64)	CloudResetPwdAgent.zip	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Windows	CloudResetPwdAgent.zip	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
CN-Hong Kong	Linux(x86_32)	CloudResetPwdAgent.zip	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(x86_64)	CloudResetPwdAgent.zip	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(aarch64)	CloudResetPwdAgent.zip	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256

Region	OS	Name	How to Obtain
	Windows	CloudResetPwdAgent.zip	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
AP-Bangkok	Linux(x86_32)	CloudResetPwdAgent.zip	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(x86_64)	CloudResetPwdAgent.zip	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Linux(aarch64)	CloudResetPwdAgent.zip	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip
		SHA256 checksum	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
	Windows	CloudResetPwdAgent.zip	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip

Region	OS	Name	How to Obtain
		SHA256 checksum	https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256

Related Operations

After obtaining the password reset plug-in, you can install the plug-in or update it as follows:

- **Installing the One-Click Password Reset Plug-in on an ECS**
- **Updating the One-Click Password Reset Plug-in for an ECS**

8.2.2 Installing the One-Click Password Reset Plug-in on an ECS

You can reset the password for logging in to an ECS with just a few clicks if you forgot the password or the password expires.

After you have created an ECS, it is a good practice to log in to it and install the password reset plug-in.

NOTE

The password reset plug-in has been installed on the ECSs created using a public image by default. To check whether the plug-in has been installed, see **Step 1**.

Notes

1. The password reset plug-in is not installed by default. You can determine whether to install it.
2. After the installation, do not uninstall the plug-in by yourself. Otherwise, the ECS password cannot be reset.
3. After you reinstall or change the OS of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
4. After you replace the system disk of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
5. The password reset plug-in cannot be installed on a CoreOS ECS.
6. To reset the password, the one-click password reset plug-in must be installed before the ECS password is lost or expires.
7. The one-click password reset plug-in can be installed only after an EIP is bound to the ECS.

Prerequisites

- The available space in drive C of a Windows ECS is greater than 300 MB, and data can be written to it.
- The available space in the root directory of a Linux ECS is greater than 300 MB, and data can be written to it.
- For Linux ECSs, **Disabling SELinux** if it has been enabled.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP is enabled in the VPC to which the ECS belongs.
- The ECS network connectivity is normal.
- Set the NIC to DHCP so that the ECS can dynamically obtain an IP address.

  NOTE

  For details about how to set the NIC to DHCP for a Linux ECS, see **Setting the NIC to DHCP**.

  For details about how to set the NIC to DHCP for a Windows ECS, see **Setting the NIC to DHCP**.
- The ECS security group rule in the outbound direction meets the following requirements:
  - **Protocol**: **TCP**
  - **Port Range**: **80**
  - **Remote End**: **169.254.0.0/16**

  If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:
  - **Protocol**: **ANY**
  - **Port Range**: **ANY**
  - **Remote End**: **0.0.0.0/16**

Installing the Password Reset Plug-in on a Linux ECS

Step 1 Use either of the following methods to check whether the password reset plug-in has been installed on the ECS:

Method 1: Use the management console for query.

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list.
   - If a dialog box is displayed, asking you to enter the new password, the password reset plug-in has been installed. No further action is required.

Figure 8-2 Information displayed if the password reset plug-in has been installed

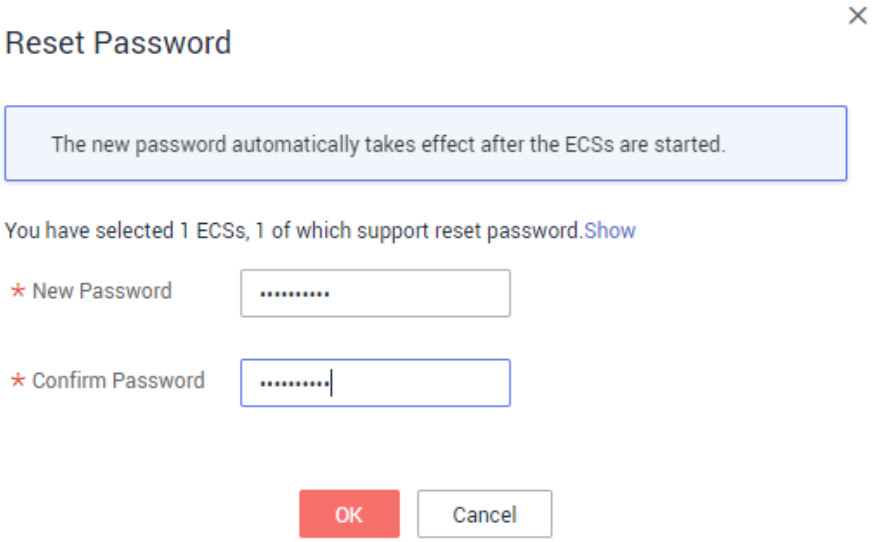

- If a dialog box is displayed, asking you to download a password reset script, the password reset plug-in has not been installed. Then, install it.

  Figure 8-3 Information displayed if the password reset plug-in has not been installed

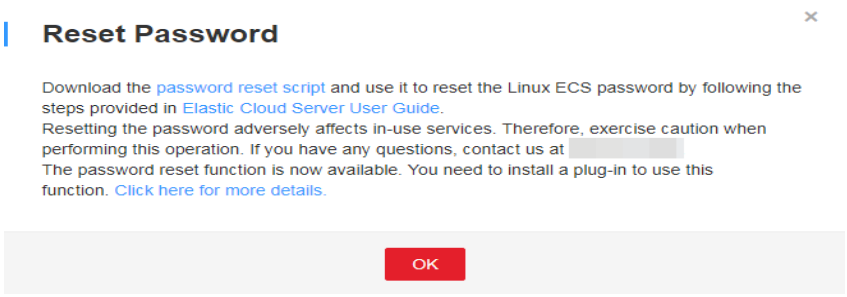

Method 2: Use the ECS for query.

1. Log in to the ECS as user **root**.
2. Run the following command to check whether CloudResetPwdAgent has been installed:

   ls -lh /Cloud*

   Figure 8-4 Checking whether the plug-in has been installed

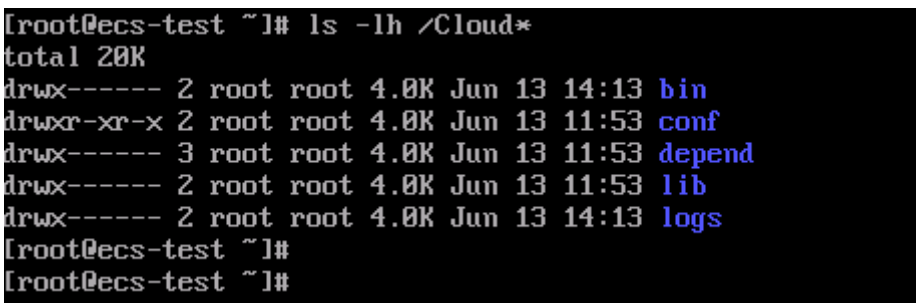

Check whether the obtained information is similar to that shown in **Figure 8-4**.

- If yes, the plug-in has been installed.
- If no, the plug-in has not been installed. Then, install it.

Step 2 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to **Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Linux)**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 3 Run the following command to decompress **CloudResetPwdAgent.zip**:

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

unzip -o -d *Decompressed directory* **CloudResetPwdAgent.zip**

An example is provided as follows:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

unzip -o -d /home/linux/test CloudResetPwdAgent.zip

Step 4 Install the one-click password reset plug-in.

1. Run the following command to open the **CloudResetPwdAgent.Linux** file:

   cd *{Plug-in decompressed directory}***/CloudResetPwdAgent/CloudResetPwdAgent.Linux**

   For example:

   If the plug-in is decompressed to **/home/linux/test**, run the following command:

   cd /home/linux/test/CloudResetPwdAgent/CloudResetPwdAgent.Linux

2. Run the following command to add the execute permission for the **setup.sh** file:

   chmod +x setup.sh

3. Run the following command to install the plug-in:

   sudo sh setup.sh

   If "cloudResetPwdAgent install successfully." is displayed and "Failed to start service cloudResetPwdAgent" is not displayed, the installation is successful.

   NOTE

   - You can also check whether the password reset plug-in has been installed using the methods provided in **Step 1**.
   - If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

Step 5 Modify the file permission of the password reset plug-in.

chmod 700 /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script

chmod 700 /CloudrResetPwdAgent/bin/wrapper

chmod 600 /CloudrResetPwdAgent/lib/*

----End

Installing the Password Reset Plug-in on a Windows ECS

Step 1 Log in to the ECS.

Step 2 Check whether the password reset plug-in CloudResetPwdAgent has been installed on the ECS. To check this, perform the following operations:

Start the **Task Manager** and check whether **cloudResetPwdAgent** is displayed on the **Services** tab. As shown in the **Figure 8-5**, the password reset plug-in has been installed on the ECS.

Figure 8-5 Successful plug-in installation

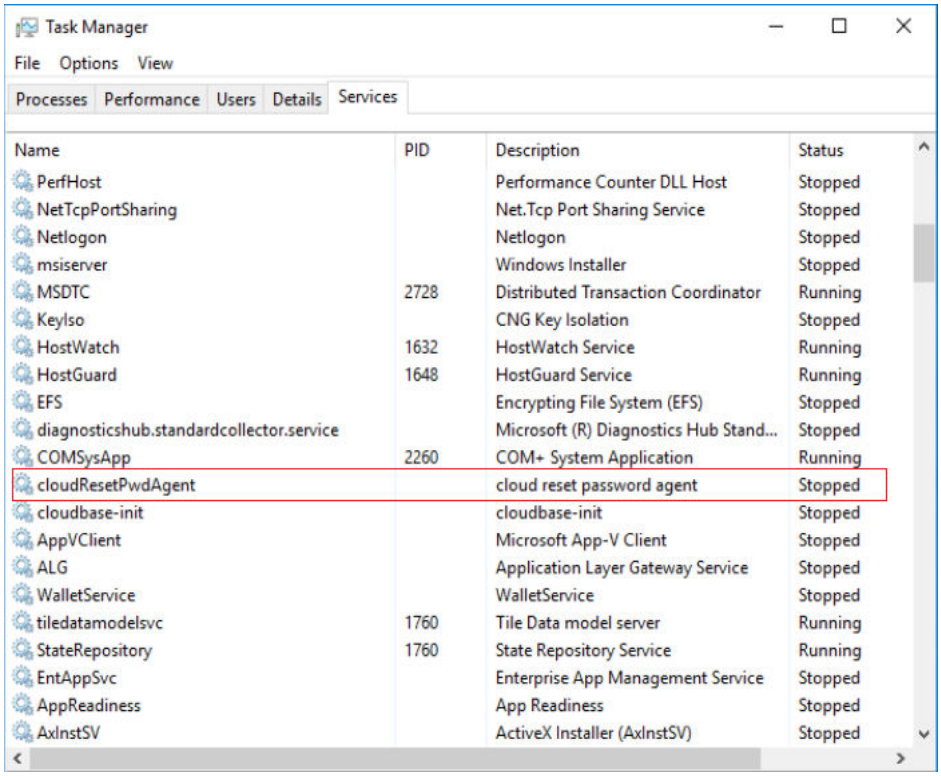

- If yes, no further action is required.
- If no, go to **Step 2**.

Step 3 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to **Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Windows)**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 4 Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

Step 5 Install the plug-in.

1. Double-click **setup.bat** in **CloudResetPwdAgent.Windows**.

   The password reset plug-in starts to be installed.

2. View the **Task Manager** and check whether the installation was successful.

   If **cloudResetPwdAgent** is displayed in the **Task Manager**, as shown in **Figure 8-6**, the installation was successful. Otherwise, the installation failed.

   Figure 8-6 Successful plug-in installation

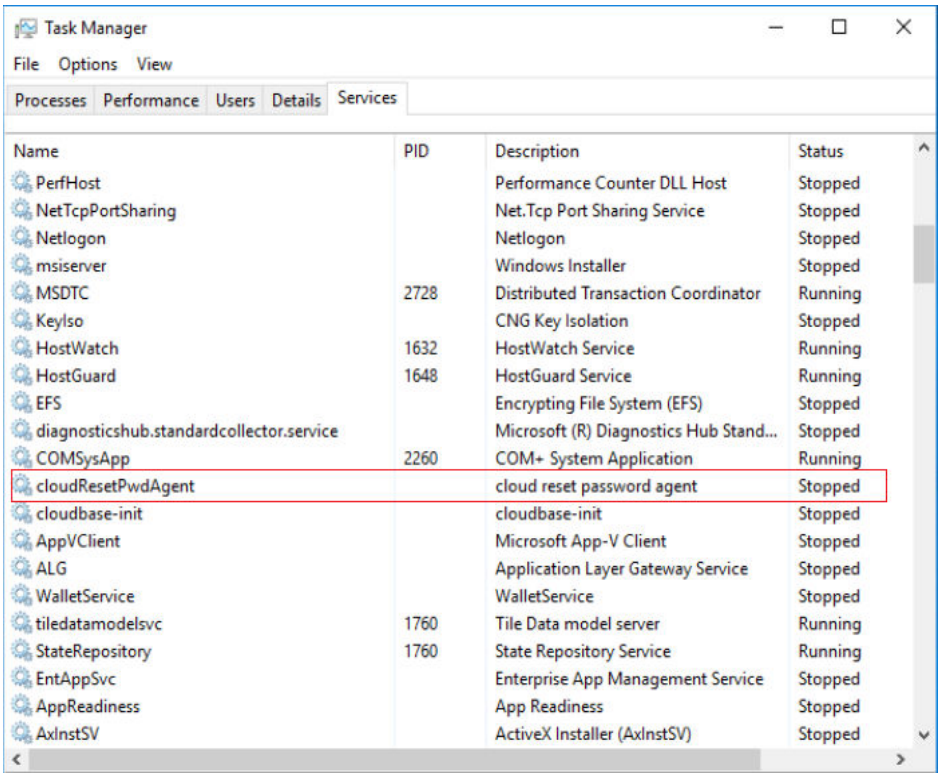

> **NOTE**
>
> If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

----End

Follow-up Procedure

- After the one-click password reset plug-in is installed, you can add it to the startup items to ensure that the plug-in automatically starts upon ECS startup. For details, see **What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?**
- After installing the one-click password reset plug-in, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.
- One-click password reset plug-in has been upgraded. New ECSs work in PIPE mode by default, preventing the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-ins randomly select idle ports with the smallest port numbers ranging from 31000 to 32999.

