Elastic Cloud Server

User Guide

Issue 01

Date 2025-05-13

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Contents

1 Using IAM to Grant Access to ECS

1.1 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 personnel based on your enterprise's organizational structure. Each IAM user has their own identity credentials for accessing ECS resources.
- Grant only the permissions required for users to perform a specific task.
- Delegate access to other Huawei Cloud accounts or cloud services for efficient O&M.

If your Huawei Cloud account does not require individual IAM users, you can 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-defined policies supported by ECS and select the policies based on service requirements.

For details about system-defined policies supported by ECS, see **ECS system-defined policies**. To grant permissions to other services, see **System-defined Permissions**.

Process Flow

Figure 1-1 Process for granting ECS permissions

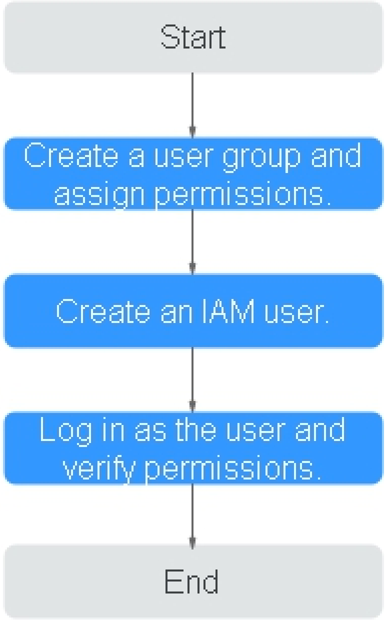
Start
Create a user group and assign permissions.
Create an IAM user.
Log in as the user and verify permissions.
End

1. Create a user group and assign permissions.

   Create a user group on the IAM console and assign the **ECS ReadOnlyAccess** permissions to the group.

2. Create a user and add the user to the user group.

   Create a user on the IAM console and add the user to the group created in step 1.

3. Log in to the management console as the created user.

   In the authorized region, perform the following operations:

   - Choose **Compute** > **Elastic Cloud Server** in the service list. On the ECS console, click **Create ECS**. If the purchase attempt failed, the **ECSReadOnlyAccess** policy has already taken effect.
   - Choose any service other than ECS in the service list. If a message appears indicating that you have insufficient permissions to access the service, the **ECSReadOnlyAccess** policy has already taken effect.

1.2 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 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 provides examples of common ECS custom policies.

Example Custom Policies

- Example 1: Only allowing users to start, stop, and restart ECSs in batches

  ```
  {
      "Version": "1.1",
      "Statement": [
          {
              "Effect": "Allow",
              "Action": [
                  "ecs:cloudServerFlavors:get",
                  "ecs:cloudServers:reboot",
                  "ecs:cloudServers:start",
                  "ecs:cloudServers:get",
                  "ecs:cloudServers:list",
                  "ecs:cloudServers:stop"
              ]
          }
      ]
  }
  ```

- Example 2: Only allowing users to stop and delete ECSs in batches

  ```
  {
      "Version": "1.1",
      "Statement": [
          {
              "Effect": "Allow",
              "Action": [
                  "ecs:cloudServers:get",
                  "ecs:cloudServers:delete",
                  "ecs:cloudServers:list",
                  "ecs:cloudServers:stop"
              ]
          }
      ]
  }
  ```

- Example 3: Only allowing VNC login

  ```
  {
      "Version": "1.1",
      "Statement": [
          {
              "Effect": "Allow",
              "Action": [
                  "ecs:cloudServerFlavors:get",
                  "ecs:cloudServers:vnc",
                   "ecs:cloudServers:get",
                  "ecs:cloudServers:list"
               ]
          }
      ]
  }
  ```