Uninstalling the Plug-in

If you do not need the password reset function anymore, perform the following operations to uninstall the plug-in:

- Linux ECS
  a. Log in to the ECS.
  b. Run the following commands to switch to the **bin** directory and delete **cloudResetPwdAgent**:
     cd /CloudrResetPwdAgent/bin
     sudo ./cloudResetPwdAgent.script remove
  c. Run the following command to delete the plug-in:
     sudo rm -rf /CloudrResetPwdAgent
     Check whether **CloudResetPwdUpdateAgent** exists. If it exists, run the following command to delete it:
     sudo rm -rf /CloudResetPwdUpdateAgent
- Windows ECS
  a. Switch to the **C:\CloudResetPwdAgent\bin** folder.
  b. Double-click **UninstallApp-NT.bat**.
  c. Delete the file in **C:\CloudResetPwdAgent**.
  d. Delete the file in **C:\CloudResetPwdUpdateAgent**.

8.2.3 Updating the One-Click Password Reset Plug-in for an ECS

You can reset the password for logging in to an ECS with just a few clicks if you forgot the password or the password expires.

This section describes how to update the one-click password reset plug-in for an ECS.

Notes

1. The one-click password reset plug-in can be updated only after an EIP is bound to the ECS.
2. By default, the one-click password reset plug-in has been installed on ECSs created using public images by default. Before updating the plug-in, uninstall it first.

Prerequisites

- The available space in drive C of a Windows ECS is greater than 300 MB, and data can be written to it.
- The available space in the root directory of a Linux ECS is greater than 300 MB, and data can be written to it.
- For Linux ECSs, **Disabling SELinux** if it has been enabled.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.

- DHCP is enabled in the VPC to which the ECS belongs.
- The ECS network connectivity is normal.
- Set the NIC to DHCP so that the ECS can dynamically obtain an IP address.

  NOTE

  For details about how to set the NIC to DHCP for a Linux ECS, see **Setting the NIC to DHCP**.

  For details about how to set the NIC to DHCP for a Windows ECS, see **Setting the NIC to DHCP**.
- The ECS security group rule in the outbound direction meets the following requirements:
  - **Protocol**: **TCP**
  - **Port Range**: **80**
  - **Remote End**: **169.254.0.0/16**

  If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:
  - **Protocol**: **ANY**
  - **Port Range**: **ANY**
  - **Remote End**: **0.0.0.0/16**

Updating the One-Click Password Reset Plug-in on a Linux ECS

Step 1 Uninstall the plug-in.

1. Log in to the ECS.
2. Run the following commands to switch to the **bin** directory and delete **cloudResetPwdAgent**:

   cd /CloudrResetPwdAgent/bin

   sudo ./cloudResetPwdAgent.script remove
3. Run the following command to delete the plug-in:

   sudo rm -rf /CloudrResetPwdAgent

   Check whether **CloudResetPwdUpdateAgent** exists. If it exists, run the following command to delete it:

   sudo rm -rf /CloudResetPwdUpdateAgent

Step 2 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to **Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Linux)**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 3 Run the following command to decompress **CloudResetPwdAgent.zip**:

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

unzip -o -d *Decompressed directory* **CloudResetPwdAgent.zip**

An example is provided as follows:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

unzip -o -d /home/linux/test CloudResetPwdAgent.zip

Step 4 Install the one-click password reset plug-in.

1. Run the following command to open the **CloudResetPwdAgent.Linux** file:

   cd *{Plug-in decompressed directory}*/**CloudResetPwdAgent/CloudResetPwdAgent.Linux**

   For example:

   If the plug-in is decompressed to **/home/linux/test**, run the following command:

   cd /home/linux/test/CloudResetPwdAgent/CloudResetPwdAgent.Linux

2. Run the following command to add the execute permission for the **setup.sh** file:

   chmod +x setup.sh

3. Run the following command to install the plug-in:

   sudo sh setup.sh

   If "cloudResetPwdAgent install successfully." is displayed and "Failed to start service cloudResetPwdAgent" is not displayed, the installation is successful.

   NOTE

   - You can also check whether the password reset plug-in has been installed using the methods provided in **Step 1**.
   - If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

Step 5 Modify the file permissions of the password reset plug-in.

chmod 700 /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script

chmod 700 /CloudrResetPwdAgent/bin/wrapper

chmod 600 /CloudrResetPwdAgent/lib/*

----End

Updating the One-Click Password Reset Plug-in on a Windows ECS

Step 1 Uninstall the plug-in.

1. Switch to the **C:\CloudResetPwdAgent\bin** folder.
2. Double-click **UninstallApp-NT.bat**.
3. Delete the file in **C:\CloudResetPwdAgent**.
4. Delete the file in **C:\CloudResetPwdUpdateAgent**.

Step 2 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to **Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Windows)**.

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 3 Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

Step 4 Install the plug-in.

1. Double-click **setup.bat** in **CloudResetPwdAgent.Windows**.

   The password reset plug-in starts to be installed.

2. View the **Task Manager** and check whether the installation was successful.

   If **cloudResetPwdAgent** is displayed in the **Task Manager**, as shown in Figure 8-7, the installation was successful. Otherwise, the installation failed.

Figure 8-7 Successful plug-in installation

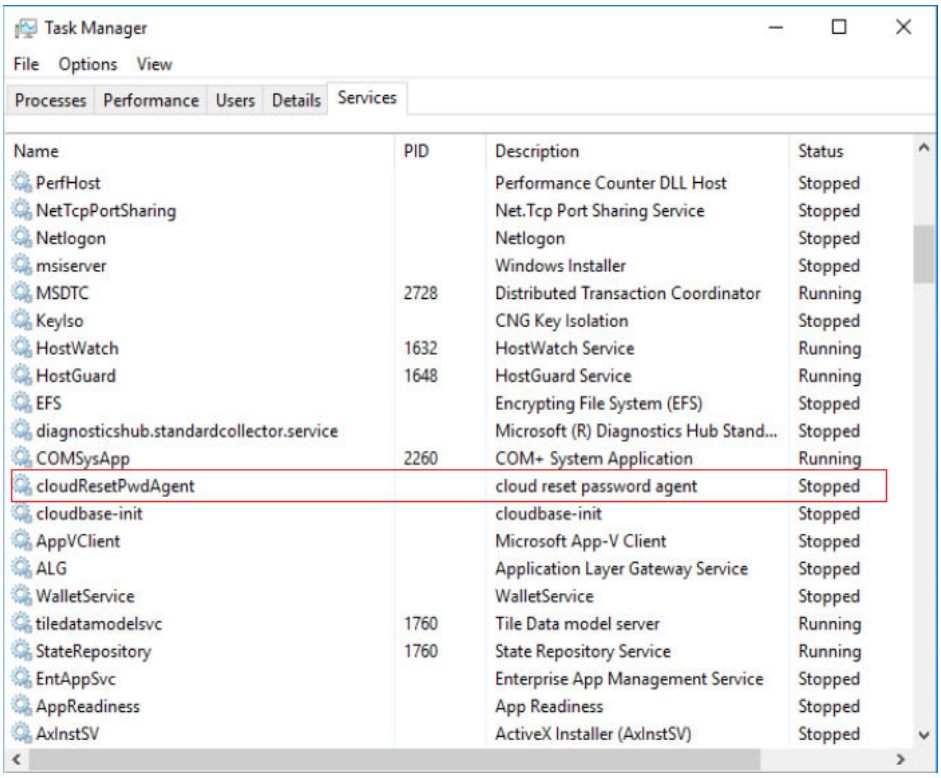

NOTE

If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

----End

Follow-up Procedure

- After the one-click password reset plug-in is updated, you can add it to the startup items to ensure that the plug-in automatically starts upon ECS startup. For details, see **What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?**
- After updating the one-click password reset plug-in, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.

- The one-click password reset plug-in has been upgraded. New ECSs work in PIPE mode by default, preventing the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in randomly selects idle ports with the smallest port numbers ranging from 31000 to 32999.

8.3 Key Pairs

8.3.1 Application Scenarios for Using Key Pairs

Key Pairs

Key pairs are a set of security credentials for identity authentication when you remotely log in to ECSs.

A key pair consists of a public key and a private key. Key Pair Service (KPS) stores the public key and you store the private key. If you have imported a public key into a Linux ECS, you can use the corresponding private key to log in to the ECS without a password. Therefore, you do not need to worry about password interception, cracking, or leakage.

You can use **Data Encryption Workshop (DEW)** to manage key pairs, including creating, importing, binding, viewing, resetting, replacing, unbinding, and deleting key pairs.

This section describes how to create and import a key pair. For details about other operations, see **Managing Key Pairs**.

Scenarios

When purchasing an ECS, you are advised to select the key pair login mode. For Windows ECSs, key pairs are required to decrypt the passwords so that you can use the decrypted password to log in.

- Logging in to a Linux ECS

  You can directly use a key pair to log in.

  - When you are creating the ECS, select the key pair login mode. For details, see "Set Login Mode" in **Step 3: Configure Advanced Settings**.
  - After the ECS is created, **bind a key pair**.
- Logging in to a Windows ECS

  You can use the key pair to obtain a password for login. The password is randomly generated and therefore is more secure.

  For details, see **Obtaining the Password for Logging In to a Windows ECS**.

Creating a Key Pair

You can create a key pair or use an existing one for remote login authentication.

- Creating a key pair

  You can create a key pair using either of the following method:

  - Follow the instructions in **(Recommended) Creating a Key Pair on the Management Console**. The public key is automatically stored in the system, and the private key is stored locally.
  - Follow the instructions in **Creating a Key Pair Using PuTTYgen**. Both the public and private keys are stored locally.

    After the key pair is created, import the key pair following the instructions provided in **Importing a Key Pair** so that you can use it.
- Using an existing key pair

  If an existing key pair (created using PuTTYgen, for example) is available, you can import the public key by referring to **Importing a Key Pair** on the management console to let the system maintain the public key for you.

  NOTE

  If the public key of the existing key pair is stored by clicking **Save public key** on PuTTY Key Generator, the public key cannot be imported to the management console.

  If you want to use this existing key pair for remote login, see **Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?**

Constraints

- Key pairs can be used to remotely log in to Linux ECSs only.
- SSH-2 key pairs created on the console support only the RSA-2048 cryptographic algorithms.
- Key pairs can be used only for ECSs in the same region.
- Imported key pairs support the following cryptographic algorithms:
  - RSA-1024
  - RSA-2048
  - RSA-4096
- Store your private key in a secure place because you need to use it to prove your identity when logging in to your ECS. The private key can be downloaded once only.

8.3.2 (Recommended) Creating a Key Pair on the Management Console

Scenarios

You can use the management console to create a key pair. ECS stores the public key and you store the private key.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.

5. On the displayed page, click **Create Key Pair**.

   NOTE

   Key pairs include private key pairs and account key pairs. Private key pairs are only available to the user itself. Account key pairs are available to all users under the account.

   You can create key pairs based on your needs.

6. Configure the following parameters:
   a. Enter a key pair name.
   b. Select a type.
   c. Select a KMS encryption algorithm.

      This option is displayed only if you select **I agree to have the private key managed by HUAWEI CLOUD**.

      NOTE

      - If you do not select **I agree to have the private key managed by HUAWEI CLOUD**, you can download the private key only once for security purposes. Please keep the private key secure.

        If the key pair is lost, you can **reset a key pair** and bind it to the ECS.
      - If you select **I agree to have the private key managed by HUAWEI CLOUD**, you can export the managed private key as required. For details, see **Exporting a Private Key**.

   d. Select **I have read and agree to the Key Pair Service Disclaimer**.
7. Click **OK**.

Related Operations

- If your private key file is lost, you can **Resetting a Key Pair**.
- If your private key file is leaked, you can **use a new key pair to replace the public key of the ECS**.

8.3.3 Creating a Key Pair Using PuTTYgen

Scenarios

You can use PuTTYgen to create a key pair and store the public key and private key locally.

NOTE

Key pairs created using puttygen.exe must be imported by referring to **Importing a Key Pair** before they are used.

Procedure

1. Download and install PuTTY and PuTTYgen.

   https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html

   NOTE

   PuTTYgen is a key generator, which is used to create a key pair that consists of a public key and a private key for PuTTY.

2. Obtain the public and private keys.
   a. Double-click **puttygen.exe** to open **PuTTY Key Generator**.

      Figure 8-8 PuTTY Key Generator

      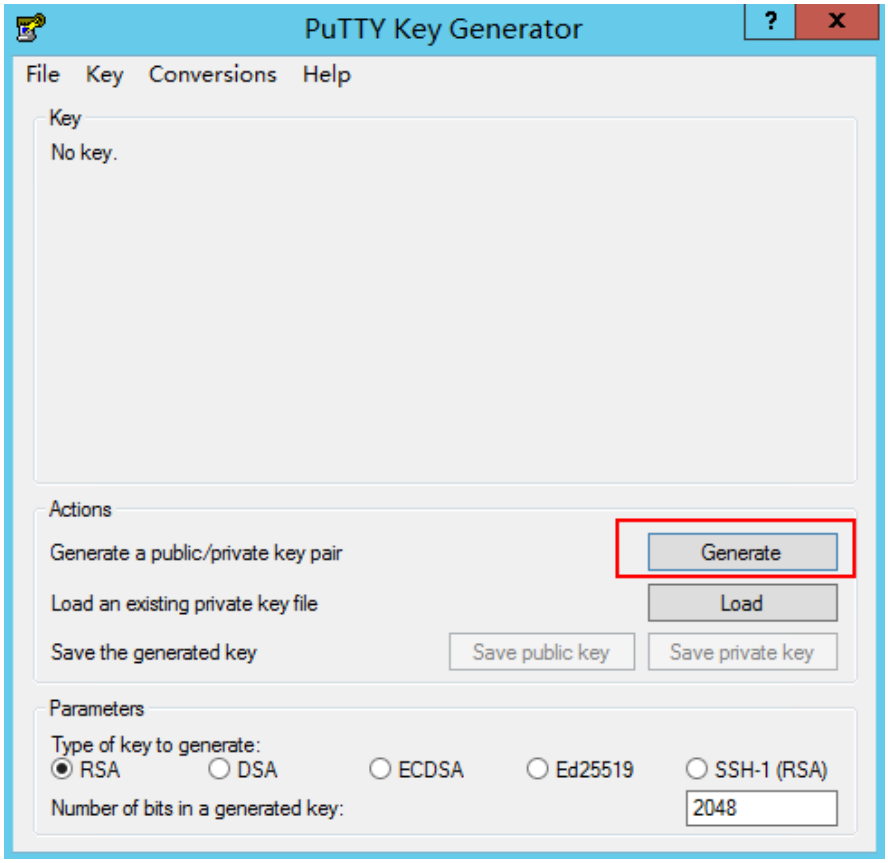

   b. Click **Generate**.

      The key generator automatically generates a key pair that consists of a public key and a private key. The content shown in the red box in **Figure 8-9** is the public key.

Figure 8-9 Generating the public and private keys

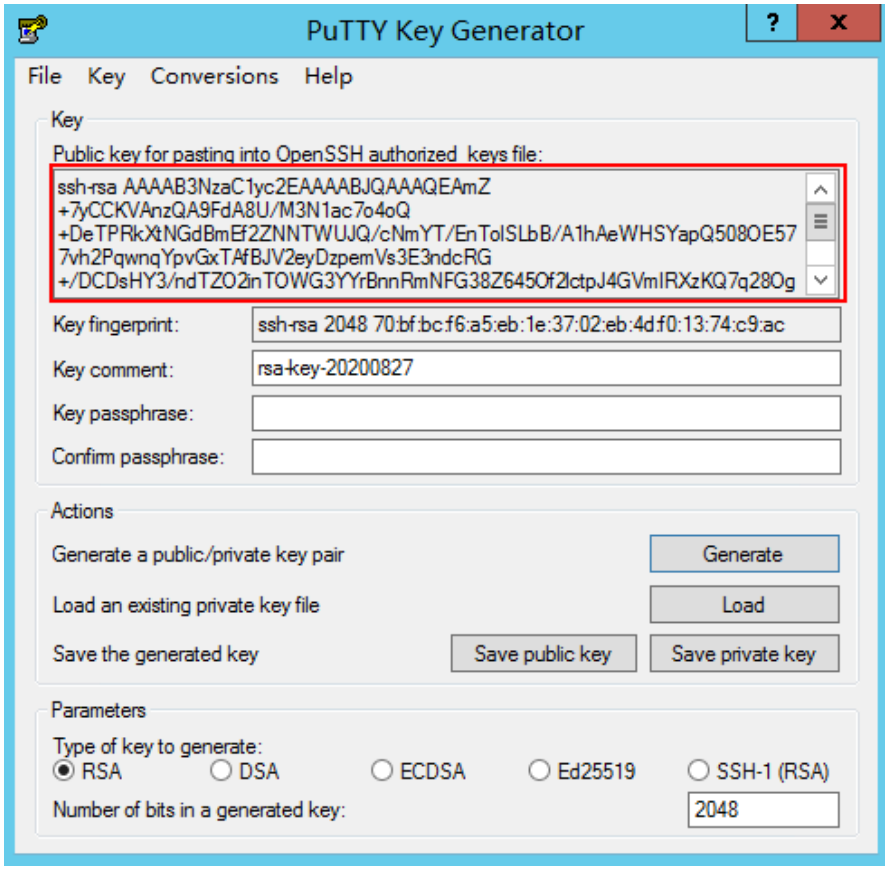

3. Copy the public key to a .txt file and save it to a local directory.

   NOTE

   Do not save the public key by clicking **Save public key** because this operation will change the format of the public key content and cause the public key to fail to be imported to the management console.

4. Save the private key and keep it secure. The private key can be downloaded only once.

   The format in which to save your private key file varies depending on application scenarios.

   - When using PuTTY to log in to a Linux ECS:

     Save the private key file in the **.ppk** format.

     i. On the **PuTTY Key Generator** page, choose **File** > **Save private key**.

Figure 8-10 Saving a private key

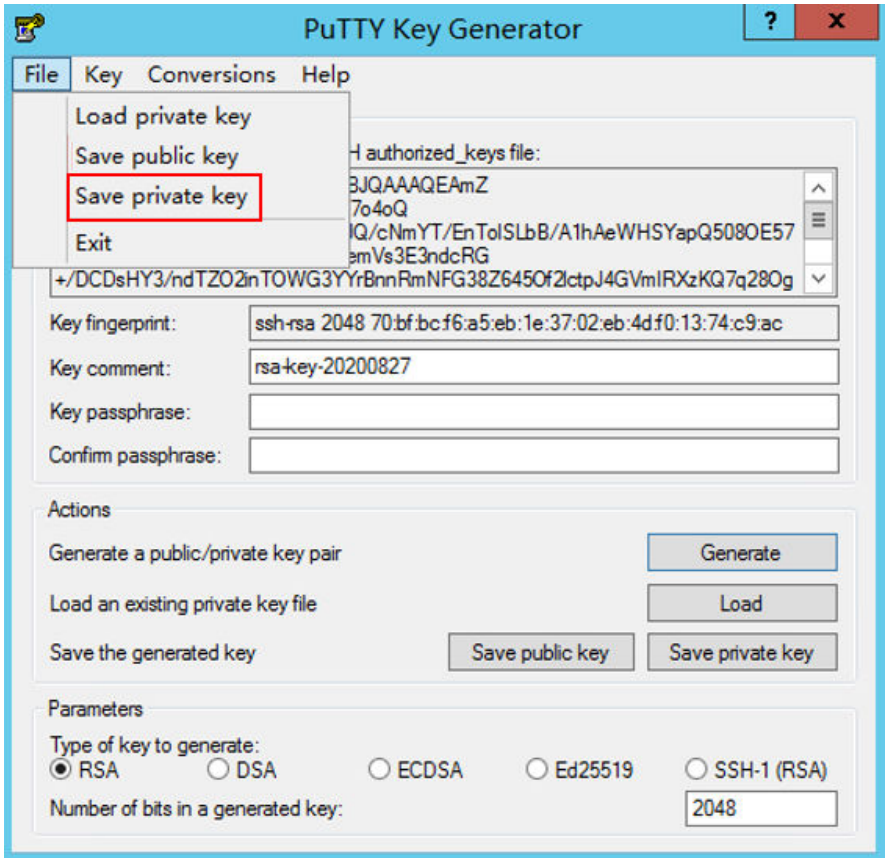

ii. Save the converted private key file, such as **kp-123.ppk**, locally.

– When using Xshell to log in to a Linux ECS or obtaining the password for logging in to a Windows ECS:

  Save the private key file in the **.pem** format.

  i. Choose **Conversions** > **Export OpenSSH key**.

  NOTE

  If you use this private file to obtain the password for logging in to a Windows ECS, do not specify **Key passphrase** for **Export OpenSSH key** so that you can obtain the password successfully.

Figure 8-11 Saving a private key

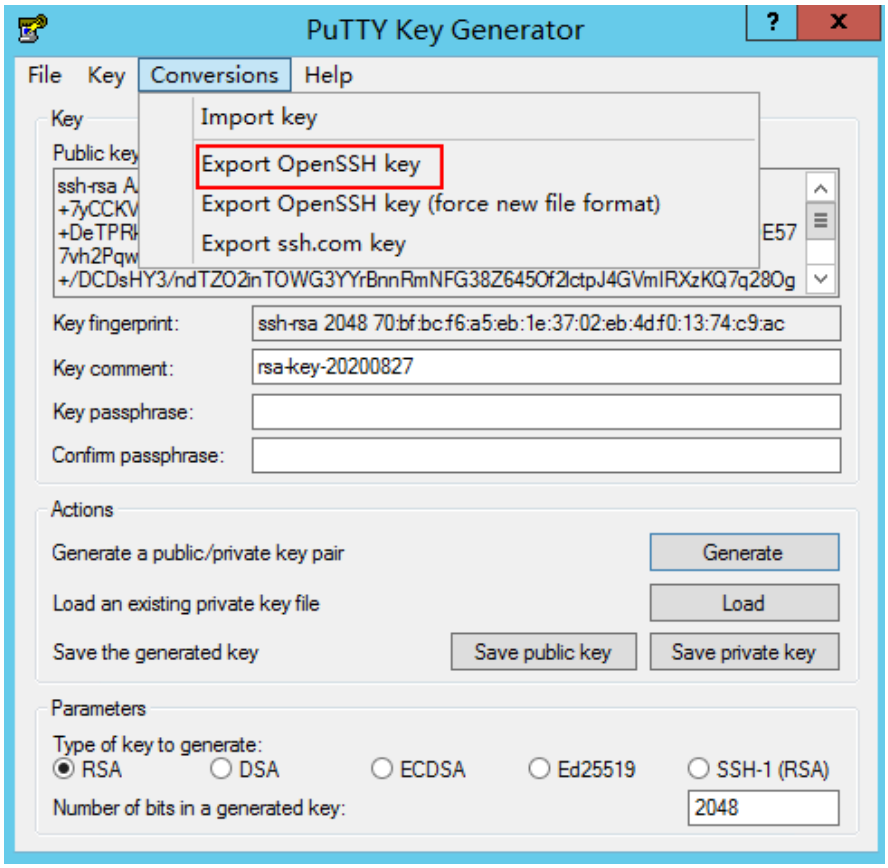

ii. Save the private key, for example, **kp-123.pem**, locally.

5. After you have saved the key pair, import your public key to the ECS by referring to **Importing a Key Pair**.

Related Operations

- If your private key file is lost, you can **Resetting a Key Pair**.
- If your private key file is leaked, you can **use a new key pair to replace the public key of the ECS**.

8.3.4 Importing a Key Pair

Scenarios

You need to import a key pair in either of the following scenarios:

- Create a key pair using PuTTYgen and import the public key to the ECS.
- Import the public key of an existing key pair to the ECS to let the system maintain your public key.

NOTE

If the public key of the existing key pair is stored by clicking **Save public key** on PuTTY Key Generator, the public key cannot be imported to the management console.

If you want to use this existing key pair for remote login, see **Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?**

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.
5. On the **Key Pair Service** page, click **Import Key Pair**.
6. Use either of the following methods to import the key pair:
   - Selecting a file
     i. In the **Import Key Pair** dialog box of the management console, click **Select File** and select the locally stored public key file (for example, the .txt file saved in **3** in **Creating a Key Pair Using PuTTYgen**).

        NOTE

        Make sure that the file to be imported is a public key file.

     ii. Click **OK**.

        After the public key is imported, you can change its name.
   - Copying the public key content
     i. Copy the public key content from the locally stored .txt file into the **Public Key Content** text box.
     ii. Click **OK**.

Helpful Links

- **What Should I Do If a Key Pair Cannot Be Imported?**
- **Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?**

8.3.5 Obtaining and Deleting the Password of a Windows ECS

8.3.5.1 Obtaining the Password for Logging In to a Windows ECS

Scenarios

Password authentication is required to log in to a Windows ECS. Therefore, you must use the key file used when you created the ECS to obtain the administrator password generated during ECS creation. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

You can obtain the initial password for logging in to a Windows ECS through the management console or APIs. For details, see this section.

Obtaining the Password Through the Management Console

1. Obtain the private key file (.pem file) used when you created the ECS.

2. Log in to the management console.
3. Click in the upper left corner and select your region and project.
4. Click . Under **Compute**, click **Elastic Cloud Server**.
5. On the **Elastic Cloud Server** page, select the target ECS.
6. In the **Operation** column, click **More** and select **Get Password**.

Figure 8-12 Obtaining a password

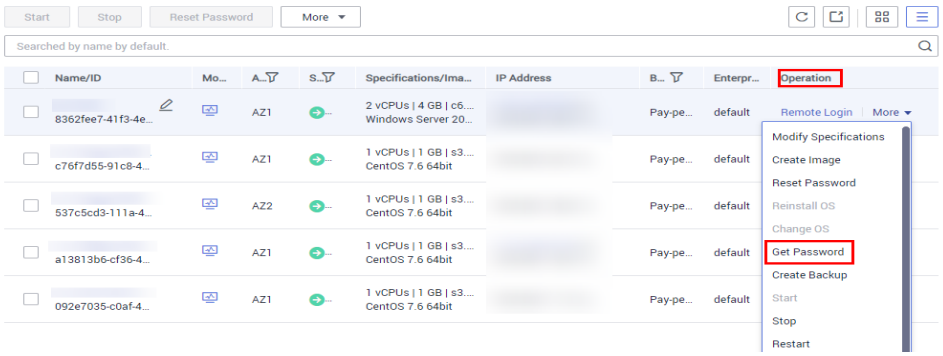

NOTE

If **Get Password** is not displayed, the one-click password reset plug-in may not be installed.

In this case, you can reset the password by referring to **Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed**.

7. Use either of the following methods to obtain the password through the key file:
   - Click **Select File** and upload the key file from a local directory.
   - Copy the key file content to the text field.
8. Click **Get Password** to obtain a random password.

Obtaining the Password Through APIs

1. Obtain the private key file (.pem file) used when you created the ECS.
2. Set up the API calling environment.
3. Call APIs. For details, see "Before You Start" in *Elastic Cloud Server API Reference*.
4. Obtain the ciphertext password.

   Call the password obtaining APIs to obtain the ciphertext password of the public key encrypted using RSA. The API URI is in the format "GET /v2/{*tenant_id*}/servers/{*server_id*}/os-server-password".

   NOTE

   For details, see "Obtaining the Password for Logging In to an ECS" in the *ECS API Reference*.
5. Decrypt the ciphertext password.

Use the private key file used when you created the ECS to decrypt the ciphertext password obtained in step **4**.

a. Run the following command to convert the ciphertext password format to ".key -nocrypt" using OpenSSL:

openssl pkcs8 -topk8 -inform PEM -outform DER -in rsa_pem.key -out pkcs8_der.key -nocrypt

b. Invoke the Java class library **org.bouncycastle.jce.provider.BouncyCastleProvider** and use the key file to edit the code decryption ciphertext.

8.3.5.2 Deleting the Initial Password for Logging In to a Windows ECS

Scenarios

After you obtain the initial password, it is a good practice to delete it to ensure system security.

Deleting the initial password does not affect ECS operation or login. Once deleted, the password cannot be retrieved. Before you delete a password, it is a good practice to record it.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. In the **Operation** column, click **More** and select **Delete Password**.

   The system displays a message, asking you whether you want to delete the password.
6. Click **OK** to delete the password.

9 Permissions Management

9.1 Creating a User and Granting ECS Permissions

Use **IAM** to implement fine-grained permissions control over your ECSs. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing ECS resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust Huawei Cloud accounts or cloud services to perform efficient O&M on your ECS resources.