- Example 4: 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 which the user belongs to. 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"
                ]
            }
      ]
}
```

2 Instances

2.1 Overview

ECS Overview

An Elastic Cloud Server (ECS) is a basic computing unit that consists of vCPUs, memory, OS, and Elastic Volume Service (EVS) disks.

After creating an ECS, you can use it like using your local computer or physical server, ensuring a secure, reliable, and efficient computing environment. ECSs support self-service creation, modification, and operation. You can create an ECS by specifying its vCPUs, memory, OS, and login authentication. After the ECS is created, you can modify its specifications as required.

There are a wide range of ECS types available to meet your compute and storage requirements. Each ECS type offers various flavors with different vCPU and memory configurations for you to choose from.

- For details about ECS types, see **ECS Types**.
- For details about ECS specifications, see **A Summary List of x86 ECS Specifications**.

Instance Configuration

You can configure parameters such as vCPUs, memory, OS, storage, and network for an ECS.

Table 2-1 Instance configurations

Item	Description	Reference
Specifications	ECS provides a range of x86 or Kunpeng ECS specifications. ECS specifications define the number of vCPUs, memory size, assured/ maximum intranet bandwidth, maximum packets per second (PPS), IPv6 support, and other configurations.	● **A Summary List of x86 ECS Specifications**
Image	An image contains the OS, application software, and initialized applications required by an ECS. You can choose from public, private, shared, and KooGallery images.	**Overview**
Storage	ECSs store data through the attached cloud disks and local disks. ● Cloud disks: are created from the dedicated distributed storage pool or EVS disks and can be used as system disks or data disks. ● Local disks: can only be used as data disks. ECSs of certain specifications have local disks by default.	**Overview**
Network	Virtual Private Cloud (VPC) allows you to create logically isolated, configurable, and manageable virtual networks for ECSs. You can configure NICs, private IP addresses, and security groups in your VPC. By default, ECSs in different VPCs cannot communicate with each other.	**Overview**

Instance Selection

Table 2-2 Instance selection

Method	Description	Reference
By Type	ECS specifications are filtered out by CPU architecture, number of vCPUs, memory size, flavor name, and instance family. You can select this mode if you want to select specific specifications of ECSs.	**Purchasing an ECS in Custom Config Mode**
By Scenario	ECS specifications are filtered out based on categories such as web applications, website applications/e-commerce, gaming, and databases as well as service volumes in different scenarios. You can select this mode if you have specific service requirements and want to select specifications based on service scenarios and volumes.	

Specifications Change

After an ECS is created, you can modify its vCPUs, memory, OS, storage, and bandwidth as required.

- To modify its vCPUs and memory, see **Modifying Specifications of Individual ECSs**.
- To modify its OS, see **Changing the OS**.
- To modify its storage, see **Adding a Disk to an ECS** and **Expanding the Capacity of an EVS Disk**.
- To modify its bandwidth, see **Modifying an EIP Bandwidth**.

2.2 Selecting an ECS Billing Mode

2.2.1 Yearly/Monthly Billing

Concept

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

For more billing information, see **Yearly/Monthly 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 anymore, unsubscribe from 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 KooGallery 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 bandwidth.

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

For details about ECS prices, see **Price Calculator**.

2.2.2 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 the 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 **Pay-per-Use 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 for each type of ECS can be obtained by dividing their hourly price by 3,600. 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 (0.68/3,600 × 30 × 60).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for $1.02 USD (0.68/3,600 × 90 × 60).

Resources Supporting Pay-per-Use Billing

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

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

  For details about ECS prices, see **Price Calculator**.

2.2.3 Spot Pricing

2.2.3.1 Spot Pricing ECSs

Concept

Huawei Cloud sells available compute 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. Paying a higher price can increase your chances of successfully purchasing a spot 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 2-1 Lifecycle of a spot ECS

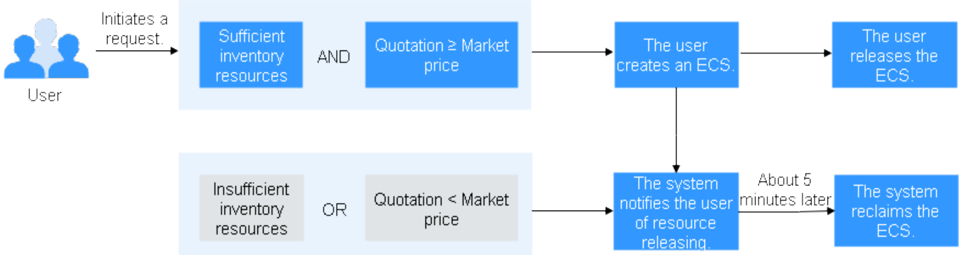
Initiates a request.
User
Sufficient inventory resources
AND
Quotation ≥ Market price
The user creates an ECS.
The user releases the ECS.
Insufficient inventory resources
OR
Quotation < Market price
The system notifies the user of resource releasing.
About 5 minutes later
The system reclaims the ECS.

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 KooGallery image.
- A spot ECS cannot be changed to a pay-per-use or yearly/monthly ECS.
- Spot ECSs do not support system disk detachment.
- 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 the data disks and EIP, set a reclaim notification. For details, see "Enabling Reclaim Notifications" in **Purchasing a Spot ECS**.

Billing Rules

See **Spot Pricing (for Spot Instances)**.

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 spot 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 instance. If during that 5 minutes, the spot price 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 **FAQs About Spot ECSs**.

2.2.4 Reserved Instances

2.2.4.1 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.

Table 2-3 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	A billing mode based on the usage frequency and duration. Pay-per-use ECSs can be created or deleted at any time.	A pay-per-use ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing a pay-per-use ECS, you can use it on the cloud.
Yearly/Monthly	A billing mode based on the required duration. This mode is ideal when the duration of ECS usage is predictable.	A yearly/monthly ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing a yearly/monthly ECS, you can use it on the cloud.
Spot pricing	A spot pricing billing mode.	A spot ECS is a basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing a spot 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 for a region and without an AZ specified. A zonal RI is purchased for an AZ.

- Attribute matching of a regional RI: The instance series, vCPU/memory ratio, and OS of a regional RI must be the same as those of a pay-per-use ECS.

  If you modify specifications or change the OS of an ECS and still want to enjoy the RI discount, you need to purchase a new RI with the same attributes as the ECS.
- Attribute matching of a zonal RI: The flavor and OS of a zonal RI must be the same as those of a pay-per-use ECS.

  If you modify specifications or change the OS of an ECS and still want to enjoy the RI discount, you need to purchase a new RI with the same attributes as the ECS.

Application Scenarios

If your ECSs will be used in a short term, it is a good practice to use the pay-per-use billing mode. 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 2-4 lists RI attributes. You can purchase your desired RIs based on these attributes.

Table 2-4 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. **NOTE** After an RI is purchased, its flavor cannot be directly changed, but you can split or combine it. For details, see **Modifying RI Attributes**.
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 instance size, can be modified during the term. However, the instance type cannot be changed.

Parameter	Description
Payment Option	No upfront

Zonal RIs

An RI purchased for a specified AZ refers to a zonal RI. It offers a billing discount for the ECSs with the same flavor and OS as the RI in that AZ.

For example, after you purchase two c3.xlarge.2 Linux RIs with a one-year term in an AZ, the RI discounts can be applied to up to two pay-per-use c3.xlarge.2 Linux ECSs running in that AZ.

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.
- Instance size flexibility: The RI discount applies to instance usage for ECSs that have the same instance series, vCPU/memory ratio, and OS as those specified in the regional RI. Instance size flexibility is determined based on the normalization factor of the instance size. Instance size flexibility does not apply to zonal RIs.

Instance size flexibility is applied from the smallest to the largest instance size within the instance series based on the normalization factor. **Table 2-5** describes the instance size within an instance type and corresponding normalization factor per hour.

NOTE

An ECS automatically benefits from the billing discount offered by a regional RI only when the instance series, vCPU/memory ratio, and OS are the same as those specified in the regional 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 2-5 Normalization factors

Instance Size	Normalization Factor
small	1
medium	1
large	2
xlarge	4
2xlarge	8
4xlarge	16

Instance Size	Normalization Factor
6xlarge	24
7xlarge	28
8xlarge	32
9xlarge	36
12xlarge	48
14xlarge	56
15xlarge	60
16xlarge	64
26xlarge	104
52xlarge	208
nxlarge	n × 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 2-2 Example RI 1

  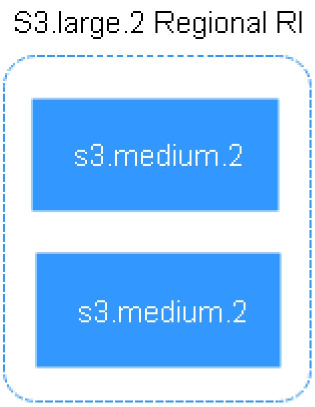

- If you have one running s3.xlarge.2 pay-per-use Linux ECS with a normalization factor of 4 in this region, the billing benefit is applied to 50% of the usage of the ECS.

Figure 2-3 Example RI 2

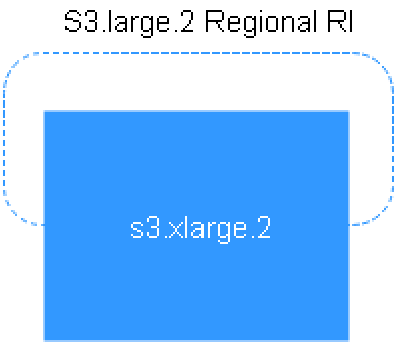

Table 2-6 Comparison between regional and zonal RIs

RI Type	AZ Flexibility	Instance Size Flexibility	Capacity Reservation
Regional RI	Supported The regional RI discount applies to any AZ in the region.	Supported The regional RI discount can be applied only when the instance series, vCPU/ memory ratio, and OS 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 recommended)	Not supported A zonal RI only applies to a specified AZ.	Not supported The instance series and OS of a zonal RI must match those of a pay-per-use instance.	Supported Resources can be reserved for creating pay-per-use ECSs.

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 instance 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). There are three running m3.xlarge.2 ECSs in your account, 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 instance 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.

2.2.4.2 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**.

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 2-7 shows specifications parameters.

Table 2-7 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. **NOTE** After an RI is purchased, its flavor cannot be directly changed, but you can split or combine it. For details, see **Modifying RI Attributes**.
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.

Parameter	Description
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 instance size, can be modified during the term. However, the instance 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 the upfront price per RI multiplied by the number of RIs. The pay-per-use price is the pay-per-use price per RI multiplied by the number of RIs.

      For details, click **Pricing details**.

14. Click **Next**.

    You can learn more about pricing details **here**.

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.**
  a. On the **Reserved Instance** page, select the target RI.
  b. Check that the selected RI is displayed at the bottom of the RI list.

Figure 2-4 Selected RI

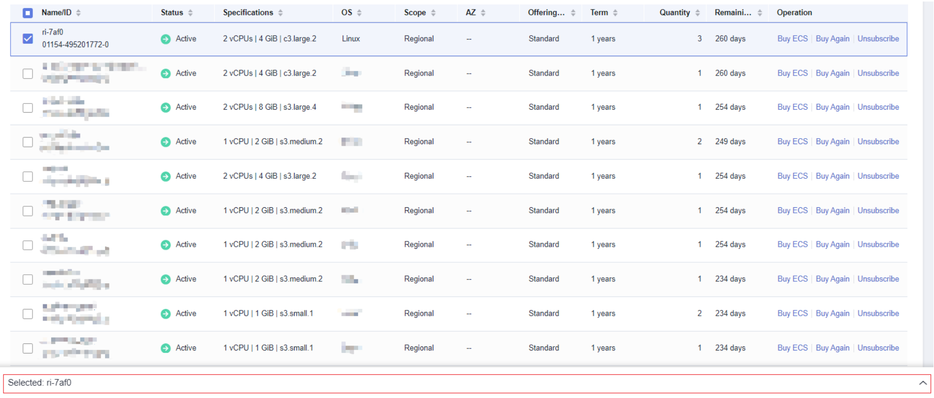

■	Name/ID	Status	Specifications	OS	Scope	AZ	Offering...	Term	Quantity	Remaini...	Operation
☑	ri-7af0 01154-495201772-0	Active	2 vCPUs | 4 GiB | c3.large.2	Linux	Regional	--	Standard	1 years	3	260 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	2 vCPUs | 4 GiB | c3.large.2	[redacted]	Regional	--	Standard	1 years	1	260 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	2 vCPUs | 8 GiB | s3.large.4	[redacted]	Regional	--	Standard	1 years	1	254 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 2 GiB | s3.medium.2	[redacted]	Regional	--	Standard	1 years	2	249 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	2 vCPUs | 4 GiB | s3.large.2	[redacted]	Regional	--	Standard	1 years	1	254 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 2 GiB | s3.medium.2	[redacted]	Regional	--	Standard	1 years	1	254 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 2 GiB | s3.medium.2	[redacted]	Regional	--	Standard	1 years	1	254 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 2 GiB | s3.medium.2	[redacted]	Regional	--	Standard	1 years	1	234 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 1 GiB | s3.small.1	[redacted]	Regional	--	Standard	1 years	2	234 days	Buy ECS | Buy Again | Unsubscribe
□	[redacted]	Active	1 vCPU | 1 GiB | s3.small.1	[redacted]	Regional	--	Standard	1 years	1	234 days	Buy ECS | Buy Again | Unsubscribe

Selected: ri-7af0

  c. Expand the RI details.

Figure 2-5 RI details

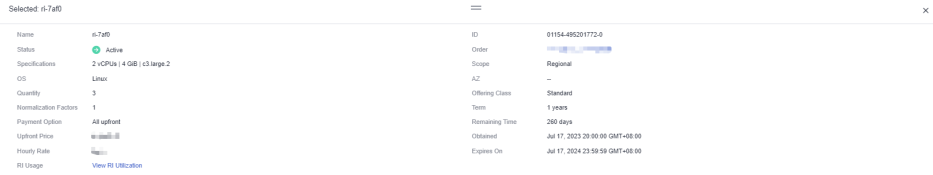

Selected: ri-7af0

Name	ri-7af0	ID	01154-495201772-0
Status	Active	Order	[redacted]
Specifications	2 vCPUs | 4 GiB | c3.large.2	Scope	Regional
OS	Linux	AZ	--
Quantity	3	Offering Class	Standard
Normalization Factors	1	Term	1 years
Payment Option	All upfront	Remaining Time	260 days
Upfront Price	[redacted]	Obtained	Jul 17, 2023 20:00:00 GMT+08:00
Hourly Rate	[redacted]	Expires On	Jul 17, 2024 23:59:59 GMT+08:00
RI Usage	View RI Utilization		

  d. In the RI details area, click **View RI Utilization** in the **RI Usage** row to go to Cost Center and view the RI utilization.

Figure 2-6 RI utilization

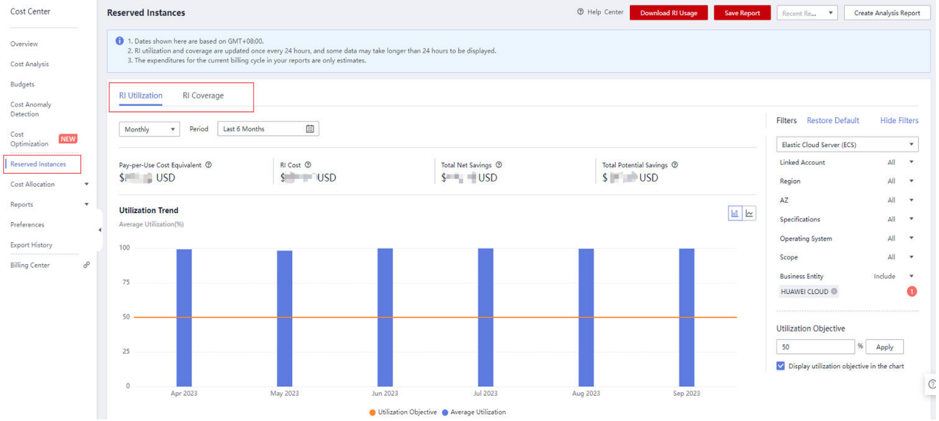

1. Dates shown here are based on GMT+08:00.
2. RI utilization and coverage are updated once every 24 hours, and some data may take longer than 24 hours to be displayed.
3. The expenditures for the current billing cycle in your reports are only estimates.

Monthly | Period: Last 6 Months

Pay-per-Use Cost Equivalent	RI Cost	Total Net Savings	Total Potential Savings
$[redacted] USD	$[redacted] USD	$[redacted] USD	$[redacted] USD

Elastic Cloud Server (ECS)
Linked Account: All
Region: All
AZ: All
Specifications: All
Operating System: All
Scope: All
Business Entity: Include
HUAWEI CLOUD

Utilization Objective: 50 %
☑ Display utilization objective in the chart

2.2.4.3 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**.

Notes and 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 RIs, but 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 2-8 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.

Allowed Operation	Description
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**.

2.2.5 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 by referring to the content in this section.

Prerequisites

- The selected ECS is billed on a pay-per-use basis.
- The target ECS must be in **Running** or **Stopped** state.

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. Above the ECS list, choose **More** > **Change to Yearly/Monthly**.
6. Confirm the ECS details, specify the required duration, and pay for the order.

2.2.6 Changing Yearly/Monthly to Pay-per-Use

Scenarios

Yearly/Monthly is a prepaid billing mode in which your ECS will be billed based on the required 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.

Constraints

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

Procedure

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

   The **Renewals** page is displayed.
3. Customize search criteria.
   - On the **Pay-per-Use After Expiration** tab, you can search for the ECSs whose billing mode has been 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).
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.
   - 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.
5. Confirm the change details and click **Change to Pay-per-Use**.

2.3 Purchasing an ECS

2.3.1 Introducing ECS Purchase Options

Huawei Cloud provides multiple options for you to purchase ECSs.

Table 2-9 Purchase options of ECSs

How to Purchase	Description
Purchasing an ECS in Custom Config Mode	Learn how to purchase an ECS in custom config mode. You can flexibly specify required parameters for your ECS, including the billing mode, instance specifications, image, storage, network, security group, and EIP.
Purchasing a Spot ECS	Learn how to purchase a spot ECS. Spot ECSs allow you to use spare ECS capacity which is available for less than the pay-per-use price. This is a good option if you want to enjoy the same performance at a lower price.
Purchasing a Spot Block ECS	Learn how to purchase a spot block ECS. Spot block ECSs allow you to use spare ECS capacity which is available for less than the pay-per-use price. This is a good option if you want to enjoy the same performance at a lower price.
Purchasing an ECS Using a Private Image	Learn how to purchase an ECS using a private image. A private image contains an OS, preinstalled public applications, and user's personal applications, saving the time for configuring the ECS repeatedly. This is a good option if you are accustomed to using certain OS and applications.
Purchasing the Same ECS	Learn how to quickly purchase ECSs with the same configurations as an existing ECS.

2.3.2 Purchasing an ECS in Custom Config Mode

Scenarios

Elastic Cloud Server (ECS) is a cloud server that provides scalable, on-demand resources, including vCPUs, memory, OS, and Elastic Volume Service (EVS) disks. After purchasing an ECS, you can use it like using your local computer or physical server.

You can create an ECS by specifying its vCPUs, memory, OS, specifications, and login mode.

This section describes how to create an ECS on the management console.

Procedure

Step	Description
Preparations	● Sign up for a HUAWEI ID, enable Huawei Cloud services, and top up your account. ● Prepare resources, such as VPCs, subnets, security groups, key pairs, Dedicated Computing Cluster (DCC), and CloudPond.
Step 1: Access the Page for Purchasing ECSs	Log in to the ECS console and open the page for purchasing ECSs.
Step 2: Specify Parameters	Specify parameters based on your service requirements.
Step 3: Confirm the Configuration and Submit the Order	Confirm the configuration details and complete the purchase.

Preparations

1. Sign up for a HUAWEI ID and complete real-name authentication.

   Before purchasing an ECS, **sign up for a HUAWEI ID and enable Huawei Cloud services** and **complete real-name authentication** first.

   If you have enabled Huawei Cloud services and completed real-name authentication, skip this step.

2. Top up your account.

   Ensure that your account has sufficient balance. If not, **top up your account**.

3. Plan network resources, such as VPCs and subnets.

   When you are purchasing an ECS, the system creates a default VPC (vpc-default) and subnet (subnet-default).

   If you do not want to use the default VPC and subnet, you can create a VPC and subnet in the corresponding region in advance.

4. Create a security group and add rules to it.

   When you are purchasing an ECS, the system creates default security groups (default, Sys-WebServer, and Sys-FullAccess). For more information about security groups and rules, see **Default Security Groups and Rules**.

   If the default security groups and rules cannot meet your service requirements, you can modify them. For details, see **Configuring Security Group Rules**.

5. Create a key pair.

   To log in to the ECS using a key pair, create one in advance. For details, see **(Recommended) Creating a Key Pair on the Management Console**.

Step 1: Access the Page for Purchasing ECSs

Log in to the management console and access the **Buy ECS** page.

Step 2: Specify Parameters

Specify parameters required for purchasing an ECS. These parameters include but are not limited to the billing mode, region, AZ, specifications, storage, and network.

Basic Configuration

1. Select a billing mode.

   You can select an appropriate billing mode based on the required duration and resource inventory to help you save costs.

Table 2-10 Billing modes

Option	Description	Scenarios and Constraints	Reference
Yearly/Monthly	Yearly/Monthly is a prepaid billing mode in which you pay for ECSs before using them. Yearly/monthly ECSs are billed by the required duration specified in the order.	This cost-effective mode is ideal when the duration of ECS usage is predictable. This billing mode is recommended for long-term users. A yearly/monthly ECS cannot be deleted. If such an ECS is not required anymore, unsubscribe from it.	**Yearly/ Monthly Billing**
Pay-per-use	Pay-per-use is a postpaid billing mode. You pay as you go and just pay for what you use. Pay-per-use ECSs are billed by the second and settled by the hour.	This mode is ideal when you want more flexibility and control on ECS usage.	**Pay-per-Use Billing**

Option	Description	Scenarios and Constraints	Reference
Spot pricing	Spot pricing is a postpaid billing mode. You pay as you go and just pay for what you use. In **Spot pricing** billing mode, your purchased ECS is billed at a lower price than that of a pay-per-use ECS with the same specifications. In **Spot pricing** billing mode, you can select **Spot** or **Spot block** for **Spot Type**. Spot ECSs and Spot block ECSs are billed by the second and settled by the hour.	Spot pricing is a good option if you want to enjoy the same performance at a lower price. In spot pricing billing mode, your purchased ECSs are not suitable for long-term workloads or workloads that require high stability.	- **Spot Pricing (for Spot Instances)** - **Spot Pricing (for Spot Block Instances)**

2. (Optional) Set **Reserved Instance**.

   This parameter is displayed only when **Billing Mode** is set to **Pay-per-use**. If you want to associate reserved instances (RIs) with your pay-per-use ECSs, select **Associate RI** and select an RI.

   NOTE

   For details, see **Reserved Instance Overview**.

   Figure 2-7 Reserved instance

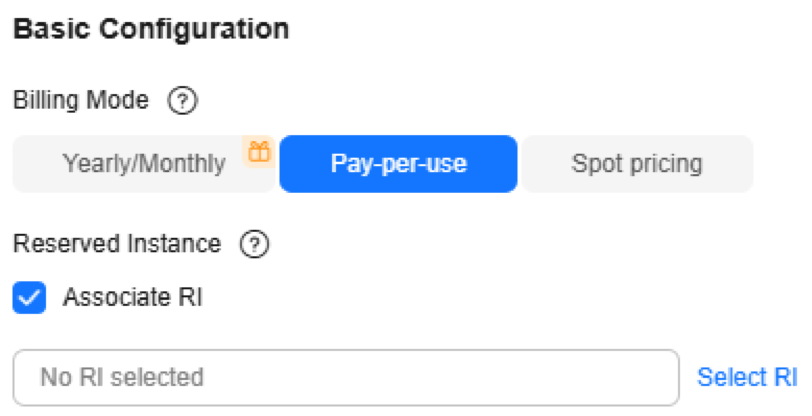

Basic Configuration

Billing Mode

Yearly/Monthly | Pay-per-use | Spot pricing

Reserved Instance

☑ Associate RI

No RI selected

Select RI

3. (Optional) Set **Spot Type**.

   This parameter is displayed only when **Billing Mode** is set to **Spot pricing**. You can select **Spot** or **Spot block** for **Spot Type**.

Figure 2-8 Spot type

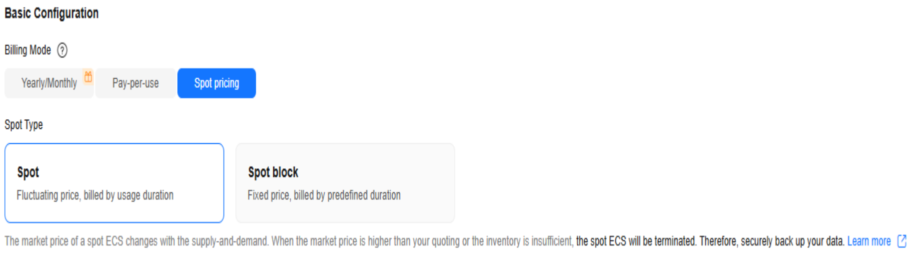

The market price of a spot ECS changes with the supply-and-demand. When the market price is higher than your quoting or the inventory is insufficient, the spot ECS will be terminated. Therefore, securely back up your data. Learn more

- **Spot**: The price of spot ECSs fluctuates with the market. For details, see **Purchasing a Spot ECS**.
- **Spot block**: The price of spot block ECSs depends on the predefined duration. For details, see **Purchasing a Spot Block ECS**.

4. Set **Region**.

   A region refers to a physical data center area where ECSs reside. For lower network latency and faster resource access, select the region nearest to your services.

   NOTE

   - ECSs in different regions cannot communicate with each other over an intranet.
   - Once ECSs are purchased, the region cannot be changed.

5. Set **AZ**.

   An AZ is a physical location that uses independent power supply and networks. AZs in the same region can communicate with each other over an intranet.

Table 2-11 Selecting an AZ

Option	Description	Scenarios and Constraints	Reference
Random	The available ECS types and flavors vary depending on AZs. The system uses hash algorithms to allocate an AZ based on your universally unique identifier (UUID) and the ECS flavor you have selected.	To view all ECS types and flavors supported by the cloud platform, select **Random** for **AZ**.	**Region and AZ**

Option	Description	Scenarios and Constraints	Reference
AZ*N*	AZs supported in the selected region. *N* indicates the sequence number of an AZ.	If you want to create an ECS in a specified AZ, select that AZ. ● For high availability (HA), create ECSs in different AZs. ● For low network latency, create ECSs in the same AZ.	

Instance

1. Select an instance selection mode.

Table 2-12 Instance parameters

Option	Description	Scenarios and Constraints	Reference
By Type	You can select ECS specifications based on different properties. ● CPU architecture: x86 or Kunpeng ● Search filters: Filter by vCPU, memory, or keyword. ● Specifications: Select specifications by ECS type and flavor.	This mode is suitable for users who are familiar with the CPU architecture, vCPUs, memory, and instance family and generation of ECSs and want to select specific specifications.	**Overview**
By Scenario	ECS specifications are recommended based on categories and sub-categories.	This mode is suitable for users who have specific service requirements.	

2. (Optional) Set **CPU Architecture**.

   This parameter is displayed only when you select **By Type**.

Table 2-13 CPU architectures

Option	Description	Scenarios and Constraints	Reference
x86	The x86 CPU architecture uses Complex Instruction Set Computer (CISC) and supports almost all general software. The execution of such an instruction is complex and time-consuming.	It is suitable for platform-dependent scenarios using Windows software and x86-compatible commercial software.	**A Summary List of x86 ECS Specifications**

3. (Optional) Set **Category** and **Sub-category**.

   These parameters are displayed only when you select **By Scenario**.

   NOTE

   The specifications vary by region and AZ. For details, see the specifications displayed on the console.

Table 2-14 Categories and sub-categories

Category	Sub-category	Description
Web applications	Traditional office	High security and reliability, suitable for traditional office scenarios like OA, ERP, and CRM with less than 200 concurrent access requests
	Enterprise websites	A balance of compute, memory, and network resources with a baseline level of vCPU performance and high cost-effectiveness
	Personal application setup	A balance of compute, memory, and network resources with a baseline level of vCPU performance and high cost-effectiveness
	Development and testing	A balance of compute, memory, and network resources with a baseline level of vCPU performance and the ability to provide burst CPU power at any time for as long as required
	Front-end servers	A balance of compute, memory, and network resources with a baseline level of vCPU performance. These ECSs can be used as front-end servers like Apache, Nginx, and IIS.

Category	Sub-category	Description
	Back-end servers	High ratio of CPUs to memory, high performance, and low latency. These ECSs are cost-effective options for back-end servers like Tomcat and JBoss.
Website applications/E-commerce	100,000 pageviews/ 1,000 active users	Cost-effective, flexible, elastic resources available anytime
	200,000 pageviews/ 2,000 active users	Suitable for e-commerce websites, which require high-performance cloud servers with fast elasticity and high stability to handle traffic bursts typical of special promotions, flash sales, and live commerce
	500,000 pageviews/ 5,000 active users	Suitable for e-commerce websites, which require high-performance cloud servers with fast elasticity and high stability to handle traffic bursts typical of special promotions, flash sales, and live commerce
Gaming	Gaming	Suitable for gaming services, which require high performance, high stability, high cost-effectiveness, and low latency
Databases	Compute	Stable, high-performance compute
	Storage	Servers that use local disks with high storage bandwidth and IOPS to provide cost-effective mass data storage
	Network	High PPS performance, high TPS throughput, and low network latency for rapid data exchange and processing
Data analytics	Management nodes	A large volume of compute resources scheduled to accelerate data processing
	Compute nodes	Balanced compute with high performance and stability
	Storage nodes	Cost-effective, high-bandwidth storage for processing large amounts of reads and writes
High performance computing	High performance computing	High-performance compute clusters with large compute and high cost-effectiveness
Image rendering	Animation rendering	CPU-accelerated rendering with high precision and stability

Category	Sub-category	Description
	Video rendering	GPU-accelerated rendering with high processing speed
AI/ Machine learning	AI training	Compatible with NVIDIA smart NICs for deep learning training, scientific computing, computational fluid dynamics, computational finance, seismic analysis, molecular modeling, and genomics.
	AI inference	Compatible with NVIDIA smart NICs for image classification and recognition, speech recognition, natural language processing, video encoding and decoding, machine learning, and lightweight training.

4. Set **Search Filters**.

   a. Select specifications.

      ▪ Select vCPUs and memory, or enter a keyword to search for ECS specifications.

        You can search for ECS flavors when you select **By Type**.

      ▪ Select ECS specifications by instance family and generation from the list.

        For details about each type, see **ECS Types**.

      NOTE

      - Sold-out vCPU and memory resources cannot be selected. You can select **Hide sold-out specifications** when purchasing ECSs.

   b. Set the scope of specifications to be displayed.

      ▪ **Only show latest generation**: After this option is selected, only newly released ECS types and specifications are displayed. If this option is not selected, all ECS types and specifications available on the cloud service platform are displayed.

      ▪ **Hide sold-out specifications**: After this option is selected, sold-out specifications are not displayed.

5. (Optional) Set **Maximum Price**.

   This parameter is available only when **Billing Mode** is set to **Spot pricing** and **Spot Type** is set to **Spot**.

   – **Automatic (Recommended)**: 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.

   For details, see **Purchasing a Spot ECS**.

6. (Optional) Set **Predefined Duration**.

   This parameter is available only when **Billing Mode** is set to **Spot pricing** and **Spot Type** is set to **Spot block**.

   – **Predefined Duration**: a duration that you specify for your spot block ECS. Prices vary depending on predefined durations.

     During the predefined duration, if your spot block ECS is automatically terminated by the system, you will not be billed for the resource usage within the predefined duration. If you delete your spot block ECS within the predefined duration, you will be billed based on the usage duration.

   – **Price for Each Spot Block ECS**: you do not need to configure this parameter.

   – (Optional) **Number of Durations**: This parameter is displayed only when **Predefined Duration** is set to **6 hours**.

   For details, see **Purchasing a Spot Block ECS**.

OS

1. Set **Image**.

   An image is an ECS template that contains an OS. It may also contain proprietary software and application software. You can use images to create ECSs.

Table 2-15 Images

Option	Description	Reference
Public image	A public image is a standard OS image which is highly stable, authorized, and visible to all users. It contains an OS and preinstalled public applications. If you need other applications or software, configure them on the new ECSs.	**Public Image Overview**
Private image	A private image is an image available only to the user who created or imported it. It contains an OS, preinstalled public applications, and the user's personal applications, saving the time for configuring the ECS repeatedly.	**Creating a Private Image** **Purchasing an ECS Using a Private Image**
Shared image	A shared image is a private image shared by another account. You can use the same image to create ECSs across accounts.	**Shared Image Overview**

2. (Optional) Set **Automatically install GPU driver**.

   This parameter is displayed only when you select a GPU-accelerated ECS type and an image with the **NO Driver** suffix.

After you select **Automatically install GPU driver** and choose the corresponding version, the GPU driver will be automatically installed for your ECSs.

NOTE

- The system can validate the specifications and GPU driver version. You can select a version from the drop-down list. For more information about GPU driver versions, see **Obtaining a Tesla Driver and CUDA Toolkit**.
- The installation takes about 5 to 10 minutes to complete. Do not stop or start the ECS during the installation, or the installation will fail.
- After the installation is complete, the ECS automatically restarts.
- Changing the OS of an existing ECS will deactivate the automatically installed GPU driver.

3. (Optional) Set **Host protection (HSS)**.

   When you select certain public images, Host Security Service (HSS) is enabled by default. HSS Basic Edition provides one-month free trial and automatically installs the HSS agent. HSS Basic Edition provides functions such as weak password and vulnerability detection.

   HSS is designed to improve the overall security for ECSs. It helps you eliminate risks, defend against intrusions and web page tampering, provide advanced defense, and manage security operations.

   NOTE

   After the one-month free trial period expires, the HSS basic edition quotas will be automatically released, and HSS will not protect your servers.

   If you want to continue using HSS or upgrade HSS security capabilities, you need to purchase HSS.

   After ECSs are purchased, you can switch between different HSS editions on the HSS console. For details about differences among different editions, see **Specifications of Different Editions**.

Storage & Backup

1. Set **System Disk**.

   A system disk stores the OS of an ECS, and is automatically created and initialized once the ECS is created.

   NOTE

   If you detach the system disk that is purchased along with a yearly/monthly ECS and want to continue using it as a system disk, you can only attach it to the original ECS. If you want to use it as a data disk, you can attach it to any ECS.

Table 2-16 System disk parameters

Parameter	Description	Scenarios and Constraints	Reference
Disk Type	Disk types are classified based on the I/O performance of disks.	Disks can be classified into the following types by I/O performance: Extreme SSD V2, Extreme SSD, General Purpose SSD V2, Ultra-high I/O, General Purpose SSD, High I/O, and Common I/O. EVS disks differ in performance and price. You can choose whichever disk type that is the best fit for your applications.	**Disk Types and Performance**
System Disk (GiB)	System disk capacity, in GiB.	EVS disks are billed by disk capacity. Select appropriate capacity based on service requirements. **NOTE** For a P1 or P2 ECS, the system disk must be greater than or equal to 15 GiB. It is recommended that the system disk be greater than 40 GiB.	
IOPS	Number of read/write operations performed by an EVS disk per second This parameter is displayed only when **General-purpose SSD V2** is selected for **Disk Type**.	You are advised to set IOPS based on the value range and service requirements.	
Throughput	Amount of data read from and written into an EVS disk per second This parameter is displayed only when **General-purpose SSD V2** is selected for **Disk Type**.	Configure a desired throughput based on the value range and your service requirements.	

2. (Optional) Set **Advanced Options** for the system disk.

   If you want to set **SCSI** and **Encryption** for the system disk, click **Advanced Options**.

Table 2-17 Advanced options

Option	Description	Scenarios and Constraints	Reference
SCSI	Specifies the SCSI device type. This parameter is selected by default.	Device types: ● **VBD**: indicates the Virtual Block Device (VBD) mode. ● **SCSI**: indicates the Small Computer System Interface (SCSI) mode. The default device type is VBD. If SCSI is selected, the disk will support transparent SCSI command transmission. **NOTE** The disk device type is configured during the purchase process. Once disks are purchased, the device type cannot be changed.	**Device Types and Usage Instructions**
Encryption	Encrypts the system disk. ● If the ECS is created from an encrypted image, the system disk of the ECS is automatically encrypted. ● If the image you selected is not encrypted, you can select **Encryption** to encrypt the system disk.	Disk encryption provides strong security protection for your data. Snapshots generated from encrypted disks and disks created using these snapshots automatically inherit the encryption attribute. For details, see **3**.	**Managing Encrypted EVS Disks**

3. (Optional) Set encryption parameters.

   This parameter is displayed only when **Encryption** is selected in **Advanced Options**.

   NOTE

   To use the encryption feature, click **Create Agency** first to grant EVS the permissions needed to obtain KMS keys for EVS disk encryption and decryption.

   If you do not have sufficient permissions to grant EVS permissions, contact the user having the **Security Administrator** permissions to grant the required permissions. For details, see

   The encryption parameters are as follows:

- **Agency Name**: specifies the name of the agency that is used to grant EVS the permissions needed to obtain KMS keys for disk encryption and decryption. When **Agency Name** is displayed as **EVSAccessKMS**, KMS permissions have been granted to EVS.

  NOTE

  To use the encryption feature, click **Create Agency** first to grant EVS the permissions needed to obtain KMS keys for EVS disk encryption and decryption.

  If you do not have sufficient permissions to grant EVS permissions, contact the user having the **Security Administrator** permissions to grant the required permissions. For details, see

- **KMS Encryption**: specifies how to obtain a KMS key.
  - **Select an existing key**: Select a KMS key from the **KMS Key Name** drop-down list.
  - **Enter a key ID**: Select a KMS key using the key ID.
- (Optional) **KMS Key Name**: specifies the name of the key used to encrypt EVS disks. This parameter is displayed only when **KMS Encryption** is set to **Select an existing key**.

  You can select an existing key pair, or click **Create KMS Key** and create a KMS key on the KMS console. The default value is **evs/default**.
- **KMS Key ID**: specifies the ID of the key used to encrypt data disks.

4. Set **Data Disk**.

   Data disks store user data. If you add data disks (click **Add Data Disk**) when purchasing ECSs, the system will automatically attach the data disks to the ECSs. If you purchase data disks after ECSs are purchased, you need to manually attach the data disks.

   NOTE

   - After you attach data disks to an ECS, you need to **initialize the disks** before using them.
   - If you detach the non-shared data disk purchased when you purchase a yearly/monthly ECS and want to attach it again, you can only attach it to the original ECS as a data disk.
   - The data disks purchased when you buy a yearly/monthly ECS does not support separate renewal, unsubscription, auto-renewal, changing to pay-per-use, and deletion.

Table 2-18 Data disk parameters

Parameter	Description	Scenarios and Constraints	Reference
Disk Type	Disk types are classified based on the I/O performance of disks.	Disks can be classified into the following types by I/O performance: Extreme SSD V2, Extreme SSD, General Purpose SSD V2, Ultra-high I/O, General Purpose SSD, High I/O, and Common I/O. EVS disks differ in performance and price. You can choose whichever disk type that is the best fit for your applications.	**Disk Types and Performance**
Data Disk (GiB)	Data disk capacity, in GiB.	EVS disks are billed by disk capacity. Select appropriate capacity based on service requirements.	
IOPS	Number of read/ write operations performed by an EVS disk per second This parameter is displayed only when **General-purpose SSD V2** is selected for **Disk Type**.	You are advised to set IOPS based on the value range and service requirements.	
Throughput	Amount of data read from and written into an EVS disk per second This parameter is displayed only when **General-purpose SSD V2** is selected for **Disk Type**.	Configure a desired throughput based on the value range and your service requirements.	
Quantity	Number of data disks.	Specify the quantity of data disks to be added as required. When creating an ECS, you can add up to 23 data disks to the ECS.	

5. (Optional) Set **Advanced Options** for data disks.

   To set **SCSI**, **Sharing**, and **Encryption** for data disks, click **Advanced Options**.

Table 2-19 Advanced options

Option	Description	Scenarios and Constraints	Reference
SCSI	Specifies the SCSI device type. This parameter is selected by default.	Device types: ● **VBD**: indicates the Virtual Block Device (VBD) mode. ● **SCSI**: indicates the Small Computer System Interface (SCSI) mode. The default device type is VBD. If SCSI is selected, the disk will support transparent SCSI command transmission. **NOTE** Disk device type is configured during purchase. It cannot be changed after the disk has been purchased.	**Device Types and Usage Instructions**
Sharing	Sharing is used to set a data disk as a shared disk.	After a data disk is configured as a shared disk, the shared disk can be attached to multiple ECSs.	**Managing Shared EVS Disks**
Encryption	**Encryption** is used to encrypt data disks.	Disk encryption provides strong security protection for your data. For details, see **3**.	**Managing Encrypted EVS Disks**
Create Disk from Data Disk Image	This option is used to create a data disk from a data disk image.	If you use a Windows or Linux image to create an ECS, you can use a data disk image to create a data disk. Select **Create Disk from Data Disk Image**. In the displayed list, select your data disk image. **NOTE** One data disk image can be used for one data disk only. This function is unavailable if you have selected a full-ECS image to create ECSs or selected **SCSI**, **Sharing**, or **Encryption** for data disks.	**Creating a Private Image**

6. (Optional) Select **Enable backup**.

   Configure this parameter only when you need to back up ECSs or EVS disks.

CBR backups can help you restore data in case of any ECS failures. To ensure data security, you are advised to enable backup.

NOTE

- For CBR pricing details, see **How Is CBR Billed?**

7. (Optional) Set CBR**Cloud Backup and Recovery**.

   This parameter is displayed only when **Enable backup** is selected.

   The following options are provided:

   – **Create new**: Set CBR parameters.

     i. Set the vault name, which consists of a maximum of 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. For example, **vault-f61e**. The default naming rule is **vault**_*xxxx*.

     ii. Set 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. The value range is from the total capacity of the ECS to 10,485,760 in the unit of GiB.

     iii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.

   – **Use existing**: Select an existing CBR vault and configure a backup policy.

     i. Select an existing cloud backup vault from the drop-down list.

     ii. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.

Network

1. Set **VPC** and **Primary NIC**.

   Virtual Private Cloud (VPC) allows you to create logically isolated, configurable, and manageable virtual networks for ECSs. You can configure security groups, Virtual Private Network (VPNs), CIDR blocks, and bandwidths in your VPC. By default, ECSs in different VPCs cannot communicate with each other.

   Figure 2-9 Network

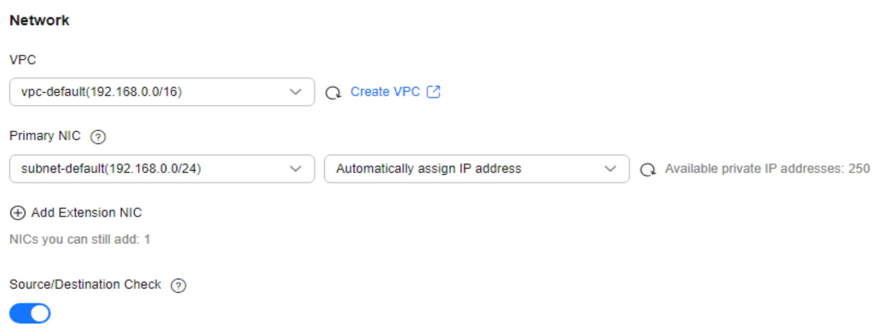

Network

VPC