If your Huawei Cloud account does not need individual IAM users, skip this section.

This section describes the procedure for granting permissions (see **Process Flow**).

Prerequisites

Before assigning permissions to user groups, you should learn about system policies supported by ECS and select the policies based on service requirements.

For more information about system policies supported by ECS, see **ECS Permissions**. For the permissions of other services, see **System Permissions**.

Process Flow

Figure 9-1 Process for granting ECS permissions

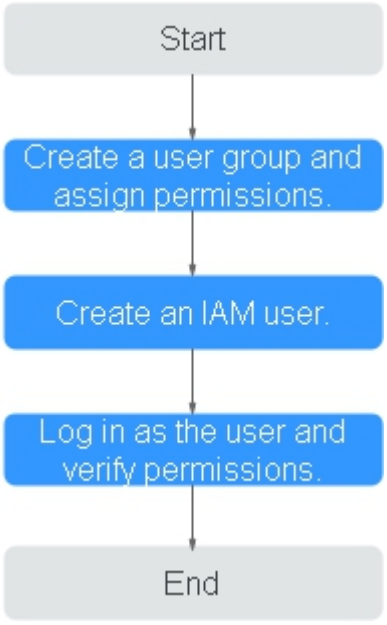

1. Create a user group and assign permissions to it.

   Create a user group on the IAM console, and attach the **ECSReadOnlyAccess** policy to the group.
2. Create an IAM user.

   Create a user on the IAM console and add the user to the group created in **1**.
3. Log in and verify permissions.

   Log in to the ECS console by using the created user, and verify that the user only has read permissions for ECS.

   - Choose **Compute** > **Elastic Cloud Server** in **Service List**. On the ECS console, click **Buy ECS**. If the creation attempt failed, the **ECSReadOnlyAccess** policy has already taken effect.
   - Choose any service other than ECS in **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **ECSReadOnlyAccess** policy has already taken effect.

9.2 ECS Custom Policies

Custom policies can be created to supplement the system-defined policies of ECS. For the actions that can be added to custom policies, see "Permissions Policies and Supported Actions" in **Elastic Cloud Server API Reference**.

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**. The following section contains examples of common ECS custom policies.

Example Custom Policies

- Example 1: Allowing users to stop and delete multiple ECSs at a time

```
{
   "Version": "1.1",
   "Statement": [
      {
         "Effect": "Allow",
         "Action": [
            "
                ecs:servers:stop
                ecs:servers:get
            "
         ]
      }
   ]
}
```

- Example 2: Denying ECS deletion

  A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  The following method can be used if you need to assign permissions of the **ECSFullAccess** policy to a user but you want to prevent the user from deleting ECSs. Create a custom policy for denying ECS deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on ECSs except deleting ECSs. The following is an example of a deny policy:

```
{
    "Version": "1.1",
    "Statement": [
        {
       "Effect": "Deny",
            "Action": [
                "ecs:cloudServers:delete"
            ]
        }
    ]
}
```


10 Launch Templates

10.1 Overview

What Is a Launch Template?

A launch template contains the configuration information to launch an ECS, for example, the ECS specifications, network settings, and a key pair (excluding the password). You can launch an ECS quickly without specifying the configuration parameters every time.

A launch template cannot be modified after it is created. However, you can create multiple versions of a template. Each version can be configured with different parameters. You can use any version of the template to create ECSs.

Creating a Launch Template

Create a launch template on the console.

For details, see **Creating a Launch Template**.

10.2 Creating a Launch Template

Scenarios

This section describes how to create a launch template on the management console.

Constraints

- Each account can have a maximum of 30 launch templates in each region.
- The parameters you can configure when you create a launch template are optional.

  However, if your launch template does not include parameters, such as the flavor and image, you need to set them when you use the template to create an ECS.

- A launch template cannot be modified after it is created. However, you can create a new version of the template to update its parameter configurations.
- Supported regions include AP-Singapore and CN-Hong Kong.

Creating a Launch Template on the Launch Templates Console

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Launch Templates**.
5. On the **Launch Templates** page, click **Create Launch Template**.
6. Configure basic settings, network, and advanced settings.

   For details about the configuration parameters, see **Purchasing an ECS**.
7. On the **Confirm** step, enter the template name and description, and click **Create Now**.

   You can view the created template on the launch template list page.

Creating a Launch Template When Buying an ECS

You can save the ECS configurations as a launch template when creating an ECS.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. Click **Buy ECS**.
5. Configure basic settings, network, and advanced settings.

   For details about the configuration parameters, see **Purchasing an ECS**.
6. On the **Confirm** step, click **Save as Launch Template**, enter the template name and description, and click **OK**.

   You can view the created template on the launch template list page.

10.3 Managing Launch Templates

Scenarios

You can:

- **Viewing Details About a Launch Template**
- **Deleting a Launch Template**

Viewing Details About a Launch Template

1. Log in to the management console.

2. Click in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Launch Templates**.
5. On the **Launch Templates** page, click the name of the launch template to view its details.

Table 10-1 Launch template details

Parameter	Description
Name	The name of the launch template.
ID	The ID of the launch template.
Created	The time when the launch template is created.
Description	The description of the launch template.
Version Information	The version information contains the configuration information about the launch template of the current version, such as the region, specifications, and image.

Deleting a Launch Template

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Launch Templates**.
4. Locate row that contains the launch template to be deleted and click **Delete** in the **Operation** column.
5. In the displayed dialog box, click **Yes**.

11 Auto Launch Groups

11.1 Overview

What Is an Auto Launch Group?

An auto launch group lets you rapidly create ECSs distributed across multiple AZs, using a combination of different types of spot and pay-per-use ECSs to meet capacity targets at the lowest price possible.

Application Scenarios

Auto launch groups are applicable to scenarios such as image rendering, stateless web services, DNA sequencing, offline analysis, function computing, batch computing, sample analysis, CI/CD, and test.

Notes

- An auto launch group can create ECSs across AZs but cannot create ECSs across regions.
- The target capacity of each auto launch group is limited.
  - If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.
  - If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.
- You can specify one launch template for each auto launch group.

Advantages

- Spot ECSs and pay-per-use ECSs

  Spot ECSs are much less expensive than regular pay-per-use ECSs, but they can be reclaimed suddenly. Spot ECSs are a great way to save money when running stateless, fault-tolerant instances that are not sensitive to interruptions. Pay-per-use ECSs can be created and deleted at any time and are a good way to save money when you are not sure about expected usage as you pay only for what you use.

An auto launch group lets you rapidly create both spot and pay-per-use ECSs to meet capacity targets at the lowest price possible.

- ECSs from different AZs

  An auto launch group can create ECSs across AZs to improve the disaster recovery capability.
- ECSs of different types

  An auto launch group can create ECSs of different types to meet your requirements of different scenarios.
- Flexible allocation strategies

  You can specify your desired target capacity and how much of that must be pay-per-use ECSs.

  You can also let your auto launch group continue to create ECSs until the total target capacity is reached or delete ECSs when the target capacity is exceeded.
- Cost effectiveness

  If your set **Optimize for** to **Lowest price**, the auto launch group will create the least expensive ECSs possible.

Pricing Details

Auto launch groups are free, but you will be billed for the ECSs created by the group.

For details, see **Elastic Cloud Server Pricing Details**.

11.2 Creating an Auto Launch Group

Scenarios

This section describes how to create an auto launch group on the management console.

Constraints

Currently, auto launch groups are supported in AP-Singapore and CN-Hong Kong regions.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Auto Launch Groups**.
5. On the **Auto Launch Groups** page, click **Create Group**.
6. Set the name of the auto launch group.

   The name can contain 2 to 64 characters, including letters, digits, underscores (_), and hyphens (-).

7. Set the total target capacity.

   You can specify the number of ECSs or vCPUs.

   If you choose to include pay-per-use ECSs, set the quantity of pay-per-use ECSs or vCPUs.

   The target capacity of each auto launch group is limited.

   - If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.
   - If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.

8. Select a launch template.

   You can select a launch template and its corresponding version as the configuration source. You can also select other required ECS configurations.

9. Set the allocation strategy.

   - **Lowest price**: The auto launch group will create the least expensive ECSs possible.
   - **Compute balancing**: The auto launch group will prioritize balancing compute loads by creating ECSs distributed across multiple AZs as evenly as possible.
   - **High specifications**: The auto launch group creates ECSs with the highest specifications possible.

     If you have configured a target number of ECSs, ECSs with more vCPUs will be prioritized and if the target is vCPUs, then that target will be met with as few ECSs as possible.

10. Select a delivery type.

    - **Single use**: The auto launch group only attempts to create ECSs to meet the target capacity when it is started, but will not create ECSs again even if the target capacity is not reached.
    - **Continuous**: The auto launch group continues to create ECSs until the total target capacity is reached.

11. Set the start time.

    Set the time when the auto launch group starts to launch ECSs. You can set both the start time and the end time to determine the validity period of the group.

    - **Immediately**: The auto launch group starts to launch ECSs immediately after the group is created.
    - **Custom**: You can specify when the auto launch group starts to launch ECSs.

12. Set the end time.

    Set the time when the auto launch group expires. You can set both the start time and the end time to determine the validity period of the group.

    - **Never expire**: The auto launch group does not expire.
    - **Custom**: You can specify when the auto launch group expires.

13. Set the global maximum price.

    Set the allowed maximum price of a single spot ECS in the auto launch group. If the market price of a spot ECS in the group exceeds the global maximum price, the spot ECS will be deleted.

If both the specific maximum price of a spot ECS and the global maximum price are set, the specific maximum price of the spot ECS will be used.

The price cannot be less than 0. If the price is set to be greater than the pay-per-use ECS price, there is no upper limit on the spot ECS price.

14. Configure ECS deletion settings.
    - **Delete ECSs When Auto Launch Group Expires**: ECSs in the auto launch group will be deleted when the group expires.
    - **Delete ECSs When Target Capacity Is Exceeded**: When the number of ECSs or vCPUs in the auto launch group exceeds the target capacity, the ECSs or vCPUs that exceed the target capacity will be deleted.

      NOTE

      If you do not select **Delete ECSs When Target Capacity Is Exceeded**, the ECSs that exceed the target capacity will be removed from the group but not deleted. The removed ECSs will be displayed in the ECS list. To avoid charges on such ECSs, manually delete them.
15. Click **Create Now**.

Execution Result

After an auto launch group is created, the group starts to create ECSs at the specified time. If you select the **Continuous** delivery mode, the auto launch group monitors the target and current capacity in real time, and automatically creates a new ECS if a spot ECS is reclaimed.

11.3 Managing Auto Launch Groups

Scenarios

You can:

- **Viewing Details About an Auto Launch Group**
- **Modifying an Auto Launch Group**
- **Deleting an Auto Launch Group**

Viewing Details About an Auto Launch Group

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. On the **Auto Launch Groups** page, click the name of the target auto launch group to view its details.

   You can view the basic information and capacity overview of the auto launch group.

   The basic information includes the name, launch template, delivery type, and allocation strategy of the group.

In the **Capacity Overview** area, you can view the total current and target capacity, and the current and target capacity of the spot ECS and pay-per-use ECS.

Modifying an Auto Launch Group

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. Locate the row that contains the target auto launch group and click **Modify** in the **Operation** column.

   You can modify the name, target capacity, quantity of pay-per-use ECSs, allowed maximum price of a spot ECS, and whether to delete ECSs when the auto launch group expires or the target capacity is exceeded.
5. Click **Yes**.

Deleting an Auto Launch Group

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. Locate the row that contains the target auto launch group and click **Delete** in the **Operation** column.
5. Determine whether to delete the ECSs in the auto launch group after the group is deleted.

   If you do not want to delete the ECSs, you can view the ECSs on the ECS list page. To avoid charges on such ECSs, manually delete them.
6. Click **Yes**.

12 Resources and Tags

12.1 Tag Management

12.1.1 Overview

Scenarios

A tag identifies an ECS. Adding tags to an ECS facilitates ECS identification and management.

You can add a tag to an ECS during the ECS creation or after the ECS is created. You can add a maximum of 10 tags to each ECS.

Basics of Tags

Tags are used to identify cloud resources. When you have many cloud resources of the same type, you can use tags to classify cloud resources by dimension (for example, use, owner, or environment).

Figure 12-1 Example tags

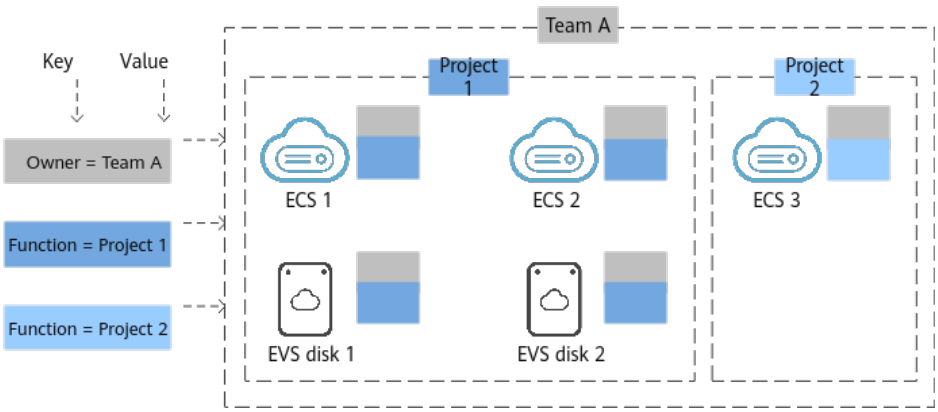

Figure 12-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Usage**. Each tag has a value.

You can quickly search for and filter specific cloud resources based on the tags added to them. For example, you can define a set of tags for cloud resources in an account to track the owner and usage of each cloud resource, making resource management easier.

Tag Naming Rules

- Each tag consists of a key-value pair.
- A maximum of 10 tags can be added to an ECS.
- For each resource, a tag key must be unique and can have only one tag value.
- A tag consists of a tag key and a tag value. Table 12-1 lists the tag key and value requirements.

Table 12-1 Tag key and value requirements

Parameter	Requirement	Example Value
Key	• Cannot be left blank. • The key value must be unique for an ECS. • Can contain a maximum of 36 characters.	Organization
Value	• Can contain a maximum of 43 characters.	Apache

12.1.2 Adding Tags

Tags are used to identify cloud resources, such as ECSs, images, and disks. If you have multiple types of cloud resources which are associated with each other, you can add tags to the resources to classify and manage them easily. For more details, see **Overview**.

You can add tags to an ECS in any of the following ways:

- **Adding Tags During ECS Creation**
- **Adding Tags on the ECS Details Page**
- **Adding Tags on the TMS Console**

For details about how to use predefined tags, see **Using Predefined Tags**.

Adding Tags During ECS Creation

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. Click **Buy ECS**.
5. Configure parameters for the ECS.

   Select **Configure now** for **Advanced Options**. Then, add a tag key and tag value. For the tag key and tag value requirements, see **Table 12-1**.

   NOTE

   - For details about other parameters, see **Purchasing an ECS**.

   Figure 12-2 Adding tags

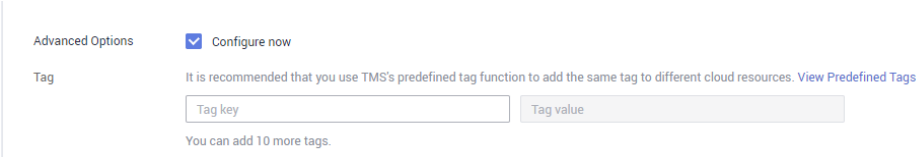

Adding Tags on the ECS Details Page

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. In the ECS list, click the name of the target ECS.

   The ECS details page is displayed.
5. Click the **Tags** tab and then **Add Tag**. In the displayed dialog box, enter the tag key and tag value. For the tag key and tag value requirements, see **Table 12-1**.

   You can change the tag value after the tag is added.

   Figure 12-3 Adding tags on the Tags tab

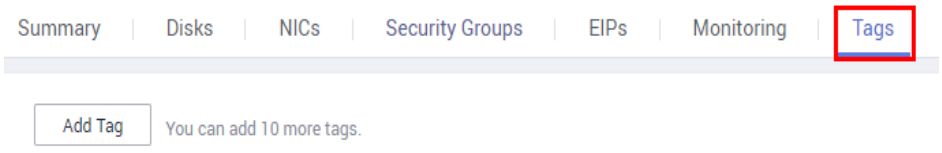

Adding Tags on the TMS Console

NOTE

This method is suitable for adding tags with the same tag key to multiple resources.

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.

Figure 12-4 Tag Management

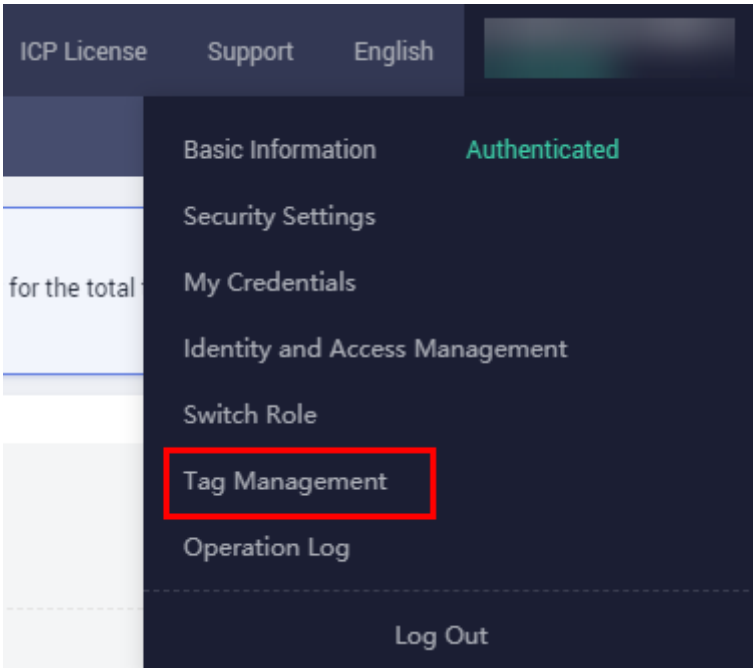

3. On the displayed **Resource Tags** page, select the region where the resource is located, select **ECS-ECS** for **Resource Type**, and click **Search**.

   All ECSs matching the search criteria are displayed.

4. In the **Search Result** area, click **Create Key**. In the displayed dialog box, enter a key (for example **project**) and click **OK**.

   After the tag is created, the tag key is added to the resource list. If the key is not displayed in the resource list, click  and select the created key from the drop-down list.

   By default, the value of the tag key is **Not tagged**. You need to set a value for the tag of each resource to associate the tag with the resource.

Figure 12-5 Resource list

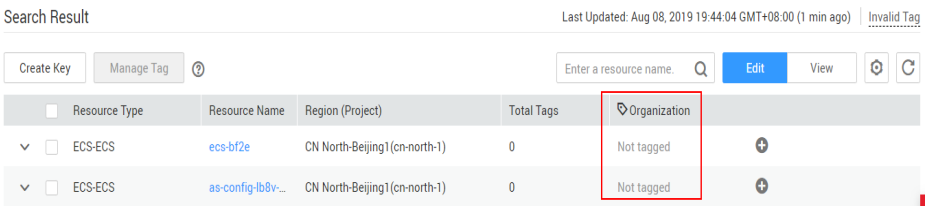

5. Click **Edit** to make the resource list editable.

6. Locate the row containing the target ECS, click  , and enter a value (for example **A**).

   After a value is set for a tag key, the number of tags is incremented by 1. Repeat the preceding steps to add tag values for other ECSs.

Figure 12-6 Setting a tag value

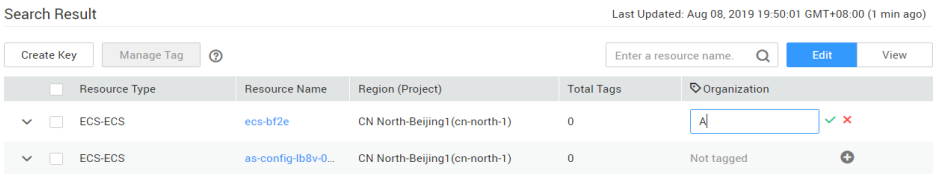

Using Predefined Tags

If you want to add the same tag to multiple ECSs or other resources, you can create a predefined tag on the TMS console and then select the tag for the ECSs or resources. This frees you from having to repeatedly enter tag keys and values. To do so, perform the following operations:

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. Choose **Predefined Tags** in the left navigation pane and click **Create Tag**. In the displayed dialog box, enter a key (for example, **project**) and a value (for example, **A**).
4. Choose **Service List** > **Compute** > **Elastic Cloud Server**, and select the predefined tag by following the procedure for adding a tag.

12.1.3 Searching for Resources by Tag

After tags are added to resources, you can search for resources by tag using either of the following methods.

Searching for ECSs by Tag

On the **Elastic Cloud Server** page, search for ECSs by tag key or value.

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. Click **Search by Tag** above the upper right corner of the ECS list to expand the search area.
5. Enter the tag of the ECS to be queried.

   Neither the tag key nor value can be empty. When the tag key or value is matched, the system automatically shows the target ECSs.
6. Add tags.

   The system supports multiple tags and uses the intersection set of all tags to search for ECSs.
7. Click **Search**.

   The system searches for ECSs based on tag keys and values.

Filtering Resources on the TMS Console

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria, including **Region**, **Resource Type**, and **Resource Tag**.
4. Click **Search**.

   All the resources that meet the search criteria will be displayed in the **Search Result** area.

12.1.4 Deleting a Tag

If you no longer need a tag, delete it in any of the following ways:

- **Deleting a Tag on the ECS Details Page**
- **Deleting a Tag on the TMS Console**
- **Batch Deleting Tags on the TMS Console**

Deleting a Tag on the ECS Details Page

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. In the ECS list, click the name of the target ECS.

   The ECS details page is displayed.
5. Click the **Tags** tab. Locate the row containing the tag to be deleted and click **Delete** in the **Operation** column. In the **Delete Tag** dialog box, click **Yes**.

Deleting a Tag on the TMS Console

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.

Figure 12-7 Tag Management

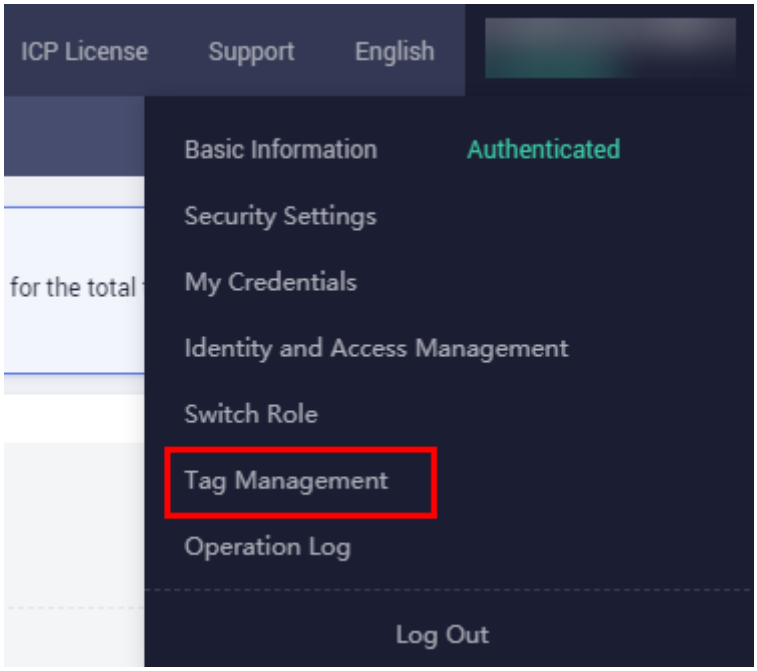

3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. In the **Search Result** area, click **Edit** to make the resource tag list editable.

   If the key of a tag you want to delete is not contained in the list, click and select the tag key from the drop-down list. It is a good practice to select at most 10 keys to display.
5. Locate the row containing the target ECS and click .
6. (Optional) Click in the upper right of the **Search Result** area.

   The resource list is refreshed and the refresh time is updated.

Batch Deleting Tags on the TMS Console

NOTICE

Exercise caution when deleting tags in a batch. After you delete the tags, they will be removed from all the associated ECSs and cannot be recovered.

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. Select the target ECSs.
5. Click **Manage Tag** in the upper left corner of the list.
6. In the displayed **Manage Tag** dialog box, click **Delete** in the **Operation** column. Click **OK**.

7. (Optional) Click in the upper right of the **Search Result** area.
   The resource list is refreshed and the refresh time is updated.

12.2 Quota Adjustment

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources** > **My Quotas**.
   The **Service Quota** page is displayed.

   Figure 12-8 My Quotas

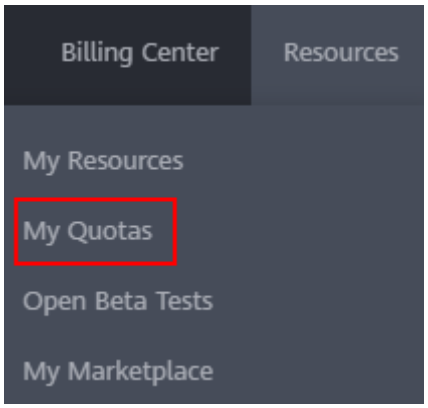

4. View the used and total quota of each type of resources on the displayed page.
   If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources** > **My Quotas**.
   The **Service Quota** page is displayed.

Figure 12-9 My Quotas

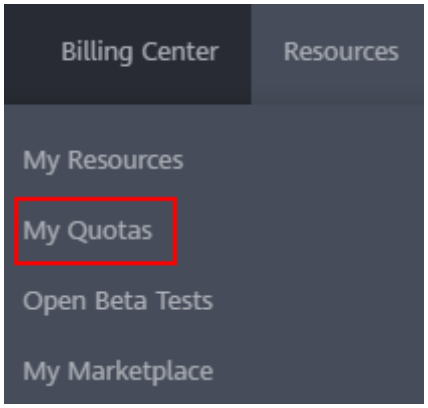

3. Click **Increase Quota** in the upper right corner of the page.

Figure 12-10 Increasing quota

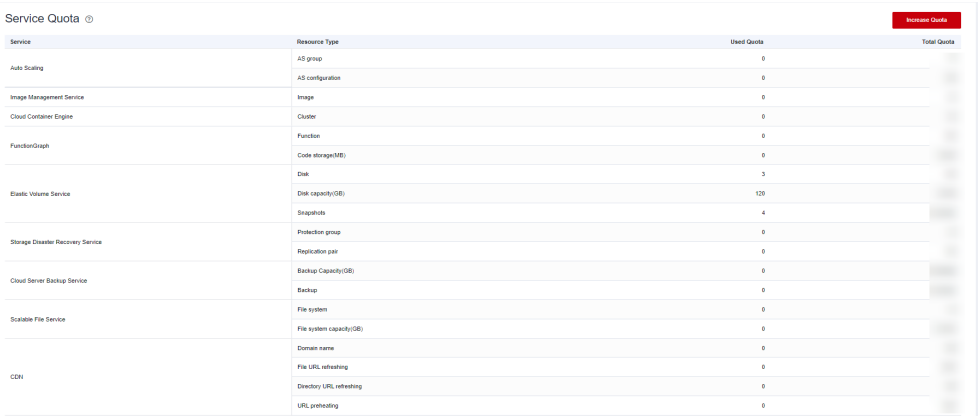

4. On the **Create Service Ticket** page, configure parameters as required.

   In the **Problem Description** area, fill in the content and reason for adjustment.

5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

13 Monitoring

13.1 Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

Server Monitoring includes **Basic Monitoring** and **OS Monitoring**.

- **Basic Monitoring** automatically reports ECS metrics to Cloud Eye.
- Using the agent installed on the target ECS, **OS Monitoring** provides system-wide, active, and fine-grained ECS monitoring.

  For instructions about how to install and configure the agent, see **Server Monitoring** in *Cloud Eye User Guide*.

This section covers the following content:

- Viewing basic ECS metrics
- Viewing OS metrics (Agent installed on ECS)
- Viewing process monitoring metrics (Agent installed on ECS)
- Customizing ECS alarm rules
- Viewing ECS running statuses for routine monitoring

One-Click Monitoring

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. The cloud platform supports automatic recovery by default. 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. For details, see **Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?**

You can enable one-click monitoring on the Cloud Eye console so that you will be notified if high availability occurs (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**.

Helpful Links

- **Why Is My Windows ECS Running Slowly?**
- **Why Is My Linux ECS Running Slowly?**

13.2 Basic ECS Metrics

Description

This section describes basic monitoring metrics reported by ECS to Cloud Eye. You can use Cloud Eye to view these metrics and alarms generated for ECSs.