vpc-default(192.168.0.0/16)

Primary NIC

subnet-default(192.168.0.0/24)

Automatically assign IP address

Available private IP addresses: 250

NICs you can still add: 1

Source/Destination Check

   You can select an available VPC from the drop-down list or create a VPC as required. By default, the system attaches a primary network interface (NIC) and specifies how a private IP address will be assigned.

   For details, see **VPC and Subnet Planning**.

NOTE

- You need to ensure that DHCP is enabled in the VPC which the ECS belongs to.

2. Set **Primary NIC**.

   The primary NIC provides the default route and cannot be deleted. It is automatically created along with an ECS. After a VPC is specified, the system attaches a primary NIC to the ECS by default and specifies how a private IP address will be assigned.

   a. Set **Primary NIC**: If there are multiple subnets in the VPC, you can select another subnet from the drop-down list as the primary NIC.

   b. Set how an IPv4 address will be assigned. **Automatically assign IP address** is selected by default.

      - **Automatically assign IP address**: The system automatically assigns a private IPv4 address to the primary NIC.

      - **Manually specify IP address**: You need to manually assign a private IPv4 address to the primary NIC. Before specifying an IP address, click **View In-Use IP Address** to avoid address conflict.

        NOTE

        If you specify a private IPv4 address when creating multiple ECSs in a batch, note the following:

        - This IP address serves as the start IP address.
        - The required IP addresses must be consecutive and available within the subnet.
        - The subnet that contains the specified IP address cannot overlap with other subnets.

      - **Use existing network interface**: This parameter is displayed only when the selected VPC has available network interfaces. You can select an existing network interface from the drop-down list as the primary network interface.

   c. (Optional) Set how an IPv6 address will be assigned. **IPv6 not required** is selected by default.

      This parameter is displayed only when the **IPv6** column of the selected ECS flavor is **Yes** and IPv6 is enabled for the subnet.

      - **IPv6 not required**: No IPv6 address is allocated to the network interface.

      - **Automatically-assigned IPv6 address**: The system automatically assigns a private IPv6 address to the network interface.

        If **Automatically-assigned IPv6 address** is selected, the system assigns IPv6 addresses. In a VPC, ECSs use IPv6 addresses to access the dual-stack intranet.

        To enable an ECS to access the Internet, you need to assign a EIP shared bandwidth, and add the ECS's IPv6 address to the shared bandwidth.

        For details, see **Adding EIPs to or Removing EIPs from a Shared Bandwidth**.

NOTE

3. (Optional) Click **Add Extension NIC**.

   Extension NICs can be separately added. If you need to attach multiple NICs to an ECS, you can add multiple extension NICs and specify their IP addresses.

   Extension NICs cannot communicate with external networks before you configure policy-based routes for them. For operation details, see **Configuring Policy-based Routes for an ECS with Multiple Network Interfaces**.

   a. Set **Extension NIC1**: If there are multiple subnets in the VPC, you can select another subnet from the drop-down list as the extension NIC.
   b. Set how an IPv4 address will be assigned. **Automatically assign IP address** is selected by default.
      - **Automatically assign IP address**: The system automatically assigns a private IPv4 address to the extension NIC.
      - **Manually specify IP address**: You need to manually assign a private IPv4 address to the extension NIC. Before specifying an IP address, click **View In-Use IP Address** to avoid address conflict.
      - **Use existing network interface**: Specify a NIC as the extension NIC. You can select a NIC from the drop-down list.
   c. (Optional) Set how an IPv6 address will be assigned. **IPv6 not required** is selected by default.

      This parameter is displayed only when the **IPv6** column of the selected ECS flavor is **Yes** and IPv6 is enabled for the subnet.

      - **IPv6 not required**: No IPv6 address is allocated to the network interface.
      - **Automatically-assigned IPv6 address**: The system automatically assigns a private IPv6 address to the network interface.

        If **Automatically-assigned IPv6 address** is selected, the system assigns IPv6 addresses. In a VPC, ECSs use IPv6 addresses to access the dual-stack intranet.

        To enable an ECS to access the Internet, you need to assign a EIP shared bandwidth, and add the ECS's IPv6 address to the shared bandwidth.

        For details, see **Adding EIPs to or Removing EIPs from a Shared Bandwidth**.

      NOTE

4. Set **Source/Destination Check**.

   When this function is enabled, source IP addresses in the outbound packets will be checked. If the IP addresses are incorrect, the packets will not be sent out. This function helps prevent spoofing packet attacks and improve security. By default, **Source/Destination Check** is enabled.

   NOTE

   The source/destination check settings apply only to the NICs created along with the ECSs.

Security Group

1. Configure a security group.

   Select an available security group from the drop-down list. You can select multiple security groups for an ECS (no more than five security groups are recommended). The access rules of all the selected security groups apply to the ECS.

   When you create an ECS for the first time, the system automatically creates the following default security groups: default, Sys-WebServer, and Sys-FullAccess. For details, see **Default Security Groups and Rules**.

   You can expand **Security Group Rules** to view details of inbound and outbound rules. Security group rules determine ECS access and usage. For details about how to configure security group rules, see **Adding a Security Group Rule**. Enable the following common ports and protocols as needed:

   - Port 80: default port for web page access through HTTP.
   - Port 443: port for web page access through HTTPS.
   - ICMP: used to ping ECSs to check their communication statuses.
   - Port 22: reserved for logging in to Linux ECS using SSH.
   - Port 3389: reserved for remote desktop login to Windows ECSs.

2. (Optional) Create a security group.

   If security groups displayed in the drop-down list do not meet your service requirements, click **Create Security Group** to create one.

Figure 2-10 Creating a security group

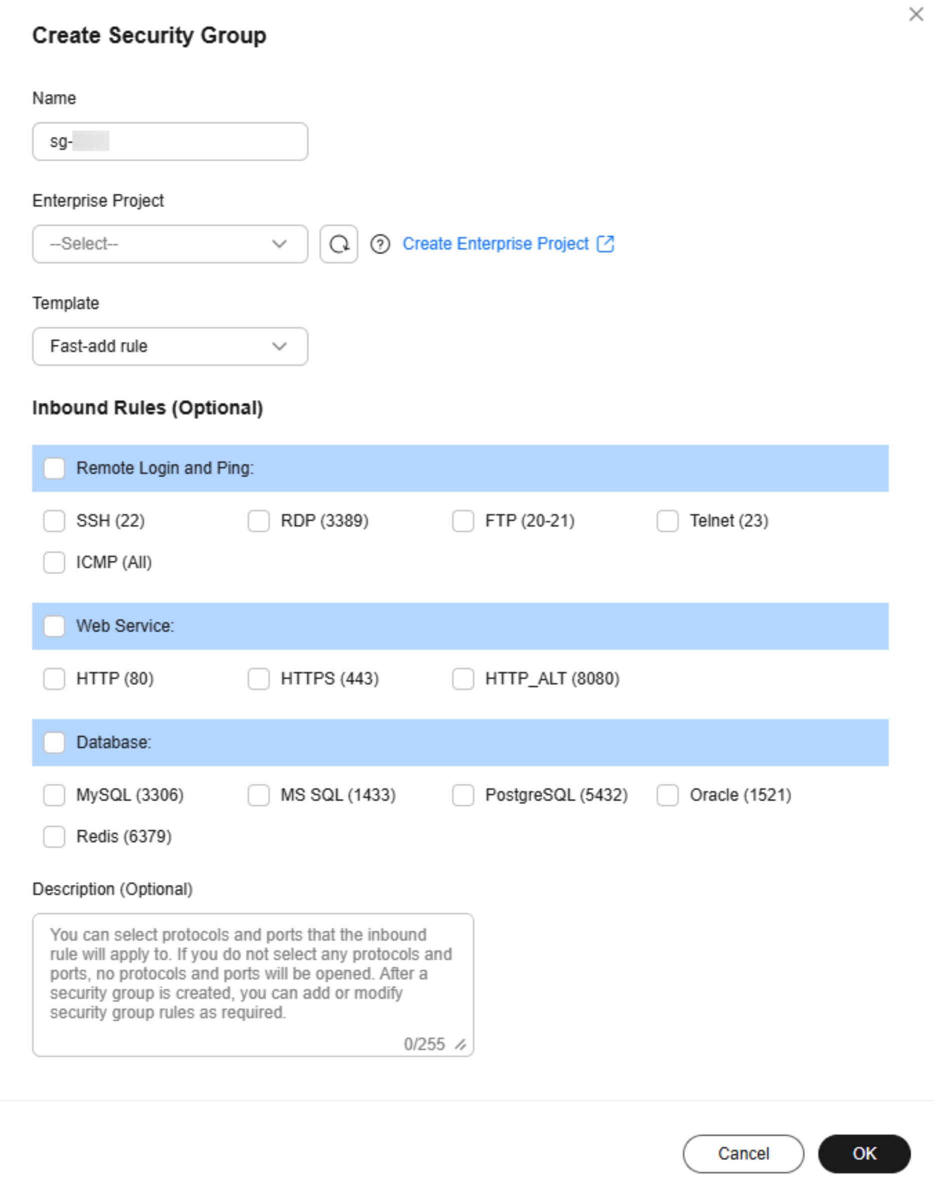

Create Security Group

Name

sg-[redacted]

Enterprise Project

--Select--

Create Enterprise Project

Template

Fast-add rule

Inbound Rules (Optional)

□ Remote Login and Ping:

□ SSH (22) □ RDP (3389) □ FTP (20-21) □ Telnet (23)

□ ICMP (All)

□ Web Service:

□ HTTP (80) □ HTTPS (443) □ HTTP_ALT (8080)

□ Database:

□ MySQL (3306) □ MS SQL (1433) □ PostgreSQL (5432) □ Oracle (1521)

□ Redis (6379)

Description (Optional)

You can select protocols and ports that the inbound rule will apply to. If you do not select any protocols and ports, no protocols and ports will be opened. After a security group is created, you can add or modify security group rules as required.

0/255

Parameters for creating a security group are as follows.

Table 2-20 Creating a security group

Parameter	Description	Example value
Name	This parameter is mandatory and specifies the name of a security group. The name: ● Can contain 1 to 64 characters. ● Can contain letters, digits, underscores (_), hyphens (-), and periods (.). **NOTE** You can change the security group name after a security group is created. It is recommended that you give each security group a different name.	sg-AB
Template	This parameter is mandatory. A template comes with default security group rules, helping you quickly create security groups. By default, **Fast-add rule** is selected. The following templates are provided: ● **General-purpose web server**: The security group that you create using this template is for general-purpose web servers and includes default rules that allow all inbound ICMP traffic and allow inbound traffic on ports 22, 80, 443, and 3389. ● **All ports open**: The security group that you create using this template includes default rules that allow traffic on all protocols and ports. **NOTE** Allowing inbound traffic on all ports poses security risks. ● **Fast-add rule**: You can select common protocols and ports that the inbound rule will apply to.	Fast-add rule
Inbound Rules	This parameter is optional. This parameter is displayed only when **Fast-add rule** is selected for **Template**. Currently, the following protocols and ports can be quickly added. Select protocols and ports as required. ● **Remote Login and Ping**: **SSH (22)**, **RDP (3389)**, **FTP (20-21)**, **Telnet (23)**, or **ICMP (All)**. ● **Web Service**: **HTTP (80)**, **HTTPS (443)**, or **HTTP_ALT (8080)**. ● **Database**: **MySQL (3306)**, **MS SQL (1433)**, **PostgreSQL (5432)**, **Oracle (1521)**, or **Redis (6379)**	-

Parameter	Description	Example value
Description	This parameter is optional and specifies supplementary information about the security group. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Show Default Rule/ Hide Default Rule	This parameter displays security group rules. You can view or hide inbound and outbound rules of the current security group.	-

3. (Optional) Configure security group rules.

   Click **Configure Security Group Rules** to modify rules of the current security group.

   For details, see **Configuring Security Group Rules**.

4. (Optional) Show or hide security group rules.

   You can click **Security Group Rules** to show or hide the security group rules.

   – Selected security groups: If there are multiple security groups, you can move a security group up or down to adjust the priority.

   – Security group rules: You can view inbound rules and outbound rules.

Public Network Access

1. Set EIP.

   An EIP is a static public IP address bound to an ECS in a VPC. Using the EIP, the ECS can provide services externally.

   The following options are provided:

   – **Auto assign**: The system automatically assigns an EIP with exclusive bandwidth for each ECS. You can set the bandwidth.

   – **Use existing**: An existing EIP will be assigned to an ECS. When using an existing EIP, you are not allowed to create ECSs in a batch.

   – **Not required**: An ECS without an EIP cannot access the Internet. It can only be used to deploy services or clusters in a private network.

   NOTE

   For a yearly/monthly ECS, **Auto assign** is unavailable for **EIP**. If an EIP is required, bind an existing EIP to the ECS. Alternatively, purchase an EIP that is billed in pay-per-use payment and then bind the EIP to the ECS.

2. Set **EIP Type**.

   – **Dynamic BGP**: If there are changes on a network using dynamic BGP, network configurations can be promptly adjusted using the specified routing protocol, ensuring network stability and optimal user experience.

   - **Static BGP** If there are changes on a network using static BGP, network configurations cannot be promptly adjusted and user experience may be affected.

3. (Optional) Set **Billed By**.

   This parameter is displayed only when **EIP** is set to **Auto assign**. Each bandwidth can be used by only one EIP.

   - **Bandwidth**: Dedicated bandwidth will be billed by size.
   - **Traffic**: Dedicated bandwidth will be billed by traffic you have actually used.
   - **Shared bandwidth**: The bandwidth can be used by multiple EIPs and you will be billed based on the shared bandwidth.

     NOTE

     - A bandwidth can be shared among a limited number of EIPs. If the number of EIPs cannot meet service requirements, switch to a higher shared bandwidth or apply for expanding the EIP quota of the existing bandwidth.
     - Yearly/monthly EIPs do not support shared bandwidths.
     - When a shared bandwidth that is billed on a yearly/monthly basis expires, the system automatically deletes the bandwidth configuration and creates a dedicated bandwidth billed by traffic for the EIPs sharing the deleted bandwidth configuration.

4. (Optional) Set **Bandwidth Size**.

   This parameter is displayed only when **EIP** is set to **Auto assign**. Select the bandwidth based on service requirements. The unit is Mbit/s.

5. (Optional) Select an EIP.

   This parameter is displayed only when **EIP** is set to **Use existing**. You can select an available EIP from the drop-down list.

Instance Management

1. Set **ECS Name**.

   The **ECS Name** will be the same as the initial hostname in the ECS OS.

   The name can contain only letters, digits, underscores (_), hyphens (-), and periods (.).

   NOTE

   The name of a Windows ECS can contain a maximum of 15 characters and must be unique, or some Windows applications may be unavailable.

   The naming rules of hostnames comply with **RFC 952** and **RFC 1123**.

   When you set the ECS name and hostname, you are advised to use letters (a-z), digits (0-9), and hyphens (-) to prevent unknown issues. In the ECS:

   - Underscores (_) will be converted to hyphens (-).
   - A combination of a hyphen and underscore (-_) will be converted to a hyphen (-).
   - Periods (.), hyphens (-), underscores (_), and non-Latin characters at the beginning of the name will be ignored.
   - For periods (.) and non-Latin characters that are not at the beginning of the name, they and any content following them will be ignored.

   - When you purchase multiple ECSs in batches, the system automatically appends numbers to the end of each ECS name. Custom naming is supported.

- Automatic naming: The system automatically appends four-digit numbers to the end of each instance name, moving up in increments of 1. For example, if you enter **ecs**, the first instance will be named as **ecs-0001**, the second as **ecs-0002**, and so on. If an ECS named **ecs-0010** already exists, the subsequently created ECSs will be automatically named from **ecs-0011**.
- Custom naming: You can create a custom naming rule using the format, "name_prefix[begin_number,bits]name_suffix", where "begin_number" is a value from 0 to 9999, and "bits" from 1 to 4. For example, if you created a custom naming rule ecs[66,3]ecs[66,3]abc and created two ECSs, the system automatically names the instances ecs066abc and ecs067abc.

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

2. Set **Login Mode**.

Login Mode specifies the method for logging in to an ECS.

Table 2-21 Login mode parameters

Option	Description	Scenarios and Constraints	Reference
Password	A username and its initial password are used for ECS login authentication. The initial password of user **root** is used for authenticating Linux ECSs.	It is recommended that you set passwords with high complexity to prevent malicious attacks. The passwords must meet the requirements described in **Table 2-22**. **NOTE** The system does not periodically change the ECS password. It is recommended that you change your password regularly for security.	**Application Scenarios for Using Passwords**
Key pair	A key pair is used for ECS login authentication. You can select an existing key pair, or click **Create Key Pair** and create a desired one.	Key pair authentication is more secure than password authentication. **NOTE** If you use an existing key pair, make sure that you have saved the key file locally. Otherwise, logging in to the ECS will fail.	**(Recommended) Creating a Key Pair on the Management Console**

Option	Description	Scenarios and Constraints	Reference
Password from image	This parameter is displayed only when you select **Private Image** for **Image** and use a Linux private image with a password configured. You can use the password of the selected private image for logging in to the ECS.	Make sure that a password has been set for the selected private image.	**Encrypting Images** Image Management Service User Guide

Table 2-22 Password complexity requirements

Parameter	Requirement
Password	● Consists of 8 to 26 characters. ● Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Linux: !@%-_=+[]:./^,{}? ● Cannot contain the username or the username spelled backwards.

3. (Optional) Set a password.

   This parameter is displayed only when **Login Mode** is set to **Password**.

   Set **Password** and **Confirm Password** by referring to **Table 2-22**. The two values entered must be the same.

4. (Optional) Set a key pair.

   This parameter is displayed only when **Login Mode** is set to **Key pair**. You can select an available key pair from the drop-down list or create a key pair by referring to **(Recommended) Creating a Key Pair on the Management Console**.

   NOTE

   If you use an existing key pair, make sure that you have saved the key file locally. Otherwise, logging in to the ECS will fail.

5. (Optional) Set **Enterprise Project**.

   This function is provided for enterprise users.

An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.

Select an enterprise project from the drop-down list. For more details, see *Enterprise Management User Guide*.

6. (Optional) Set **Tag**.

   You can add tags to ECSs.

   Tags help you easily identify and manage your ECSs. You can add up to 10 tags to an ECS.

   For details, see **Overview**.

   NOTE

   Tags added during ECS creation will also be added to the created EIP and EVS disks (including the system disk and data disks) of the ECS. If the instance uses an existing EIP, the tags will not be added to that EIP.

   After creating the ECS, you can view the tags on the pages providing details about the ECS, EIP, and EVS disks.

Advanced Settings

1. Set **Detailed monitoring**.

   If you select certain public images, it is a good practice to use the host monitoring function. Host monitoring collects ECS OS metrics, such as CPU usage, memory usage, and network status, so that you can use these metrics to monitor resource utilization or locate a fault.

   After you enable detailed monitoring, an agent will be installed on the ECS to provide 1-minute fine-grained monitoring of ECS metrics, such as vCPUs, memory, network, disks, and processes.

   For details about the monitoring metrics after the agent is installed, see **OS Monitoring Metrics Supported by ECSs with the Agent Installed**.

2. (Optional) Set **ECS group**.

   An ECS group applies the anti-affinity policy to the ECSs in it so that the ECSs are automatically allocated to different hosts. For instructions about how to create an ECS group, see section **Managing ECS Groups**.

   NOTE

   An existing ECS attached with a local disk cannot be added to an ECS group. To use ECS group functions, select an ECS group when creating an ECS.

3. Add an ECS description.

4. Set **User Data**.

   You can inject user data, for example, inject the OS initialization script, to ECSs during ECS creation. During the first startup of the ECSs, the data will be automatically injected.

   - **As text**: allows you to enter the user data in the text box.
   - **As file**: enables the text to automatically inject a script file or other files into a specified directory on an ECS when you create the ECS.

   For details, see **Injecting User Data**.

5. Set **Agency**.

When your ECS resources need to be shared with other accounts, or your ECS is delegated to professional personnel or team for management, the tenant administrator creates an agency in IAM and grants the ECS management permissions to the personnel or team. The delegated account can log in to the cloud system and switch to your account to manage resources. You do not need to share security credentials (such as passwords) with other accounts, ensuring the security of your account.

If you have created an agency in IAM, select the agency from the drop-down list. For more information about agencies, see **Account Delegation**.

6. Set **CPU Options**.
   - To configure hyper-threading for an ECS, select **Specify CPU options**.

     For details about hyper-threading, see **Enabling or Disabling Hyper-Threading**.
   - Set **Threads per Core**.

     This parameter is displayed when **Specify CPU options** is selected. You can select a parameter value from the drop-down list.
     - **1**: one thread per core, which means hyper-threading is disabled.
     - **2** (default value): two threads per core, which means hyper-threading is enabled.

Purchase Details

1. (Optional) Select the required duration for ECSs.

   This parameter is displayed only when **Billing Mode** is set to **Yearly/Monthly**. The duration can be from 1 month to .
2. (Optional) Set **Auto-renew**.

   This parameter is displayed only when **Billing Mode** is set to **Yearly/Monthly**.

   You can select **Auto-renew** to automatically renew yearly/monthly resources when they expire.
   - Monthly: Your subscription will be automatically renewed each month.
   - Yearly: Your subscription will be automatically renewed each year.

   For details about auto-renewal, see **Auto-Renewal Rules**.
3. (Optional) Determine whether to select **Set scheduled deletion time** for **Required Duration**.

   This parameter is displayed only when **Billing Mode** is set to **Pay-per-use** or **Spot pricing**. If you select **Set scheduled deletion time** and set a time, the ECS will be automatically deleted when the time is reached.

   NOTICE

   After you set a scheduled deletion time, the system automatically deletes the ECS at the specified time. Back up data in advance.

   The scheduled deletion time must be at least 1 hour from the current time but not more than 3 years from now. You can change the scheduled deletion time before the instance is deleted.

The system executes the scheduled deletion task every 5 minutes and stops the billing after the ECS is deleted.

4. Set **Quantity**.

   You can set how many ECSs to be created in a batch. ECSs created in a batch have the same configurations.

   The remaining number of ECSs you are allowed to create is displayed. If the number of ECSs you want to create exceeds the quota,**increase the quota**.

   NOTE

   You can set the following at the bottom of the purchase page:

   - When **Billing Mode** is set to **Yearly/Monthly**, you can set quantity and required duration.
   - When **Billing Mode** is set to **Pay-per-use** or **Spot pricing**, you can set quantity.

   After the setting is complete, you can hover over the price to view billing items. If you have any questions about the price, click **Pricing details** to learn more.

Step 3: Confirm the Configuration and Submit the Order

1. In the **Configuration Summary** panel on the right side, confirm the ECS details.

   Mandatory fields that are not configured are displayed in red. You need to set them in the parameter configuration area.

2. Read and select the agreement, and click **Submit**.

   After an ECS is created, it will start by default.

2.3.3 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 spot ECSs, see **Spot Pricing (for Spot Instances)**.

Notes and 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 KooGallery image.

- Spot ECSs cannot be switched to yearly/monthly ECSs.
- Spot ECSs do not support system disk detachment.
- 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.

Purchasing a Spot ECS

Follow the instructions provided in **Purchasing an ECS**, **Login Overview (Windows)**, or **Login Overview (Linux)** 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.

(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 & Deployment**, 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 2-23**.

Table 2-23 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 Notification	Select **Yes**.
	SMN Topic	Select a topic from the drop-down list. If there are no proper SMN topics, create one. 1. Click **Create Topic** to switch to the SMN console. 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 & Deployment**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in **Table 2-24** and search for traces as needed.

Table 2-24 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

2.3.4 Purchasing a Spot Block ECS

Scenarios

A spot block ECS is billed in spot pricing mode. You can purchase and use such ECSs at a discount price. A spot block ECS performs as well as the ECSs with the same specifications in other billing modes. If inventory resources are insufficient, the spot block ECS will be reclaimed.

Compared with pay-per-use and yearly/monthly ECSs, spot block ECSs offer the same performance at a lower price.

For more information about spot block ECSs, see **Spot Pricing (for Spot Block Instances)**.

Notes and Constraints

- Only general computing-plus ECSs support spot block pricing payments.
- Huawei Cloud provides the utmost efforts to ensure the proper running of your spot block ECSs. However, when system resources are insufficient or in other extreme cases, the spot block ECSs will be released. Back up data in advance.
- Spot block ECSs are only supported in some regions and for some specifications. For details, see the information displayed on the management console.

- The price of spot block ECSs varies by the predefined duration.
- Spot block ECSs cannot be switched to pay-per-use or yearly/monthly ECSs.
- Spot block ECSs do not support specifications modification.
- Spot block ECSs do not support OS change.
- Spot block ECSs do not support automatic recovery.
- Spot block ECSs do not support system disk detachment.
- When a spot block 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 block ECS will not be released after it is reclaimed. If you want to be notified when a spot block ECS is reclaimed so that you can determine whether to manually release data disks and EIP, set a reclaim notification.

Purchasing a Spot Block ECS

Follow the instructions provided in **Purchasing an ECS**, **Login Overview (Windows)**, or **Login Overview (Linux)** to buy and log in to spot block ECSs. Pay attention to the following settings:

When purchasing a spot block 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, when inventory resources are insufficient, the spot block ECSs will be released. Back up data in advance.
- Set **Spot Type** to **Spot block**.
- Specify **Predefined Duration** and **Number of Durations**.

  Predefined duration is a duration that you specify for your spot block ECS. During the predefined duration, if your spot block ECS is automatically terminated by the system, you will not be billed for the resource usage within the predefined duration. If you delete your spot block ECS within the predefined duration, you will be billed based on the usage duration.
- Click **Next**, confirm that the specifications and price are correct, agree to the service agreement, and click **Submit**.

  NOTE

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

Enabling Reclaim Notifications

After purchasing a spot block ECS, you can use it like using the ECSs in other billing modes. However, the system will reclaim the instance when the predefined duration is reached or system resources are insufficient. Use the following method to enable notifications:

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 & Deployment**, 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 2-25**.

Table 2-25 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 Notification	Select **Yes**.
	SMN Topic	Select a topic from the drop-down list. If there are no proper SMN topics, create one. 1. Click **Create Topic** to switch to the SMN console. 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 & Deployment**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in **Table 2-26** and search for traces as needed.

Table 2-26 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

2.3.5 Purchasing an ECS Using a Private Image

Scenarios

A private image is a personal image created or imported by a user and is visible only to the user who created or imported it. It contains an OS, preinstalled public applications, and the user's personal applications, saving the time for configuring an ECS repeatedly.

The difference between ECSs created using public and private images is as follows:

- ECSs created using public images contain only the OS and pre-installed public applications. You need to install your personal applications if required.
- ECSs created using private or shared images contain the OS, pre-installed public applications, and your personal applications.

NOTE

You can also use an encrypted image to create ECSs. For details, see **Encrypting Images**.

Notes and Constraints

- If you use a full-ECS image to create an ECS, the EVS disks associated with the full-ECS image do not support the function of creating disks from a data disk image.
- If a full-ECS image is in **Normal** state and the system displays message "Available in AZ*x*", the full-ECS image can be used to create ECSs in this AZ

only, and the encryption attributes of the system and data disks of the created ECSs are the same as those of the system and data disks specified in the full-ECS image. The SCSI, encryption, and sharing attribute settings of the system and data disks cannot be modified during ECS creation.

- If a full-ECS image is in **Normal** state and the system does not display message "Available in AZ*x*", the full-ECS image can be used to create ECSs in the entire region, and the encryption attributes of the system and data disks of the created ECSs are the same as those of the system and data disks specified in the full-ECS image. The SCSI, encryption, and sharing attribute settings of data disks can be modified during ECS creation.
- An ISO image created from an ISO file cannot be used to create ECSs. It can only be used to create a temporary ECS. You need to install an OS and drivers on the temporary ECS and use the temporary ECS to create a system disk image first.
- You are advised to use ECSs created from ISO images only for OS installation because such ECSs do not support some functions, such as disk attachment.
- To ensure that NIC multi-queue is enabled on an ECS created using a private image, configure NIC multi-queue when creating such a private image. NIC multi-queue assigns interrupts for queues to different vCPUs for higher network packets per second (PPS) and bandwidth.

  For details, see **How Do I Enable NIC Multi-Queue for an Image?**

Procedure (on the ECS Console)

1. Log in to the management console and go to the **ECS console**.
2. Click **Buy ECS** and specify parameters in **Basic Configuration** (**Billing Mode**, **Region**, and **AZ**) and **Instance**.
3. In the **OS** module, choose **Private image** for **Image**.
4. Select a private image from the drop-down list.
5. Configure other parameters and complete the ECS purchase.

   For details, see **Purchasing an ECS in Custom Config Mode**.

Procedure (on the IMS Console)

1. Log in to the IMS console and go to the **private image** list.
2. Locate the row that contains the target image and click **Apply for Server** in the **Operation** column. The **Buy ECS** page is displayed.
3. On the displayed **Buy ECS** page, set required parameters.

   For details, see **Purchasing an ECS in Custom Config Mode**.

   NOTE

   - When creating an ECS using a system disk image, you can reset the ECS specifications and system disk type. The ECS's system disk size must be greater than that of the image.
   - If you use a private full-ECS image that contains one or more data disks to create an ECS, the system automatically sets the data disk parameters for the ECS. The capacity of the system and data disks can only be expanded. It cannot be reduced.
   - If a full-ECS image contains multiple data disks, it takes some time to load and display the disk information.

2.3.6 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 and Constraints

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

Procedure

1. Log in to the management console and access the **Elastic Cloud Server** page.
2. Locate the row containing the target ECS and choose **More** > **Buy Same ECS** in the **Operation** column.

   The system switches to the **Buy ECS** page and automatically copies the parameter settings of the selected ECS.
3. Adjust the settings of the new ECSs as needed.

   For details about parameter settings, see **Purchasing an ECS**.
4. Confirm the ECS details, select the agreement, and click **Create**.

   NOTE

   For security purposes, you need to 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**.

2.4 Logging In to a Windows ECS

2.4.1 Login Overview (Windows)

Constraints

- Only a running ECS can be logged in to.
- The username for logging in to a Windows ECS is **Administrator**.
- 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 advised to log in to such 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 2-27 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 **Logging In to a Windows ECS Using MSTSC**.	The target ECS needs to have 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 macOS Server**.	
	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**.	

ECS OS	Local OS	Connection Method	Requirement
	Windows	Through the management console. For details, see **Logging In to a Windows ECS Using VNC**.	No EIP is required.

Helpful Links

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

2.4.2 Logging In to a Windows ECS Using VNC

Scenarios

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

If you cannot use the MSTSC or other remote login tools to log in to an ECS, you can use the VNC login mode. This login mode is mainly used in emergency O&M scenarios for you to view and perform maintenance operations.

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.

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.
5. In the **Operation** column of the target ECS, click **Remote Login**.
6. In the **Logging In to a Windows ECS** dialog box, expand **Other Login Modes** and click **Log In** in the **VNC Login** area.
7. (Optional) When the system displays "Press Ctrl+Alt+Delete to unlock", click **Ctrl+Alt+Del** in the upper part of the remote login page to log in to the ECS.

Figure 2-11 Ctrl+Alt+Del

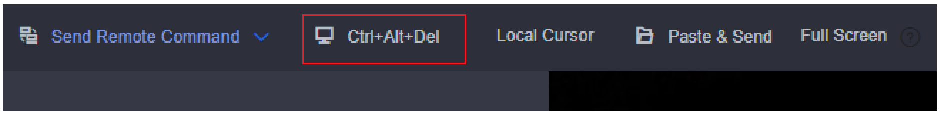

8. Enter the ECS password as prompted.

Helpful Links

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

2.4.3 Logging In to a Windows ECS 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.
- 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 which the ECS belongs to. 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 2-12 Logging in to an ECS using MSTSC

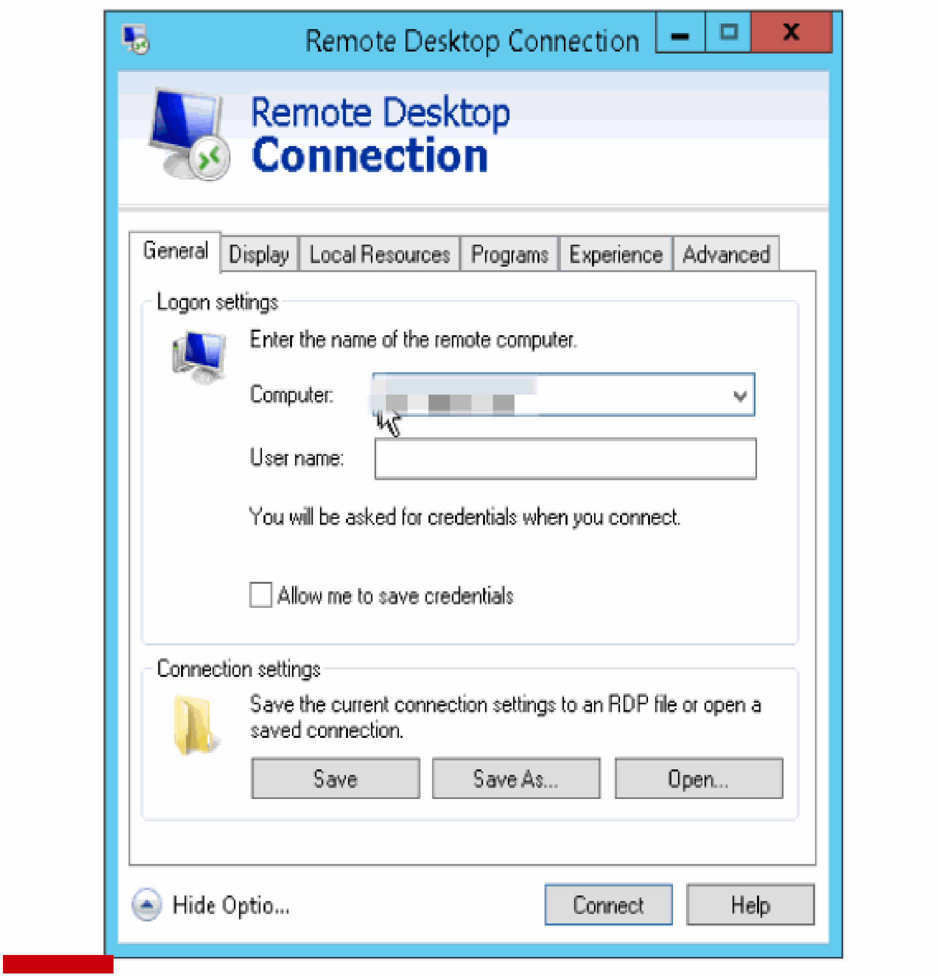

Enter the name of the remote computer.

Computer: [redacted]

User name:

You will be asked for credentials when you connect.

□ Allow me to save credentials

Save the current connection settings to an RDP file or open a saved connection.

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 2-13 Showing options

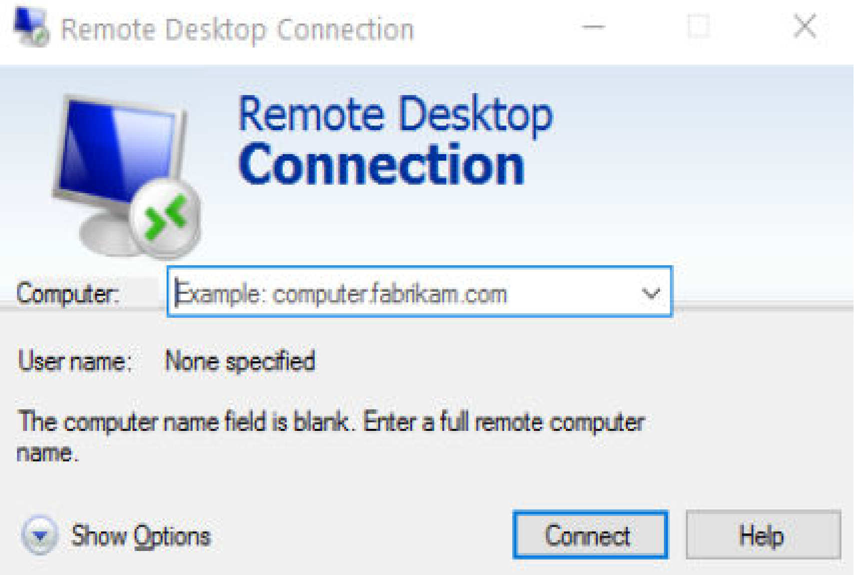

Computer: Example: computer.fabrikam.com

User name: None specified

The computer name field is blank. Enter a full remote computer name.

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 2-14 Remote Desktop Connection

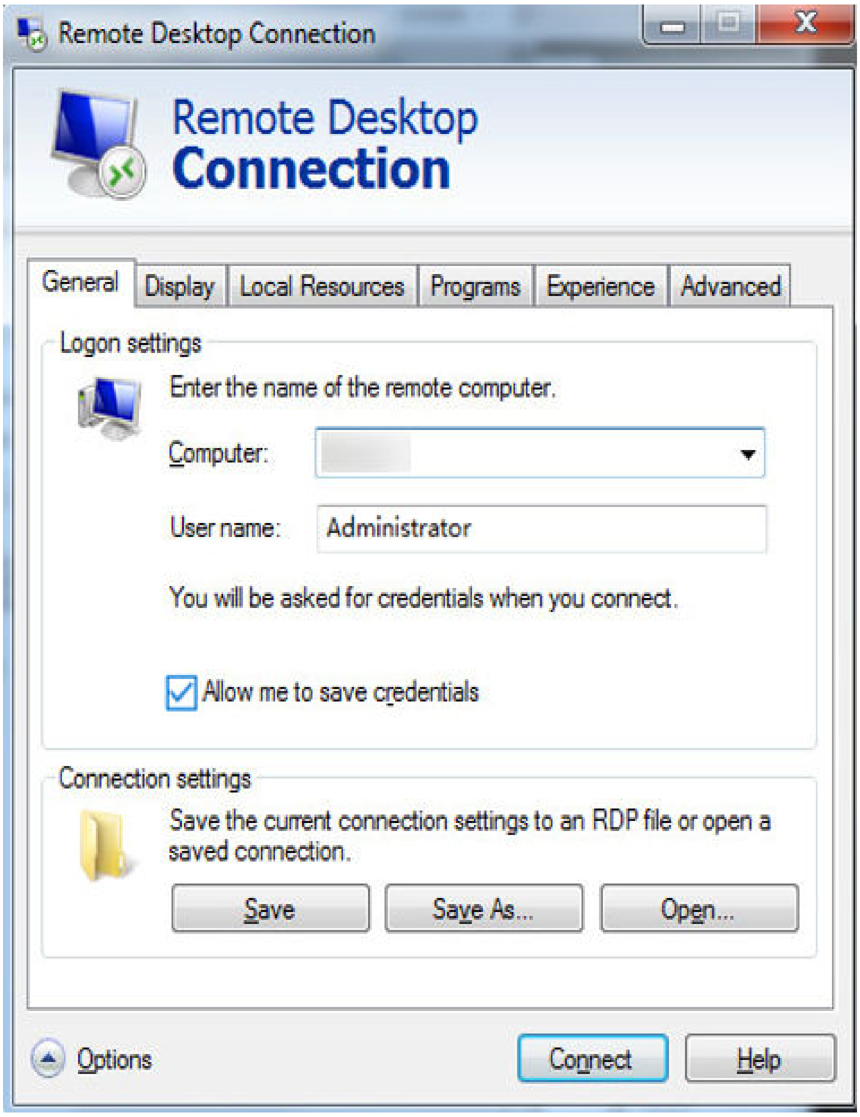

Logon settings

Enter the name of the remote computer.

Computer:

User name: Administrator

You will be asked for credentials when you connect.

☑ Allow me to save credentials

Connection settings

Save the current connection settings to an RDP file or open a saved 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 2-15 Clipboard

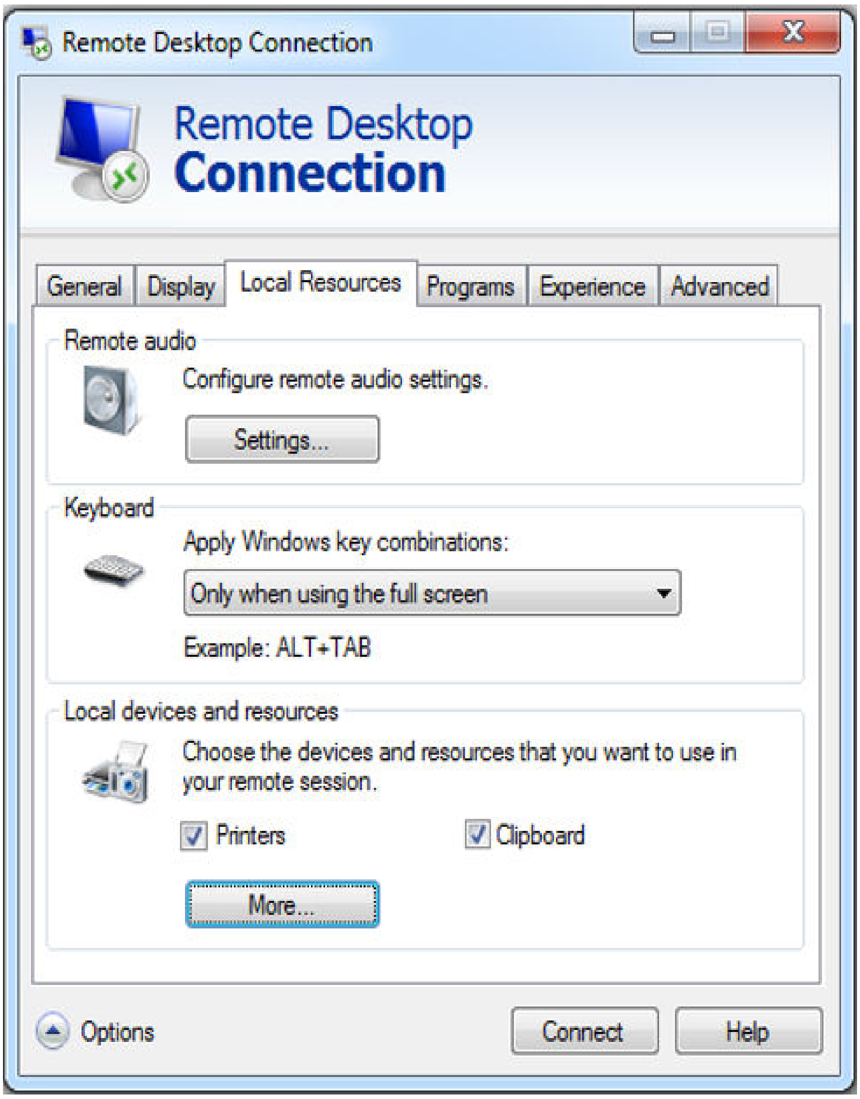

Configure remote audio settings.

Apply Windows key combinations: Only when using the full screen

Example: ALT+TAB

Choose the devices and resources that you want to use in your remote session.

☑ Printers ☑ Clipboard

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

Figure 2-16 Drives

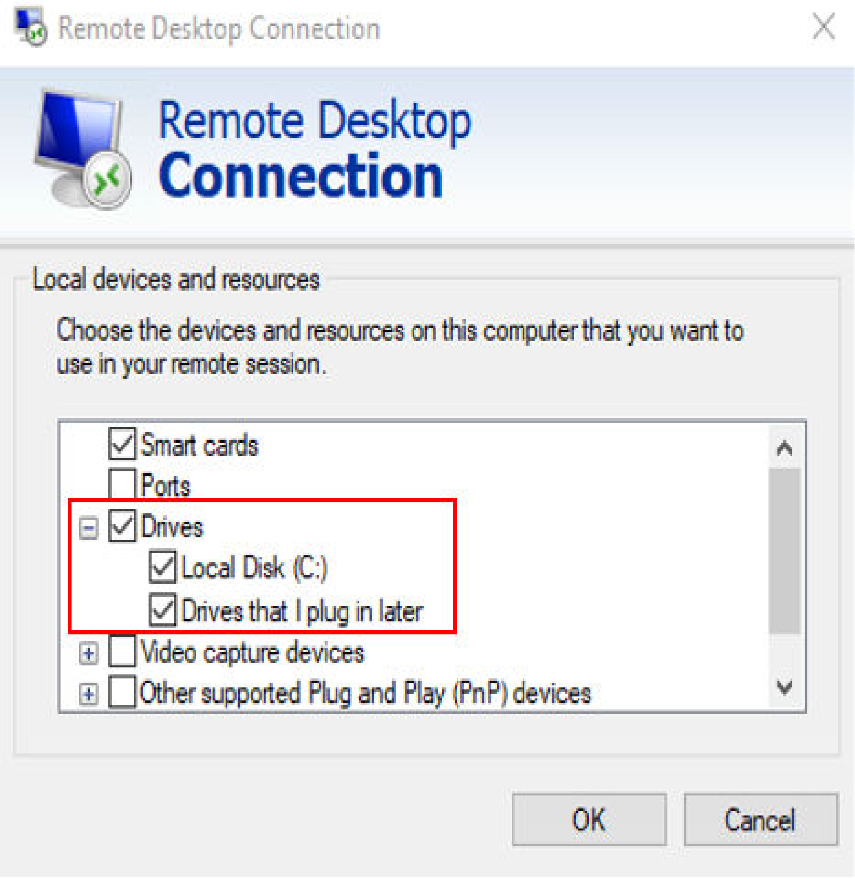

Remote Desktop Connection

Remote Desktop Connection

Local devices and resources

Choose the devices and resources on this computer that you want to use in your remote session.

☑ Smart cards
☐ Ports
☑ Drives
☑ Local Disk (C:)
☑ Drives that I plug in later
☐ Video capture devices
☐ Other supported Plug and Play (PnP) devices

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

Figure 2-17 Adjusting the size of the desktop

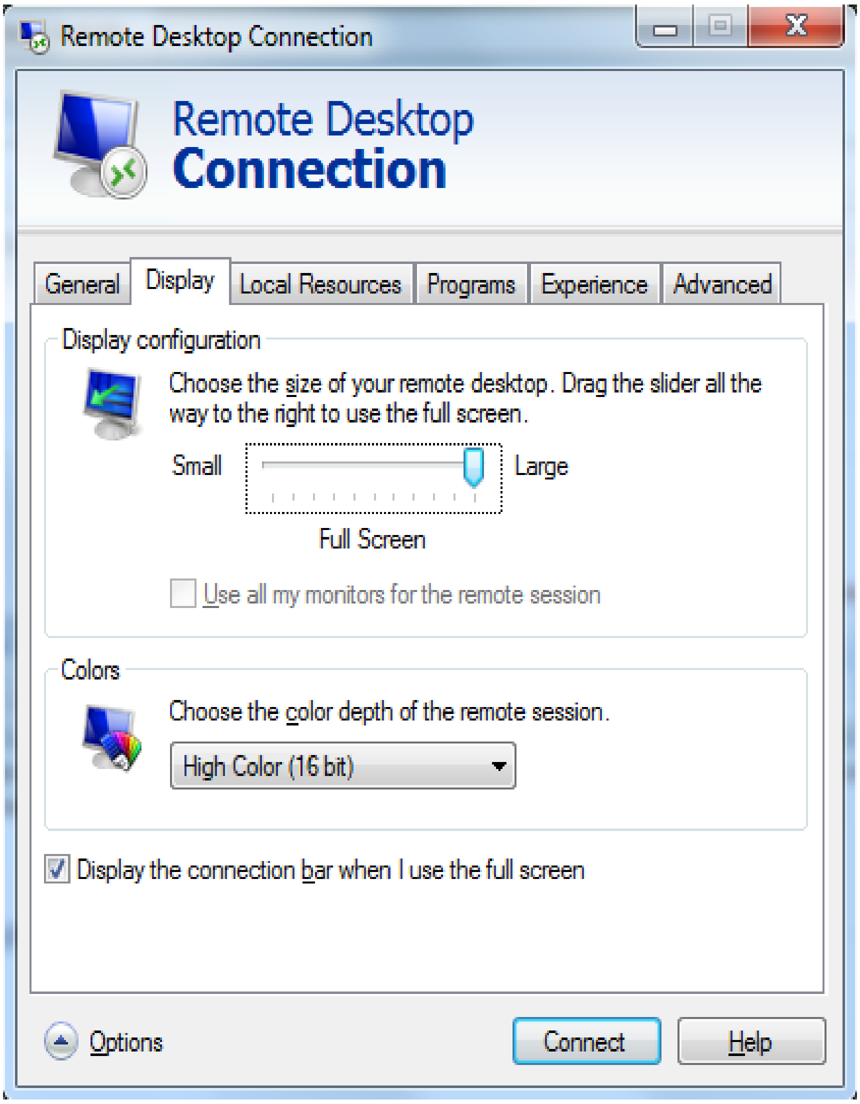

Display configuration

Choose the size of your remote desktop. Drag the slider all the way to the right to use the full screen.

Small Large

Full Screen

□ Use all my monitors for the remote session

Colors

Choose the color depth of the remote session.

High Color (16 bit)

☑ Display the connection bar when I use the full screen

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 **Logging In to a Windows ECS 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 2-18 System Properties

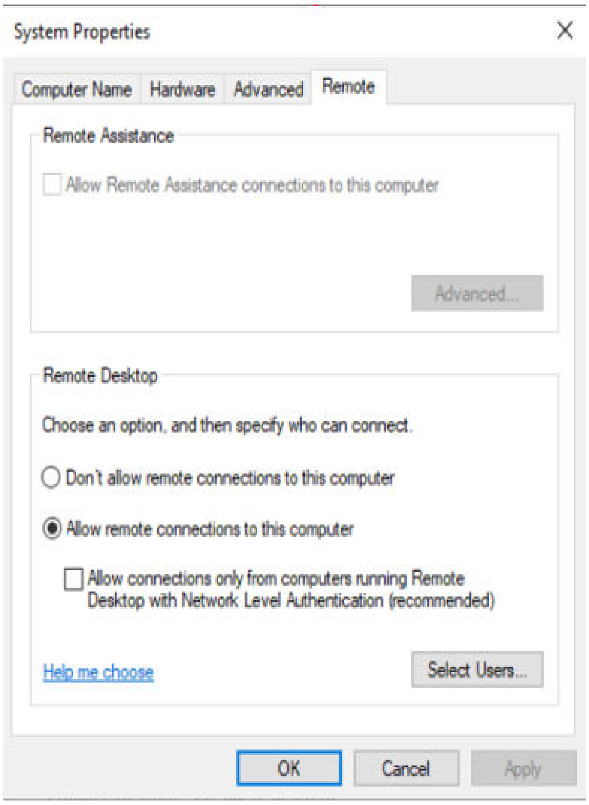

System Properties ×

Computer Name | Hardware | Advanced | Remote

Remote Assistance

□ Allow Remote Assistance connections to this computer

Advanced...

Remote Desktop

Choose an option, and then specify who can connect.

○ Don't allow remote connections to this computer

● Allow remote connections to this computer

□ Allow connections only from computers running Remote Desktop with Network Level Authentication (recommended)

Help me choose

Select Users...

OK | Cancel | Apply

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

Helpful Links

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

2.4.4 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.
- The ECS must have an EIP bound.

  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 which the ECS belongs to.
- 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 2-28 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 **Logging In to a Windows ECS 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 2-19 System Properties

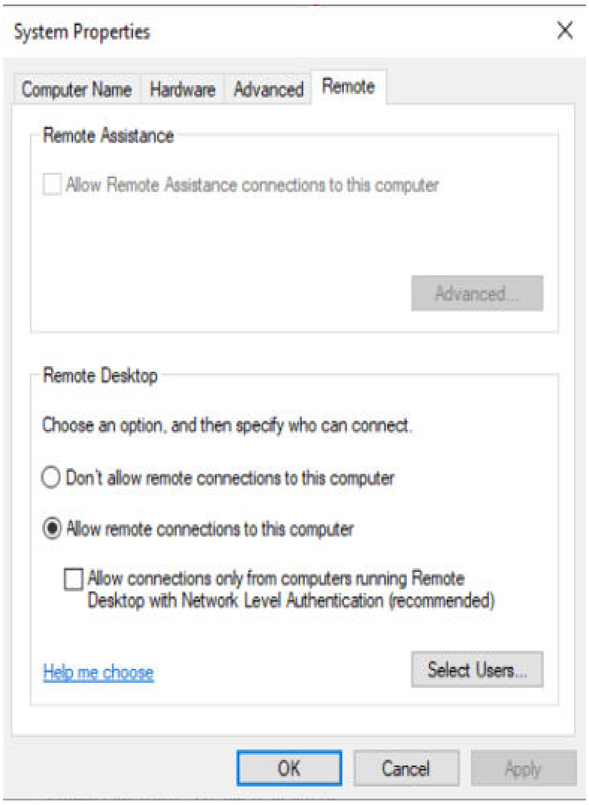

Remote Assistance

□ Allow Remote Assistance connections to this computer

Remote Desktop

Choose an option, and then specify who can connect.

○ Don't allow remote connections to this computer

● Allow remote connections to this computer

□ Allow connections only from computers running Remote Desktop with Network Level Authentication (recommended)

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

2.4.5 Logging In to a Windows ECS from a macOS Server

Scenarios

This section describes how to use a remote login tool to log in to a Windows ECS from a macOS server. 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.
- 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 which the ECS belongs to. For details, see **Configuring Security Group Rules**.

- The remote access tool supported by Mac, such as Microsoft Remote Desktop for Mac has been installed.

  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 2-20 Add Desktop

   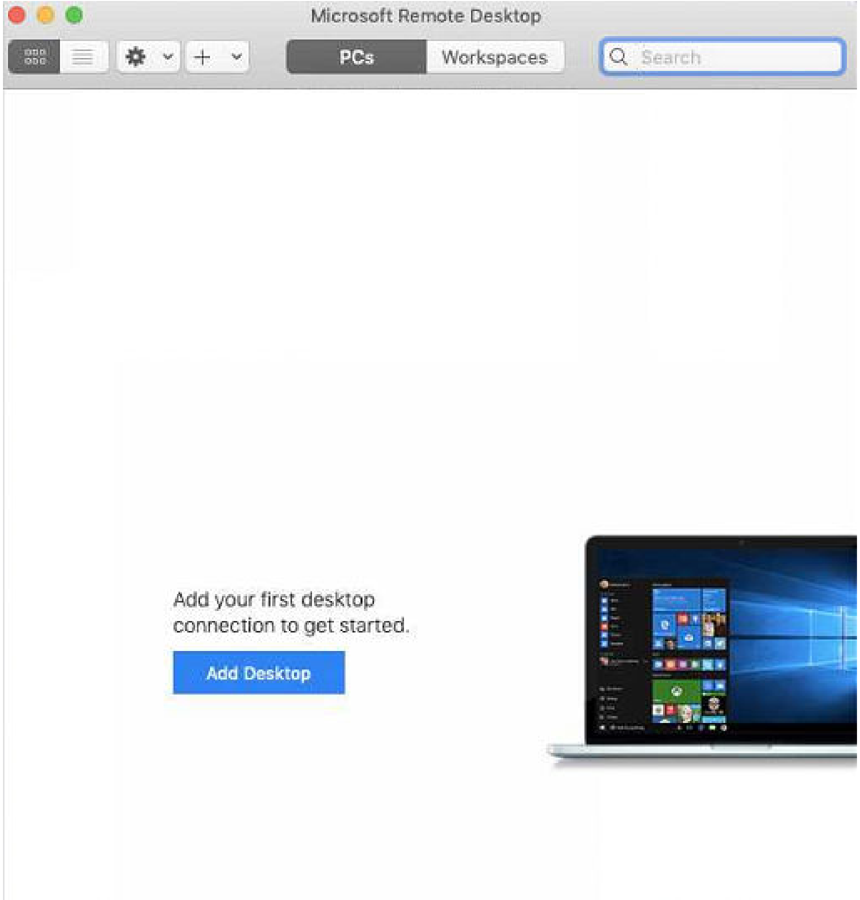

3. On the **Add PC** page, set login information.
   - **PC name**: Enter the EIP bound to the target Windows ECS.
   - **User account**: Select **Add user account** from the drop-down list.

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

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

Figure 2-21 Add user account

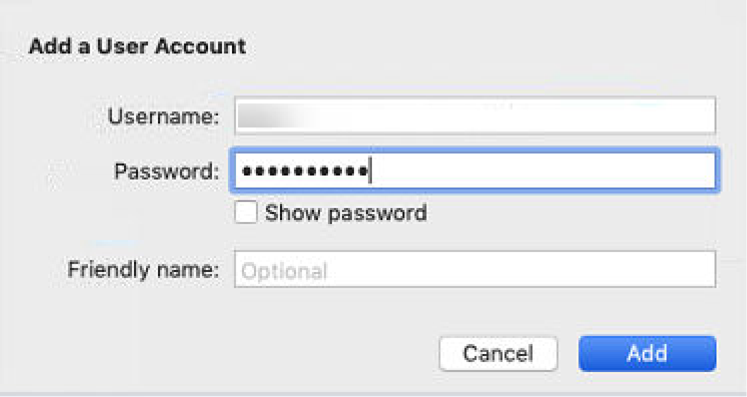

Add a User Account

Username:

Password: ••••••••••••

□ Show password

Friendly name: Optional

Figure 2-22 Add PC

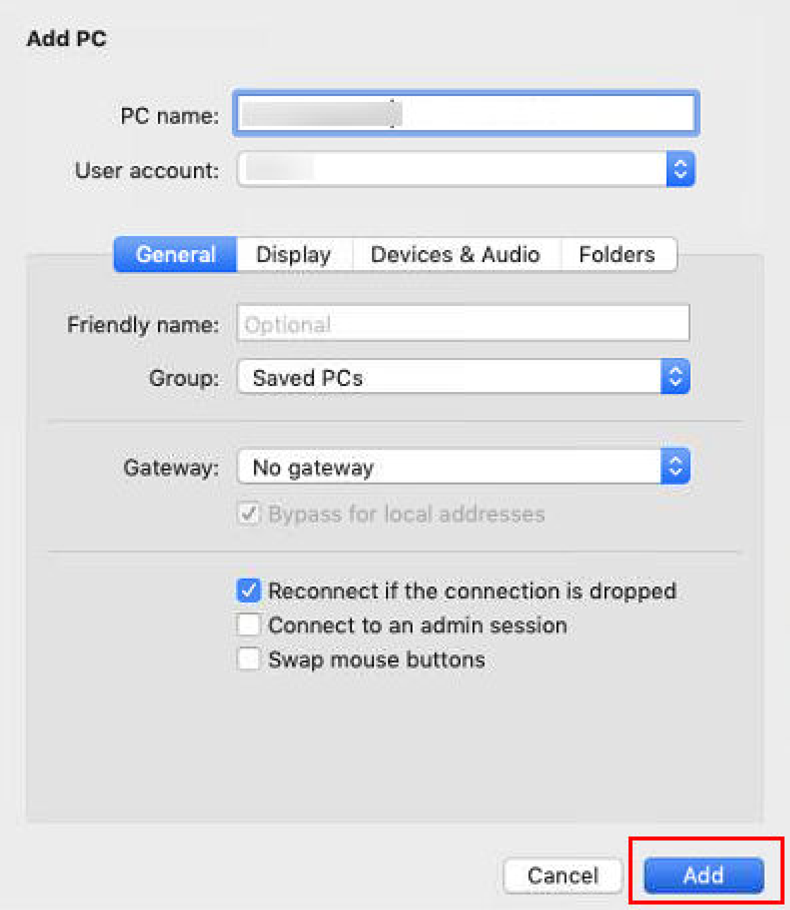

Add PC

PC name:

User account:

General | Display | Devices & Audio | Folders

Friendly name: Optional

Group: Saved PCs

Gateway: No gateway

☑ Bypass for local addresses

☑ Reconnect if the connection is dropped

□ Connect to an admin session

□ Swap mouse buttons

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

Figure 2-23 Double-click for login

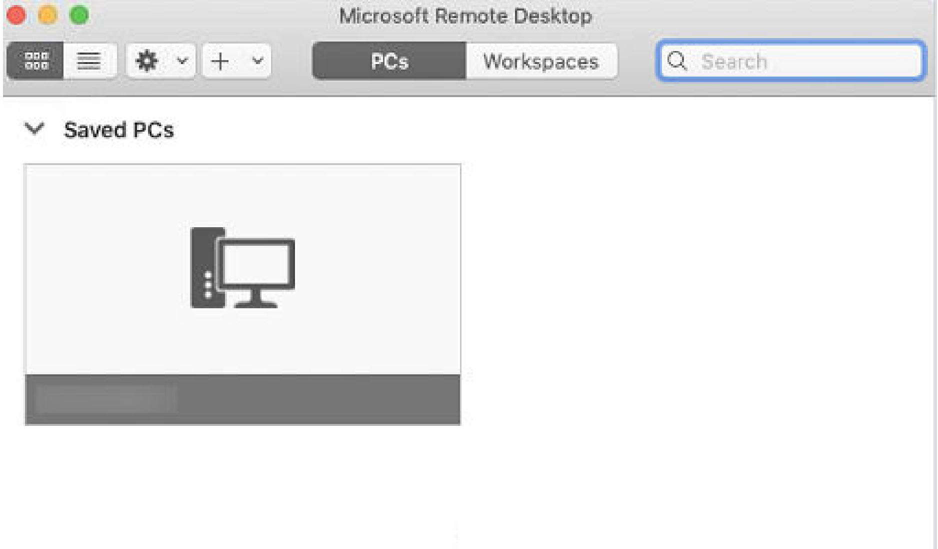

5. Confirm the information and click **Continue**.

   You have logged in to the Windows ECS.

Figure 2-24 Successful login

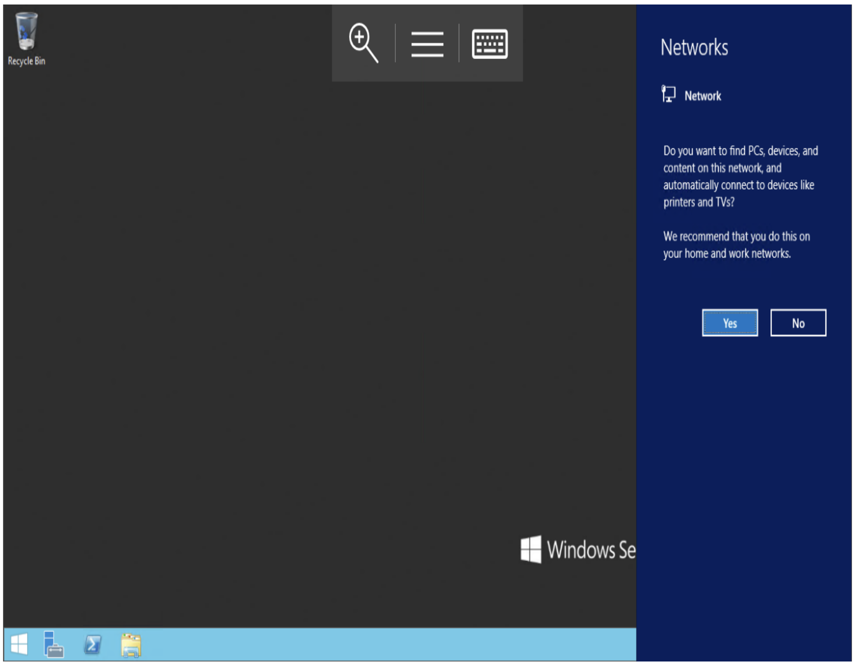

Do you want to find PCs, devices, and content on this network, and automatically connect to devices like printers and TVs?

We recommend that you do this on your home and work networks.

2.4.6 Logging In to a Windows ECS from a Mobile Terminal

Scenarios

If you want to manage Windows ECSs on the cloud anytime, anywhere, you can log in to them from a remote desktop application on your mobile device.

This section uses Remote Desktop released by Microsoft as an example to describe how to log in to a Windows ECS from an Android mobile device.

NOTE

The supported remote desktop applications may vary depending on the OS type of the mobile device. For details, see the operation guide of the corresponding application.

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- 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 which the ECS belongs to. For details, see **Configuring Security Group Rules**.
- Microsoft Remote Desktop has been installed on the mobile terminal.

Procedure

1. Start Remote Desktop on the mobile device.
2. In the upper right corner of Remote Desktop, click + and select **Desktop**.

Figure 2-25 Remote desktop

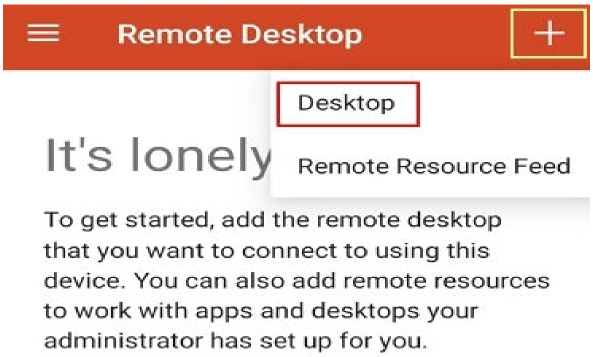

It's lonely

To get started, add the remote desktop that you want to connect to using this device. You can also add remote resources to work with apps and desktops your administrator has set up for you.

3. In the **Add desktop** dialog box, enter the Windows ECS hostname or EIP for **PC name**.
4. Click **SAVE**.
5. On the **Remote Desktop** page, click the icon of the Windows ECS to be logged in to.

Figure 2-26 Logging in to the Windows ECS

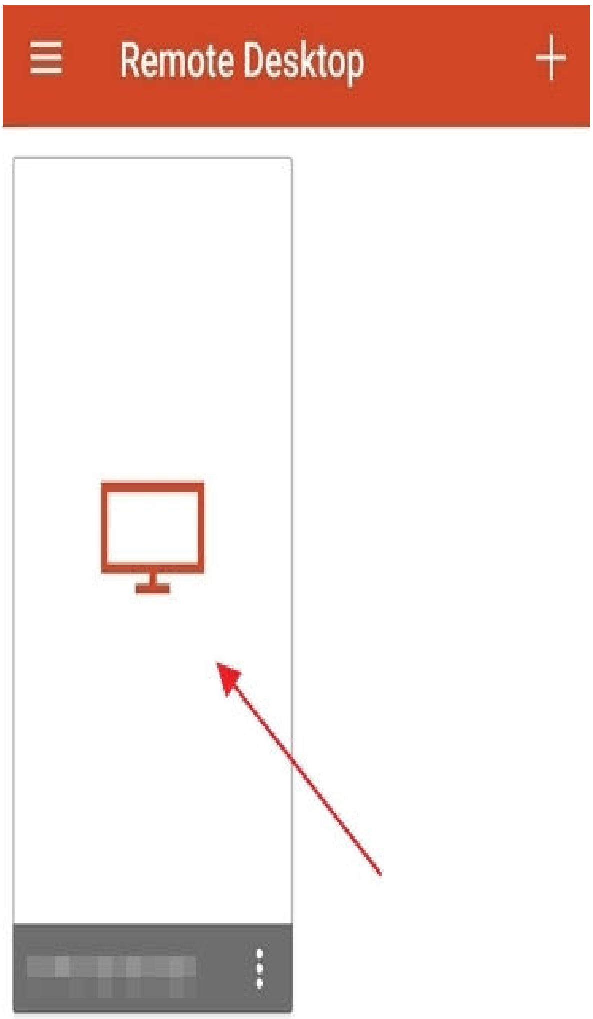

6. If the message "Certificate can't be verified. Do you want to connect anyway?" is displayed, confirming the information and click **CONNECT**.

Figure 2-27 Confirmation

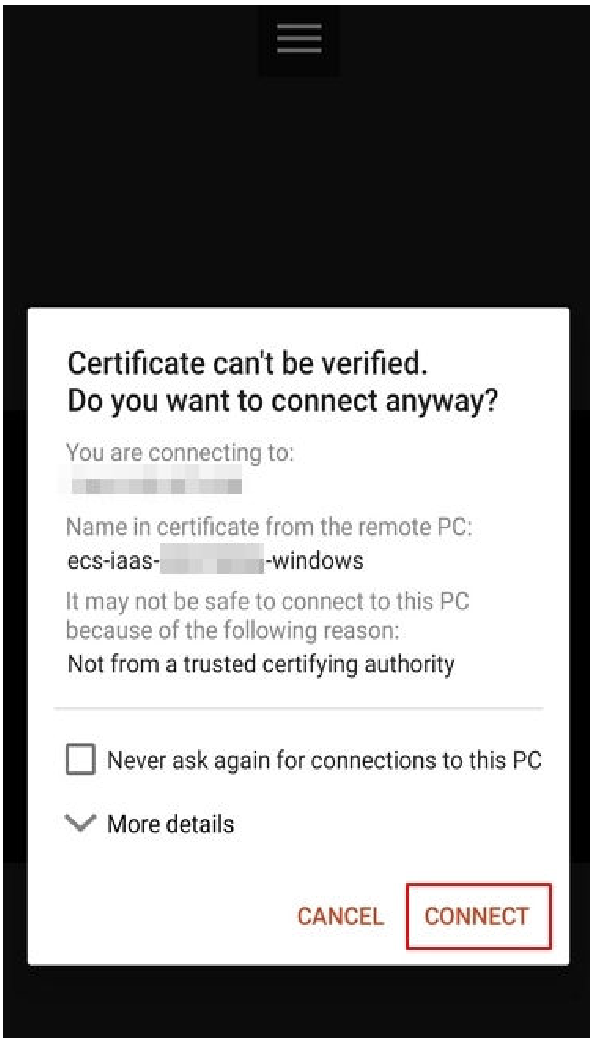

Certificate can't be verified.
Do you want to connect anyway?

You are connecting to:
[redacted]

Name in certificate from the remote PC:
ecs-iaas-[redacted]-windows

It may not be safe to connect to this PC because of the following reason:
Not from a trusted certifying authority

□ Never ask again for connections to this PC

7. On the sign-in page, enter the username (such as administrator) and password, and click **CONNECT**.

   You have logged in to the Windows ECS.

Figure 2-28 Successful login

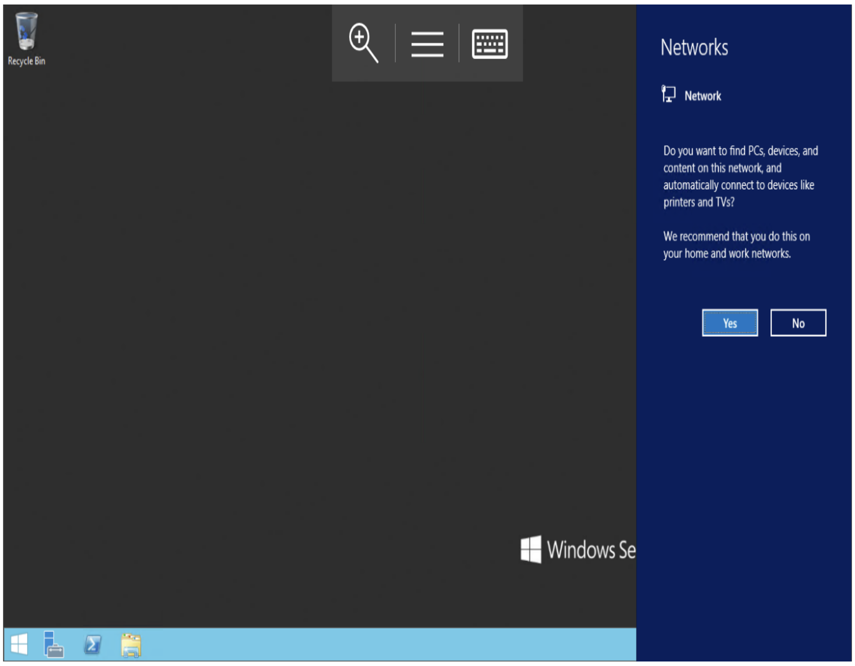
Networks

Network

Do you want to find PCs, devices, and content on this network, and automatically connect to devices like printers and TVs?

We recommend that you do this on your home and work networks.

2.5 Logging In to a Linux ECS

2.5.1 Login Overview (Linux)

Constraints

- Only a running ECS can be logged in to.
- The username for logging in to a Linux ECS is **root**.

Login Modes

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

Table 2-29 Linux ECS login modes

ECS OS	Local OS	Connection Method	Requirement
Linux	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**	The target ECS needs to have 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.)

ECS OS	Local OS	Connection Method	Requirement
	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 **Logging In to a Linux ECS 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?**

2.5.2 Logging In to a Linux ECS Using VNC

Scenarios

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

If you cannot use other remote login tools to log in to an ECS, you can use the VNC login mode. This login mode is mainly used in emergency O&M scenarios for you to view and perform maintenance operations.

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 **Logging In to a Linux ECS Using an SSH Key Pair** and **Logging In to a Linux ECS 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**.
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 2-29 Username (root as an example) and password

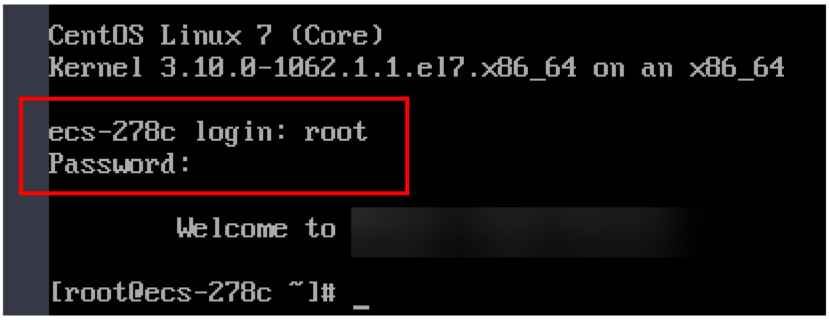