Namespace

SYS.ECS

Basic ECS Metrics

Basic ECS metrics vary depending on ECS OSs and types. For details, see**Table 13-1**. √ indicates that the metric is supported, and x indicates that the metric is not supported.

NOTE

- Certain ECS metrics require the installation of UVP VMTools on the image from which the ECS is created. For details about how to install UVP VMTools, see **https://github.com/UVP-Tools/UVP-Tools/**.
- Certain ECS metrics require the installation of the Agent on the ECS. After the Agent is installed, log in to the management console and choose **Cloud Eye** under **Management & Deployment**. On the Cloud Eye console, choose **Server Monitoring** > **Elastic Cloud Server** from the left navigation pane to view ECS metrics, such as **AGT. User Space CPU Usage**. For details, see **OS Monitoring Metrics Supported by ECSs with the Agent Installed**.
  - For details about how to install the Agent on a Windows ECS, see "Installing and Configuring the Agent (Windows)" in *Cloud Eye User Guide*.
  - For details about how to install the Agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in *Cloud Eye User Guide*.

Table 13-1 Basic ECS metrics

Metric ID	Metric	Windows		Linux	
-	None	Xen	KVM	Xen	KVM
cpu_util	CPU Usage	Supported	Supported	Supported	Supported

Metric ID	Metric	Windows		Linux	
mem_util	Memory Usage	Supported	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
disk_util_inband	Disk Usage	Supported	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
disk_read_bytes_rate	Disk Read Bandwidth	Supported	Supported	Supported	Supported
disk_write_bytes_rate	Disk Write Bandwidth	Supported	Supported	Supported	Supported
disk_read_requests_rate	Disk Read IOPS	Supported	Supported	Supported	Supported
disk_write_requests_rate	Disk Write IOPS	Supported	Supported	Supported	Supported
network_incoming_bytes_rate_inband	Inband Incoming Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported

Metric ID	Metric	Windows		Linux	
network_outgoing_bytes_rate_inband	Inband Outgoing Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Supported	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
network_incoming_bytes_aggregate_rate	Outband Incoming Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
network_outgoing_bytes_aggregate_rate	Outband Outgoing Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
network_vm_bandwidth_in	Inbound Bandwidth	Not supported	Supported	Not supported	Supported
network_vm_bandwidth_out	Outbound Bandwidth	Not supported	Supported	Not supported	Supported
network_vm_pps_in	Inbound PPS	Not supported	Supported	Not supported	Supported

Metric ID	Metric	Windows		Linux	
network_vm_pps_out	Outbound PPS	Not supported	Supported	Not supported	Supported
network_vm_newconnections	New Connections	Not supported	Supported	Not supported	Supported

Table 13-2 describes these basic ECS metrics.

The monitoring intervals for the following ECSs with raw monitoring metrics are as follows:

- Xen ECSs: 4 minutes
- KVM and QingTian ECSs: 5 minutes

Table 13-2 Basic metric description

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
cpu_util	CPU Usage	CPU usage of an ECS This metric is used to show the CPU usage of the physical server accommodating the monitored ECS, which is not accurate as the CPU usage obtained on the monitored ECS. For details, see **Why Is Basic Monitoring Data Inconsistent with Data Monitored by the OS?** Unit: Percent Formula: CPU usage of an ECS/Number of vCPUs in the ECS	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
mem_util	Memory Usage	Memory usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Unit: Percent Formula: Used memory of an ECS/ Total memory of the ECS **NOTE** The memory usage of QingTian ECSs cannot be monitored.	≥ 0	ECS	5 minutes
disk_util_inband	Disk Usage	Disk usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Unit: Percent Formula: Used capacity of an ECS-attached disk/Total capacity of the ECS-attached disk	≥ 0	ECS	5 minutes
disk_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes read from an ECS-attached disk/ Monitoring interval byte_out = (rd_bytes - last_rd_bytes)/Time difference	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
disk_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes written to an ECS-attached disk/ Monitoring interval	≥ 0	ECS	5 minutes
disk_read_requests_rate	Disk Read IOPS	Number of read requests sent to an ECS-attached disk per second Unit: request/s Formula: Total number of read requests sent to an ECS-attached disk/ Monitoring interval req_out = (rd_req - last_rd_req)/Time difference	≥ 0	ECS	5 minutes
disk_write_requests_rate	Disk Write IOPS	Number of write requests sent to an ECS-attached disk per second Unit: request/s Formula: Total number of write requests sent to an ECS-attached disk/ Monitoring interval req_in = (wr_req - last_wr_req)/Time difference	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_incoming_bytes_rate_inband	Inband Incoming Rate	Number of incoming bytes on an ECS per second Unit: byte/s Formula: Total number of inband incoming bytes on an ECS/Monitoring interval	≥ 0	ECS	5 minutes
network_outgoing_bytes_rate_inband	Inband Outgoing Rate	Number of outgoing bytes on an ECS per second Unit: byte/s Formula: Total number of inband outgoing bytes on an ECS/Monitoring interval	≥ 0	ECS	5 minutes
network_incoming_bytes_aggregate_rate	Outband Incoming Rate	Number of incoming bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband incoming bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_outgoing_bytes_aggregate_rate	Outband Outgoing Rate	Number of outgoing bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband outgoing bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	ECS	5 minutes
network_vm_connections	Network Connections	Total number of TCP and UDP connections to an ECS Unit: count **NOTE** This metric is collected out-of-band and its value may be greater than the number of network connections queried in the OS.	≥ 0	ECS	5 minutes
network_vm_bandwidth_in	Inbound Bandwidth	Number of public and private bits received by the ECS per second Unit: byte/s	≥ 0	ECS	5 minutes
network_vm_bandwidth_out	Outbound Bandwidth	Number of public and private bits sent by the ECS per second Unit: byte/s	≥ 0	ECS	5 minutes
network_vm_pps_in	Inbound PPS	Number of public and private packets received by the ECS per second Unit: packet/s	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_vm_pps_out	Outbound PPS	Number of public and private packets sent by the ECS per second Unit: packet/s	≥ 0	ECS	5 minutes
network_vm_newconnections	New Connections	Number of new connections (including TCP, UDP, and ICMP) created on the ECS Unit: count	≥ 0	ECS	5 minutes

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

13.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed

Description

OS monitoring provides system-level, proactive, and fine-grained monitoring. It requires the Agent to be installed on the ECSs to be monitored. This section describes OS monitoring metrics reported to Cloud Eye.

OS monitoring supports metrics about CPU, CPU load, memory, disk, disk I/O, file system, GPU, NIC, NTP, and TCP.

After the Agent is installed, you can view monitoring metrics of ECSs running different OSs. Monitoring data is collected every 1 minute.

OS Metrics: CPU

Table 13-3 CPU metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object Unit: percent ● Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s)** value. ● Windows: Obtain the metric value using the Windows API **GetSystemTimes**.	0-100	ECS	1 minute
cpu_usage_idle	(Agent) Idle CPU Usage	Percentage of time that CPU is idle Unit: percent ● Linux: Check metric value changes in file **/proc/stat** in a collection period. ● Windows: Obtain the metric value using the Windows API **GetSystemTimes**.	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_user	(Agent) User Space CPU Usage	Percentage of time that the CPU is used by user space Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) us** value. • Windows: Obtain the metric value using the Windows API **GetSystemTimes**.	0-100	ECS	1 minute
cpu_usage_system	(Agent) Kernel Space CPU Usage	Percentage of time that the CPU is used by kernel space Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) sy** value. • Windows: Obtain the metric value using the Windows API **GetSystemTimes**.	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_other	(Agent) Other Process CPU Usage	Percentage of time that the CPU is used by other processes Unit: percent • Linux: **Other Process CPU Usage** = 1- **Idle CPU Usage** - **Kernel Space CPU Usage** - **User Space CPU Usage** • Windows: **Other Process CPU Usage** = 1- **Idle CPU Usage** - **Kernel Space CPU Usage** - **User Space CPU Usage**	0-100	ECS	1 minute
cpu_usage_nice	(Agent) Nice Process CPU Usage	Percentage of time that the CPU is in user mode with low-priority processes which can easily be interrupted by higher-priority processes Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) ni** value. • Windows is not supported currently.	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_iowait	(Agent) iowait Process CPU Usage	Percentage of time that the CPU is waiting for I/O operations to complete Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) wa** value. • Windows is not supported currently.	0-100	ECS	1 minute
cpu_usage_irq	(Agent) CPU Interrupt Time	Percentage of time that the CPU is servicing interrupts Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) hi** value. • Windows is not supported currently.	0-100	ECS	1 minute
cpu_usage_softirq	(Agent) CPU Software Interrupt Time	Percentage of time that the CPU is servicing software interrupts Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s) si** value. • Windows is not supported currently.	0-100	ECS	1 minute

OS Metric: CPU Load

Table 13-4 CPU load metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
load_average1	(Agent) 1-Minute Load Average	CPU load averaged from the last 1 minute Linux: Obtain the metric value from the number of logic CPUs in **load1/** in file **/proc/loadavg**. Run the **top** command to check the **load1** value.	≥ 0	ECS	1 minute
load_average5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes Linux: Obtain the metric value from the number of logic CPUs in **load5/** in file **/proc/loadavg**. Run the **top** command to check the **load5** value.	≥ 0	ECS	1 minute
load_average15	(Agent) 15-Minute Load Average	CPU load averaged from the last 15 minutes Linux: Obtain the metric value from the number of logic CPUs in **load15/** in file **/proc/loadavg**. Run the **top** command to check the **load15** value.	≥ 0	ECS	1 minute

NOTE

The Windows OS does not support the CPU load metrics.

OS Metric: Memory

Table 13-5 Memory metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_available	(Agent) Available Memory	Amount of memory that is available and can be given instantly to processes Unit: GB ● Linux: Obtain the metric value from **/proc/meminfo**. – If **MemAvailable** is displayed in **/proc/meminfo**, obtain the value. – If **MemAvailable** is not displayed in **/proc/meminfo**, **MemAvailable** = **MemFree** + **Buffers+Cached** ● Windows: The metric value is calculated by available memory minuses used memory. The value is obtained by calling the Windows API GlobalMemoryStatusEx.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object Unit: percent ● Linux: Obtain the metric value from the **/proc/meminfo** file: (**MemTotal** - **MemAvailable**)/ **MemTotal** – If **MemAvailable** is displayed in **/proc/meminfo**, **MemUsedPercent** = (**MemTotal**-**MemAvailable**)/ **MemTotal** – If **MemAvailable** is not displayed in **/proc/meminfo**, **MemUsedPercent** = (**MemTotal** – **MemFree** – **Buffers** – **Cached**)/ **MemTotal** ● Windows: The calculation formula is as follows: Used memory size/Total memory size*100%.	0-100	ECS	1 minute
mem_free	(Agent) Idle Memory	Amount of memory that is not being used Unit: GB ● Linux: Obtain the metric value from **/proc/meminfo**. ● Windows is not supported currently.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_buffers	(Agent) Buffer	Amount of memory that is being used for buffers Unit: GB • Linux: Obtain the metric value from **/proc/meminfo**. Run the **top** command to check the **KiB Mem:buffers** value. • Windows is not supported currently.	≥ 0	ECS	1 minute
mem_cached	(Agent) Cache	Amount of memory that is being used for file caches Unit: GB • Linux: Obtain the metric value from **/proc/meminfo**. Run the **top** command to check the **KiB Swap:cached Mem** value. • Windows is not supported currently.	≥ 0	ECS	1 minute
total_open_files	(Agent) Total File Handles	Total handles used by all processes Unit: count • Linux: Use the **/proc/{pid}/fd** file to summarize the handles used by all processes. • Windows is not supported currently.	≥0	ECS	1 minute

OS Metric: Disk

Table 13-6 Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks Unit: GB • Linux: Run the **df -h** command to check the value in the **Avail** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	Total space on the disks, including used and free Unit: GB • Linux: Run the **df -h** command to check the value in the **Size** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	Used space on the disks Unit: GB • Linux: Run the **df -h** command to check the value in the **Used** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_usedPercent	(Agent) Disk Usage	Percentage of total disk space that is used, which is calculated as follows: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity** Unit: percent • Linux: It is calculated as follows: Used/Size. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Use the WMI interface to call GetDiskFreeSpaceExW API to obtain disk space data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS - Mount point	1 minute

OS Metric: Disk I/O

Table 13-7 Disk I/O metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_read_bytes_rate	(Agent) Disks Read Rate	Number of bytes read from the monitored disk per second Unit: byte/s • Linux: The disk read rate is calculated based on the data changes in the sixth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: – Use Win32_PerfFormattedData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. – The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0 bytes/s	• ECS - Disk • ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		– When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data.			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_read_requests_rate	(Agent) Disks Read Requests	Number of read requests sent to the monitored disk per second Unit: request/s • Linux: The disk read requests are calculated based on the data changes in the fourth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: – Use Win32_PerfFormattedData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. – The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). – When the CPU usage is high, monitoring data obtaining timeout	≥ 0 requests/s	• ECS - Disk • ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		may occur and result in the failure of obtaining monitoring data.			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_write_bytes_rate	(Agent) Disks Write Rate	Number of bytes written to the monitored disk per second Unit: byte/s • Linux: The disk write rate is calculated based on the data changes in the tenth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: – Use Win32_PerfFormattedData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. – The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). – When the CPU usage is high, monitoring data obtaining timeout	≥ 0 bytes/s	• ECS - Disk • ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		may occur and result in the failure of obtaining monitoring data.			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_write_requests_rate	(Agent) Disks Write Requests	Number of write requests sent to the monitored disk per second Unit: request/s • Linux: The disk write requests are calculated based on the data changes in the eighth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: – Use Win32_PerfFormattedData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. – The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). – When the CPU usage is high, monitoring data obtaining timeout	≥ 0 requests/s	• ECS - Disk • ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		may occur and result in the failure of obtaining monitoring data.			
disk_readTime	(Agent) Average Read Request Time	Average amount of time that read requests have waited on the disks Unit: ms/count • Linux: The average read request time is calculated based on the data changes in the seventh column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0 ms/Count	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_writeTime	(Agent) Average Write Request Time	Average amount of time that write requests have waited on the disks Unit: ms/count • Linux: The average write request time is calculated based on the data changes in the eleventh column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0 ms/Count	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_ioUtils	(Agent) Disk I/O Usage	Percentage of the time that the disk has had I/O requests queued to the total disk operation time Unit: percent • Linux: The disk I/O usage is calculated based on the data changes in the thirteenth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	0-100	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_queue_length	(Agent) Disk Queue Length	Average number of read or write requests queued up for completion for the monitored disk in the monitoring period Unit: count • Linux: The average disk queue length is calculated based on the data changes in the fourteenth column of the corresponding device in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_write_bytes_per_operation	(Agent) Average Disk Write Size	Average number of bytes in an I/O write for the monitored disk in the monitoring period Unit: byte/op • Linux: The average disk write size is calculated based on the data changes in the tenth column of the corresponding device to divide that of the eighth column in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0 bytes/op	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_read_bytes_per_operation	(Agent) Average Disk Read Size	Average number of bytes in an I/O read for the monitored disk in the monitoring period Unit: byte/op • Linux: The average disk read size is calculated based on the data changes in the sixth column of the corresponding device to divide that of the fourth column in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0 bytes/op	• ECS - Disk • ECS - Mountpoint	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_io_svctm	(Agent) Disk I/O Service Time	Average time in an I/O read or write for the monitored disk in the monitoring period Unit: ms/op • Linux: The average disk I/O service time is calculated based on the data changes in the thirteenth column of the corresponding device to divide the sum of data changes in the fourth and eighth columns in file **/proc/diskstats** in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows is not supported currently.	≥ 0	• ECS - Disk • ECS - Mountpoint	1 minute

OS Metric: File System

Table 13-8 File system metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_fs_rwstate	(Agent) File System Read/ Write Status	Read and write status of the mounted file system of the monitored object Possible values are **0** (read and write) and **1** (read only). Linux: Check file system information in the fourth column in file **/proc/mounts**.	• **0**: readable and writable • **1**: read-only	ECS - Mount point	1
disk_inodesTotal	(Agent) Disk inode Total	Total number of index nodes on the disk Linux: Run the **df -i** command to check the value in the **Inodes** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute
disk_inodesUsed	(Agent) Total inode Used	Number of used index nodes on the disk Linux: Run the **df -i** command to check the value in the **IUsed** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_inodesUsedPercent	(Agent) Percentage of Total inode Used	Number of used index nodes on the disk Unit: percent Linux: Run the **df -i** command to check the value in the **IUse%** column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS - Mount point	1 minute

NOTE

The Windows OS does not support the file system metrics.

OS Metric: NIC

Table 13-9 NIC metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_bitRecv	(Agent) Outbound Bandwidth	Number of bits sent by this NIC per second Unit: bit/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 bit/s	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_bitSent	(Agent) Inbound Bandwidth	Number of bits received by this NIC per second Unit: bit/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 bit/s	ECS	1 minute
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Unit: count/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 Counts/s	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Unit: count/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Use the MibIfRow object in the WMI to obtain network metric data.	≥ 0 Counts/s	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_errin	(Agent) Receive Error Rate	Percentage of receive errors detected by this NIC per second Unit: percent • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows is not supported currently.	0-100	ECS	1 minute
net_errout	(Agent) Transmit Error Rate	Percentage of transmit errors detected by this NIC per second Unit: percent • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows is not supported currently.	0-100	ECS	1 minute
net_dropin	(Agent) Received Packet Drop Rate	Percentage of packets received by this NIC which were dropped per second Unit: percent • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows is not supported currently.	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_dropout	(Agent) Transmitted Packet Drop Rate	Percentage of packets transmitted by this NIC which were dropped per second Unit: percent • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows is not supported currently.	0-100	ECS	1 minute

OS Metric: NTP

Table 13-10 NTP metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
ntp_offset	(Agent) NTP Offset	NTP offset of the monitored object Unit: ms Linux: Run the **nvidia-smi** command to check the value in the **Perf** column.	≥ 0 ms	ECS	1 minute

OS Metric: TCP

Table 13-11 TCP metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_total	(Agent) TCP TOTAL	Total number of TCP connections Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_established	(Agent) TCP ESTABLISHED	Number of ESTABLISHED TCP connections Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_sys_sent	(Agent) TCP SYS_SENT	Number of TCP connections that are being requested by the client Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_sys_recv	(Agent) TCP SYS_RECV	Number of pending TCP connections received by the server Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_fin_wait1	(Agent) TCP FIN_WAIT1	Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_fin_wait2	(Agent) TCP FIN_WAIT2	Number of TCP connections in the FIN_WAIT2 state Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_close	(Agent) TCP CLOSE	Number of closed TCP connections Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_last_ack	(Agent) TCP LAST_ACK	Number of TCP connections waiting for ACK packets when the connections are being passively closed by the client Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_listen	(Agent) TCP LISTEN	Number of TCP connections in the LISTEN state Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_closing	(Agent) TCP CLOSING	Number of TCP connections to be automatically closed by the server and the client at the same time Unit: count • Linux: Obtain TCP connections in all states from the **/proc/net/tcp** file, and then collect the number of connections in each state. • Windows: Obtain the metric value using WindowsAPI GetTcpTable2.	≥ 0	ECS	1 minute
net_tcp_retrans	(Agent) TCP Retransmission Rate	Percentage of packets that are resent Unit: percent • Linux: Obtain the metric value from the **/proc/net/snmp** file. The value is the ratio of the number of sent packets to the number of retransmitted packages in a collection period. • Windows: Obtain the metric value using WindowsAPI GetTcpStatistics.	0-100%	ECS	1 minute

OS Metric: GPU

Table 13-12 GPU metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_status	GPU Health Status	Overall measurement of the GPU health Unit: none • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	• 0: The GPU is healthy. • 1: The GPU is subhealthy. • 2: The GPU is faulty.	• ECS • ECS - GPU	1 minute
gpu_usage_encoder	Encoding Usage	Encoding capability usage on the GPU Unit: percent • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	0-100%	• ECS • ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_usage_decoder	Decoding Usage	Decoding capability usage on the GPU Unit: percent • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	0-100%	• ECS • ECS - GPU	1 minute
gpu_volatile_correctable	Volatile Correctable ECC Errors	Number of correctable ECC errors since the GPU is reset. The value is reset to **0** each time the GPU is reset. Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_volatile_uncorrectable	Volatile Uncorrectable ECC Errors	Number of uncorrectable ECC errors since the GPU is reset. The value is reset to **0** each time the GPU is reset. Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute
gpu_aggregate_correctable	Aggregate Correctable ECC Errors	Aggregate correctable ECC errors on the GPU Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute
gpu_aggregate_uncorrectable	Aggregate Uncorrectable ECC Errors	Aggregate uncorrectable ECC Errors on the GPU Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_retired_page_single_bit	Retired Page Single Bit Errors	Number of retired page single bit errors, which indicates the number of single-bit pages blocked by the graphics card Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute
gpu_retired_page_double_bit	Retired Page Double Bit Errors	Number of retired page double bit errors, which indicates the number of double-bit pages blocked by the graphics card Unit: count • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	≥ 0	• ECS • ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_performance_state	(Agent) Performance Status	GPU performance of the monitored object Unit: none • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	P0-P15, P32 • **P0**: indicates the maximum performance status. • **P15**: indicates the minimum performance status. • **P32**: indicates the unknown performance status.	ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_usage_mem	(Agent) GPU Memory Usage	GPU memory usage of the monitored object Unit: percent • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	0-100	ECS - GPU	1 minute
gpu_usage_gpu	(Agent) GPU Usage	GPU usage of the monitored object Unit: percent • Linux: Obtain the metric value using the **libnvidia-ml.so.1** library file of the graphics card. • Windows: Obtain the metric value using the **nvml.dll** library of the graphics card.	0-100	ECS - GPU	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.
ECS - Disk	disk	Specifies the disks attached to an ECS.
ECS - Mount point	mount_point	Specifies the mount point of a disk.
ECS - GPU	gpu	Specifies the graphics card of an ECS.

13.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed

Description

Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. By default, Cloud Eye collects CPU usage, memory usage, and number of opened files of active processes.

This section describes process monitoring metrics reported to Cloud Eye.

Process Metrics

After the agent is installed, you can view the default process metrics listed in the following table, regardless of ECS types and OSs.

Table 13-13 Process metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_cpu	CPU Usage	CPU consumed by a process. **pHashId** (process name and process ID) is the value of **md5**. Unit: percent • Linux: Check metric value changes in file **/proc/pid/stat**. • Windows: Call the Windows API GetProcessTimes to obtain the CPU usage of the process.	0-100%	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_mem	Memory Usage	Memory consumed by a process. **pHashId** (process name and process ID) is the value of **md5**. Unit: percent • Linux: RSS*PAGESIZE/MemTotal Obtain the **RSS** value by checking the second column of file **/proc/pid/statm**. Obtain the **PAGESIZE** value by running the **getconf PAGESIZE** command. Obtain the **MemTotal** value by checking file **/proc/meminfo**. • Windows: Call the Windows API procGlobalMemoryStatusEx to obtain the total memory size. Call GetProcessMemoryInfo to obtain the used memory size. Use the used memory size to divide the total memory size to get the memory usage.	0-100%	ECS	1 minute
proc_pHashId_file	Open Files	Number of files opened by a process. **pHashId** (process name and process ID) is the value of **md5**. • Linux: Run the **ls -l /proc/pid/fd** command to view the number of opened files. • Windows is not supported currently.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_running_count	(Agent) Running Processes	Number of processes that are running ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows is not supported currently.	≥0	ECS	1 minute
proc_idle_count	(Agent) Idle Processes	Number of processes that are idle ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows is not supported currently.	≥0	ECS	1 minute
proc_zombie_count	(Agent) Zombie Processes	Number of zombie processes ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows is not supported currently.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_blocked_count	(Agent) Blocked Processes	Number of processes that are blocked ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows is not supported currently.	≥0	ECS	1 minute
proc_sleeping_count	(Agent) Sleeping Processes	Number of processes that are sleeping ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows is not supported currently.	≥0	ECS	1 minute
proc_total_count	(Agent) Total Processes	Total number of processes on the monitored object ● Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. ● Windows: Obtain the total number of processes by using the system process status support module **psapi.dll**.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_specified_count	(Agent) Specified Processes	Number of specified processes • Linux: You can obtain the state of each process by checking the **Status** value in the **/proc/pid/status** file, and then collect the total number of processes in each state. • Windows: Obtain the total number of processes by using the system process status support module **psapi.dll**.	≥0	• ECS • ECS - Process	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.
ECS - Process	proc	Specifies the ECS process.

13.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed and Using Simplified Monitoring Metrics

Description

This section describes the OS metrics supported by ECSs. In the following regions, the agent of the latest version is used with simplified monitoring metrics:

CN East-Shanghai1, CN East-Shanghai2, CN North-Beijing1, CN North-Beijing4, CN South-Guangzhou, CN-Hong Kong, AP-Bangkok, AP-Singapore, and AF-Johannesburg.

After installing the agent on an ECS, you can view its OS monitoring metrics. Monitoring data is collected every 1 minute.

OS Monitoring Metrics

Table 13-14 OS monitoring metrics

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object Unit: percent • Linux: Check metric value changes in file **/proc/stat** in a collection period. Run the **top** command to check the **%Cpu(s)** value. • Windows: Obtain the metric value using the Windows API GetSystemTimes.	0-100	ECS	1 minute
load_average5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes • Linux: Obtain the metric value from the number of logic CPUs in **load5/** in file **/proc/loadavg**. Run the **top** command to check the **load5** value. • Windows does not support this metric.	≥ 0	ECS	1 minute
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object Unit: percent • Linux: Obtain the metric value from the **/proc/meminfo** file: (**MemTotal** - **MemAvailable**)/**MemTotal** • Windows: Obtain the value using the following formula: Used memory size/Total memory size x 100%	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mountPointPrefix_disk_free	(Agent) Available Disk Space	Free disk space Unit: GB • Linux: Run the **df -h** command to check the value in the **Avail** column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mountPointPrefix_disk_usedPercent	(Agent) Disk Usage	Percentage of total disk space that is used Unit: percent • Linux: Obtain the metric value using following formula: **Disk Usage** = **Used Disk Space**/**Disk Storage Capacity**. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows: Obtain the metric value using the WMI API GetDiskFreeSpaceExW. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mountPointPrefix_disk_ioUtils and volumePrefix_disk_ioUtils	(Agent) Disk I/O Usage	Percentage of the time that the disk has had I/O requests queued to the total disk operation time Unit: percent • Linux: Obtain the metric value by calculating the data changes in the thirteenth column of the monitored object in file **/proc/diskstats** in a collection period. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows does not support this metric.	0-100	ECS	1 minute
mountPointPrefix_disk_inodesUsedPercen	(Agent) Percentage of Total inode Used	Number of used index nodes on the disk Unit: percent • Linux: Run the **df -i** command to check the value in the **IUse%** column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~). • Windows does not support this metric.	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
net_bitSent	(Agent) Inbound Bandwidth	Number of bits received by the monitored object per second Unit: bit/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Obtain the metric value using the WMI MibIfRow object.	≥ 0	ECS	1 minute
net_bitRecv	(Agent) Outbound Bandwidth	Number of bits sent by the target NIC per second Unit: bit/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Obtain the metric value using the WMI MibIfRow object.	≥ 0	ECS	1 minute
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Unit: count/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Obtain the metric value using the WMI MibIfRow object.	≥ 0	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Unit: count/s • Linux: Check metric value changes in file **/proc/net/dev** in a collection period. • Windows: Obtain the metric value using the WMI MibIfRow object.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
net_tcp_total	(Agent) Total Number of TCP Connections	Total number of TCP connections of this NIC	≥ 0	ECS	1 minute
net_tcp_established	(Agent) Number of ESTABLISHED TCP Connections	Number of ESTABLISHED TCP connections of this NIC	≥ 0	ECS	1 minute

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

13.6 Setting Alarm Rules

Scenarios

Setting ECS alarm rules allows you to customize the monitored objects and notification policies so that you can closely monitor your ECSs.

This section describes how to set ECS alarm rules, including alarm rule names, monitoring objects, monitoring metrics, alarm thresholds, monitoring intervals, and notifications.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Management & Governance**, choose **Cloud Eye**.
4. In the navigation pane on the left, choose **Alarm Management** > **Alarm Rules**.
5. On the **Alarm Rules** page, click **Create Alarm Rule** to create an alarm rule, or modify an existing alarm rule.

The following uses modifying an existing alarm rule as an example.

a. Click the target alarm rule.
b. Click **Modify** in the upper right corner of the page.
c. On the **Modify Alarm Rule** page, set parameters as prompted.
d. Click **Modify**.