```
CentOS Linux 7 (Core)
Kernel 3.10.0-1062.1.1.el7.x86_64 on an x86_64

ecs-278c login: root
Password:

        Welcome to [redacted]

[root@ecs-278c ~]# _
```

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 2-30 Paste & Send

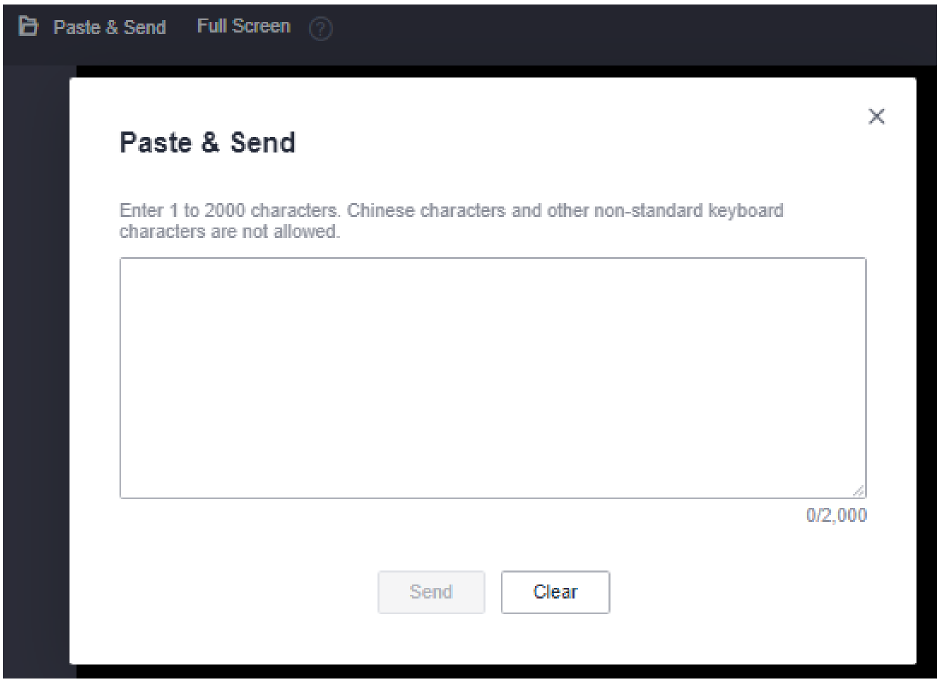

Enter 1 to 2000 characters. Chinese characters and other non-standard keyboard characters are not allowed.

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 Paste & Send 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 Paste & Send 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?**

2.5.3 Logging In to a Linux ECS Using an SSH Key Pair

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**.

- 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 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.

Figure 2-31 Importing the private key file

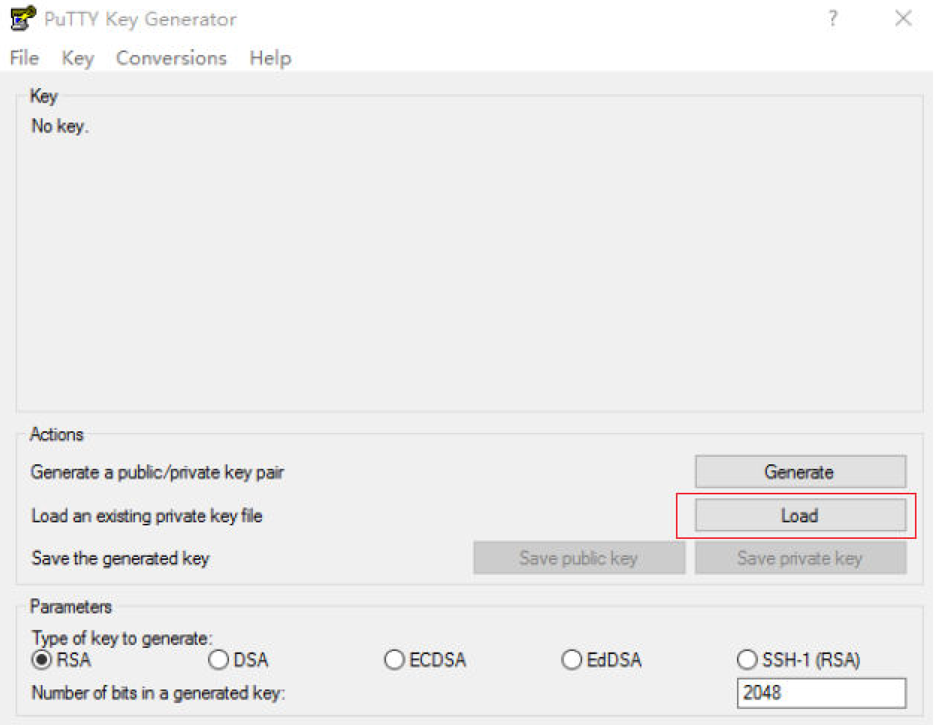

No key.

Type of key to generate:
● RSA ○ DSA ○ ECDSA ○ EdDSA ○ SSH-1 (RSA)

Number of bits in a generated key: 2048

5. In the **Actions** area, 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 2-32 Configuring the EIP

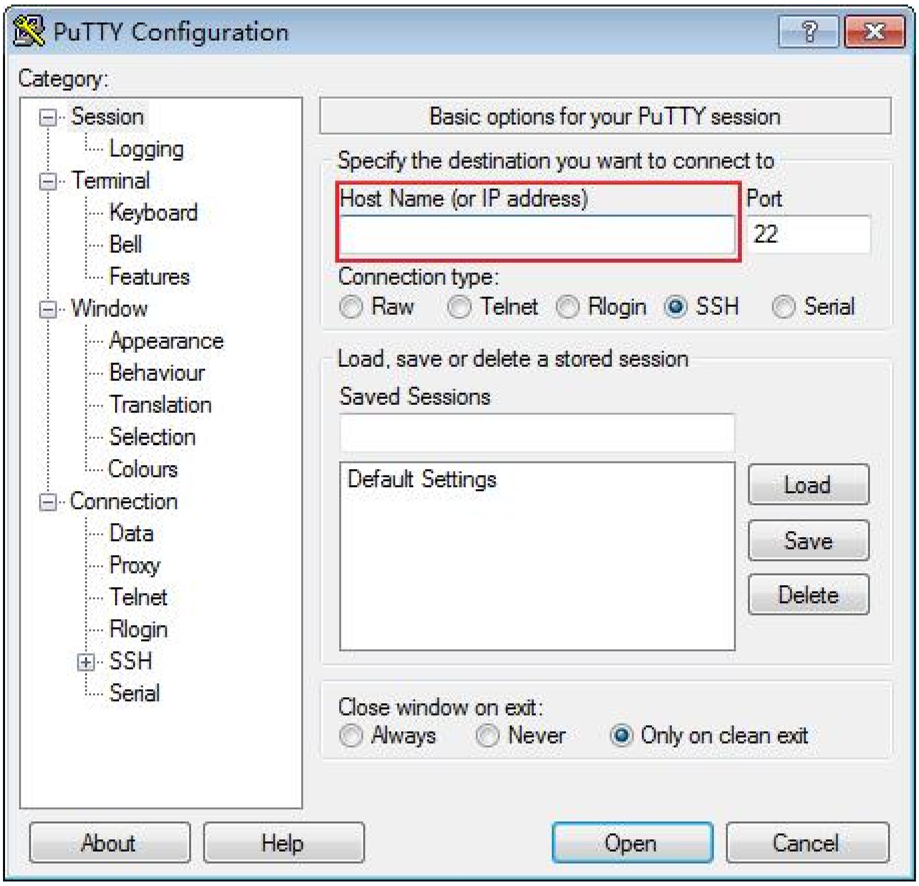

Basic options for your PuTTY session

Specify the destination you want to connect to

Host Name (or IP address)

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Load, save or delete a stored session

Saved Sessions

Default Settings

Close window on exit: ○ Always ○ Never ● Only on clean exit

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

Figure 2-33 Entering the username

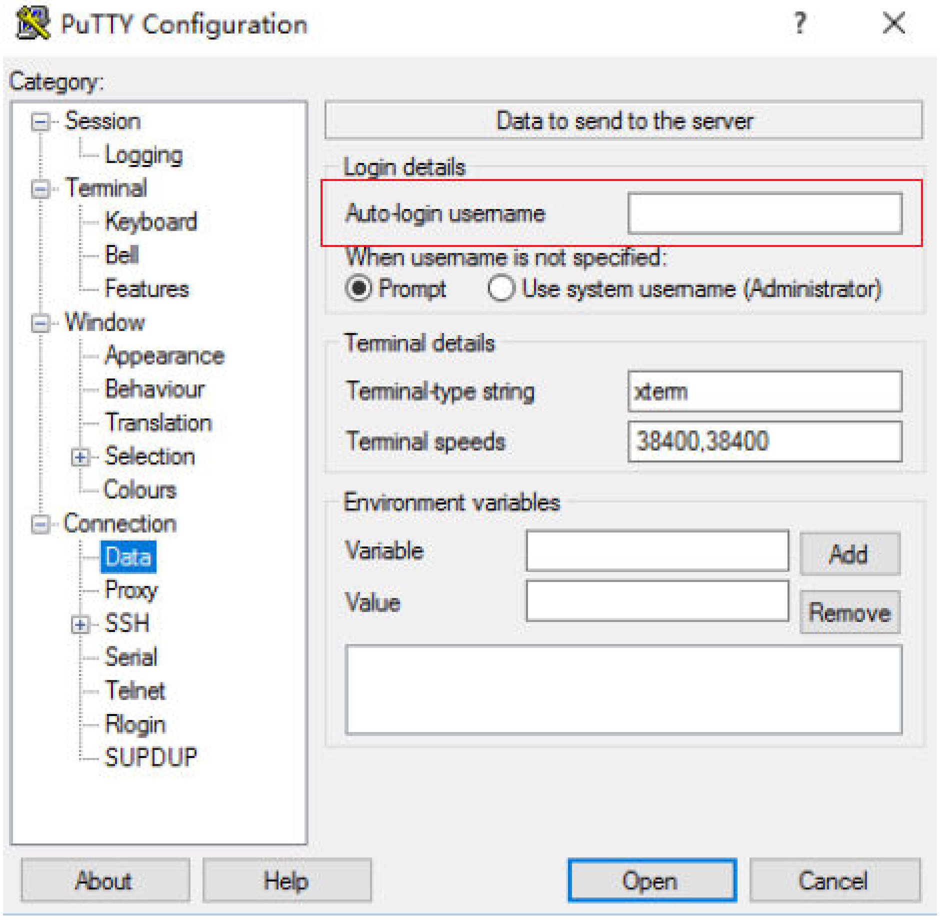

PuTTY Configuration

Category:

- Session
  - Logging
- Terminal
  - Keyboard
  - Bell
  - Features
- Window
  - Appearance
  - Behaviour
  - Translation
  - Selection
  - Colours
- Connection
  - Data
  - Proxy
  - SSH
  - Serial
  - Telnet
  - Rlogin
  - SUPDUP

Data to send to the server

Login details

Auto-login username

When username is not specified:

● Prompt ○ Use system username (Administrator)

Terminal details

Terminal-type string: xterm

Terminal speeds: 38400,38400

Environment variables

Variable

Value

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** > **Credentials**. In the configuration item **Private key file for authentication**, click **Browse** and select the private key converted in step 6.

Figure 2-34 Importing the private key file

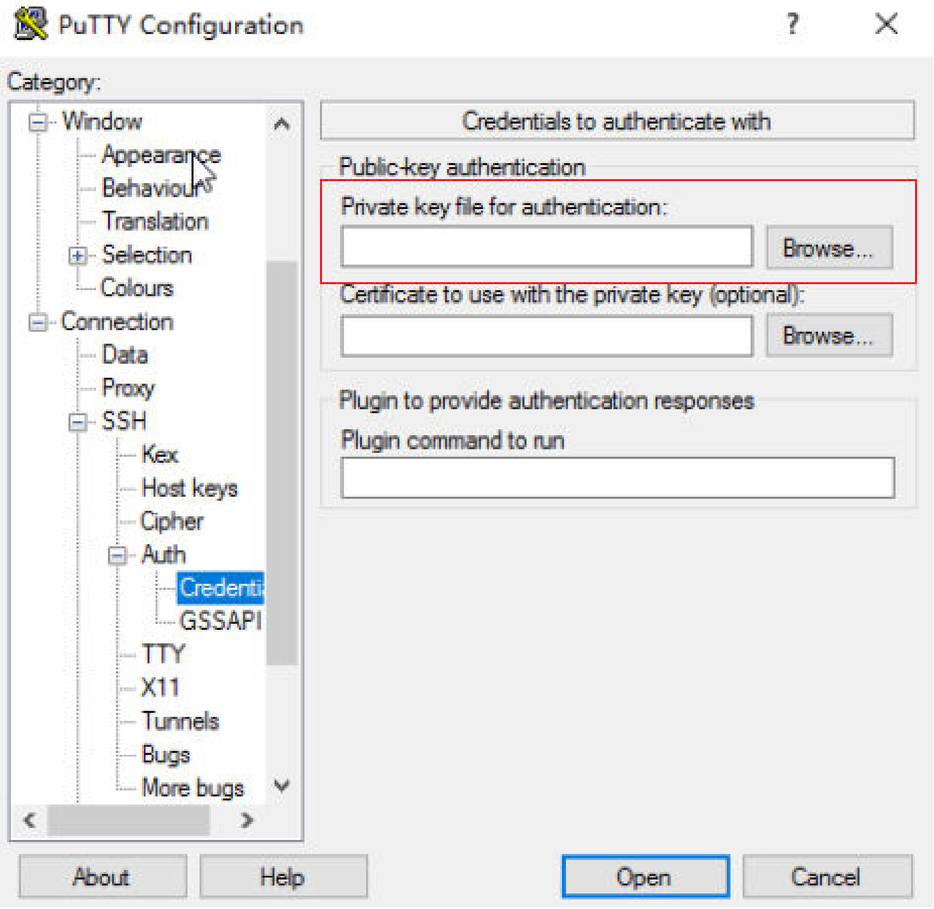

11. Click **Open** to log in to the ECS.

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 2-35 SSH Security Warning

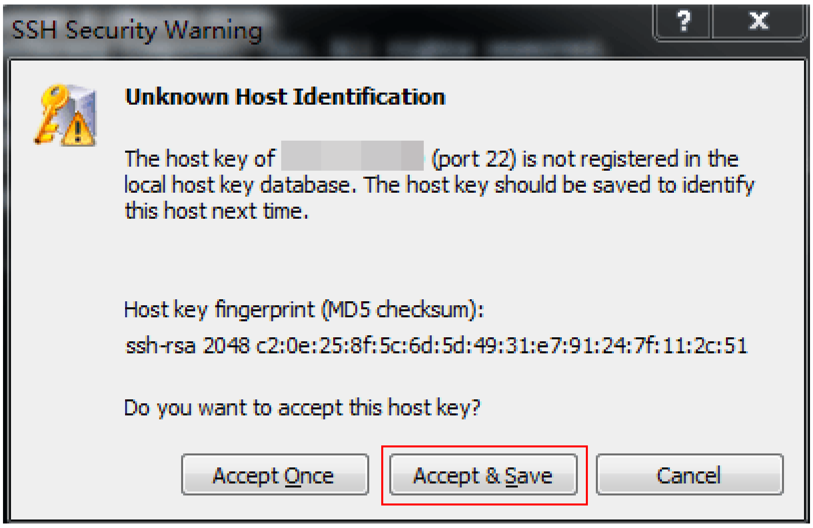

Unknown Host Identification

The host key of [redacted] (port 22) is not registered in the local host key database. The host key should be saved to identify this host next time.

Host key fingerprint (MD5 checksum):
ssh-rsa 2048 c2:0e:25:8f:5c:6d:5d:49:31:e7:91:24:7f:11:2c:51

Do you want to accept this host key?

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?**

2.5.4 Logging In to a Linux ECS 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 which the ECS belongs to. 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 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 2-36 Session

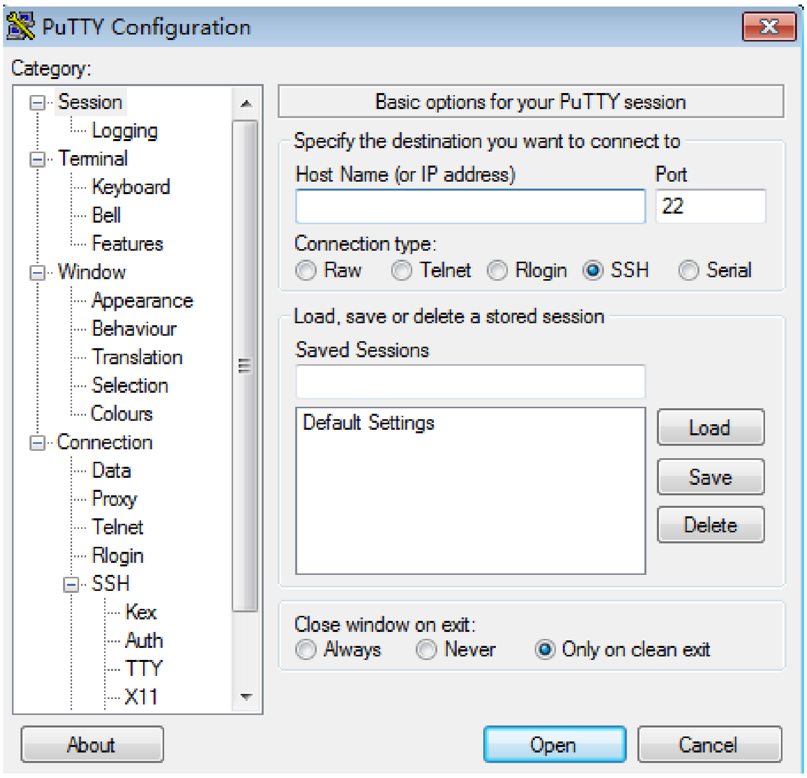

Host Name (or IP address):

Port: 22

Connection type: ○ Raw ○ Telnet ○ Rlogin ● SSH ○ Serial

Close window on exit: ○ Always ○ Never ● Only on clean exit

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

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?**

2.5.5 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.
- If you choose the key pair-based SSH login, the **Login Mode** has been set to **Key pair** during the ECS purchase process and the private key file of the Linux ECS has been obtained.

  The private key file of the Linux ECS is generated during the key pair creation. If the private key file is lost, you can **reset the key pair** to assign a new key pair to the ECS. If you select **I agree to host the private key of the key pair**, you can export the managed private key as required. For details, see **Exporting a Private Key**.
- 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 which the ECS belongs to. 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 (including CoreOS), 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

The private key file of the Linux ECS is generated during the key pair creation. If the private key file is lost, you can **reset the key pair** to assign a new key pair to the ECS. If you select **I agree to host the private key of the key pair**, you can export the managed private key as required. For details, see **Exporting a Private Key**.

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 Linux ECS using VNC.

2.5.6 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 Termius, 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 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 which the ECS belongs to. 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 2-37 New Host

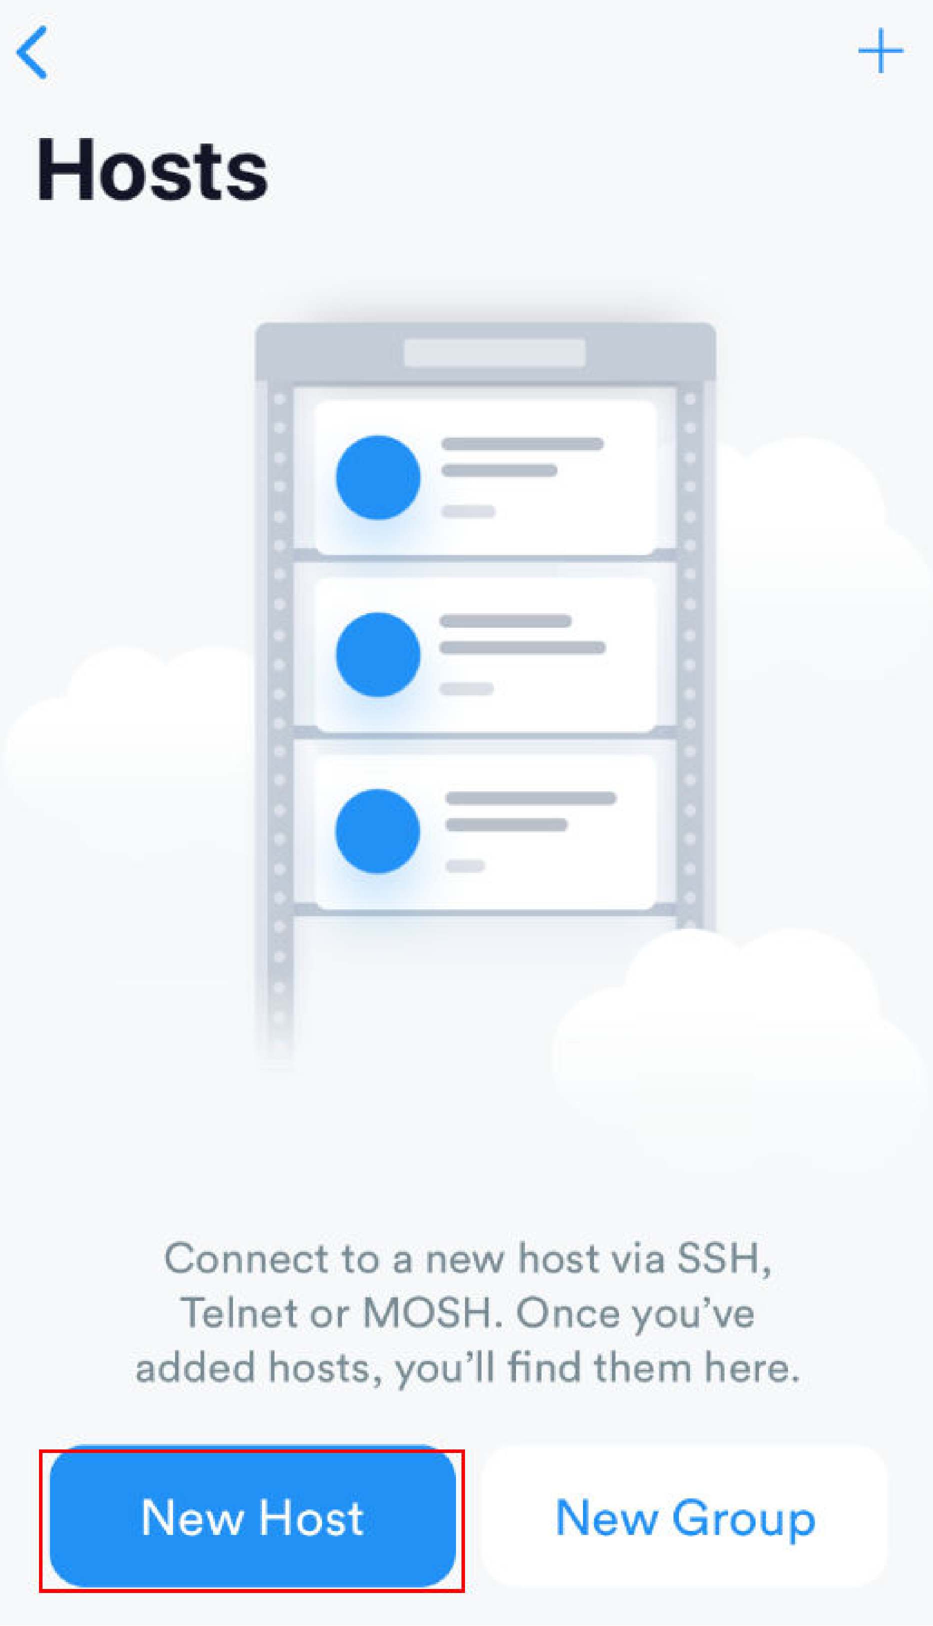
Hosts
Connect to a new host via SSH, Telnet or MOSH. Once you've added hosts, you'll find them here.
New Host
New Group

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 2-38 Setting parameters

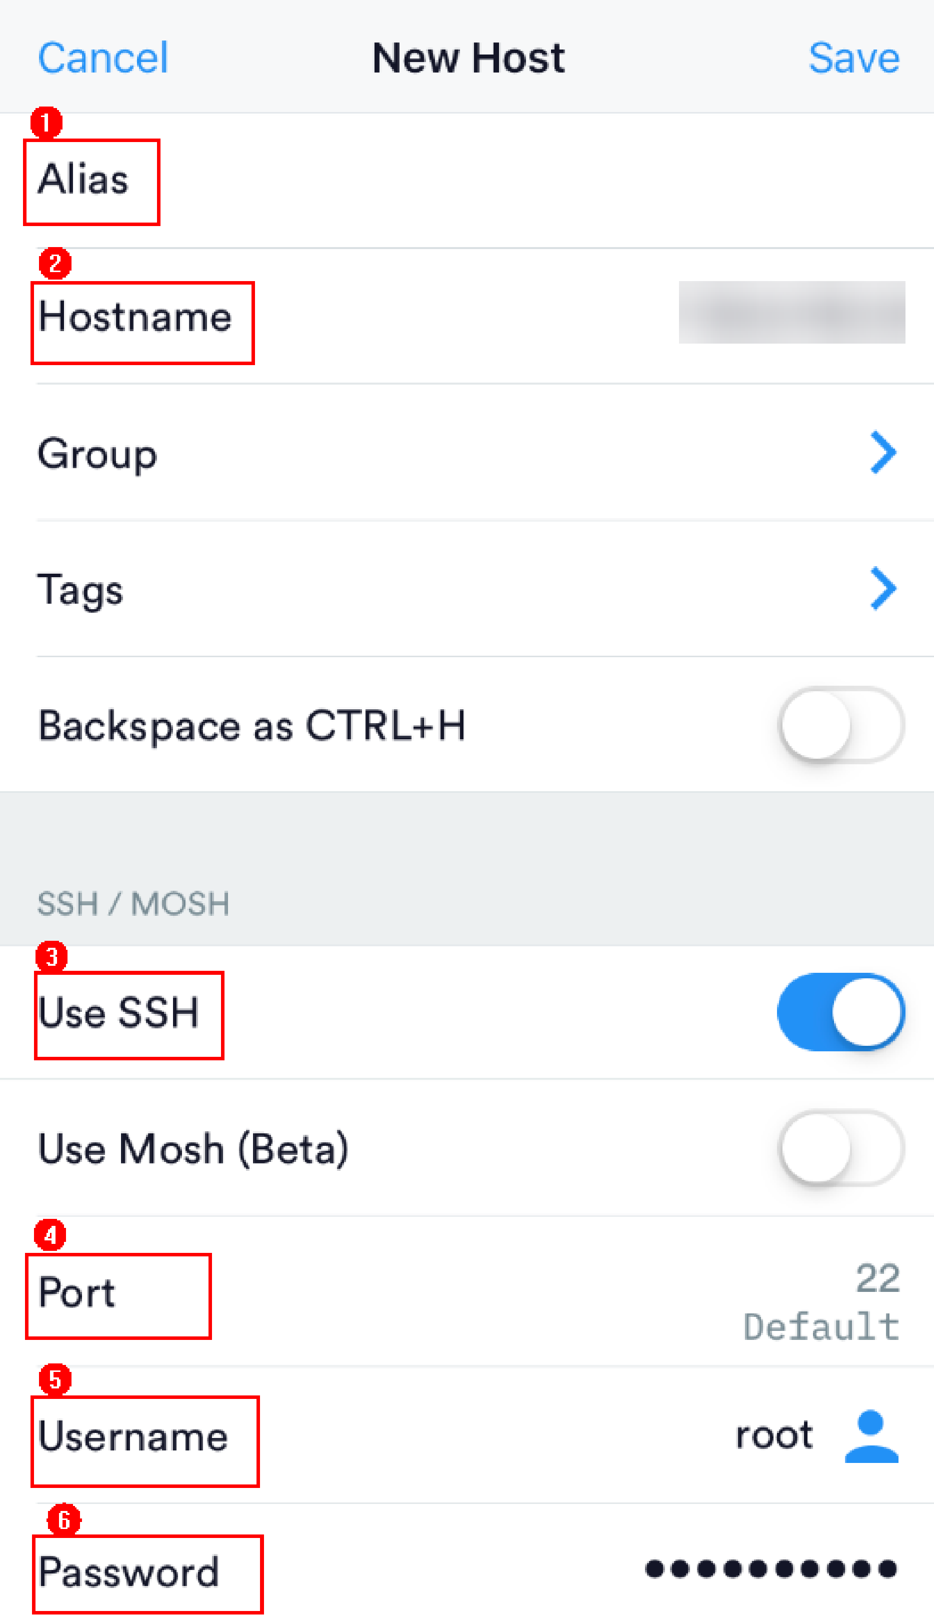

Alias

Hostname [redacted]

Group

Tags

Backspace as CTRL+H ○

SSH / MOSH

Use SSH ●

Use Mosh (Beta) ○

Port 22 Default

Username root

Password ●●●●●●●●●●

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 2-39 Login information

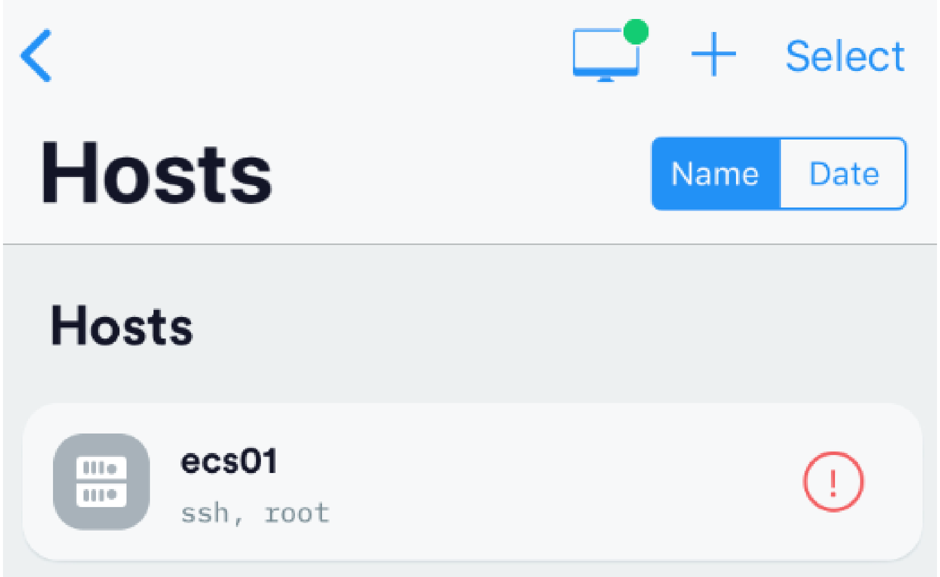

If the following page is displayed, you have connected to the Linux ECS.

Figure 2-40 Connected

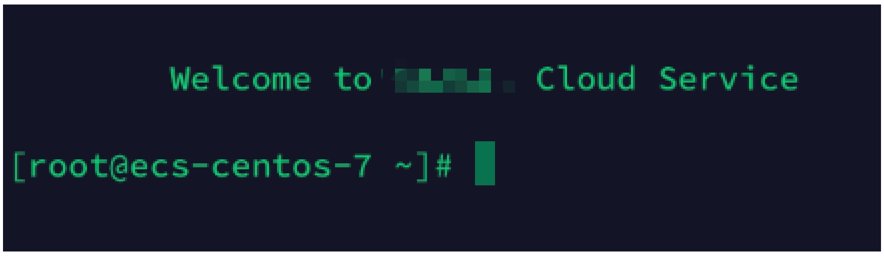

```
        Welcome to [redacted] Cloud Service

[root@ecs-centos-7 ~]# 
```

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 2-41 Starting JuiceSSH

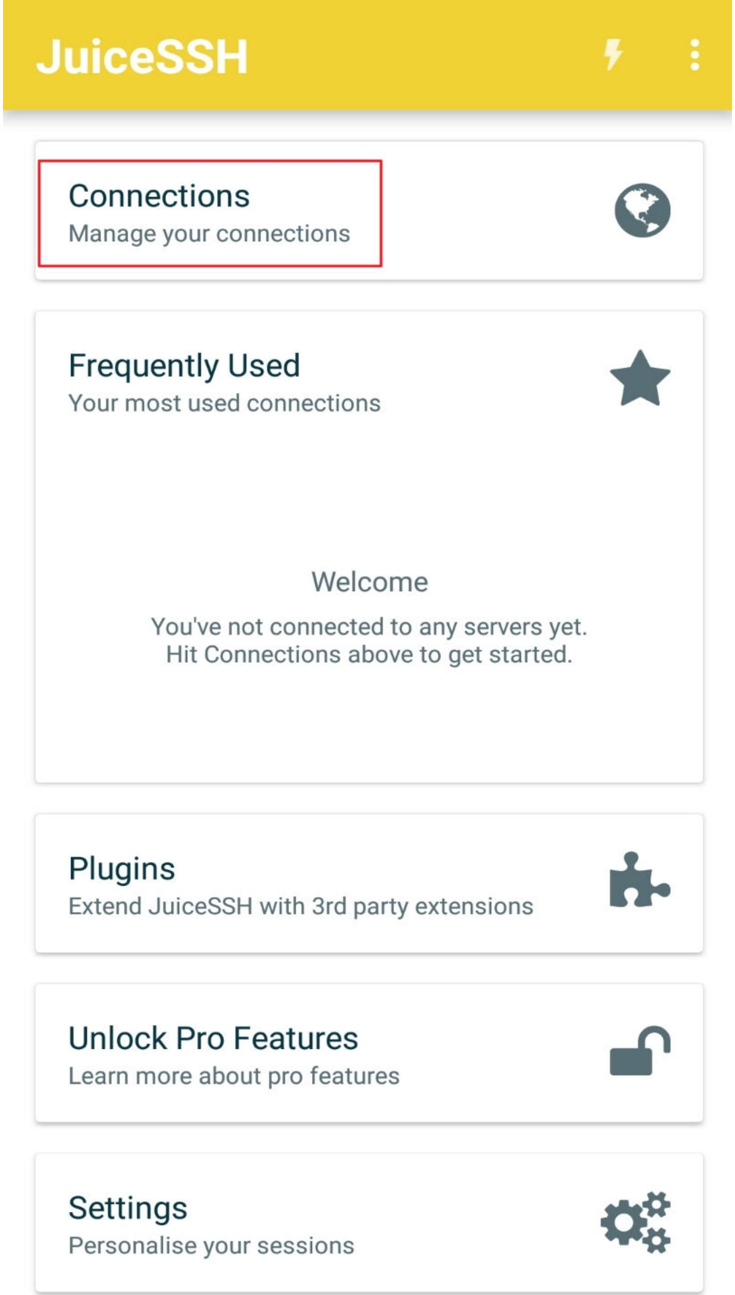

JuiceSSH

Connections
Manage your connections

Frequently Used
Your most used connections

Welcome

You've not connected to any servers yet.
Hit Connections above to get started.

Plugins
Extend JuiceSSH with 3rd party extensions

Unlock Pro Features
Learn more about pro features

Settings
Personalise your sessions

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

Figure 2-42 Connections

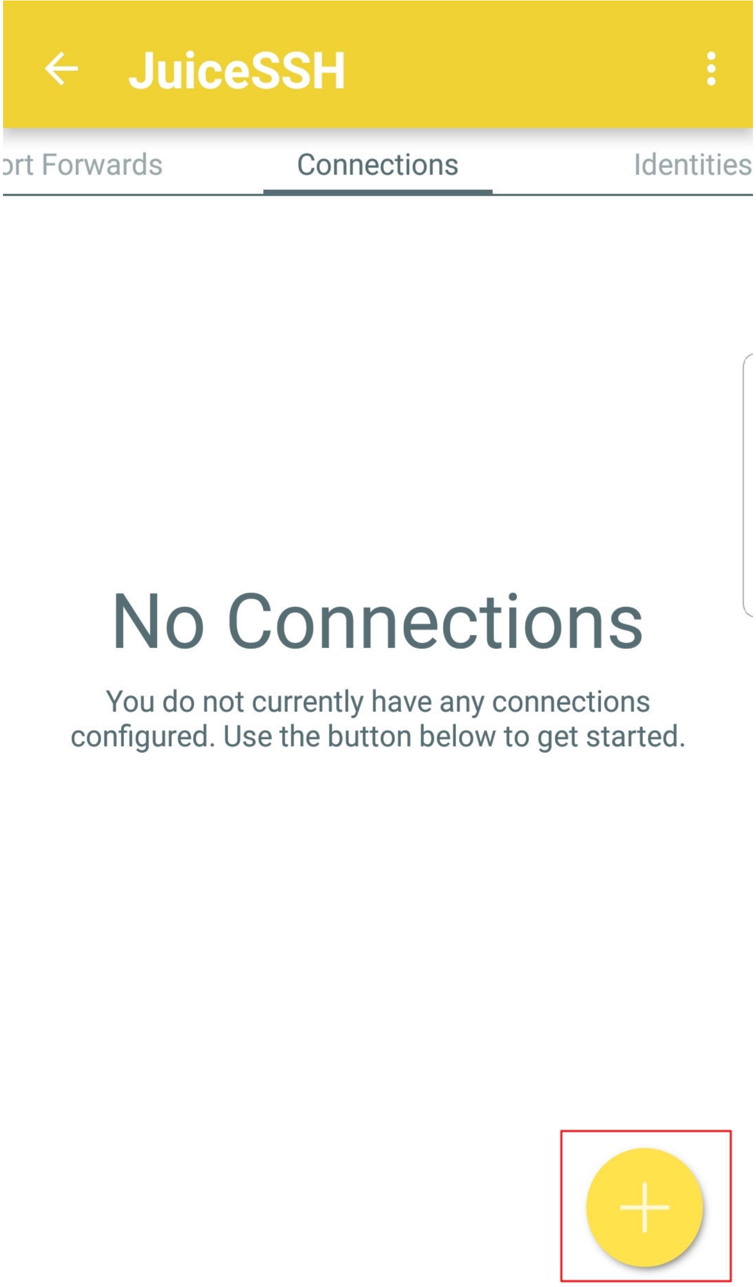

No Connections

You do not currently have any connections configured. Use the button below to get started.

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 2-43 New Identity

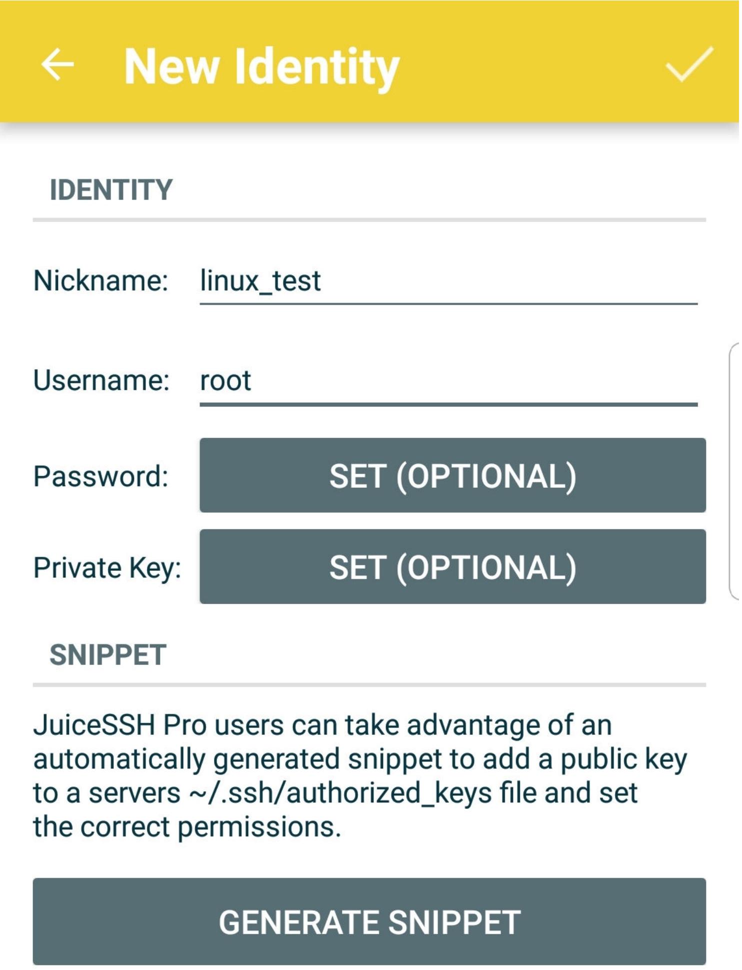

New Identity

IDENTITY

Nickname: linux_test

Username: root

Password: SET (OPTIONAL)

Private Key: SET (OPTIONAL)

SNIPPET

JuiceSSH Pro users can take advantage of an automatically generated snippet to add a public key to a servers ~/.ssh/authorized_keys file and set the correct permissions.

– **Port**: Enter port number **22**.

Figure 2-44 Port

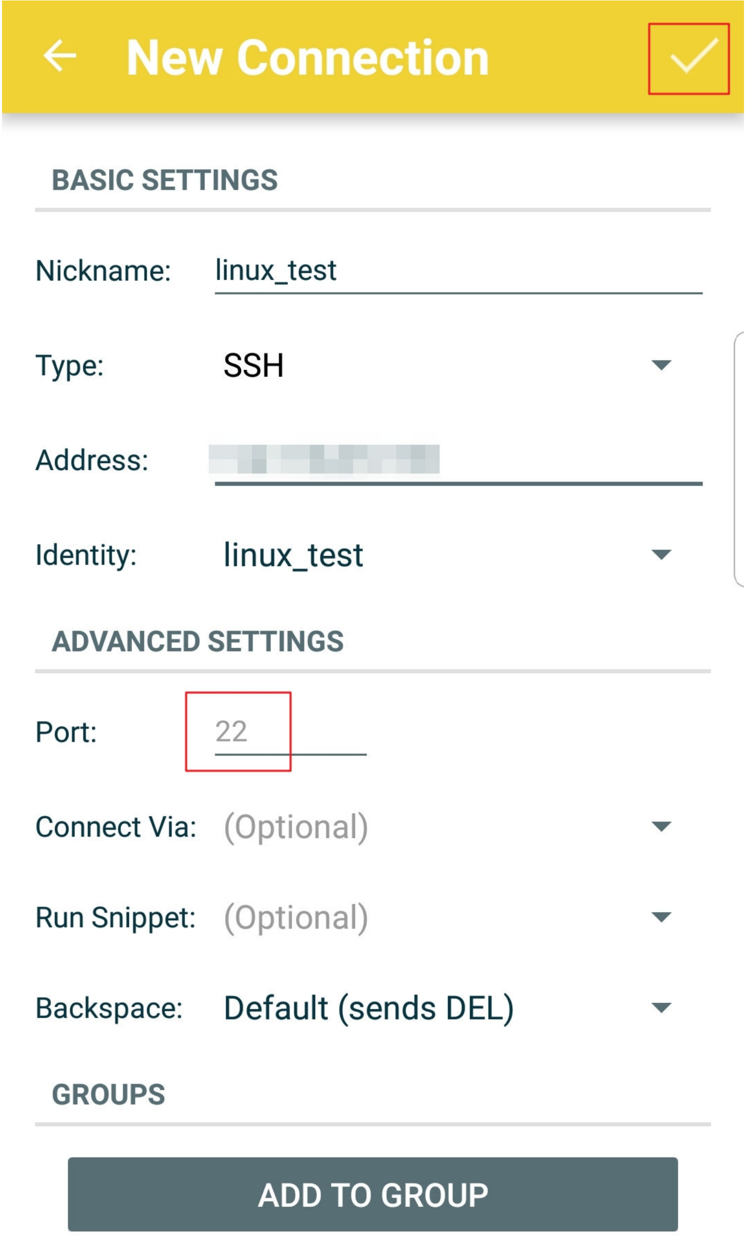

Nickname: linux_test

Type: SSH

Address: [redacted]

Identity: linux_test

Port: 22

Connect Via: (Optional)

Run Snippet: (Optional)

Backspace: Default (sends DEL)

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

Figure 2-45 Connections

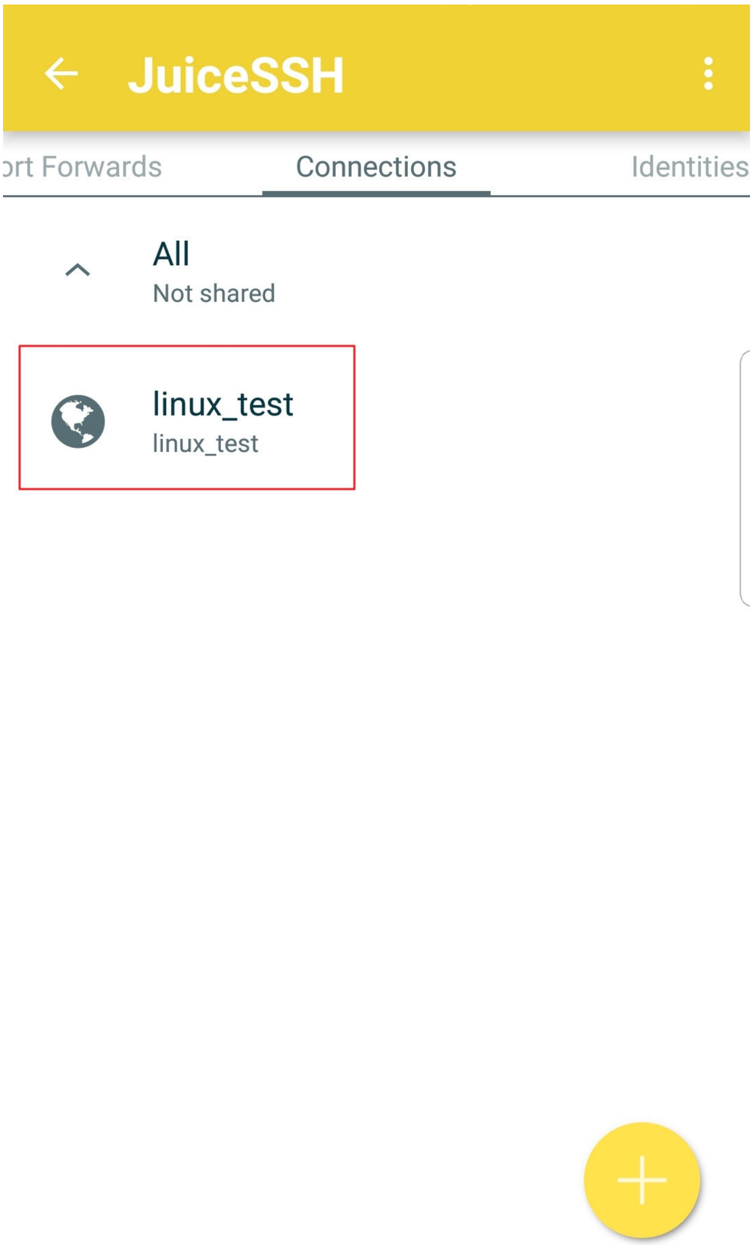

5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 2-46 Confirming the information

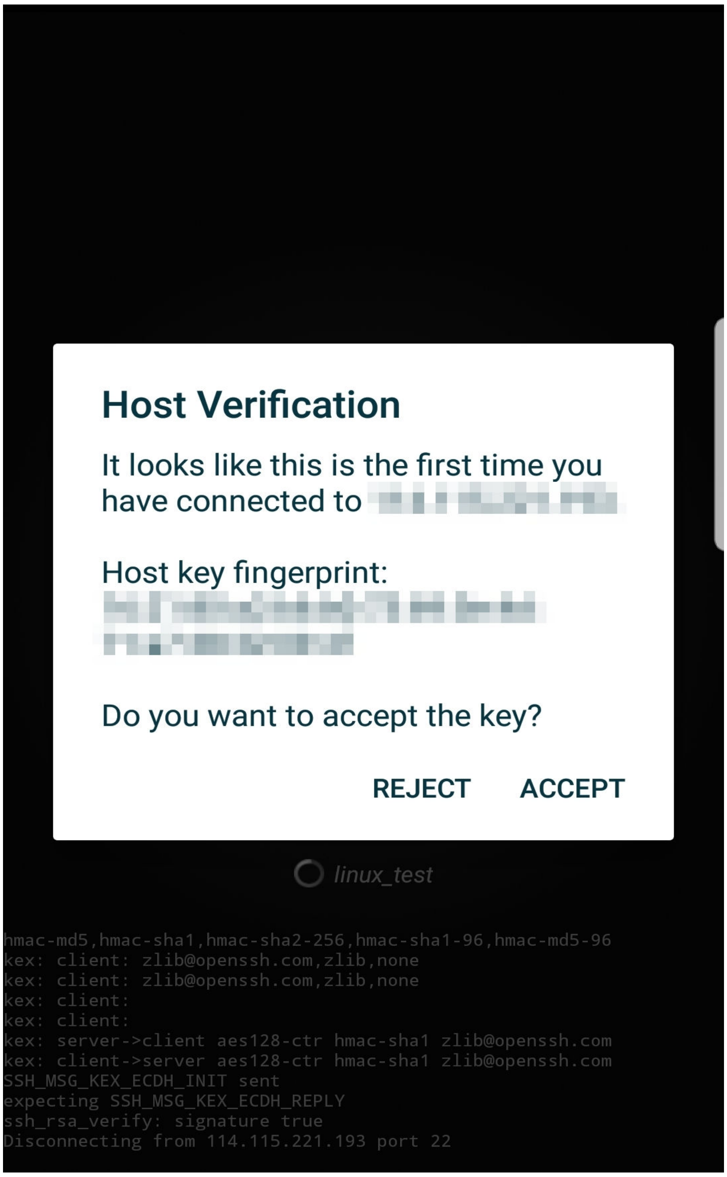

Host Verification

It looks like this is the first time you have connected to [redacted]

Host key fingerprint:
[redacted]

Do you want to accept the key?

```
hmac-md5,hmac-sha1,hmac-sha2-256,hmac-sha1-96,hmac-md5-96
kex: client: zlib@openssh.com,zlib,none
kex: client: zlib@openssh.com,zlib,none
kex: client:
kex: client:
kex: server->client aes128-ctr hmac-sha1 zlib@openssh.com
kex: client->server aes128-ctr hmac-sha1 zlib@openssh.com
SSH_MSG_KEX_ECDH_INIT sent
expecting SSH_MSG_KEX_ECDH_REPLY
ssh_rsa_verify: signature true
Disconnecting from 114.115.221.193 port 22
```

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 2-47 Tutorial

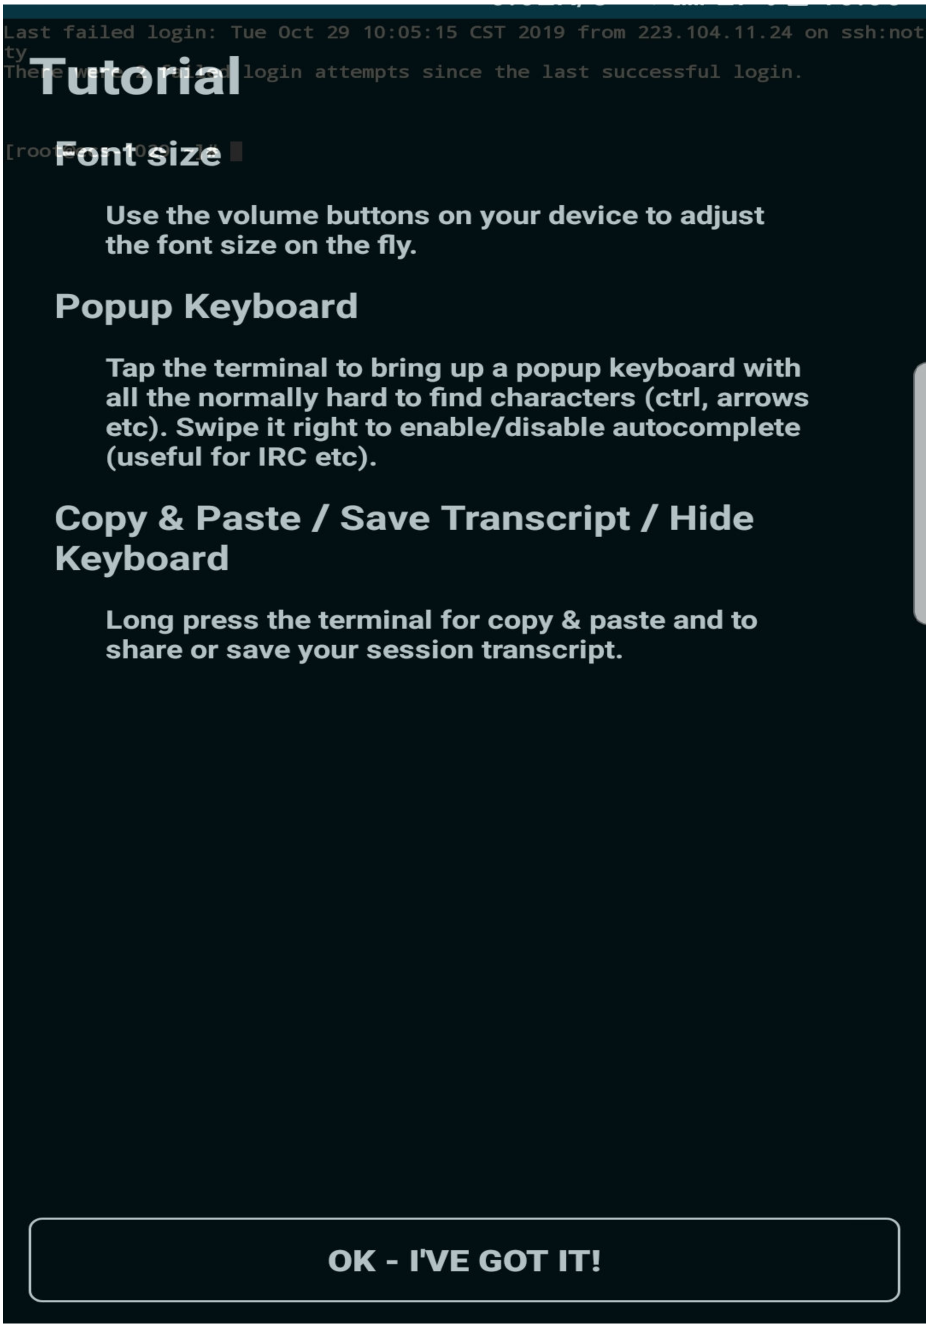

```
Last failed login: Tue Oct 29 10:05:15 CST 2019 from 223.104.11.24 on ssh:not
ty
There were    failed login attempts since the last successful login.

[root@         ~]# 
```

Tutorial

Font size

Use the volume buttons on your device to adjust the font size on the fly.

Popup Keyboard

Tap the terminal to bring up a popup keyboard with all the normally hard to find characters (ctrl, arrows etc). Swipe it right to enable/disable autocomplete (useful for IRC etc).

Copy & Paste / Save Transcript / Hide Keyboard

Long press the terminal for copy & paste and to share or save your session transcript.

You have logged in to the Linux ECS.

Figure 2-48 Successful login

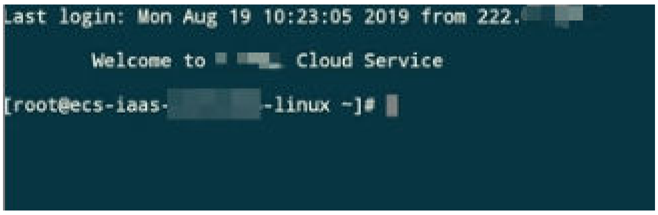

```
Last login: Mon Aug 19 10:23:05 2019 from 222.[redacted]

       Welcome to [redacted] Cloud Service

[root@ecs-iaas-[redacted]-linux ~]# 
```

2.6 Managing GPU Drivers of GPU-accelerated ECSs

2.6.1 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 **Manually 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 **Manually Installing a Tesla Driver on a GPU-accelerated ECS**.

Table 2-30 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.

Driver	License	CUDA	OpenGL	DirectX	Vulkan	Application Scenario	Description
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.

2.6.2 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 **Manually Installing a Tesla Driver on a GPU-accelerated ECS**.

Downloading a Tesla Driver

Download a driver based on your ECS type.

Table 2-31 Mapping between Tesla drivers and ECS types

ECS Type	Driver	Product Series	Product
Pi2	Tesla	T	T4

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

There is a mapping between the driver version and CUDA Toolkit version. If the versions do not match, the driver may be unavailable.

For mapping details, see **Downloading the Official NVIDIA Drivers**.

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 2-49 Selecting the CUDA Toolkit version

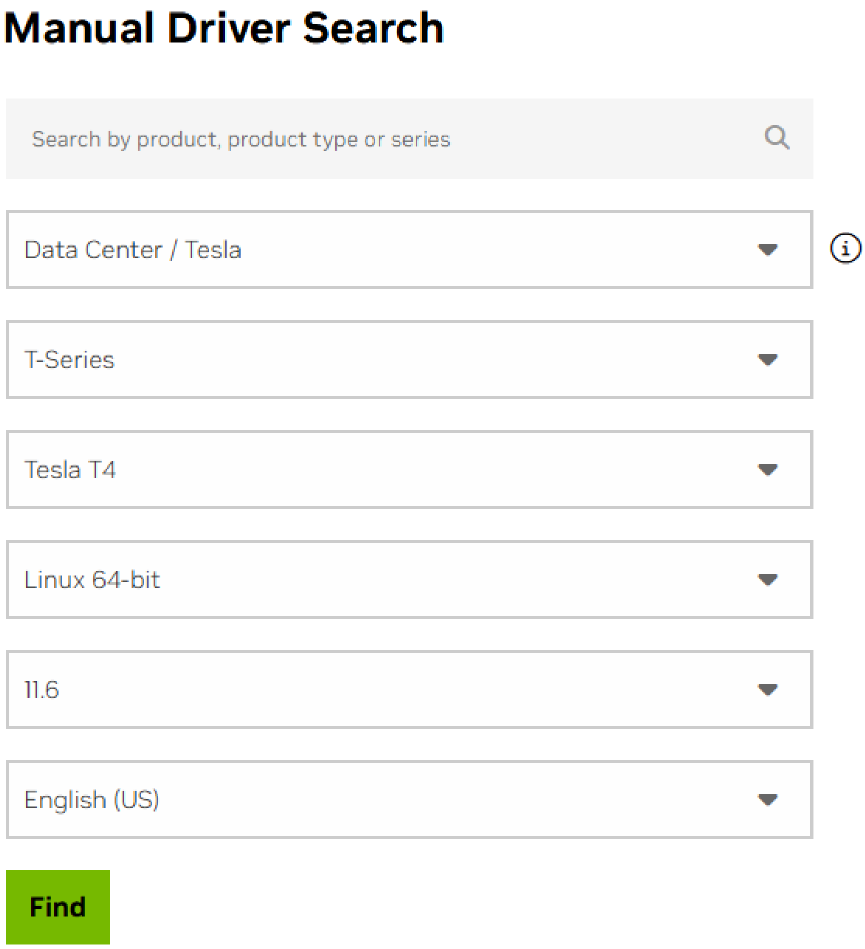

Manual Driver Search

Search by product, product type or series

Data Center / Tesla

T-Series

Tesla T4

Linux 64-bit

11.6

English (US)

2. Select your desired version and download the package.

2.6.3 Manually 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**

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 sign up 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 2-50 Applying for a trial license

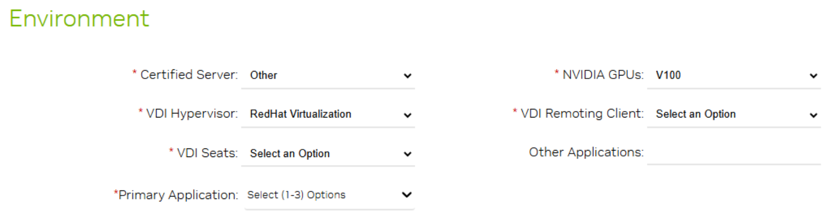

Environment

* Certified Server: Other

* VDI Hypervisor: RedHat Virtualization

* VDI Seats: Select an Option

*Primary Application: Select (1-3) Options

* NVIDIA GPUs: V100

* VDI Remoting Client: Select an Option

Other Applications:

Downloading GRID Driver and Software License Packages

1. Obtain the driver installation package required for an OS. For details, see **Table 2-32**.

   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 2-32 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
PI2	GPU passthrough	● CentOS 8.2 64bit ● CentOS 7.6 64bit ● Ubuntu 20.04 Server 64bit ● Ubuntu 18.04 Server 64bit	Select a version as needed.	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. Refer to **Figure 2-51** to obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 2-51 PAK

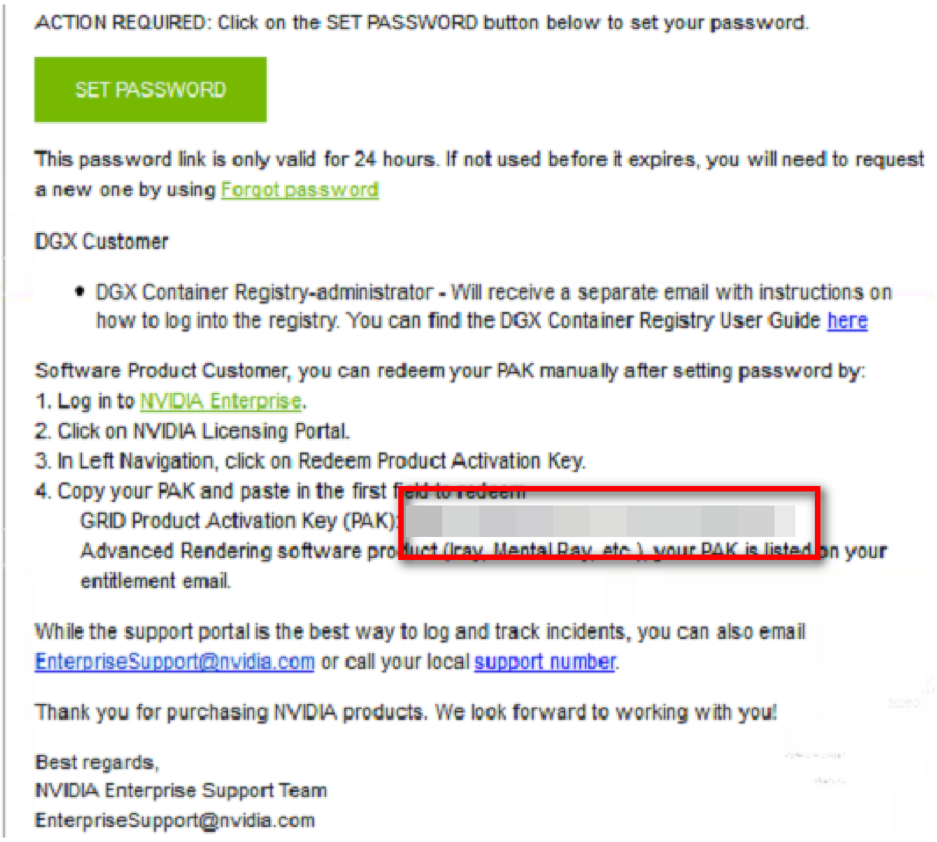

ACTION REQUIRED: Click on the SET PASSWORD button below to set your password.

SET PASSWORD

This password link is only valid for 24 hours. If not used before it expires, you will need to request a new one by using Forgot password

DGX Customer

- DGX Container Registry-administrator - Will receive a separate email with instructions on how to log into the registry. You can find the DGX Container Registry User Guide here

Software Product Customer, you can redeem your PAK manually after setting password by:
1. Log in to NVIDIA Enterprise.
2. Click on NVIDIA Licensing Portal.
3. In Left Navigation, click on Redeem Product Activation Key.
4. Copy your PAK and paste in the first field to redeem
   GRID Product Activation Key (PAK): [redacted]
   Advanced Rendering software product (Iray, Mental Ray, etc.), your PAK is listed on your entitlement email.

While the support portal is the best way to log and track incidents, you can also email EnterpriseSupport@nvidia.com or call your local support number.

Thank you for purchasing NVIDIA products. We look forward to working with you!

Best regards,
NVIDIA Enterprise Support Team
EnterpriseSupport@nvidia.com

5. Enter the PAK obtained in step **4** on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 2-52 Redeem Product Activation Keys

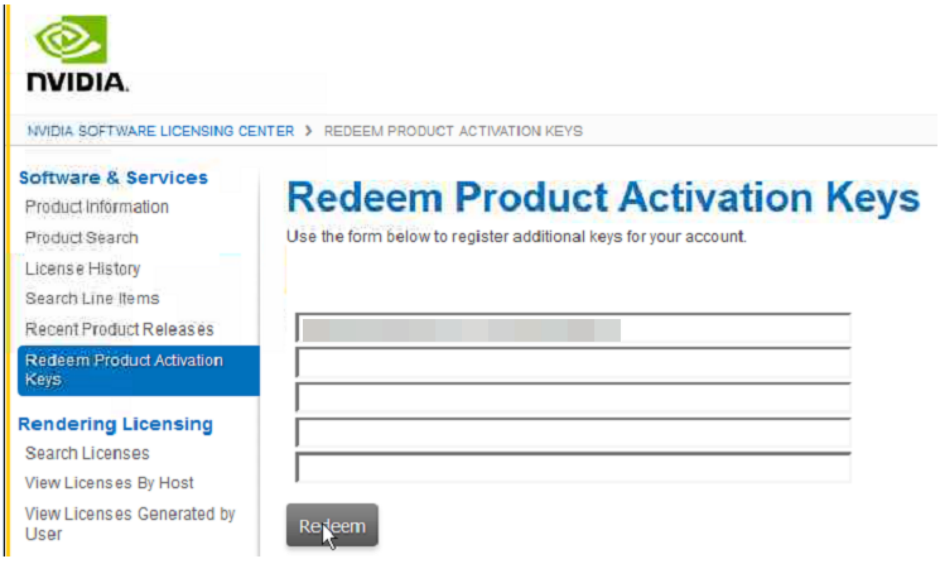

NVIDIA SOFTWARE LICENSING CENTER > REDEEM PRODUCT ACTIVATION KEYS

Redeem Product Activation Keys

Use the form below to register additional keys for your account.

6. Specify **Username** and **Password** and click **LOGIN**.

Figure 2-53 Logging in to the official NVIDIA website

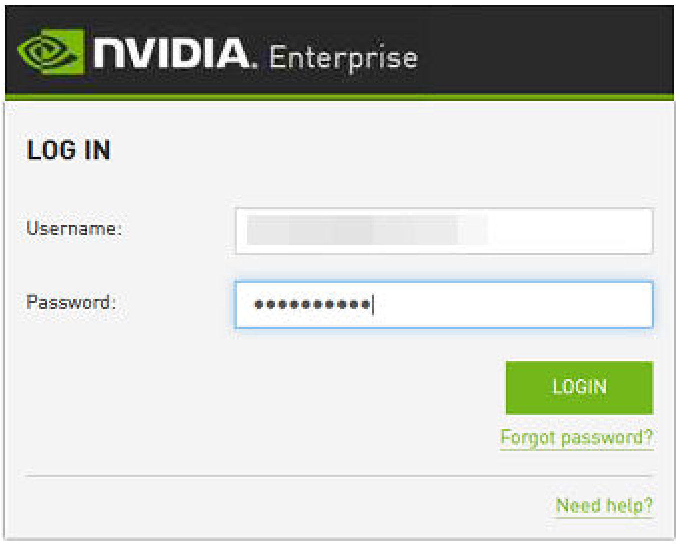

LOG IN

Username:

Password: ●●●●●●●●●●●

7. Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.

Figure 2-54 SOFTWARE DOWNLOADS page

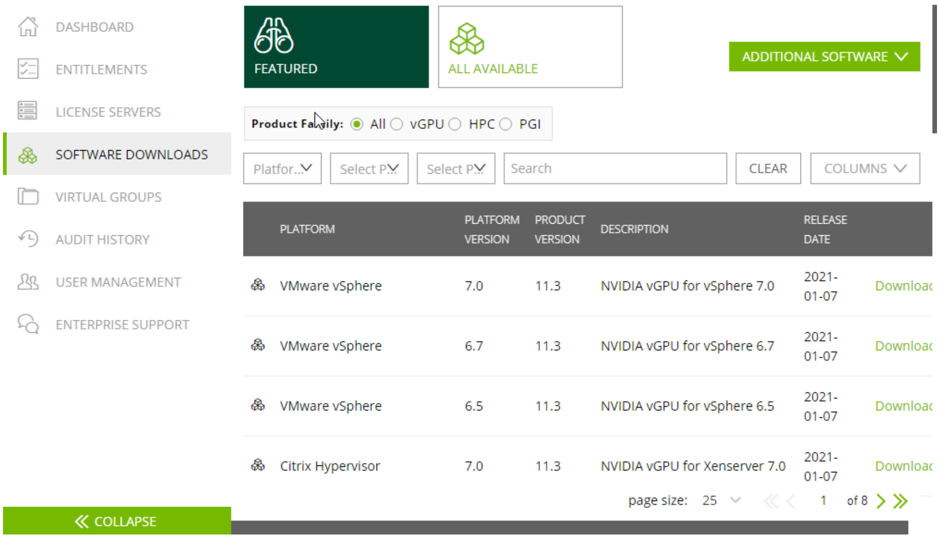

PLATFORM	PLATFORM VERSION	PRODUCT VERSION	DESCRIPTION	RELEASE DATE	
VMware vSphere	7.0	11.3	NVIDIA vGPU for vSphere 7.0	2021-01-07	Downloac
VMware vSphere	6.7	11.3	NVIDIA vGPU for vSphere 6.7	2021-01-07	Downloac
VMware vSphere	6.5	11.3	NVIDIA vGPU for vSphere 6.5	2021-01-07	Downloac
Citrix Hypervisor	7.0	11.3	NVIDIA vGPU for Xenserver 7.0	2021-01-07	Downloac

8. Download the GRID driver of the required version. For details, see **Table 2-32**.
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 2-55 ADDITIONAL SOFTWARE

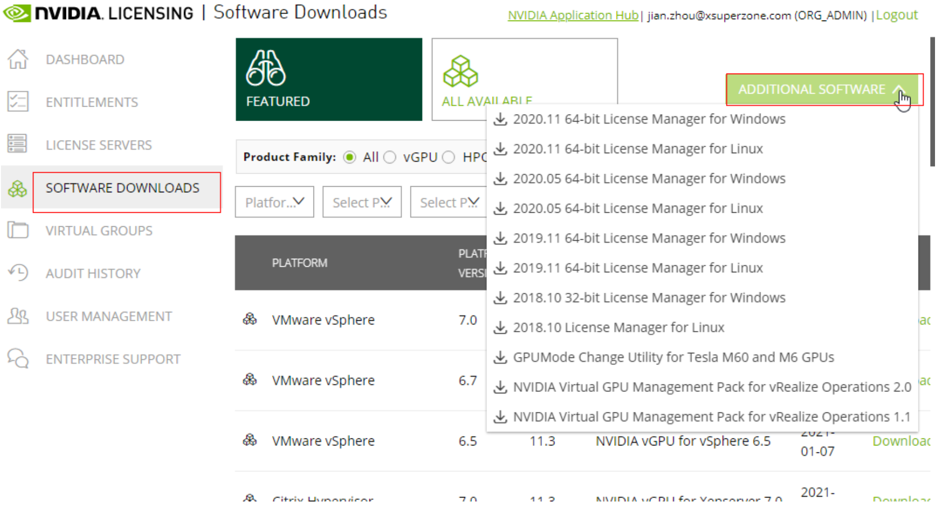
NVIDIA. LICENSING | Software Downloads
NVIDIA Application Hub | jian.zhou@xsuperzone.com (ORG_ADMIN) | Logout
DASHBOARD
ENTITLEMENTS
LICENSE SERVERS
SOFTWARE DOWNLOADS
VIRTUAL GROUPS
AUDIT HISTORY
USER MANAGEMENT
ENTERPRISE SUPPORT
FEATURED
ALL AVAILABLE
ADDITIONAL SOFTWARE
Product Family: ● All ○ vGPU ○ HPC
Platfor...
Select P...
Select P...
2020.11 64-bit License Manager for Windows
2020.11 64-bit License Manager for Linux
2020.05 64-bit License Manager for Windows
2020.05 64-bit License Manager for Linux
2019.11 64-bit License Manager for Windows
2019.11 64-bit License Manager for Linux
2018.10 32-bit License Manager for Windows
2018.10 License Manager for Linux
GPUMode Change Utility for Tesla M60 and M6 GPUs
NVIDIA Virtual GPU Management Pack for vRealize Operations 2.0
NVIDIA Virtual GPU Management Pack for vRealize Operations 1.1
PLATFORM
VMware vSphere 7.0
VMware vSphere 6.7
VMware vSphere 6.5 11.3 NVIDIA vGPU for vSphere 6.5 01-07 Downloac

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 which the license server belongs to 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 step **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.

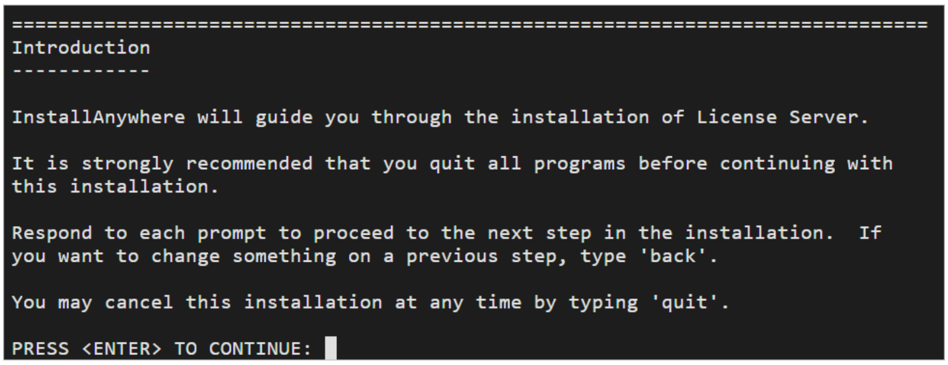

```
===============================================================================
Introduction
------------

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**.

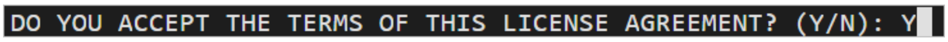

```
DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N): Y
```

   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**.

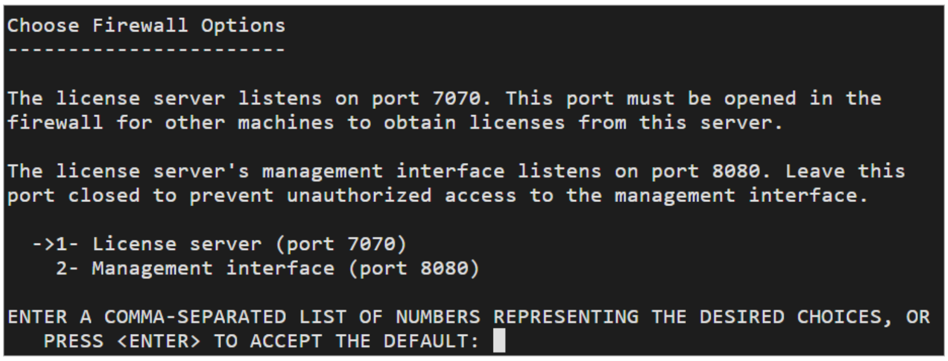

```
Choose Firewall Options
-----------------------

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.

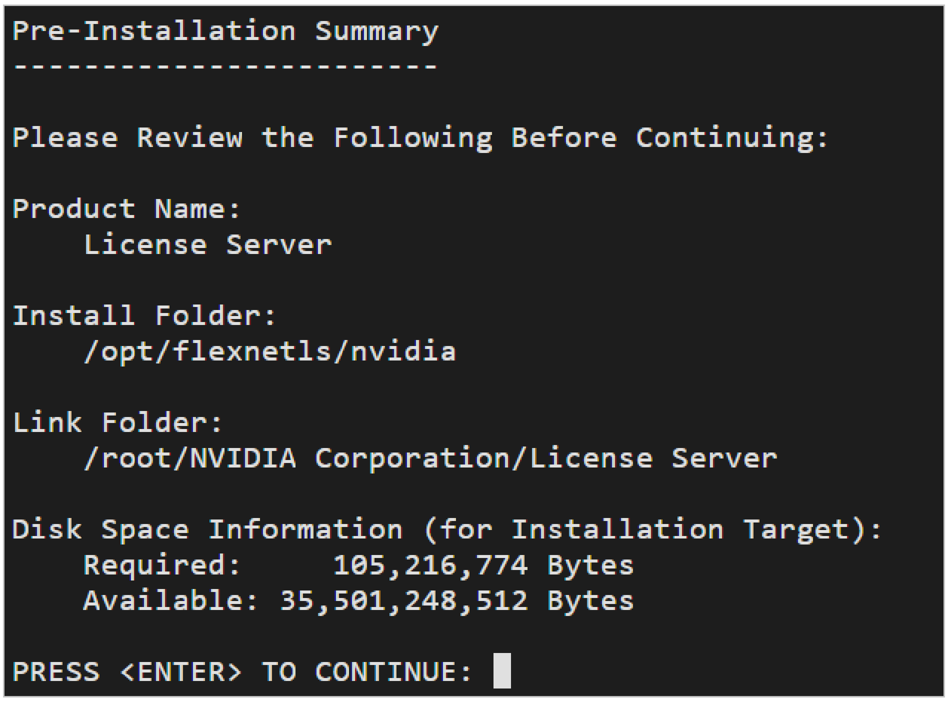

```
Pre-Installation Summary
------------------------

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.