   After an alarm rule is modified, the system automatically notifies you of an alarm when the alarm complying with the alarm rule is generated.

NOTE

For more information about ECS alarm rules, see **Cloud Eye User Guide**.

13.7 Viewing ECS Metrics

Scenarios

The cloud platform provides Cloud Eye, which monitors the running statuses of your ECSs. You can obtain the monitoring metrics of each ECS on the management console.

There a short time delay between transmission and display of monitoring data. The status of an ECS displayed on Cloud Eye is the status obtained 5 to 10 minutes before. If an ECS is just created, wait for 5 to 10 minutes to view the real-time monitoring data.

Prerequisites

- The ECS is running properly.

  Cloud Eye does not display the monitoring data for a stopped, faulty, or deleted ECS. After such an ECS restarts or recovers, the monitoring data is available in Cloud Eye.

  NOTE

  Cloud Eye discontinues monitoring ECSs that remain in **Stopped** or **Faulty** state for 24 hours and removes them from the monitoring list. However, the alarm rules for such ECSs are not automatically deleted.

- Alarm rules have been configured in Cloud Eye for the target ECS.

  The monitoring data is unavailable for the ECSs without alarm rules configured in Cloud Eye. For details, see **Setting Alarm Rules**.

- The target ECS has been properly running for at least 10 minutes.

  The monitoring data and graphics are available for a new ECS after the ECS runs for at least 10 minutes.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click  . Under **Compute**, click **Elastic Cloud Server**.

4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS. The page providing details about the ECS is displayed.
6. Click the **Monitoring** tab to view the monitoring data.
7. In the ECS monitoring area, select a duration to view the monitoring data.

   You can view the monitoring data of the ECS in the last 1 hour, last 3 hours, last 12 hours, last 1 day, or last 7 days.

14 CTS

14.1 Key Operations Supported by CTS

Scenarios

Cloud Trace Service (CTS) records user operations performed on ECSs and related resources for further query, auditing, and backtracking.

Prerequisites

CTS has been provisioned.

Key ECS Operations Recorded by CTS

Table 14-1 ECS operations recorded by CTS

Operation	Resource Type	Event Name
Creating an ECS	ecs	createServer createServerV2 createServerV21
Deleting an ECS	ecs	deleteServer deleteServerV2 deleteServerV21
Starting an ECS	ecs	startServer
Restarting an ECS	ecs	rebootServer
Stopping an ECS	ecs	stopServer
Adding an ECS NIC	ecs	addNic
Deleting an ECS NIC	ecs	deleteNic delNic

Operation	Resource Type	Event Name
Attaching a disk	ecs	attachVolume attachVolumeV2
Attaching a disk (on the EVS console)	ecs	attachVolume2
Detaching a disk	ecs	detachVolume
Reinstalling an OS	ecs	reinstallOs
Changing an OS	ecs	changeOs
Modifying specifications	ecs	resizeServer
Enabling automatic recovery on an ECS	ecs	addAutoRecovery
Disabling automatic recovery on an ECS	ecs	deleteAutoRecovery
Updating metadata/Setting metadata of a specified key	ecs	updateMetadata
Logging in to an ECS using VNC	ecs	remoteConsole
Modifying ECS information	ecs	updateServer

14.2 Viewing Audit Logs

Scenarios

CTS starts to record ECS operations after it is provisioned. You can view the operation records of the last seven days on the management console.

This section describes how to view the operation records.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click **Service List**. Under **Management & Governance**, choose **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Trace List**.
5. Specify filter criteria as needed. The following filter criteria are available:
   - **Trace Type**, **Trace Source**, **Resource Type**, and **Search By**:

     Select a filter criterion from the drop-down list.

     If you select **Trace name** for **Search By**, you need to select a specific trace name.

If you select **Resource ID** for **Search By**, you need to select or enter a specific resource ID.

When you select **Resource name** for **Search By**, you need to select or enter a specific resource name.

- **Operator**: Select a specific operator (which is a user rather than the tenant).
- **Trace Status**: Available options include **All trace statuses**, **Normal**, **Warning**, and **Incident**. You can only select one of them.
- **Time Range**: In the upper-right corner, you can select any time range of the last seven days to view traces generated during that period.

6. Expand the trace for details.

Figure 14-1 Expanding trace details

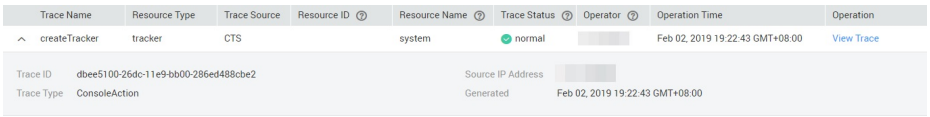

7. Click **View Trace**. A dialog box is displayed, in which the trace structure details are displayed.

   For more information about CTS, see **Cloud Trace Service User Guide**.

A Change History

Released On	Description
2023-07-31	This issue is the forty-second official release. Updated the procedure for installing the one-click password reset plug-in in the following sections: • **Installing the One-Click Password Reset Plug-in on an ECS** • **Updating the One-Click Password Reset Plug-in for an ECS**
2023-06-28	This issue is the forty-first official release. Added **Login Using CloudShell**. Added descriptions about CloudShell login in **Login Overview**.
2023-05-31	This issue is the fortieth official release. Modified the following content: • Modified billing rules for stopped ECSs in **Pay-per-Use Billing**. • Added key operations supported by CTS in **Key Operations Supported by CTS**. • Added dimensions in **OS Monitoring Metrics Supported by ECSs with the Agent Installed** and **Process Monitoring Metrics Supported by ECSs with the Agent Installed**.
2023-01-13	This issue is the thirty-ninth official release. Added the following content: • **Starting and Stopping ECSs** • **Uninstalling a GPU Driver from a GPU-accelerated ECS** Modified the following content: Added billing conversion rules for associated resources when ECS's billing mode is changed from pay-per-use to yearly/monthly in **Changing Pay-per-Use to Yearly/Monthly**.

Released On	Description
2022-08-29	This issue is the thirty-eighth official release. Added **Obtaining the One-Click Password Reset Plug-in**.
2022-05-16	This issue is the thirty-seventh official release. Modified content in the following (Launch Template and Auto Launch Group were commercially used): • **Overview** • **Overview**
2022-04-20	This issue is the thirty-sixth official release. Optimized the content in "Passwords and Key Pairs". Added the following content: • **Application Scenarios for Using Passwords** • **Application Scenarios for Using Key Pairs** • **Creating a Key Pair Using PuTTYgen** • **Importing a Key Pair** Modified the following content: **(Recommended) Creating a Key Pair on the Management Console**
2022-03-28	This issue is the thirty-fifth official release. Added the prerequisites for disabling SELinux in **Installing the One-Click Password Reset Plug-in on an ECS**.
2022-01-07	This issue is the thirty-fourth official release. Added **Updating the One-Click Password Reset Plug-in for an ECS**.
2021-11-12	This issue is the thirty-third official release. • Modified the content in **Enabling and Purchasing a Reserved Instance**.
2021-09-25	This issue is the thirty-second official release. • Modified the procedure on Linux proxy ECS in **Enabling Internet Connectivity for an ECS Without an EIP**. • Optimized the description of scenarios and constraints in **Changing the OS**.
2021-09-02	This issue is the thirty-first official release. • Added the differences between backups, snapshots, and images in section **Overview**. • Modified the procedure for uninstalling plug-ins in Linux OS in **Installing the One-Click Password Reset Plug-in on an ECS**.

Released On	Description
2021-06-30	This issue is the thirtieth official release. • Added the procedure for changing the file permission of the password reset plug-in in **Installing the One-Click Password Reset Plug-in on an ECS**. • Modified the IDs of the (Agent) Disk I/O Usage and (Agent) Percentage of Total inode Used in **OS Monitoring Metrics Supported by ECSs with the Agent Installed and Using Simplified Monitoring Metrics**.
2021-06-11	This issue is the twenty-ninth official release. Added the following content: • **Migrating an ECS**
2021-05-20	This issue is the twenty-eighth official release. Added the following content: • **Launch Templates** • **Auto Launch Groups**
2021-03-10	This issue is the twenty-seventh official release. Modified the following content: • Added the URL for downloading the PV driver and UVP VMTools in **Changing a Xen ECS to a KVM ECS (Windows)**.
2021-02-20	This issue is the twenty-sixth official release. Modified the following content: • Modified steps in **Changing Yearly/Monthly to Pay-per-Use**.
2020-12-17	This issue is the twenty-fifth official release. Added the following content: • **OS Monitoring Metrics Supported by ECSs with the Agent Installed and Using Simplified Monitoring Metrics**
2020-09-29	This issue is the twenty-fourth official release. Added the following content: • **Changing a Xen ECS to a KVM ECS (Windows)** • **Automatically Changing a Xen ECS to a KVM ECS (Linux)** • **Manually Changing a Xen ECS to a KVM ECS (Linux)**
2020-08-31	This issue is the twenty-third official release. Added the following content: • Added notes and firewall configuration examples in **Obtaining Metadata**. • Deleted the example of using a plaintext password in **Passing User Data to ECSs**.

Released On	Description
2020-06-09	This issue is the twenty-second official release. Added the following content: • **Dynamically Assigning IPv6 Addresses**
2020-05-06	This issue is the twenty-first official release. Added the following content: • **Adding a Disk to an ECS** • **Unbinding an EIP** Modified the following content: • Optimized the operations for installing a Tesla driver and CUDA toolkit in **Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS**. • Added **GPU Driver** to describe GRID and Tesla drivers. • Modified operations in **Reinstalling the OS** and **Changing the OS**. • **Binding an EIP**
2020-02-18	This issue is the twentieth official release. • Added **Logging In to a Windows ECS from a Mac**. • Added **Overview**. • Added **Overview**.
2019-11-28	This issue is the nineteenth official release. • Added constraints in **Managing ECS Groups**. • Modified description in **Resetting the Password for Logging In to an ECS on the Management Console** because login passwords can be reset when the target ECSs are running. • Moved "Changing the Login Password on an ECS" to FAQs. • Moved "Resetting the Forgotten Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed" to FAQs. • Moved "Resetting the Password for Logging In to a Linux ECS Without the Password Reset Plug-in Installed" to FAQs. • Added an example for setting a time zone in **Changing the Time Zone for an ECS**.

Released On	Description
2019-10-28	This issue is the eighteenth official release. Added the following content: • **Installing a GRID Driver on a GPU-accelerated ECS** • "Installing a NVIDIA GPU Driver and CUDA Toolkit (Recommended)" • **Installing a Tesla Driver and CUDA Toolkit on a GPU-accelerated ECS** • Modified **Managing ECS Groups**. • Modified NIC multi-queue supported by Linux ECSs in **Enabling NIC Multi-Queue**.
2019-10-25	This issue is the seventeenth official release. Modified the following content: • **Basic ECS Metrics** • **Viewing ECS Creation Statuses** Added the following content: • **OS Monitoring Metrics Supported by ECSs with the Agent Installed**
2019-09-23	This issue is the sixteenth official release. Modified the following content: • Modified application scenarios in **Changing an EIP Bandwidth**. • Modified application scenarios in **Changing an EIP**. • Added notes for adding an ECS to an ECS group in **Managing ECS Groups**.
2019-09-06	This issue is the fifteenth official release. Added the following content: • **Spot Pricing** • **Purchasing a Spot ECS** Modified the following content: • **Configuring Security Group Rules** • **Resetting the Password for Logging In to an ECS on the Management Console**
2019-07-30	This issue is the fourteenth official release. Added the following content: • **Logging In to a Windows ECS from a Mobile Terminal** • **Logging In to a Linux ECS from a Mobile Terminal**

Released On	Description
2019-06-30	This issue is the thirteenth official release. Added the following content: • **Yearly/Monthly Billing** • **Pay-per-Use Billing** • **Changing Yearly/Monthly to Pay-per-Use** Modified the following content: • Optimized the document structure. • Modified the character set for resetting a login password.
2019-05-30	This issue is the twelfth official release. Modified the following content: • Added the example of user data to **Passing User Data to ECSs**. • Deleted "Troubleshooting". • Added constraints in **General Operations**. • **Enabling NIC Multi-Queue** • **Login Using MSTSC**
2019-04-03	This issue is the eleventh official release. Added the following content: • **Changing Pay-per-Use to Yearly/Monthly**
2019-03-04	This issue is the tenth official release. Added the following content: • Purchasing an RI • Modifying an RI Modified the following content: • Modified ECS metadata types in **Obtaining Metadata**. • Added use restrictions in **Passing User Data to ECSs**.
2019-02-28	This issue is the ninth official release. Modified the following content: • Modified ECS metrics in **Basic ECS Metrics**.
2018-11-19	This issue is the eighth official release. Modified the following content: • Updated supported images in **Enabling NIC Multi-Queue**.

Released On	Description
2018-07-30	This issue is the seventh official release. Added the following content: • Changing the Login Password on an ECS Modified the following content: • **Installing the One-Click Password Reset Plug-in on an ECS** • Discarded "6.1 Application Scenarios" for resetting the password for logging in to an ECS. • Modified description in **Changing a Security Group**, allowing you to change the security group in the **Operation** column.
2018-06-30	This issue is the sixth official release. Added the following content: • **CTS** Modified the following content: • Allowed to export certain ECSs in **Exporting ECS Information**. • Modified prerequisites in **Changing the OS**, allowing you to change the OS of an ECS on which reinstalling the OS failed.
2018-05-30	This issue is the fifth official release. Modified the following content: • Changed the term EIP in Chinese. • Modified the description of **local-ipv4** and **public-ipv4** in **Obtaining Metadata**. • Modified **Basic ECS Metrics** because the monitoring metric **System Status Check Failed** has been terminated. • Added introduction to user data scripts in **Passing User Data to ECSs**.
2018-04-30	This issue is the fourth official release. Added the following content: • **Creating an Image** • **Viewing Failed Tasks** Modified the following content: • Added description in **Login Using VNC**, indicating that data can be copied and pasted on VNC pages.
2018-01-30	This issue is the third official release. Modified the following content: • Added "Follow-up Procedure" in **General Operations**.

Released On	Description
2018-01-04	This issue is the second official release. Added the following content: • Added one-click password reset in **Installing the One-Click Password Reset Plug-in on an ECS**. • **Resetting the Password for Logging In to an ECS on the Management Console**
2017-12-31	This issue is the first official release.