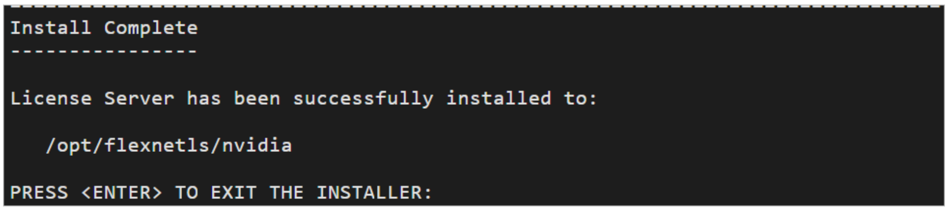

```
Install Complete
----------------

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 2-56 LICENSE SERVERS

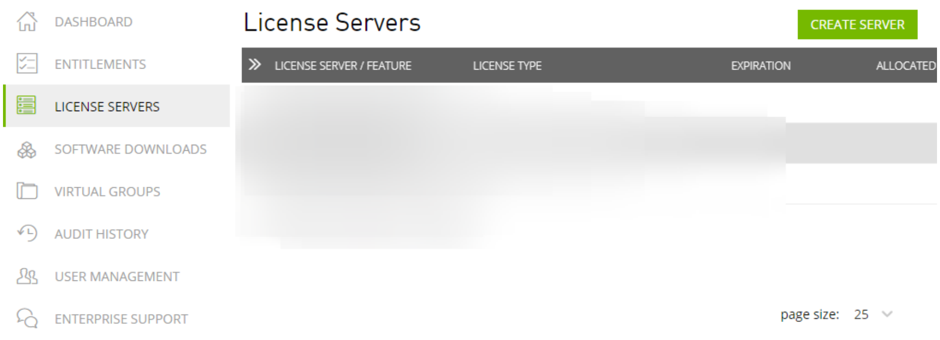

b. Click **CREATE SERVER**.

c. On the displayed **Create License Server** page, configure parameters.

Figure 2-57 Create License Server

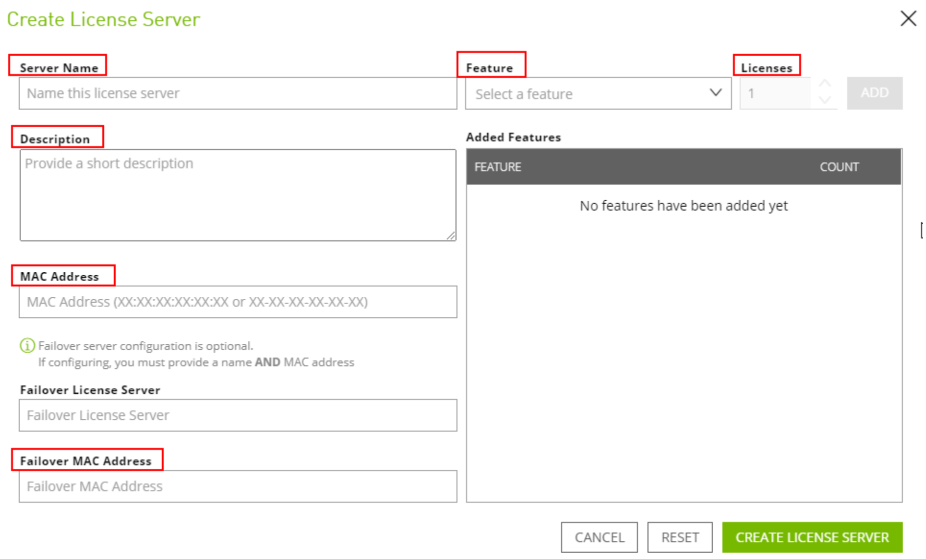

Table 2-33 Parameters for creating a license server

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.

Parameter	Description
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 2-58 Downloading the license file

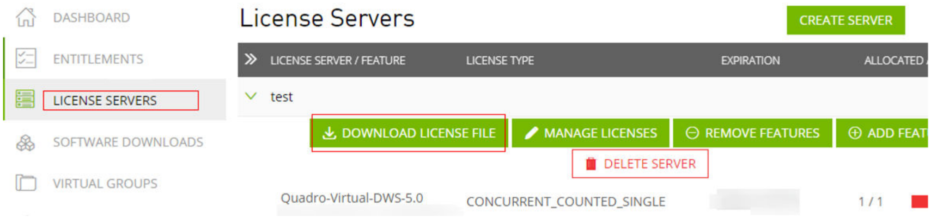

LICENSE SERVER / FEATURE	LICENSE TYPE	EXPIRATION	ALLOCATED
Quadro-Virtual-DWS-5.0	CONCURRENT_COUNTED_SINGLE	[redacted]	1 / 1

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 2-59 Uploading a license file

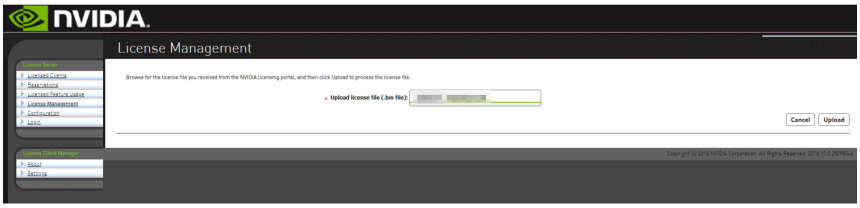

License Management

Browse for the license file you received from the NVIDIA licensing portal, and then click Upload to process the license file.

• Upload license file (.bin file): [redacted]

Copyright (c) 2018 NVIDIA Corporation. All Rights Reserved. 2018.10.0.25098346

2.6.4 Manually Installing a Tesla Driver 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 ECS must have an EIP bound.
- Check whether the CUDA toolkit and Tesla driver have been installed on the ECS.

NOTE

- If the CUDA toolkit has not been installed, download it 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 CUDA Toolkit
  - **Installing the CUDA Toolkit on a Linux ECS**

Installing a Tesla Driver on a Linux ECS

The following uses Ubuntu 20.04 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 at **NVIDIA Driver Downloads** based on the ECS type.

   Figure 2-60 Selecting a NVIDIA driver version

   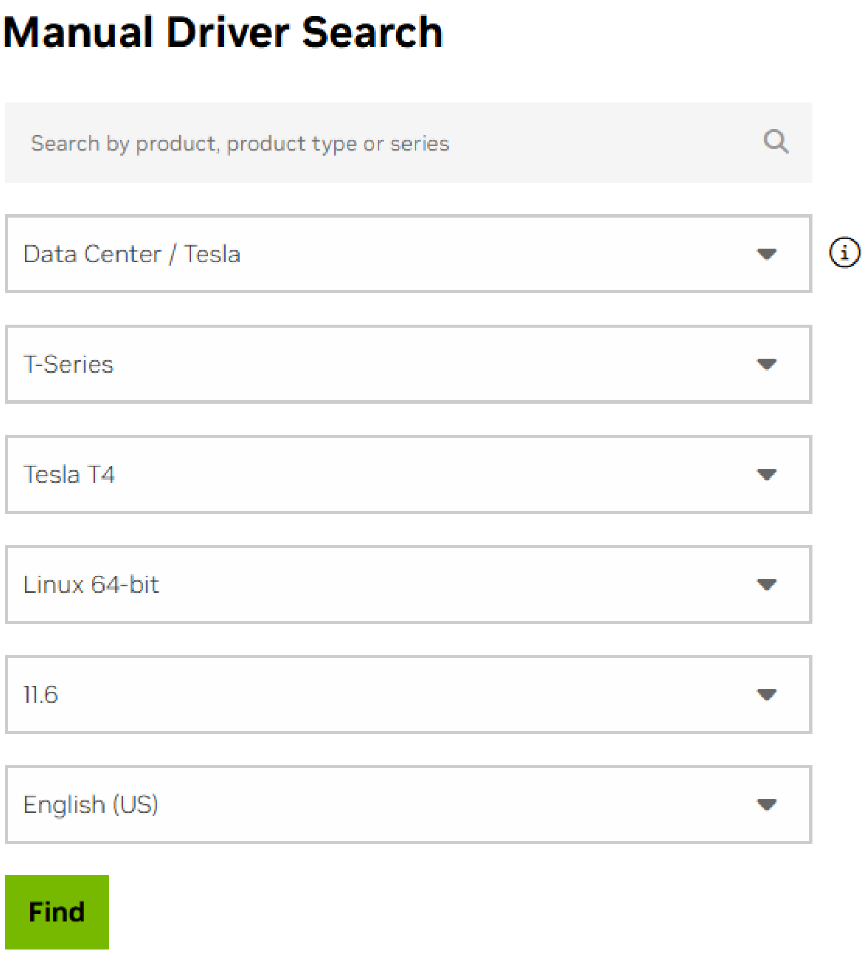

   Manual Driver Search

   Search by product, product type or series

   Data Center / Tesla

   T-Series

   Tesla T4

   Linux 64-bit

   11.6

   English (US)

4. Select a driver version as required. The following uses Tesla 418.67 as an example.

   Figure 2-61 Selecting a driver version

   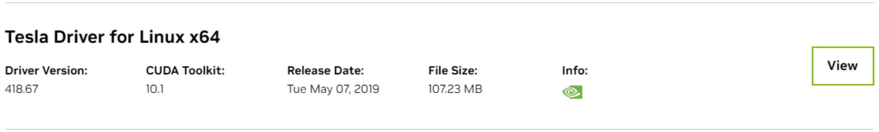

   Tesla Driver for Linux x64

Driver Version:	CUDA Toolkit:	Release Date:	File Size:	Info:
418.67	10.1	Tue May 07, 2019	107.23 MB	

5. Click **View** in the row containing the driver to be downloaded.
6. Right-click **Download** and copy 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 2-62 Obtaining the installation package

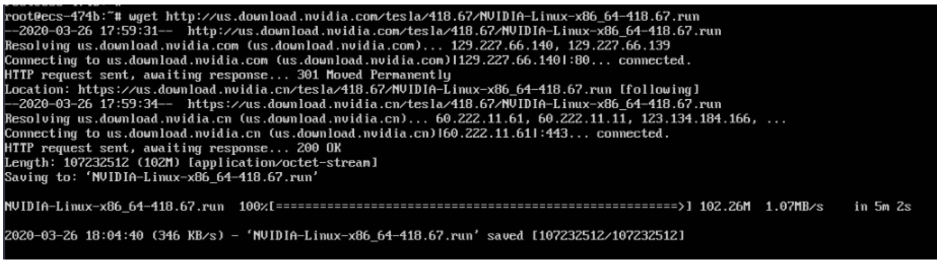

```
root@ecs-474b:~# wget http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
--2020-03-26 17:59:31--  http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
Resolving us.download.nvidia.com (us.download.nvidia.com)... 129.227.66.140, 129.227.66.139
Connecting to us.download.nvidia.com (us.download.nvidia.com)|129.227.66.140|:80... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: https://us.download.nvidia.cn/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run [following]
--2020-03-26 17:59:34--  https://us.download.nvidia.cn/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run
Resolving us.download.nvidia.cn (us.download.nvidia.cn)... 60.222.11.61, 60.222.11.11, 123.134.184.166, ...
Connecting to us.download.nvidia.cn (us.download.nvidia.cn)|60.222.11.61|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 107232512 (102M) [application/octet-stream]
Saving to: 'NVIDIA-Linux-x86_64-418.67.run'

NVIDIA-Linux-x86_64-418.67.run  100%[=====================================================>] 102.26M  1.07MB/s    in 5m 2s

2020-03-26 18:04:40 (346 KB/s) - 'NVIDIA-Linux-x86_64-418.67.run' saved [107232512/107232512]
```

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 2-63 Disabling the Nouveau driver

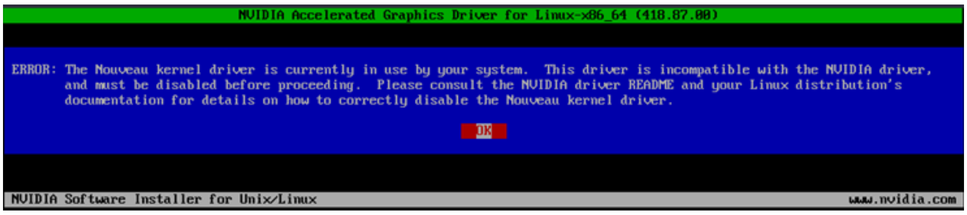

```
NVIDIA Accelerated Graphics Driver for Linux-x86_64 (418.87.00)

ERROR: The Nouveau kernel driver is currently in use by your system.  This driver is incompatible with the NVIDIA driver,
       and must be disabled before proceeding.  Please consult the NVIDIA driver README and your Linux distribution's
       documentation for details on how to correctly disable the Nouveau kernel driver.

                                                    OK

NVIDIA Software Installer for Unix/Linux                                                                  www.nvidia.com
```

   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 2-64 Completing the NVIDIA driver installation

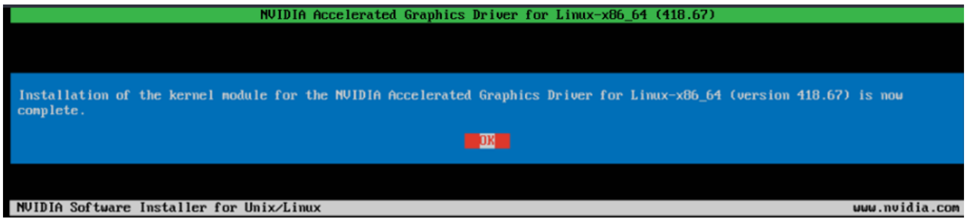

```
NVIDIA Accelerated Graphics Driver for Linux-x86_64 (418.67)

Installation of the kernel module for the NVIDIA Accelerated Graphics Driver for Linux-x86_64 (version 418.67) is now
complete.

NVIDIA Software Installer for Unix/Linux                                                         www.nvidia.com
```

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 2-65 Viewing the NVIDIA driver version

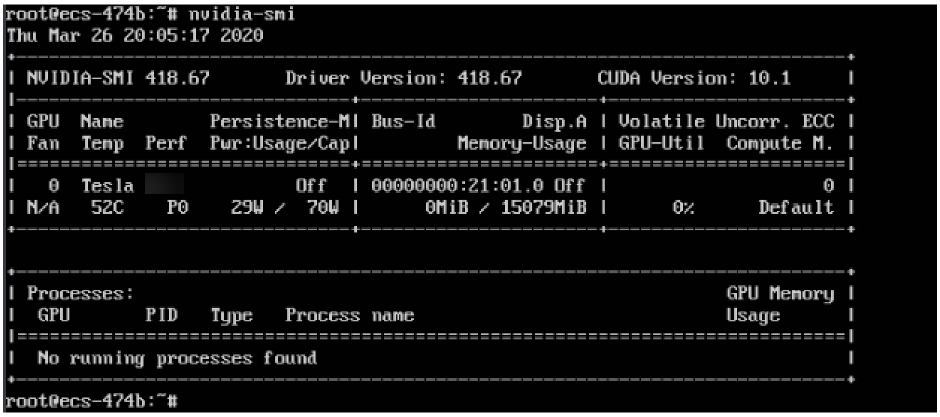

```
root@ecs-474b:~# nvidia-smi
Thu Mar 26 20:05:17 2020
+-----------------------------------------------------------------------------+
| NVIDIA-SMI 418.67       Driver Version: 418.67       CUDA Version: 10.1     |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla [redacted]    Off  | 00000000:21:01.0 Off |                    0 |
| N/A   52C    P0    29W /  70W |      0MiB / 15079MiB |      0%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Processes:                                                       GPU Memory |
|  GPU       PID   Type   Process name                             Usage      |
|=============================================================================|
|  No running processes found                                                 |
+-----------------------------------------------------------------------------+
root@ecs-474b:~#
```

Installing the CUDA Toolkit on a Linux ECS

The following uses Ubuntu 20.04 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 2-66 Selecting a CUDA version

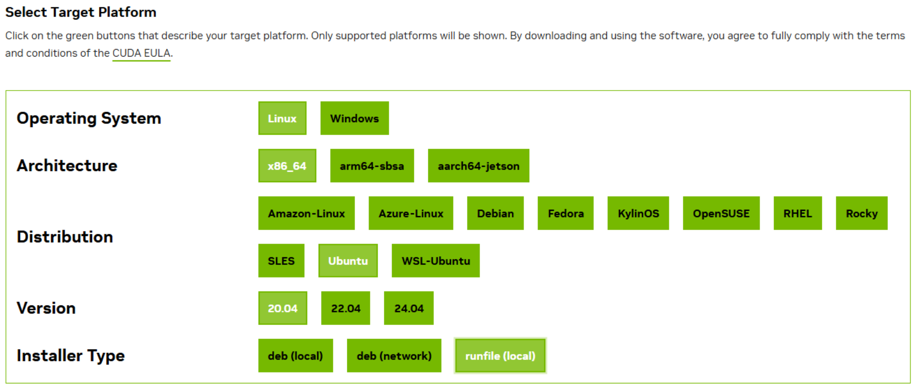

Select Target Platform

Click on the green buttons that describe your target platform. Only supported platforms will be shown. By downloading and using the software, you agree to fully comply with the terms and conditions of the CUDA EULA.

Operating System	Linux Windows
Architecture	x86_64 arm64-sbsa aarch64-jetson
Distribution	Amazon-Linux Azure-Linux Debian Fedora KylinOS OpenSUSE RHEL Rocky SLES Ubuntu WSL-Ubuntu
Version	20.04 22.04 24.04
Installer Type	deb (local) deb (network) runfile (local)

4. Find the link for downloading CUDA 10.1 corresponding to Ubuntu 20.04 64bit and copy the link.

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 2-67 Downloading CUDA

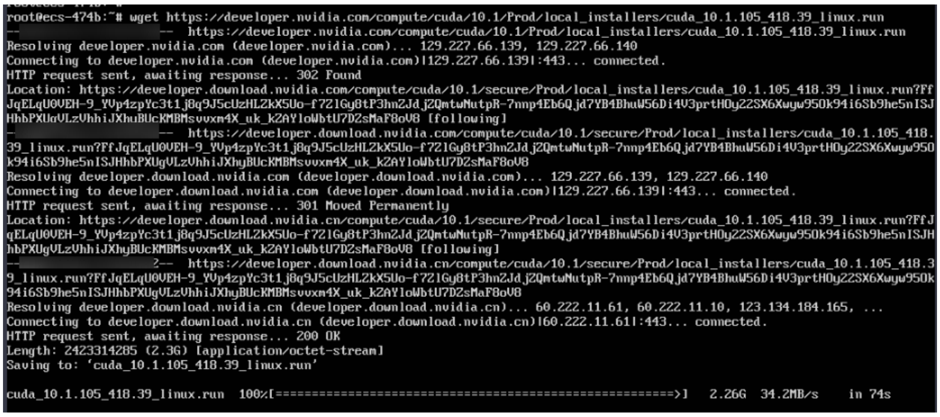

```
root@ecs-474b:~# wget https://developer.nvidia.com/compute/cuda/10.1/Prod/local_installers/cuda_10.1.105_418.39_linux.run
--[redacted]--  https://developer.nvidia.com/compute/cuda/10.1/Prod/local_installers/cuda_10.1.105_418.39_linux.run
Resolving developer.nvidia.com (developer.nvidia.com)... 129.227.66.139, 129.227.66.140
Connecting to developer.nvidia.com (developer.nvidia.com)|129.227.66.139|:443... connected.
HTTP request sent, awaiting response... 302 Found
Location: https://developer.download.nvidia.com/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.39_linux.run?Ff
JqELqU0VEH-9_YVp4zpYc3t1j8q9J5cUzHLZkX5Uo-f7Z1Gy8tP3hnZJdjZQmtwNutpR-7nnp4Eb6Qjd7YB4BhuW56Di4V3prtHOy2ZSX6Xwyw950k94i6Sb9he5nISJ
HhbPXUgVLzVhhiJXhyBUcKMBMsvvxm4X_uk_kZAYloWbtU7DZsMaF8oV8 [following]
--[redacted]--  https://developer.download.nvidia.com/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.
39_linux.run?FfJqELqU0VEH-9_YVp4zpYc3t1j8q9J5cUzHLZkX5Uo-f7Z1Gy8tP3hnZJdjZQmtwNutpR-7nnp4Eb6Qjd7YB4BhuW56Di4V3prtHOy2ZSX6Xwyw950
k94i6Sb9he5nISJHhbPXUgVLzVhhiJXhyBUcKMBMsvvxm4X_uk_kZAYloWbtU7DZsMaF8oV8
Resolving developer.download.nvidia.com (developer.download.nvidia.com)... 129.227.66.139, 129.227.66.140
Connecting to developer.download.nvidia.com (developer.download.nvidia.com)|129.227.66.139|:443... connected.
HTTP request sent, awaiting response... 301 Moved Permanently
Location: https://developer.download.nvidia.cn/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.39_linux.run?FfJ
qELqU0VEH-9_YVp4zpYc3t1j8q9J5cUzHLZkX5Uo-f7Z1Gy8tP3hnZJdjZQmtwNutpR-7nnp4Eb6Qjd7YB4BhuW56Di4V3prtHOy2ZSX6Xwyw950k94i6Sb9he5nISJH
hbPXUgVLzVhhiJXhyBUcKMBMsvvxm4X_uk_kZAYloWbtU7DZsMaF8oV8 [following]
--[redacted]2--  https://developer.download.nvidia.cn/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.3
9_linux.run?FfJqELqU0VEH-9_YVp4zpYc3t1j8q9J5cUzHLZkX5Uo-f7Z1Gy8tP3hnZJdjZQmtwNutpR-7nnp4Eb6Qjd7YB4BhuW56Di4V3prtHOy2ZSX6Xwyw950k
94i6Sb9he5nISJHhbPXUgVLzVhhiJXhyBUcKMBMsvvxm4X_uk_kZAYloWbtU7DZsMaF8oV8
Resolving developer.download.nvidia.cn (developer.download.nvidia.cn)... 60.222.11.61, 60.222.11.10, 123.134.184.165, ...
Connecting to developer.download.nvidia.cn (developer.download.nvidia.cn)|60.222.11.61|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 2423314285 (2.3G) [application/octet-stream]
Saving to: 'cuda_10.1.105_418.39_linux.run'

cuda_10.1.105_418.39_linux.run  100%[=====================================================>]   2.26G  34.2MB/s    in 74s
```

6. Install CUDA.

   Follow the instructions provided on the official NVIDIA website.

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 2-68 Installing CUDA_1

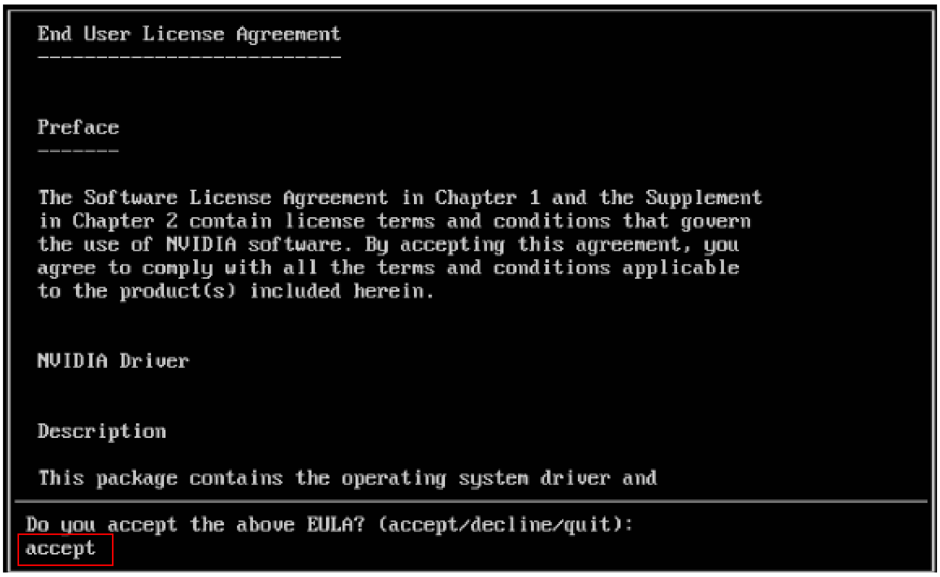

```
  End User License Agreement
  --------------------------


  Preface
  -------

  The Software License Agreement in Chapter 1 and the Supplement
  in Chapter 2 contain license terms and conditions that govern
  the use of NVIDIA software. By accepting this agreement, you
  agree to comply with all the terms and conditions applicable
  to the product(s) included herein.


  NVIDIA Driver


  Description

  This package contains the operating system driver and

 Do you accept the above EULA? (accept/decline/quit):
 accept
```

9. Select **Install** and press **Enter** to start the installation.

Figure 2-69 Installing CUDA_2

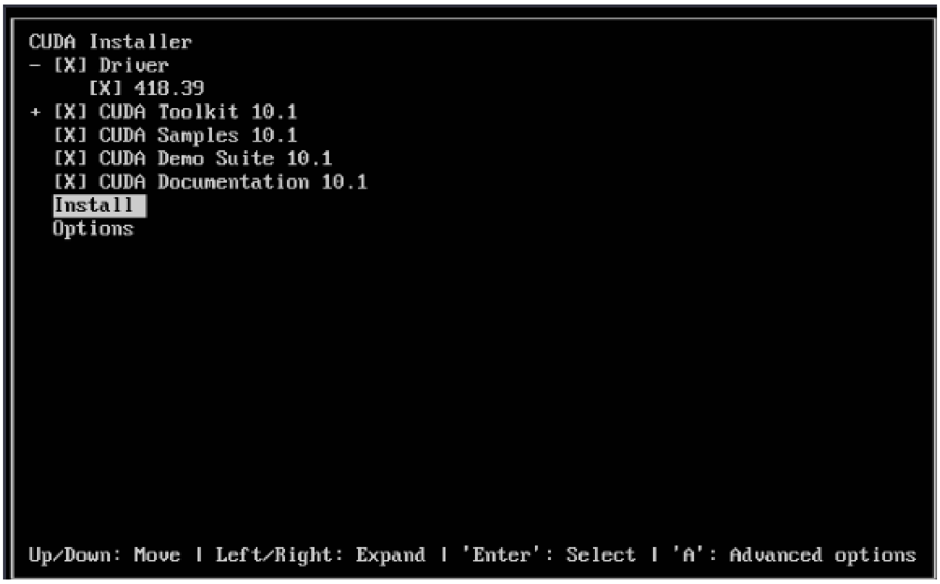

```
CUDA Installer
- [X] Driver
     [X] 418.39
+ [X] CUDA Toolkit 10.1
  [X] CUDA Samples 10.1
  [X] CUDA Demo Suite 10.1
  [X] CUDA Documentation 10.1
  Install
  Options










Up/Down: Move | Left/Right: Expand | 'Enter': Select | 'A': Advanced options
```

Figure 2-70 Completing the installation

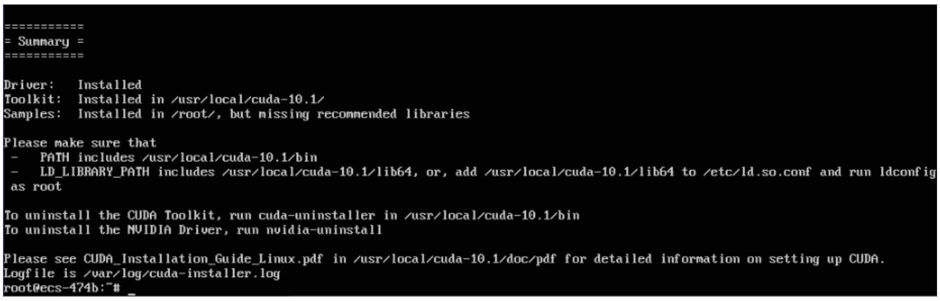

```
===========
= Summary =
===========

Driver:   Installed
Toolkit:  Installed in /usr/local/cuda-10.1/
Samples:  Installed in /root/, but missing recommended libraries

Please make sure that
 -   PATH includes /usr/local/cuda-10.1/bin
 -   LD_LIBRARY_PATH includes /usr/local/cuda-10.1/lib64, or, add /usr/local/cuda-10.1/lib64 to /etc/ld.so.conf and run ldconfig
 as root

To uninstall the CUDA Toolkit, run cuda-uninstaller in /usr/local/cuda-10.1/bin
To uninstall the NVIDIA Driver, run nvidia-uninstall

Please see CUDA_Installation_Guide_Linux.pdf in /usr/local/cuda-10.1/doc/pdf for detailed information on setting up CUDA.
Logfile is /var/log/cuda-installer.log
root@ecs-474b:~# _
```

10. (Optional) Check whether CUDA has been installed.

If the CUDA version is 11.5 or earlier, perform the following operations to check whether CUDA has been installed: If the CUDA version is 11.6 or later, skip this step.

a. 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

b. Run the **make** command to automatically compile the deviceQuery program.

c. 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 2-71 deviceQuery common output

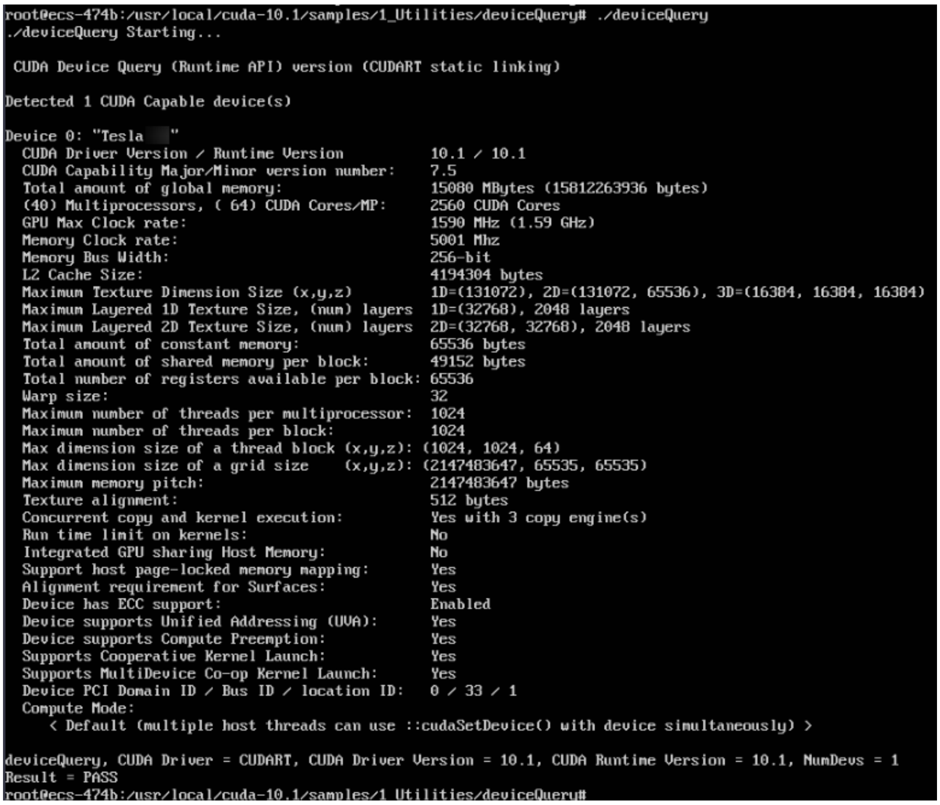

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery# ./deviceQuery
./deviceQuery Starting...

 CUDA Device Query (Runtime API) version (CUDART static linking)

Detected 1 CUDA Capable device(s)

Device 0: "Tesla [redacted]"
  CUDA Driver Version / Runtime Version          10.1 / 10.1
  CUDA Capability Major/Minor version number:    7.5
  Total amount of global memory:                 15080 MBytes (15812263936 bytes)
  (40) Multiprocessors, ( 64) CUDA Cores/MP:     2560 CUDA Cores
  GPU Max Clock rate:                            1590 MHz (1.59 GHz)
  Memory Clock rate:                             5001 Mhz
  Memory Bus Width:                              256-bit
  L2 Cache Size:                                 4194304 bytes
  Maximum Texture Dimension Size (x,y,z)         1D=(131072), 2D=(131072, 65536), 3D=(16384, 16384, 16384)
  Maximum Layered 1D Texture Size, (num) layers  1D=(32768), 2048 layers
  Maximum Layered 2D Texture Size, (num) layers  2D=(32768, 32768), 2048 layers
  Total amount of constant memory:               65536 bytes
  Total amount of shared memory per block:       49152 bytes
  Total number of registers available per block: 65536
  Warp size:                                     32
  Maximum number of threads per multiprocessor:  1024
  Maximum number of threads per block:           1024
  Max dimension size of a thread block (x,y,z): (1024, 1024, 64)
  Max dimension size of a grid size    (x,y,z): (2147483647, 65535, 65535)
  Maximum memory pitch:                          2147483647 bytes
  Texture alignment:                             512 bytes
  Concurrent copy and kernel execution:          Yes with 3 copy engine(s)
  Run time limit on kernels:                     No
  Integrated GPU sharing Host Memory:            No
  Support host page-locked memory mapping:       Yes
  Alignment requirement for Surfaces:            Yes
  Device has ECC support:                        Enabled
  Device supports Unified Addressing (UVA):      Yes
  Device supports Compute Preemption:            Yes
  Supports Cooperative Kernel Launch:            Yes
  Supports MultiDevice Co-op Kernel Launch:      Yes
  Device PCI Domain ID / Bus ID / location ID:   0 / 33 / 1
  Compute Mode:
     < Default (multiple host threads can use ::cudaSetDevice() with device simultaneously) >

deviceQuery, CUDA Driver = CUDART, CUDA Driver Version = 10.1, CUDA Runtime Version = 10.1, NumDevs = 1
Result = PASS
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Utilities/deviceQuery#
```

11. Check the CUDA version.

    /usr/local/cuda/bin/nvcc -V

Figure 2-72 Checking the CUDA version

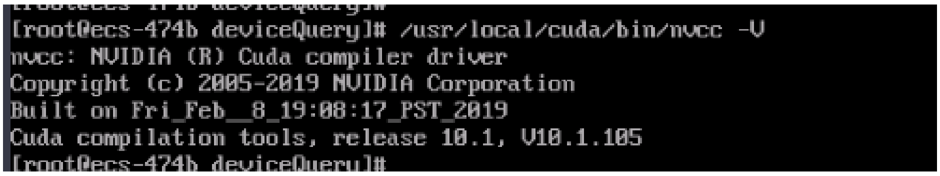

```
[root@ecs-474b deviceQuery]# /usr/local/cuda/bin/nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri_Feb__8_19:08:17_PST_2019
Cuda compilation tools, release 10.1, V10.1.105
[root@ecs-474b deviceQuery]#
```

12. 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.

2.7 Managing ECS Configurations

2.7.1 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.
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.
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
    ```

2.7.2 Enabling or Disabling Hyper-Threading

Scenarios

When purchasing an x86 ECS, you can enable or disable hyper-threading by specifying CPU options. If you do not specify it, hyper-threading is enabled by default.

When you purchase x86 ECSs, you can determine whether to enable hyper-threading based on your service scenarios:

- If you require CPU cores to concurrently process a large amount of data and background tasks, enabling hyper-threading can greatly improve computing performance.
- For compute-intensive or high-performance computing (HPC) applications, such as computational materials science, disabling hyper-threading is a better choice.

You can enable or disable hyper-threading when purchasing x86 ECSs or modifying their specifications.

Background

The processors of x86 ECSs support hyper-threading, which enables two threads to run concurrently on each CPU core. Each thread is represented as a virtual CPU (vCPU). A vCPU is a virtual logical core. After hyper-threading is enabled, a CPU core contains two vCPUs.

A flavor defines the number of vCPUs. You can query the number of vCPUs that an x86 ECS has by referring to **Querying the Number of vCPUs of an ECS**.

Hyper-threading is enabled for most x86 ECSs by default. If hyper-threading is disabled during the x86 ECS creation or specification modification, the number of vCPUs queried from the x86 ECS is half of the number of vCPUs defined by the ECS flavor.

Notes and Constraints

- After an x86 ECS is , you cannot directly change its hyper-threading status. To do so, you can change its flavor to change the hyper-threading status.
- Enabling or disabling hyper-threading is free of charge.
- For details about ECS flavors that support hyper-threading, see **A Summary List of x86 ECS Specifications**.

Enabling or Disabling Hyper-Threading (During ECS Purchase)

1. Log in to the management console and access the **Buy ECS** page.

   Configure basic, network, and advanced settings for ECSs based on service requirements. For details, see **Purchasing an ECS**.
2. Select **Configure now** to configure advanced options.
3. Select **Specify CPU options**.

   Figure 2-73 Specifying CPU options

   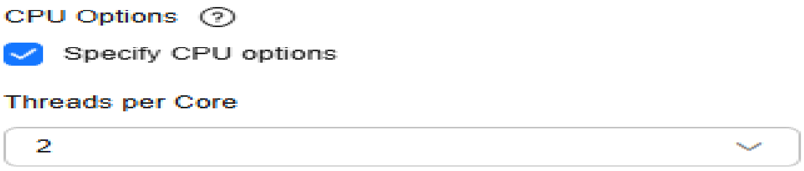

   CPU Options

   ☑ Specify CPU options

   Threads per Core

   2
4. Set **Threads per Core**.

   This parameter is displayed when **Specify CPU options** is selected. You can select a parameter value from the drop-down list.
   - **1**: one thread per core, which means hyper-threading is disabled.
   - **2** (default value): two threads per core, which means hyper-threading is enabled.
5. Click **Next: Confirm** to confirm the settings and complete the ECS purchase.

   After purchasing an ECS, you can query the hyper-threading status by referring to the **cpu_options** parameter in **Querying Details About an ECS**.

Enabling or Disabling Hyper-Threading (During ECS Flavor Change)

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, locate the row containing the target ECS and choose **More** > **Modify Specifications** in the **Operation** column.

   The **Modify ECS Specifications** page is displayed.
5. Select a new ECS type and flavor.

Figure 2-74 Modifying ECS specifications

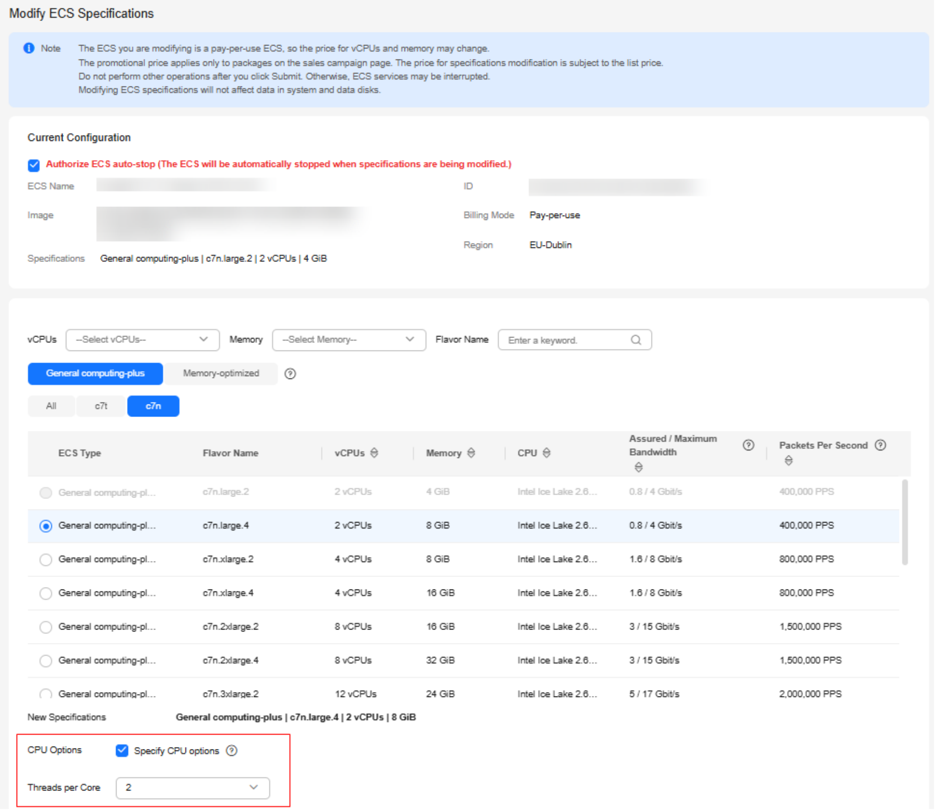

Modify ECS Specifications

Note The ECS you are modifying is a pay-per-use ECS, so the price for vCPUs and memory may change.
The promotional price applies only to packages on the sales campaign page. The price for specifications modification is subject to the list price.
Do not perform other operations after you click Submit. Otherwise, ECS services may be interrupted.
Modifying ECS specifications will not affect data in system and data disks.

Current Configuration

☑ Authorize ECS auto-stop (The ECS will be automatically stopped when specifications are being modified.)

ECS Name	[redacted]	ID	[redacted]
Image	[redacted]	Billing Mode	Pay-per-use
		Region	EU-Dublin
Specifications	General computing-plus | c7n.large.2 | 2 vCPUs | 4 GiB		

vCPUs --Select vCPUs-- Memory --Select Memory-- Flavor Name Enter a keyword.

ECS Type	Flavor Name	vCPUs	Memory	CPU	Assured / Maximum Bandwidth	Packets Per Second
○ General computing-pl...	c7n.large.2	2 vCPUs	4 GiB	Intel Ice Lake 2.6...	0.8 / 4 Gbit/s	400,000 PPS
● General computing-pl...	c7n.large.4	2 vCPUs	8 GiB	Intel Ice Lake 2.6...	0.8 / 4 Gbit/s	400,000 PPS
○ General computing-pl...	c7n.xlarge.2	4 vCPUs	8 GiB	Intel Ice Lake 2.6...	1.6 / 8 Gbit/s	800,000 PPS
○ General computing-pl...	c7n.xlarge.4	4 vCPUs	16 GiB	Intel Ice Lake 2.6...	1.6 / 8 Gbit/s	800,000 PPS
○ General computing-pl...	c7n.2xlarge.2	8 vCPUs	16 GiB	Intel Ice Lake 2.6...	3 / 15 Gbit/s	1,500,000 PPS
○ General computing-pl...	c7n.2xlarge.4	8 vCPUs	32 GiB	Intel Ice Lake 2.6...	3 / 15 Gbit/s	1,500,000 PPS
○ General computing-pl...	c7n.3xlarge.2	12 vCPUs	24 GiB	Intel Ice Lake 2.6...	5 / 17 Gbit/s	2,000,000 PPS

New Specifications General computing-plus | c7n.large.4 | 2 vCPUs | 8 GiB

CPU Options ☑ Specify CPU options

Threads per Core 2

6. Click **Next**.
7. Confirm the settings, read and select the agreement, and then click **Submit Application**.

   After modifying ECS specifications, you can query the hyper-threading status by referring to the **cpu_options** parameter in **Querying Details About an ECS**.

Querying the Number of vCPUs of an ECS

You can log in to an ECS and view the number of its vCPUs.

- For Linux ECSs:
  a. **Log in to a Linux ECS.**
  b. Run the following command to view the number of logical cores of the ECS:

     lscpu

     As shown in **Figure 2-75**, **CPU(s)** indicates the number of logical cores.

Figure 2-75 Viewing the result

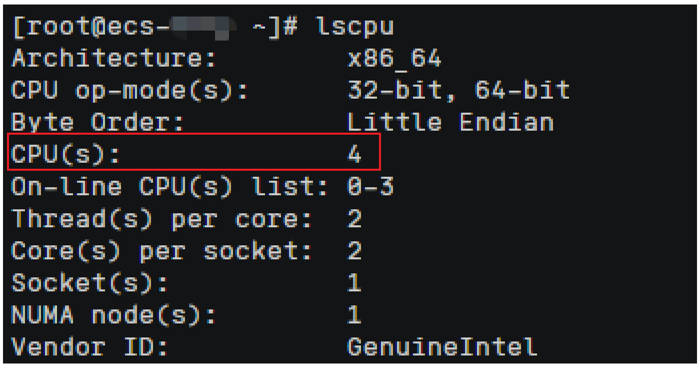

```
[root@ecs-[redacted] ~]# lscpu
Architecture:          x86_64
CPU op-mode(s):        32-bit, 64-bit
Byte Order:            Little Endian
CPU(s):                4
On-line CPU(s) list: 0-3
Thread(s) per core:    2
Core(s) per socket:    2
Socket(s):             1
NUMA node(s):          1
Vendor ID:             GenuineIntel
```

- For Windows ECSs:
  a. **Log in to a Windows ECS.**
  b. Choose **Control Panel** > **Device Manager** and expand **Processors** to view the number of logical cores (threads) of the ECS.

Figure 2-76 Viewing the result

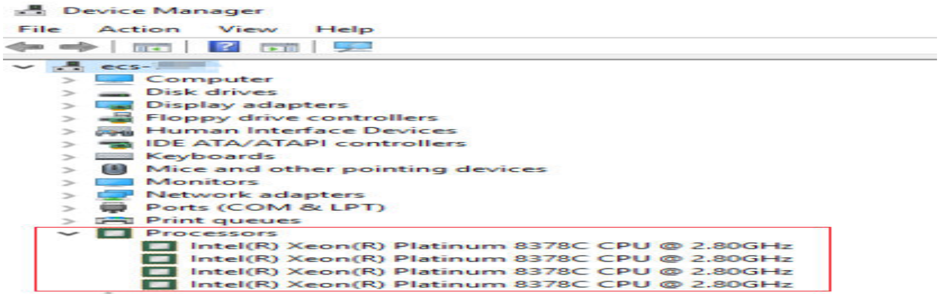

```
ecs-[redacted]
  > Computer
  > Disk drives
  > Display adapters
  > Floppy drive controllers
  > Human Interface Devices
  > IDE ATA/ATAPI controllers
  > Keyboards
  > Mice and other pointing devices
  > Monitors
  > Network adapters
  > Ports (COM & LPT)
  > Print queues
  v Processors
      Intel(R) Xeon(R) Platinum 8378C CPU @ 2.80GHz
      Intel(R) Xeon(R) Platinum 8378C CPU @ 2.80GHz
      Intel(R) Xeon(R) Platinum 8378C CPU @ 2.80GHz
      Intel(R) Xeon(R) Platinum 8378C CPU @ 2.80GHz
```

2.7.3 Obtaining Metadata and Passing User Data

2.7.3.1 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 2-34**. 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:

- 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 2-34 does not contain the following EC2-compatible 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 2-34 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 2-35**.
OpenStack	/password	Displays the password for logging in to an ECS.
OpenStack	/user_data	Displays ECS user data. This metadata allows you to specify scripts and configuration files for initializing ECSs. For details, see **Injecting User Data**. 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, authorize agency permissions to the **op_svc_ecs** account and ECSs in IAM. **NOTE** You can determine what permissions are granted to the agency based on the principal of least privilege (PoLP). ECSs will not use agencies to perform operations on resources.
OpenStack	/spot/instance-action	Queries the prompt of stopping a spot ECS.

Metadata Type	Metadata Item	Description
EC2-compatible	/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-compatible	/meta-data/local-hostname	The meaning of this field is the same as that of hostname.
EC2-compatible	/meta-data/public-hostname	The meaning of this field is the same as that of hostname.
EC2-compatible	/meta-data/instance-type	Displays an ECS flavor.
EC2-compatible	/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-compatible	/meta-data/placement/availability-zone	Displays the AZ accommodating an ECS.
EC2-compatible	/meta-data/public-ipv4	Displays the EIP bound to the ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.
EC2-compatible	/meta-data/public-keys/0/openssh-key	Displays the public key of an ECS.
EC2-compatible	/user-data	Displays ECS user data.
EC2-compatible	/meta-data/security-groups	Displays the security group of an ECS.

Table 2-35 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.

Parameter	Type	Description
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. For details about the default security group rules for the outbound direction, see **Default Security Groups and Rules**.

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

```
{
    "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

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

NOTE

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

Figure 2-77 404 Not Found

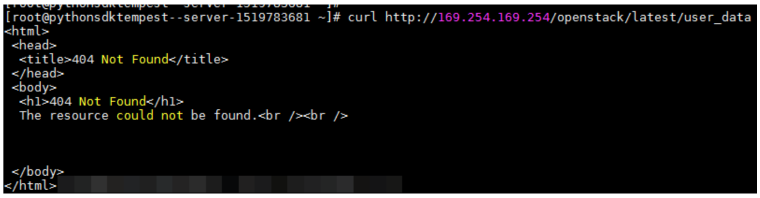

```
[root@pythonsdktempest--server-1519783681 ~]#
[root@pythonsdktempest--server-1519783681 ~]# curl http://169.254.169.254/openstack/latest/user_data
<html>
 <head>
  <title>404 Not Found</title>
 </head>
 <body>
  <h1>404 Not Found</h1>
  The resource could not be found.<br /><br />



 </body>
</html>[redacted]
```

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

```
{
    "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 a temporary AK/SK.

NOTE

- If an ECS needs to obtain a temporary AK/SK, you need to create and authorize an agency on the IAM console and then go to the ECS details page to configure **Agency** for the ECS in the **Management Information** area.

  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

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

```
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

```
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

```
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

```
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

  ```
  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

  ```
  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

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

Helpful Links

Why Can't My Linux ECS Obtain Metadata?

2.7.3.2 Injecting User Data

Scenarios

Specify **User Data** to inject user data into ECSs to:

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

Constraints

- 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.

Injecting 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 inject 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.
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 2-36 Linux ECS script types

Item	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.
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 injected into 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 2-78 Viewing operating results

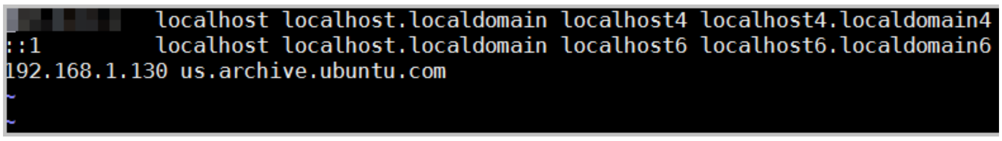

```
[redacted]      localhost localhost.localdomain localhost4 localhost4.localdomain4
::1             localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.1.130 us.archive.ubuntu.com
~
~
```

Case 1

This case illustrates how to inject 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 injection 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 2-37**.

Table 2-37 Password complexity requirements

Parameter	Requirement
Password	● Consists of 8 to 26 characters. ● Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Linux: !@%-_=+[]:./^,{}? ● Cannot contain the username or the username spelled backwards.

User data example:

Using a ciphertext password (recommended)

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

In the preceding command output, **$6$V6azyeLwcD3CHlpY $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 injection 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 2-38**.

Table 2-38 Password complexity requirements

Parameter	Requirement
Password	● Consists of 8 to 26 characters. ● Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Linux: !@%-_=+[]:./^,{}? ● Cannot contain the username or the username spelled backwards.

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 5

This case illustrates how to use the user data injection 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 inject the user data to assign user **root** permissions for remotely logging in to a Linux ECS. After injecting the file into 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 injection cases, visit the official Cloud-init/Cloudbase-init website:

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

2.7.4 Changing ECS Names

Scenarios

After an ECS is created, you can change its name as needed.

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 ECS details page, click ✎ next to the ECS name and edit the name.

   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 ✓.

Changing the Names of Multiple ECSs in a Batch

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. Select the target ECSs.
5. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.
6. Enter a new name.
7. Click **OK**.

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

2.7.5 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.
- In an ECS group associated with an anti-affinity policy, ECSs are deployed on physical servers.
- 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.

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

- ECSs of specific types must be stopped before being added to an ECS group. Stop these ECSs as prompted when adding them to an ECS group.
- 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.

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. 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 ECS group 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 **OK**.

   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 **OK**.

2.7.6 Automatically Recovering ECSs

Scenarios

ECSs run on physical servers. Although there are multiple mechanisms to ensure system reliability, error tolerance, and high availability, server hardware might be damaged or power failure might occur. If physical servers cannot be powered on or restarted due to damage, CPU and memory data will be lost, and the ECSs cannot recover through live migration.

The cloud platform provides auto recovery to restart ECSs through cold migration, ensuring high availability and top-performing dynamic migration capability of ECSs. You can enable auto recovery during or after ECS creation. If a physical server accommodating ECSs breaks down, the ECSs with auto recovery enabled will automatically be migrated to a functional server to minimize the impact on your services. During this process, the ECSs will restart.

Notes

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

Procedure

1. Log in to the management console.
2. Click 📍 in the upper left corner and select your region and project.

3. Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.

   The page providing details about the ECS is displayed.
5. Set **Auto Recovery** to **Enable** or **Disable**.

   Auto recovery is enabled by default.
   - If a physical server accommodating ECSs breaks down, the ECSs with auto recovery enabled will automatically be migrated to a functional server to minimize the impact on your services. During this process, the ECSs will restart.
   - If **Auto Recovery** is disabled, you must wait for the system administrator to recover ECSs when hardware becomes faulty.

2.7.7 Configuring Mapping Between Hostnames and IP Addresses in the Same VPC

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

2.8 Modifying ECS Specifications (vCPUs and Memory)

2.8.1 Modifying Specifications of Individual ECSs

Scenarios

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

Notes

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- During the specification modification, do not perform any operation on the ECS, such as stopping or restarting the ECS. Otherwise, the modification will fail.
- When modifying the specifications of an ECS, sold-out vCPU and memory resources cannot be selected.

- Downgrading ECS specifications (vCPUs or memory) will reduce performance.
- Certain ECS types do not support specifications modification. For details about ECS types and functions, see **ECS Types**. For details about constraints on using different types of ECSs, see their notes.
- When the disk status is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.
- Before modifying specifications, make sure that the ECS has been stopped.
- For yearly/monthly ECSs that use paid images, the instance specifications cannot be downgraded. This means you cannot change the specifications to lower-cost ones.

Preparations

After ECS specifications are modified, network interface 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 network interface 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
  ```

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

   If the ECS is not in **Stopped** state, choose **More** > **Stop** in the **Operation** column.
5. Choose **More** > **Modify Specifications** in the **Operation** column.

   The **Modify ECS Specifications** page is displayed.
6. Select the new ECS type, vCPUs, and memory as prompted.
7. Click **Next**.
8. Confirm the modified configurations and click **Submit**.
9. 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 **9.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.

- 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 & Deployment**, 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.

2.9 Reinstalling or Changing the OS

2.9.1 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 system disk, including the system partition. 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.
- 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 or select **Stop the ECS** in the **Reinstall OS** dialog box.
5. Select the login mode.

   If the target ECS uses key pair authentication, you can replace the original key pair.
6. Click **OK**.
7. In the **Reinstall OS** dialog box, 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

   During the reinstallation process, a temporary ECS is created. After the reinstallation is complete, 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.

2.9.2 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

- The OS change takes about 1 to 4 minutes During this process, the ECS status is **Changing OS**.
- Do not perform any operations on the ECS before the system injects the password or key, or the login will fail.
- The target ECS must have a system disk attached.
- The EVS disk quota must be greater than 0.
- The system disk type cannot be changed.
- 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.

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.
- Changing the OS clears the data in all partitions of the system disk, including the system partition. Back up data before changing the OS. For details, see **Backing Up an ECS**.
- 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.
- An OS change takes about 1 to 4 minutes to complete. During this process, the ECS is in **Changing OS** state.

Prerequisites

- The data is backed up.

  For details, see **Cloud Backup and Recovery**.

- If you want to change the login authentication mode from password to key pair during the OS change, create a key file in advance.

  For details, see **(Recommended) Creating a Key Pair on the Management Console**.

- 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** > **Change OS** in the **Operation** column.

   Before changing the OS, stop the ECS first or select **Stop the ECS** in the **Change OS** dialog box.
5. Select the target image.
6. Configure the login mode.

   If the target ECS uses key pair authentication, you can replace the original key pair.
7. Click **OK**.
8. In the **Change OS** dialog box, confirm the specifications, and click **OK**.

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

   NOTE

   During the OS change process, a temporary ECS is created. After the OS change 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.

2.10 Viewing ECS Information

2.10.1 Viewing the ECS Creation Status

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. After creating an ECS, view the creation status above the ECS list beside the common operations (**Start**, **Stop**, **Restart**, and **More**).
5. Click the number displayed above **Creating** and view task details.

   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 has been created successfully and displayed in the ECS list, see **Why Does the Failures Area Show an ECS Creation Failure But the ECS List Displays the Created ECS?**

2.10.2 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 2-39 lists the types of failures that can be recorded in the **Failures** area.

Table 2-39 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**. ● Automatic recovery enabled during ECS creation Automatic recovery is enabled during ECS creation. After the ECS is created, if the system fails to enable automatic recovery, 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.

   NOTE

   If **Failures** is not displayed on the management console, the following tasks have been successfully executed:

   - The ECS specifications are modified.
   - Automatic recovery is enabled during ECS creation.
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.

2.10.3 Viewing ECS Details

Scenarios

After obtaining ECSs, you can view and manage them on the management console. This section describes how to view ECS configuration details, including its name, image, system disk, data disks, VPC, network interfaces, security group, EIP address, and bandwidth.

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 the ECSs you have purchased as well as their details such as the 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 press **Enter**.
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**.

2.10.4 Exporting ECS Information

Scenarios

The information of all ECSs in your account can be exported to an XLSX file locally. 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.

2.10.5 Searching for ECSs

Scenarios

After purchasing an ECS, you can use the search function on the management console to search for ECSs in the current region. You can search for ECSs by name, ID, AZ, status, flavor name, image ID, EIP, private IP address, creation time, billing mode, VPC ID, enterprise project, or resource tag.

Search Syntax

A variety of ECS search types are available. For details, see **Table 2-40**.

NOTE

- For certain properties, if you enter complete values, the system can automatically identity the property type and search for them.
- The following properties only support exact search and you need to enter complete values: ID, image ID, EIP, VPC ID, and enterprise project.
- When you use multiple properties including the private IP address for search, you need to enter complete values of the IP address.
- The private IP address must be in the following CIDR blocks: 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24.
- You can search by tag key or key-value pair. You can add one or more tags. If keys are different, the tags are automatically joined with AND. If the keys are the same but the values are different, the tags are also automatically joined with AND.
- Tags do not support multi-value search if no property is selected.
- You cannot use both the private IP address and EIP for a combination of search.

Table 2-40 Search syntax

Search Type	Supported Property	Format	Example	Description
Property value Automatic property matching	ID Flavor name EIP Private IP address	Complete property value	ID: 4a79dfec-f0d8-4181-9bef-495b8b7220e1 Flavor Name: s2.xlarge.4 Private IP Address: 192.168.99.231	When you search by keyword, enter a complete property value instead of selecting a property. The system can automatically match the property type for search. Separate every two values with a comma (,). Otherwise, only the last value will be used for the search. Multiple property values are in OR relationship.
Property value Fuzzy search	Name Private IP address Flavor name	Incomplete property value	Name: ecs-c Flavor name: s7n Private IP address: 192.168.0	Select a property, and enter or choose the corresponding property value.
Single property	All properties on the console	Property: Value	Private IP Address: 192.168.99.231	Select a property, and enter or choose the corresponding property value. The following properties only support exact search and you need to enter complete values for them: ID, EIP, image ID, and VPC ID.
Multiple properties	All properties on the console	Property 1: Value Property 2: Value	Private IP Address: 192.168.99.231 Name: ecs-c	You can search by multiple properties and the properties are in AND relationship. The following properties only support exact search and you need to enter complete values for them: ID, EIP, image ID, and VPC ID.

Search Type	Supported Property	Format	Example	Description
Single property with multiple values	ID Flavor name EIP	Property: Value 1,Value 2	ID: 624eda28-6bd9-402a-934b-26c8969f7169,bf6c0281-f749-42d7-b732-23ac69d80ebe Flavor Name: s2,s3	Select a property and enter or choose multiple values. The values are in OR relationship. The ID only supports exact search. You need to enter a complete ECS ID for search.
	Status Billing mode	Property: Value 1 Property: Value 2	Status: Running Status: Stopped	Select a property and choose multiple values. The values are in OR relationship.
Multiple properties with multiple values	ID, flavor name, status, EIP, billing mode, and private IP address	Property1: Value 1,Value 2 Property 2: Value 1,Value 2	ID: 624eda28-6bd9-402a-934b-26c8969f7169,bf6c0281-f749-42d7-b732-23ac69d80ebe Flavor Name: s2,s3	Multi-property and multi-value search ● Multiple properties are in AND relationship. ● Multiple values of the same property are in OR relationship. The ID only supports exact search. You need to enter a complete ECS ID for search. You can use vertical bars (|) to separate values of the following properties: status and billing mode. Alternatively, you can directly select the properties and corresponding values.

Procedure

1. Log in to the management console.
2. Click ☰ . Under **Compute**, choose **Elastic Cloud Server**.

   The ECS console is displayed.

3. In the search box, specify search criteria.

   You can use either of the following methods:

   - Directly enter a property value for search.
   - Select a property first and specify the property value for search.
     i. Click the search box and select a property.

        Figure 2-79 Selecting a property

        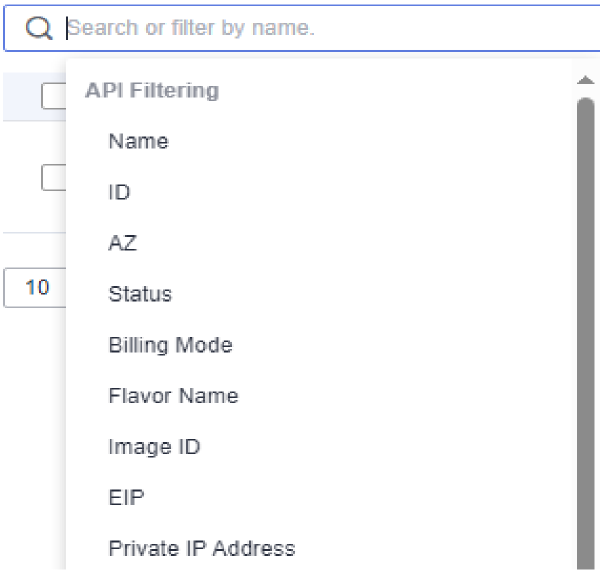

     ii. Specify a property value and press **Enter** to search.

Example 1: Searching by Property Value

After you enter a complete property value, the system automatically identifies the property type and search for it. Separate every two values with a comma (,). Otherwise, only the last value is used for the search. Multiple property values are in OR relationship.

- Searching by a single value

  Enter a complete ECS ID and press **Enter** to search.

  Figure 2-80 Entering a complete ECS ID

  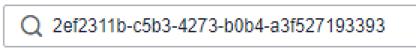
  2ef2311b-c5b3-4273-b0b4-a3f527193393

  Figure 2-81 Automatic property matching

  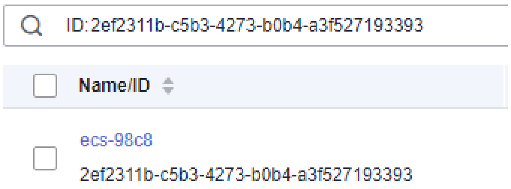
  ID:2ef2311b-c5b3-4273-b0b4-a3f527193393

□	Name/ID
□	ecs-98c8 2ef2311b-c5b3-4273-b0b4-a3f527193393

- Searching by multiple values

Enter multiple complete flavor names and press **Enter** to search.

Figure 2-82 Entering multiple complete flavor names

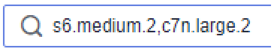

s6.medium.2,c7n.large.2

Figure 2-83 Automatic property matching for search

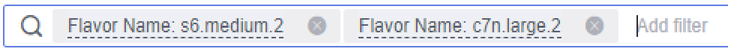

Flavor Name: s6.medium.2

Flavor Name: c7n.large.2

Example 2: Searching by a Single Property

Select a property, and enter or choose the corresponding property value.

The following properties only support exact search and you need to enter complete values for them: ID, EIP, image ID, and VPC ID.

- Fuzzy search by private IP address
  a. Select **Private IP Address** in the search box.
  b. Enter a private IP address and press **Enter** to search. Private IP addresses can be used for fuzzy search. For example, you can enter **192.168.0** to search for all ECSs that use the 192.168.0 IP address.

Figure 2-84 Searching by private IP address

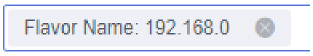

Flavor Name: 192.168.0

- Exact search by ID
  a. Select **ID** in the search box.
  b. Enter a complete ECS ID and press **Enter** to search.

Figure 2-85 Searching for an ECS by ID

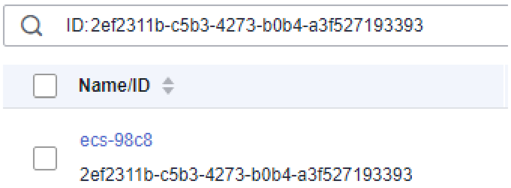

ID:2ef2311b-c5b3-4273-b0b4-a3f527193393

□	Name/ID
□	ecs-98c8 2ef2311b-c5b3-4273-b0b4-a3f527193393

Example 3: Searching by Multiple Properties

You can search by multiple properties and the properties are in AND relationship.

The following properties only support exact search and you need to enter complete values for them: ID, EIP, image ID, and VPC ID.

In this example, use **Name** and **Private IP Address** for a combination of search.

1. Select **Name** and enter an ECS name in the search box for fuzzy search.
2. Select **Private IP Address** and enter a private IP address for fuzzy search.

Figure 2-86 Searching by name and private IP address

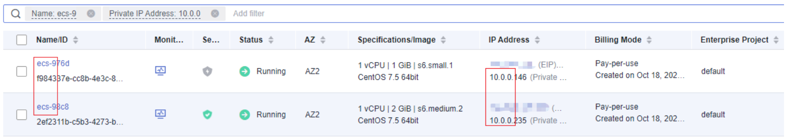

☐	Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address	Billing Mode	Enterprise Project
☐	ecs-976d f984337e-cc8b-4e3c-8...			Running	AZ2	1 vCPU | 1 GiB | s6.small.1 CentOS 7.5 64bit	[redacted] (EIP)... 10.0.0.146 (Private ...	Pay-per-use Created on Oct 18, 202...	default
☐	ecs-98c8 2ef2311b-c5b3-4273-b...			Running	AZ2	1 vCPU | 2 GiB | s6.medium.2 CentOS 7.5 64bit	[redacted] (... 10.0.0.235 (Private ...	Pay-per-use Created on Oct 18, 202...	default

Example 4: Searching by a Single Property with Multiple Values

You can choose from the following properties: status, ID, flavor name, private IP address, EIP, billing mode, and tag.

The following properties only support exact search and you need to enter complete values for them: ID, private IP address, and EIP.

Select a property and enter multiple values. The values are in OR relationship.

- Fuzzy search
  a. Select **Flavor Name** in the search box.
  b. Enter multiple flavor names and separate them with commas (,).

Figure 2-87 Searching by flavor name

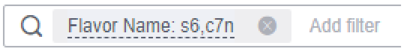

- Exact search
  a. Select **Private IP Address** in the search box.
  b. Enter multiple complete private IP addresses and separate them with commas (,).

Figure 2-88 Searching by private IP address

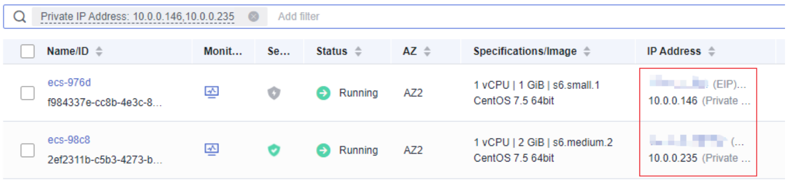

☐	Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
☐	ecs-976d f984337e-cc8b-4e3c-8...			Running	AZ2	1 vCPU | 1 GiB | s6.small.1 CentOS 7.5 64bit	[redacted] (EIP)... 10.0.0.146 (Private ...
☐	ecs-98c8 2ef2311b-c5b3-4273-b...			Running	AZ2	1 vCPU | 2 GiB | s6.medium.2 CentOS 7.5 64bit	[redacted] (... 10.0.0.235 (Private ...

Example 5: Searching by Multiple Properties with Multiple Values

You can choose from the following properties to search by properties with multiple values: ID, flavor name, status, billing mode, and EIP.

Multi-property, multi-value search

- Multiple properties are in AND relationship.
- Multiple values of the same property are in OR relationship.

The ID only supports exact search. You need to enter a complete ECS ID for search.

- Fuzzy search
  a. Select **Status** in the search box and choose **Running** and **Stopped**.
  b. Add **Flavor Name**, enter multiple flavors, and separate them with commas (,).

Figure 2-89 Searching by status and flavor name

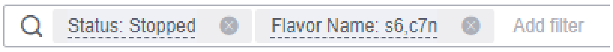

Status: Stopped ⊗ Flavor Name: s6,c7n ⊗ Add filter

- Exact search
  a. Select **Status** in the search box and choose **Running** and **Stopped**.
  b. Add **ID**, enter multiple IDs, and separate them with commas (,).

Figure 2-90 Searching by status and ID

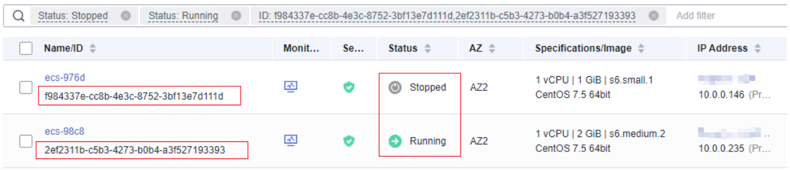

Status: Stopped ⊗ Status: Running ⊗ ID: f984337e-cc8b-4e3c-8752-3bf13e7d111d,2ef2311b-c5b3-4273-b0b4-a3f527193393 ⊗ Add filter

□	Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
□	ecs-976d f984337e-cc8b-4e3c-8752-3bf13e7d111d		✓	Stopped	AZ2	1 vCPU | 1 GiB | s6.small.1 CentOS 7.5 64bit	[redacted] 10.0.0.146 (Pr...
□	ecs-98c8 2ef2311b-c5b3-4273-b0b4-a3f527193393		✓	Running	AZ2	1 vCPU | 2 GiB | s6.medium.2 CentOS 7.5 64bit	[redacted] .. 10.0.0.235 (Pr...

Example 6: Searching by Tags

You can search by tag key or key-value pair.

You can add one or more tags. If keys are different, the tags are automatically joined with AND.

If the keys are the same but the values are different, the tags are also automatically joined with AND.

- Searching by a single tag
  In the search box, select a tag key under **Resource Tag** and then select a tag value for auto search.

Figure 2-91 Searching by tag

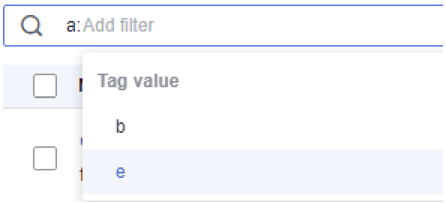

a: Add filter

- Searching by multiple tags
  In the search box, select multiple tag key-value pairs for auto search.
  If you search by multiple tags, the tags are in the AND relationship.

Figure 2-92 Searching by tag

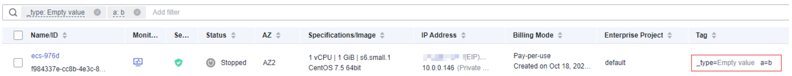

_type: Empty value ⊗ a: b ⊗ Add filter

□	Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address	Billing Mode	Enterprise Project	Tag
□	ecs-976d f984337e-cc8b-4e3c-8...		✓	Stopped	AZ2	1 vCPU | 1 GiB | s6.small.1 CentOS 7.5 64bit	[redacted] (EIP)... 10.0.0.146 (Private ...	Pay-per-use Created on Oct 18, 202...	default	_type=Empty value a=b

3 Images

3.1 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 3-1 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 KooGallery 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.
- A full-ECS image is shareable only when it is created from a CBR backup or from an ECS that has never had a CSBS backup. Full-ECS images created using other methods cannot be shared.

KooGallery Image

A KooGallery image is a third-party image that has an OS, application environment, and software preinstalled. You can use such an image for website setup, application development, and visualized management with just a few clicks. No additional configurations are required.

A KooGallery image can be free or paid, based on the image service providers. When you use a paid image to create an ECS, you need to pay for the KooGallery image and ECS.

Helpful Links

- **Creating a Private Image**

3.2 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**.

Prerequisites

Before creating an image, ensure that you have completed required configurations.

For details, see **How Do I Configure an ECS, BMS, or Image File Before I Use It to Create an Image?**

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, choose **More** > **Manage Image/Disk/Backup** > **Create Image** in the **Operation** column.
5. Configure the following information:

Table 3-2 and **Table 3-3** list the parameters in the **Image Type and Source** and **Image Information** areas, respectively.

Table 3-2 Image type and source

Parameter	Description
Image Type	Select **System disk image**.
Source	Click the **ECS** tab and select an ECS with required configurations.

Table 3-3 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 submit the request.

4 Disks

4.1 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 User Guide***.

Device Types

EVS disks have two device types, Virtual Block Device (VBD) and Small Computer System Interface (SCSI).

- VBD:

  When you create an EVS disk on the management console, **Device Type** of the EVS disk is VBD by default. VBD EVS disks support only simple SCSI read/write commands.

- SCSI:

  You can create EVS disks whose **Device Type** is SCSI on the management console. These EVS disks support transparent SCSI command transmission, allowing ECS OS to directly access underlying storage media. SCSI EVS disks support both basic and advanced SCSI commands.

NOTE

For more information about how to use SCSI EVS disks, for example, how to install a driver for SCSI EVS disks, see **Do I Need to Install a Driver for SCSI EVS Disks?**

Helpful Links

- **Attaching an EVS Disk to an ECS**
- **Initializing an EVS Data Disk**
- **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?**

4.2 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:

- During the ECS purchase. Data disks added in this way are automatically attached to the ECS.
- After the ECS is purchased. Data disks added in this way must be manually attached to the ECS.

This section describes how to add a data disk after an ECS is purchased.

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 buying disks 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 purchased, it is attached to the ECS by default.

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**.

4.3 Attaching a 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 (choosing **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.
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 details about constraints 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**.

4.4 Initializing Data Disks

After you attach a new data disk to a server, you must initialize the disk including creating partitions, creating file systems, and mounting the partitions before you can use the disk.

Scenario

- **System disk**

  When a server is created, a system disk is automatically initialized with master boot record (MBR).

- **New data disk**
  - If a data disk is created together with a server, EVS automatically attaches it to the server. You only need to initialize it to make it available for use.
  - If a data disk is created explicitly, you need to first attach it to a server and then initialize it.

  For detailed operation instructions, see **Table 4-1**.

- **Existing data disk**

  An existing data disk is a disk created from a snapshot, a backup, or an image, or a disk detached from another server.

  - You can choose not to initialize the disk and use the disk existing partitions.
    - In Linux, mount the partitions on desired mount points and configure auto mount at system start.

      For details, see **Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)**.

    - In Windows, no further action is required. You can simply use the existing partitions.
  - You can also re-initialize the data disk.

    Re-partitioning a disk will erase all the existing data on the disk, so you are advised to use snapshots to back up the disk data first.

    - In Linux, unmount the partitions, delete them (by running **fdisk** *<Disk name>*, entering **d** and the partition number, and entering **w**), and then re-initialize the disk.

- In Windows, delete the partitions (using the volume deletion tool) and then re-initialize the disk.

For detailed initialization operations, see **Table 4-1**.

NOTE

Initializing a disk does not delete the snapshots created for the disk, so you can still use snapshots to roll back data to the source disk after the disk is initialized.

Operation Instructions

Table 4-1 Disk initialization instructions

Disk Capacity	Partition Style	Partition Type	Operating System	Reference
Capacity ≤ 2 TiB	GPT/MBR	● GPT partitions are not classified, and there is no limit on the number of GPT partitions. ● MBR partitions can be: – Four primary partitions – Three primary partitions and one extended partition The number of logical partitions allowed in the extended partition is not limited, so theoretically you can create as many logical partitions as you want. If you need five or more partitions, use the "primary partitions + one extended partition" model and then create logical partitions in the extended partition.	Linux	**Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)**
			Windows	**Initializing a Windows Data Disk**
Capacity > 2 TiB	GPT	GPT partitions are not classified, and there is no limit on the number of GPT partitions.	Linux	**Initializing a Linux Data Disk (Greater Than 2 TiB)**
			Windows	**Initializing a Windows Data Disk**

NOTICE

- The maximum disk size that MBR supports is 2 TiB, and that GPT supports is 18 EiB. If your disk is greater than 2 TiB or you may expand it to over 2 TiB later, use GPT when initializing disks.
- If you change the partition style of a disk, data on the disk will be erased. Select an appropriate partition style when initializing disks.
- In Linux, you can use either fdisk or parted to create MBR partitions, and use only parted to create GPT partitions.

4.5 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 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

- 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 4-2** lists the second part of supported OSs.

Table 4-2 OSs supporting EVS disk detachment from a running ECS

OS	Version
CentOS	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Ubuntu Server	16.04 64bit
	14.04 64bit
	14.04.4 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**.

4.6 Expanding the Capacity of an EVS Disk

Scenarios

You can expand the disk capacity if the disk space is insufficient. The capacities of both system disks and data disks can be expanded.

Procedure

The capacity of an EVS disk can be expanded in either of the following ways:

- Create 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 about how to expand the capacity of an EVS disk, see **Disk Capacity Expansion**.

For details, see **Expansion Overview**.

NOTE

After the disk capacity is expanded, only the storage capacity of the EVS disk is expanded. To use the added storage space, you also need to log in to the ECS and extend the partition and file system.

4.7 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 created in **1** to the disk-intensive ECS for which you want to expand the capacity.
3. Back up local disk data.

   Back up the data stored in the local disks to the EVS disk that is newly attached to the disk-intensive ECS.

   - For Windows ECSs, directly copy the data to be backed up to the EVS disk.
   - For Linux ECSs, run the cp command to copy the data to be backed up to the EVS disk.
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 to go to the ECS details page.
   c. Click the **Disks** tab. Locate the row containing the EVS data disk and click **Detach** to detach the disk from the ECS.
5. Prepare a new disk-intensive ECS with higher specifications and larger capacity than the original one.

   Ensure that the local disk capacity can 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 in step **7** to the local disks of the new disk-intensive ECS.

5 Elastic Network Interfaces

5.1 Overview

An elastic network interface (referred to as a network interface in this documentation) is a virtual network card. You can create and configure network interfaces and attach them to your instances (ECSs and BMSs) to obtain flexible and highly available network configurations.

Network Interface Types

- A primary network interface is created together with an instance by default, and cannot be detached from the instance.
- An extended network interface can be created on the **Network Interfaces** tab, and can be attached to or detached from an instance.

Application Scenarios

- Flexible migration

  You can detach a network interface from an instance and then attach it to another instance. The network interface retains its private IP address, EIP, and security group rules. In this way, service traffic on the faulty instance can be quickly migrated to the standby instance, implementing quick service recovery.
- Traffic management

  You can attach multiple network interfaces that belong to different subnets in a VPC to the same instance, and configure the network interfaces to carry the private network traffic, public network traffic, and management network traffic of the instance. You can configure access control policies and routing policies for each subnet, and configure security group rules for each network interface to isolate networks and service traffic.

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 ECS to which you want to attach a network interface.

   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.
     1. (Optional) Search for the network interface by name, ID, or private IP address.
     2. 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

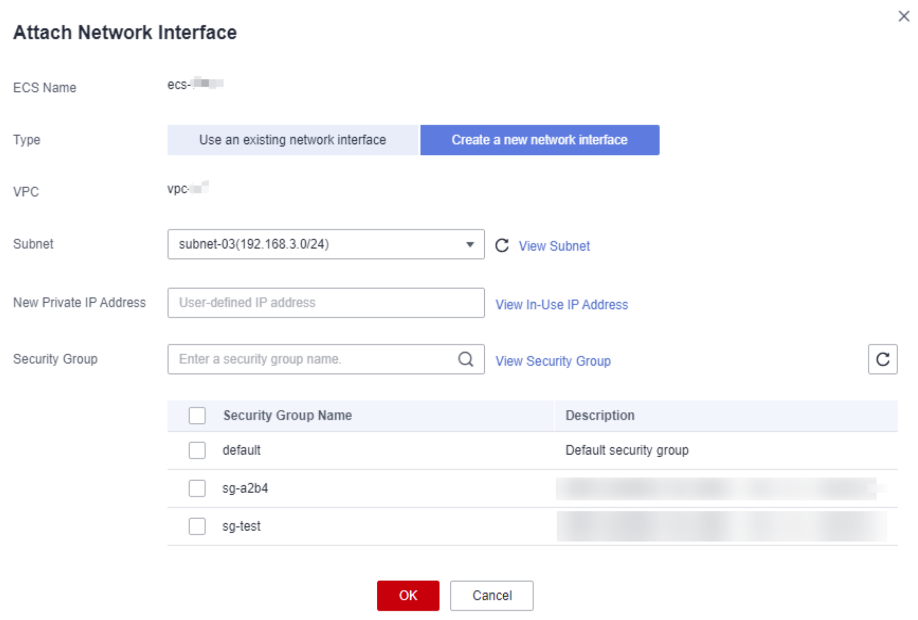

Attach Network Interface

ECS Name: ecs-[redacted]

Type: Use an existing network interface | Create a new network interface

VPC: vpc-[redacted]

Subnet: subnet-03(192.168.3.0/24)

New Private IP Address: User-defined IP address

Security Group: Enter a security group name.

	Security Group Name	Description
□	default	Default security group
□	sg-a2b4	[redacted]
□	sg-test	[redacted]

- **Subnet**: the subnet that the network interface belongs to.
- **New 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 access rules of all the selected security groups will apply to the ECS.

7. Click **OK**.

   To ensure that extension NICs can communicate with external networks, you need to configure policy-based routes for the NICs on the ECS after they are added.

   For details, see **How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?**

Follow-up Procedure

Some OSs cannot identify newly attached network interfaces. In this case, you need to manually activate the network interfaces. The following uses Ubuntu as an example to show how to activate network interfaces. Operations may vary depending on the operating system. You can refer to the corresponding OS documentation for assistance.

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 eth*X*** 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 serial number of the network interface, 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. 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 **ifconfig** again 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 a maximum of 12 network interfaces, including a primary network interface that cannot be detached. This section describes how to detach an extension network interface.

⚠ CAUTION

Detaching a networking interface may cause network interruptions. Evaluate the impact in advance and exercise caution when performing this operation.

Procedure

1. Log in to the management console.
2. Click ☰. Under **Compute**, click **Elastic Cloud Server**.
3. In the ECS list, click the name of the ECS from which you want to detach a network interface

   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 detach the primary network interface (the first one displayed in the network interface list).
5. In the displayed dialog box, click **OK**.

   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

- Only running or stopped ECSs support VPC change.
- 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 VPC switchover, do not bind, unbind, or replace the EIP. Otherwise, a message indicating insufficient permissions will be displayed, but you do not need to take any action.
- If an ECS NIC has an IPv6 address, the VPC of the ECS cannot be changed.

Notes

- A VPC can be changed on a running ECS, but the ECS network connection will be interrupted during the change process.

  NOTE

  If you intend to change the VPC for a running ECS, the VPC change may fail when traffic is routed to the ECS NIC. In this case, you are advised to try again later or stop the ECS first and then try to change the VPC.
- 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, suc h as ELB, VPN, NAT, traffic mirroring, and DNS.

Prerequisites

The target VPC, subnet, private IP address, and security group are available.

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 and choose **More** > **Manage Network** > **Change VPC** in the **Operation** column.

   The **Change VPC** dialog box is displayed.

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 will apply to 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.
- To change the private IP address for a backend server of a load balancer, remove the backend server from the backend server group first.

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.

Binding a Virtual IP Address

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**.

Configuring a Virtual IP Address for an ECS

After you bind one or more virtual IP addresses to an ECS on the console, you must log in to the ECS to manually configure these virtual IP address.

The following OSs are used as examples here. For other OSs, see the help documents on their official websites.

- Linux: CentOS 7.2 64bit and Ubuntu 22.04 server 64bit
- Windows: Windows Server

Linux (CentOS)

The following uses CentOS 7.2 64bit as an example.

1. Obtain the network interface that the virtual IP address is to be bound and the connection of the network interface:

   nmcli connection

   Information similar to the following is displayed:

   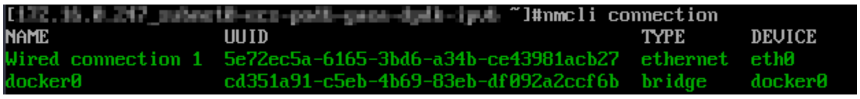

   ```
   [[redacted] ~]#nmcli connection
   NAME                UUID                                  TYPE      DEVICE
   Wired connection 1  5e72ec5a-6165-3bd6-a34b-ce43981acb27  ethernet  eth0
   docker0             cd351a91-c5eb-4b69-83eb-df092a2ccf6b  bridge    docker0
   ```

   The command output in this example is described as follows:

   - **eth0** in the **DEVICE** column indicates the network interface that the virtual IP address is to be bound.

   - **Wired connection 1** in the **NAME** column indicates the connection of the network interface.

2. Add the virtual IP address for the connection:

   nmcli connection modify "*<connection-name-of-the-network-interface>***" +ipv4.addresses** *<virtual-IP-address>*

   Configure the parameters as follows:

   - *connection-name-of-the-network-interface*: The connection name of the network interface obtained in **1**. In this example, the connection name is **Wired connection 1**.
   - *virtual-IP-address*: Enter the virtual IP address to be added. If you add multiple virtual IP addresses at a time, separate every two with a comma (,).

   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**

3. Make the configuration in **2** take effect:

   nmcli connection up "*<connection-name-of-the-network-interface>***"**

   In this example, run the following command:

   nmcli connection up "Wired connection 1"

   Information similar to the following is displayed:

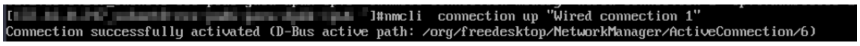

```
[[redacted] ]#nmcli  connection up "Wired connection 1"
Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/6)
```

4. 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 network interface eth0.

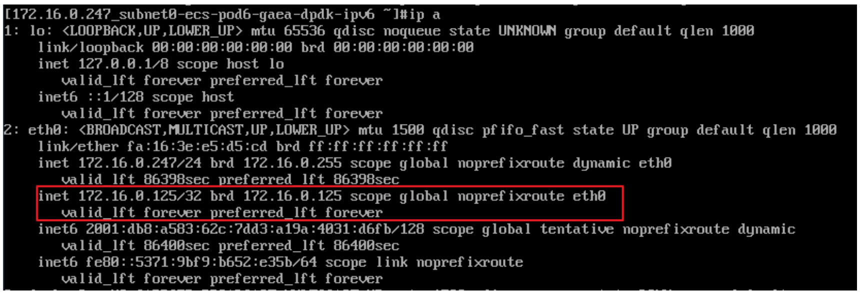

```
[172.16.0.247_subnet0-ecs-pod6-gaea-dpdk-ipv6 ~]#ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether fa:16:3e:e5:d5:cd brd ff:ff:ff:ff:ff:ff
    inet 172.16.0.247/24 brd 172.16.0.255 scope global noprefixroute dynamic eth0
       valid_lft 86398sec preferred_lft 86398sec
    inet 172.16.0.125/32 brd 172.16.0.125 scope global noprefixroute eth0
       valid_lft forever preferred_lft forever
    inet6 2001:db8:a583:62c:7dd3:a19a:4031:d6fb/128 scope global tentative noprefixroute dynamic
       valid_lft 86400sec preferred_lft 86400sec
    inet6 fe80::5371:9bf9:b652:e35b/64 scope link noprefixroute
       valid_lft forever preferred_lft forever
```

NOTE

To delete an added virtual IP address, perform the following steps:

1. Delete the virtual IP address from the connection of the network interface:

   nmcli connection modify "*<connection-name-of-the-network-interface>***" -ipv4.addresses** *<virtual-IP-address>*

   To delete multiple virtual IP addresses at a time, separate every two with a comma (,). Example commands are as follows:

   - Deleting a single virtual IP address: **nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125**
   - Deleting multiple virtual IP addresses: **nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125,172.16.0.126**

2. Make the deletion take effect by referring to **3**.

Linux (Ubuntu)

The following uses Ubuntu 22.04 server 64bit as an example. If the ECS runs **Ubuntu 22** or **Ubuntu 20**, perform the following operations:

1. Obtain the network interface that the virtual IP address is to be bound:

   ifconfig

   Information similar to the following is displayed. In this example, the network interface bound to the virtual IP address is **eth0**.

```
root@ecs-X-ubantu:~# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 172.16.0.210  netmask 255.255.255.0  broadcast 172.16.0.255
        inet6 fe80::f816:3eff:fe01:f1c3  prefixlen 64  scopeid 0x20<link>
        ether fa:16:3e:01:f1:c3  txqueuelen 1000  (Ethernet)
        RX packets 43915  bytes 63606486 (63.6 MB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 3364  bytes 455617 (455.6 KB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0
...
```

2. Switch to the **/etc/netplan** directory:

   cd /etc/netplan

3. Add a virtual IP address to the network interface.

   a. Open the configuration file **01-netcfg.yaml**:

      vim 01-netcfg.yaml

   b. Press **i** to enter the editing mode.

   c. In the network interface configuration area, add a virtual IP address.

      In this example, add a virtual IP address for **eth0**:

      addresses:

      - 172.16.0.26/32

      The file content is as follows:

```
network:
    version: 2
    renderer: NetworkManager
    ethernets:
        eth0:
            dhcp4: true
            addresses:
            - 172.16.0.26/32
        eth1:
            dhcp4: true
    eth2:
      dhcp4: true
    eth3:
      dhcp4: true
    eth4:
      dhcp4: true
```

d. Press **Esc**, enter **:wq!**, save the configuration, and exit.

4. Make the configuration in **3** take effect:

   netplan apply

5. Check whether the virtual IP address has been bound:

   ip a

   Information similar to the following is displayed. In the command output, virtual IP address 172.16.0.26 is bound to network interface eth0.

```
root@ecs-X-ubantu:/etc/netplan# ip a
...
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default
qlen 1000
    link/ether fa:16:3e:01:f1:c3 brd ff:ff:ff:ff:ff:ff
    altname enp0s3
    altname ens3
    inet 172.16.0.26/32 scope global noprefixroute eth0
       valid_lft forever preferred_lft forever
    inet 172.16.0.210/24 brd 172.16.0.255 scope global dynamic noprefixroute eth0
       valid_lft 107999971sec preferred_lft 107999971sec
    inet6 fe80::f816:3eff:fe01:f1c3/64 scope link
       valid_lft forever preferred_lft forever
```

NOTE

To delete an added virtual IP address, perform the following steps:

1. Open the configuration file **01-netcfg.yaml** and delete the virtual IP address of the corresponding network interface by referring to **3**.
2. Make the deletion take effect by referring to **4**.

Windows OS

The following operations use Windows Server as an example.

1. In **Control Panel**, click **Network and Sharing Center**, and click the corresponding local connection.
2. On the displayed page, click **Properties**.
3. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.
4. Click **Properties**.
5. 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 5-2 Configuring private IP address

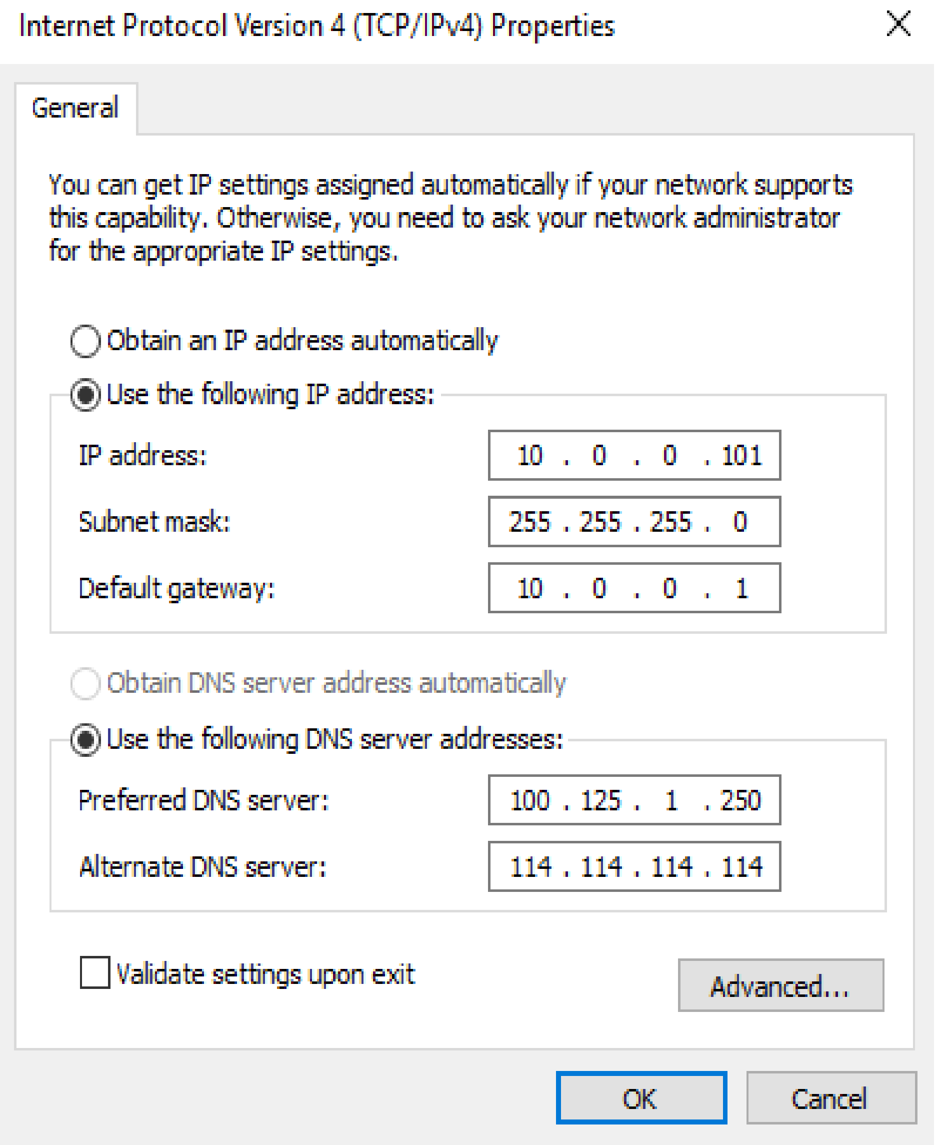

You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings.

○ Obtain an IP address automatically

● Use the following IP address:

IP address:	10 . 0 . 0 . 101
Subnet mask:	255 . 255 . 255 . 0
Default gateway:	10 . 0 . 0 . 1

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server:	100 . 125 . 1 . 250
Alternate DNS server:	114 . 114 . 114 . 114

□ Validate settings upon exit

6. Click **Advanced**.
7. On the **IP Settings** tab, click **Add** in the **IP addresses** area.

   Add the virtual IP address, for example, 10.0.0.154.

Figure 5-3 Configuring virtual IP address

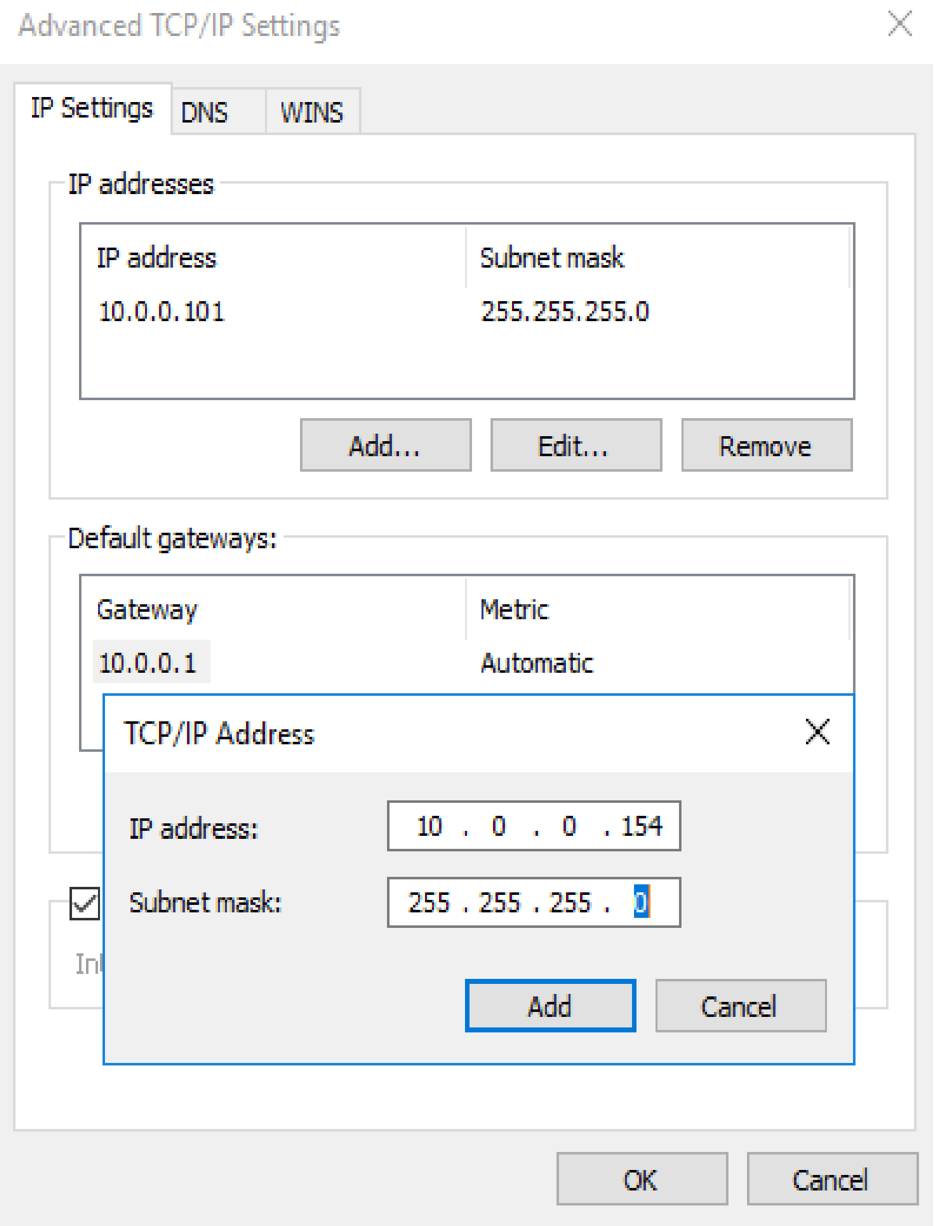

IP address	Subnet mask
10.0.0.101	255.255.255.0

Gateway	Metric
10.0.0.1	Automatic

IP address: 10 . 0 . 0 . 154

Subnet mask: 255 . 255 . 255 . 0

8. Click **OK**.
9. 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's network interface has been correctly configured.

5.7 Enabling NIC Multi-Queue

Scenarios

With the increase of network I/O bandwidth, single-core CPUs face bottlenecks in handling network interrupts. NIC multi-queue assigns interrupts to different CPUs for higher packets per second (PPS) and bandwidth.

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 5-1** support NIC multi-queue.

  NOTE

  - 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 Linux ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Ubuntu 20.04 server 64bit	Yes	Yes
CentOS 6.*/7.* 64bit	Yes	Yes
EulerOS 2.9 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

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 5-4 shows an example request body for modifying the NIC multi-queue attribute.

Figure 5-4 Example request body

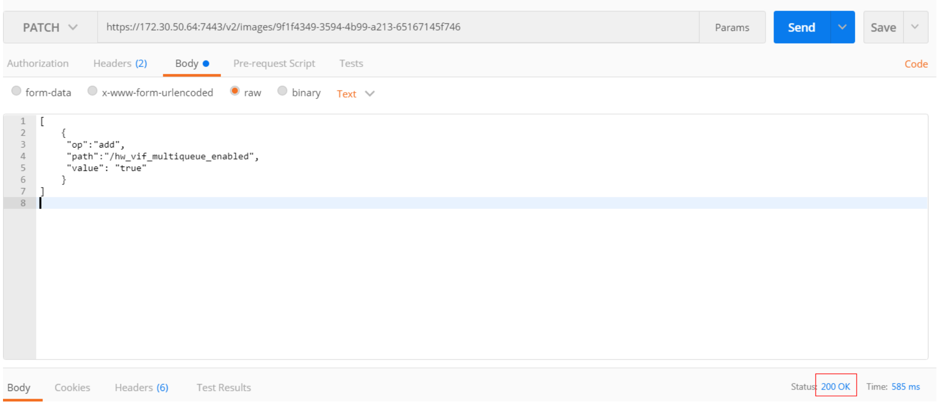

PATCH https://172.30.50.64:7443/v2/images/9f1f4349-3594-4b99-a213-65167145f746

```
[
    {
     "op":"add",
     "path":"/hw_vif_multiqueue_enabled",
     "value": "true"
    }
]
```

Status: 200 OK Time: 585 ms

Creating an ECS Using a Private Image

When using a registered private image to create an ECS, note the following parameter settings:

- **Region**: Select the region where the private image is located.
- **Image**: Click **Private image** and then select the desired image from the drop-down list.

Running the Script for Configuring NIC Multi-Queue

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.

NOTE

The script for automatically enabling NIC multi-queue only supports eth0 NICs.

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 output:

```
[root@localhost ~]# ethtool -l eth0   # 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 NIC Queues

The following uses a Linux ECS as an example to describe how to view the number of NIC queues.

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 Enabling IPv6 for a Network Interface

Scenarios

The IPv4/IPv6 dual-stack network provides IPv4 and IPv6 addresses for ECSs.

When purchasing ECSs, you can select a flavor that supports IPv6 and a primary network interface with IPv6 enabled. Then the ECSs can have both IPv4 and IPv6 addresses.

If IPv6 is not enabled when you purchase an ECS, the ECS has only an IPv4 address. In this case, you can enable IPv6 for the created ECS.

This section describes how to enable IPv6 for a created ECS.

Notes

- Ensure that the ECS flavor allows IPv6 to be enabled after the ECS is created.

  Currently, IPv6 can be enabled only for flavors of v7 and later versions, such as C7n and M7n

  On the **Network Interfaces** tab of the ECS details page, if **Enable IPv6** is displayed in the upper right corner of the row containing the target network interface, IPv6 can be enabled after the ECS is created.
- Only one IPv6 address can be bound to a network interface.

(Optional) Step 1: Enabling IPv6 for a Subnet

NOTE

- After IPv6 is enabled for the subnet that an ECS belongs to, an IPv6 CIDR block is automatically assigned to the subnet. IPv6 cannot be disabled once it is enabled for a subnet.
- If you have selected **Enable** for **IPv6 CIDR Block** when creating a subnet, an IPv6 CIDR block will be automatically assigned to the subnet. There is no need to perform the following steps.

1. Log in to the management console and go to the **Elastic Cloud Server** page.
2. Click the name of the ECS for which IPv6 is to be enabled. The ECS details page is displayed.
3. In the **NICs** area of the **Summary** tab, click the subnet name. The subnet details page is displayed.
4. On the **Summary** tab of the subnet details page, click **Enable IPv6**.
5. Click **OK** to enable IPv6 for the subnet.

Step 2: Enabling IPv6 for a Network Interface

1. Access the **Elastic Cloud Server** page.
2. Click the name of the ECS for which IPv6 is to be enabled. The ECS details page is displayed.
3. On the **Network Interfaces** tab, click **Enable IPv6** in the upper right corner of the row containing the target network interface.

   NOTE

   - **Enable IPv6** or **Disable IPv6** is displayed only when the ECS flavor supports IPv6 after the ECS is created.

     You can check whether IPv6 can be enabled after an ECS is created based on **Notes**.
   - If no IPv6 address is assigned after IPv6 is enabled, restart the ECS and check again. Alternatively, configure the IPv6 address by referring to **Dynamically Assigning IPv6 Addresses**.
   - You can disable IPv6 on this page if it is no longer used. After IPv6 is disabled, no IPv6 address is displayed for the network interface.
   - If you want to enable IPv6 again after it is disabled, you need to restart the ECS, log in to the ECS and manually clear the IPv6 cache, and request an IPv6 address again.
4. Click **OK** to enable IPv6 for the ECS network interface.

5.9 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.

Constraints

- Ensure that IPv6 has been enabled on the subnet where the ECS works.

  If IPv6 is not enabled, enable it by referring to **Enabling IPv6 for a Network Interface**. Once enabled, IPv6 cannot be disabled.
- 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.

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 **Automatically-assigned IPv6 address** is selected during ECS creation.
- After the 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.
- Check that the ECS network configuration is correct.

  For details about how to check the network configuration, see **Checking the ECS Network Configuration**.

  If the network configuration is incorrect, **submit a service ticket**.

Procedure

- Windows: Windows Server 2012 is used as an example to describe how to enable dynamic assignment of IPv6 addresses in Windows, as shown in **Table 5-2**.
- Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually, as shown in **Table 5-2**.

  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. You can set the timeout duration for assigning IPv6 addresses by referring to **Setting the Timeout Duration for IPv6 Address Assignment**.

Table 5-2 Enabling dynamic assignment of IPv6 addresses for different OSs

OS	Auto/Manual	Reference
Windows Server 2012	Auto	**Windows Server 2012**
Windows Server 2008	Auto	**Windows Server 2008**
Linux	Auto (recommended)	**Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)**

OS	Auto/Manual	Reference
Linux	Manual	**Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)**

Windows Server 2012

Step 1 Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window:

ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

Figure 5-5 Querying the IPv6 address

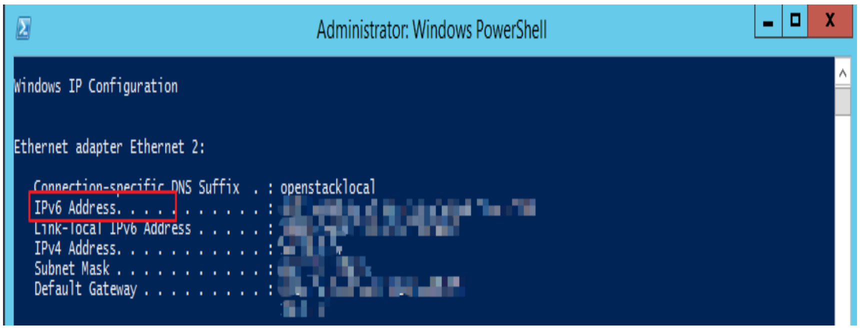

```
Windows IP Configuration

Ethernet adapter Ethernet 2:

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

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 5-6 Link-local IPv6 address

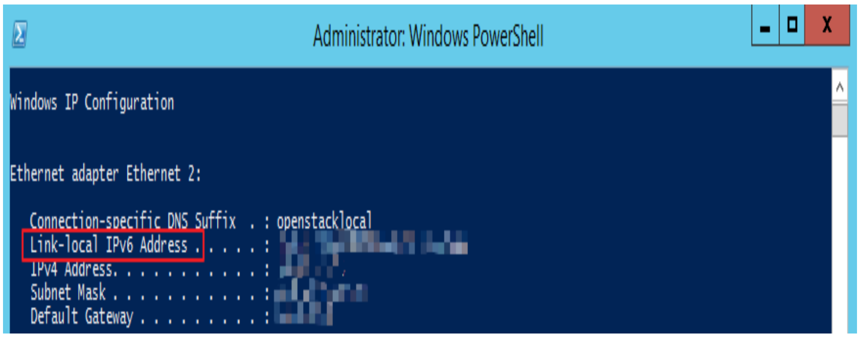

```
Windows IP Configuration

Ethernet adapter Ethernet 2:

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

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to **Step 3**.

Figure 5-7 IPv6 disabled

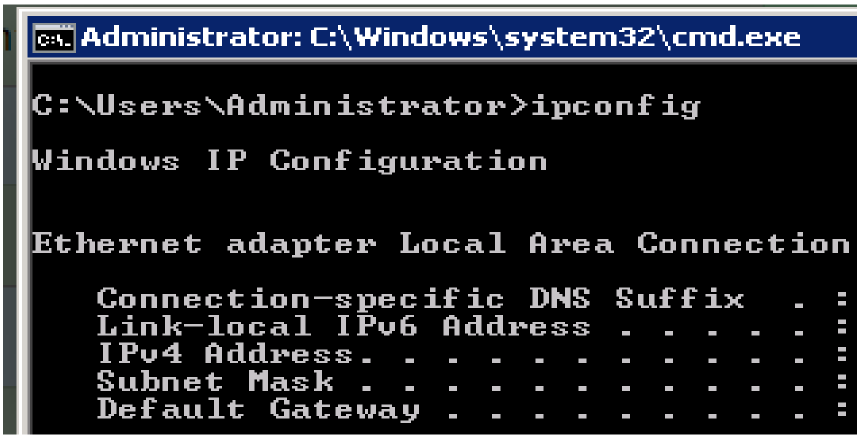

Administrator: C:\Windows\system32\cmd.exe

```
C:\Users\Administrator>ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection

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

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 5-5**. 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 5-8 Ethernet connection

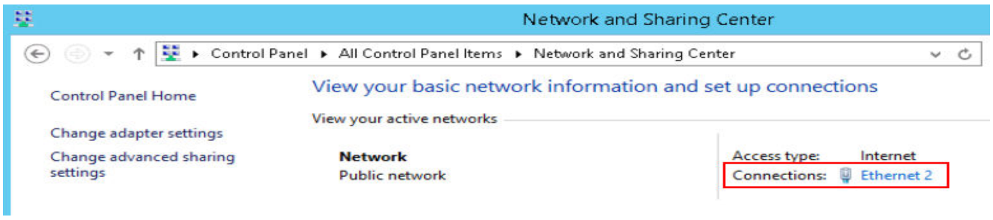

Network and Sharing Center
Control Panel ▸ All Control Panel Items ▸ Network and Sharing Center
Control Panel Home
Change adapter settings
Change advanced sharing settings

View your basic network information and set up connections

View your active networks

Network
Public network

Access type: Internet
Connections: Ethernet 2

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 5-9 Configuring dynamic IPv6 address assignment

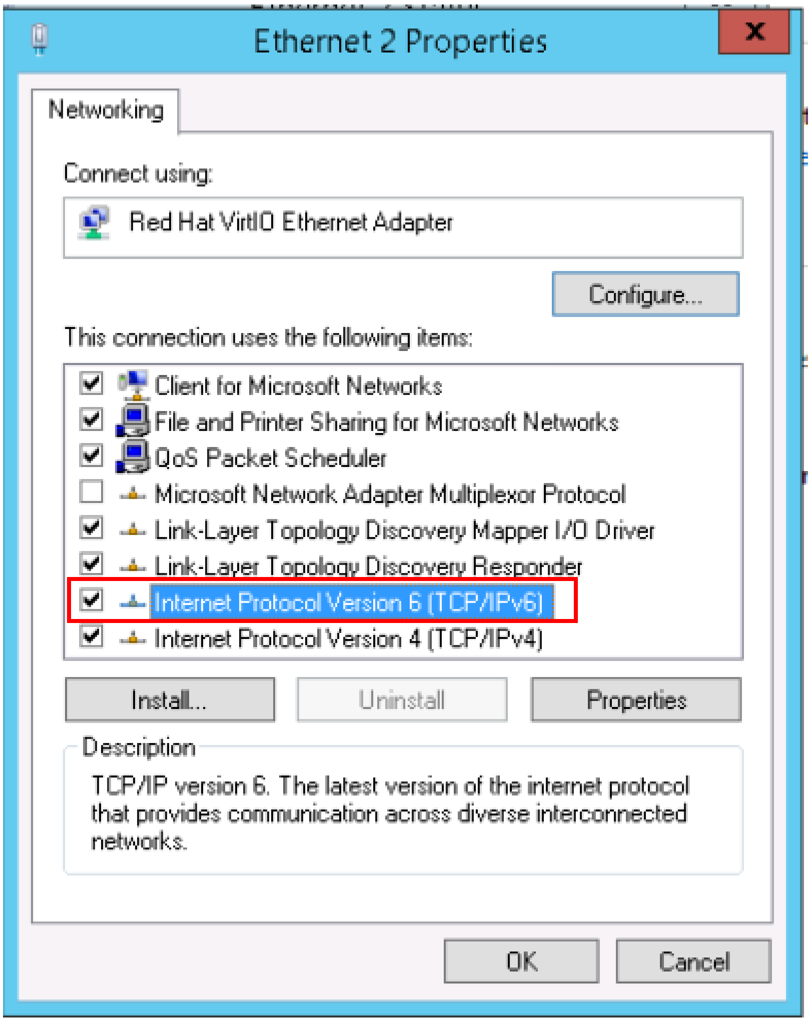

Ethernet 2 Properties

Networking

Connect using:

Red Hat VirtIO Ethernet Adapter

This connection uses the following items:

☑ Client for Microsoft Networks
☑ File and Printer Sharing for Microsoft Networks
☑ QoS Packet Scheduler
☐ Microsoft Network Adapter Multiplexor Protocol
☑ Link-Layer Topology Discovery Mapper I/O Driver
☑ Link-Layer Topology Discovery Responder
☑ Internet Protocol Version 6 (TCP/IPv6)
☑ Internet Protocol Version 4 (TCP/IPv4)

Description

TCP/IP version 6. The latest version of the internet protocol that provides communication across diverse interconnected networks.

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 5-10 Configuring an IPv6 address and a DNS server address

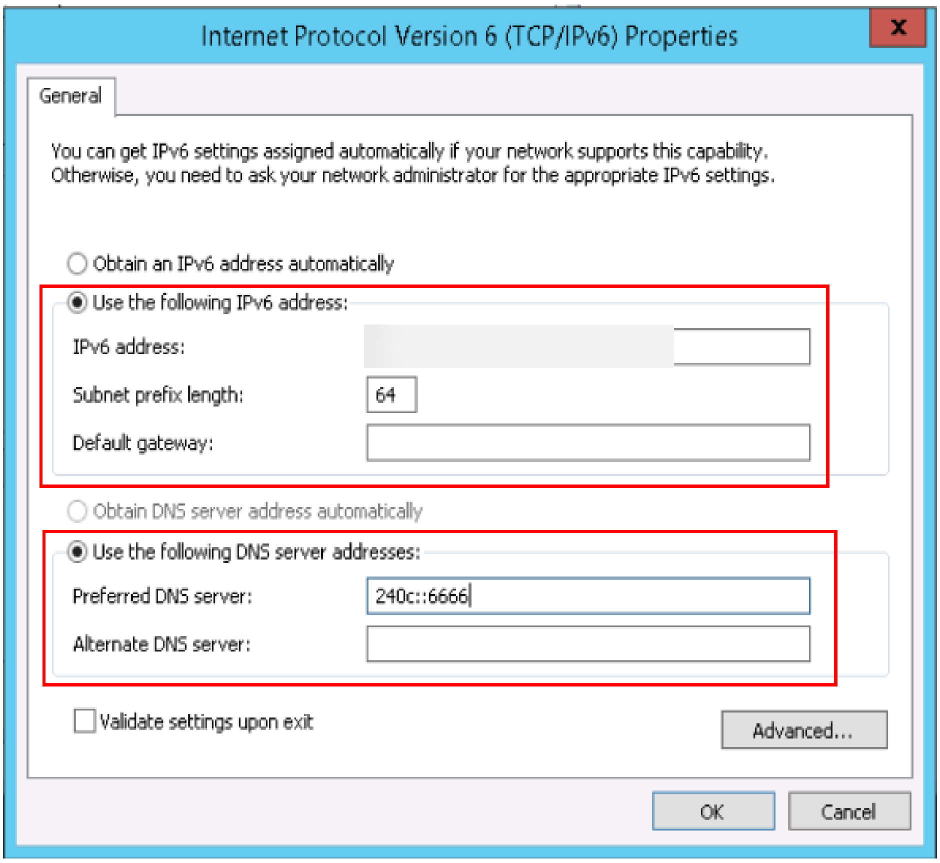

Internet Protocol Version 6 (TCP/IPv6) Properties

General

You can get IPv6 settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IPv6 settings.

○ Obtain an IPv6 address automatically

● Use the following IPv6 address:

IPv6 address:

Subnet prefix length: 64

Default gateway:

○ Obtain DNS server address automatically

● Use the following DNS server addresses:

Preferred DNS server: 240c::6666

Alternate DNS server:

□ Validate settings upon exit

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:

ipconfig

- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

Figure 5-11 Querying the IPv6 address

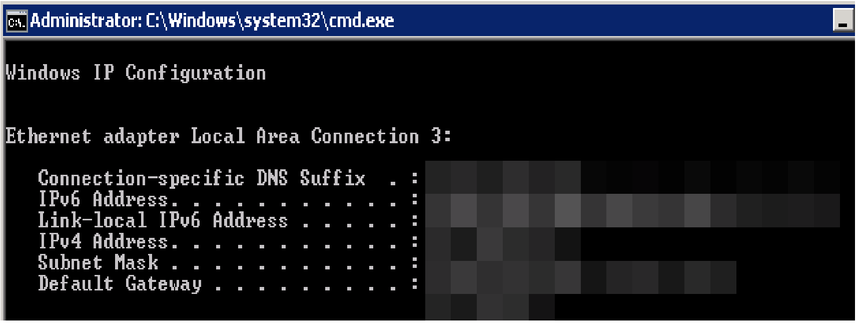

Administrator: C:\Windows\system32\cmd.exe

```
Windows IP Configuration

Ethernet adapter Local Area Connection 3:

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

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to **Step 2**.

Figure 5-12 Link-local IPv6 address

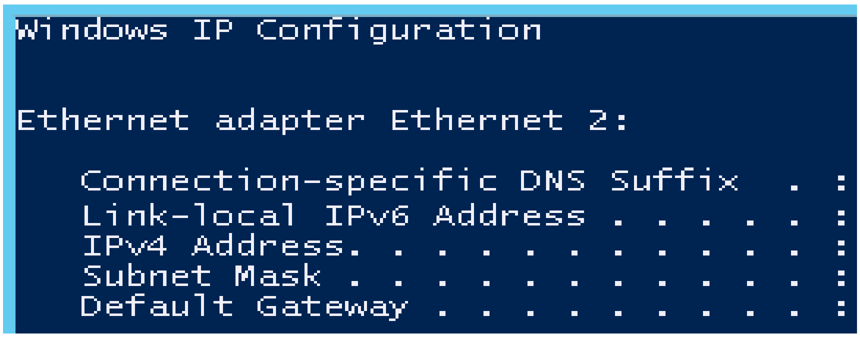

```
Windows IP Configuration

Ethernet adapter Ethernet 2:

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

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to **Step 3**.

Figure 5-13 IPv6 disabled

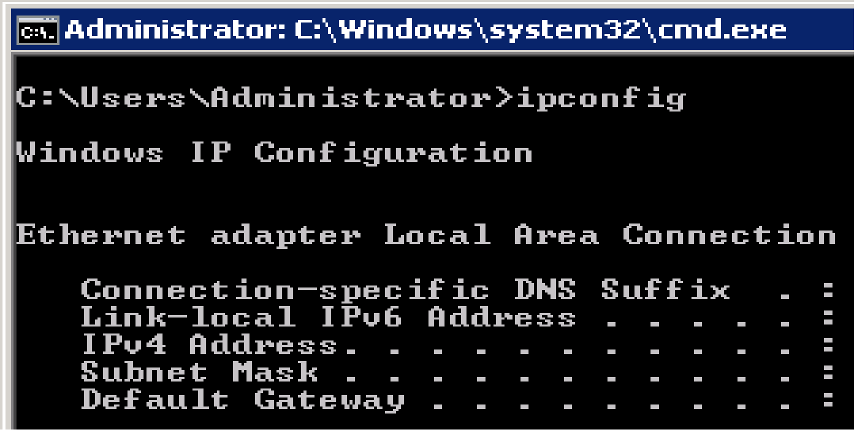

Administrator: C:\Windows\system32\cmd.exe

```
C:\Users\Administrator>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection

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

NOTE

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in **Figure 5-11**. 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 5-14 Configuring dynamic IPv6 address assignment

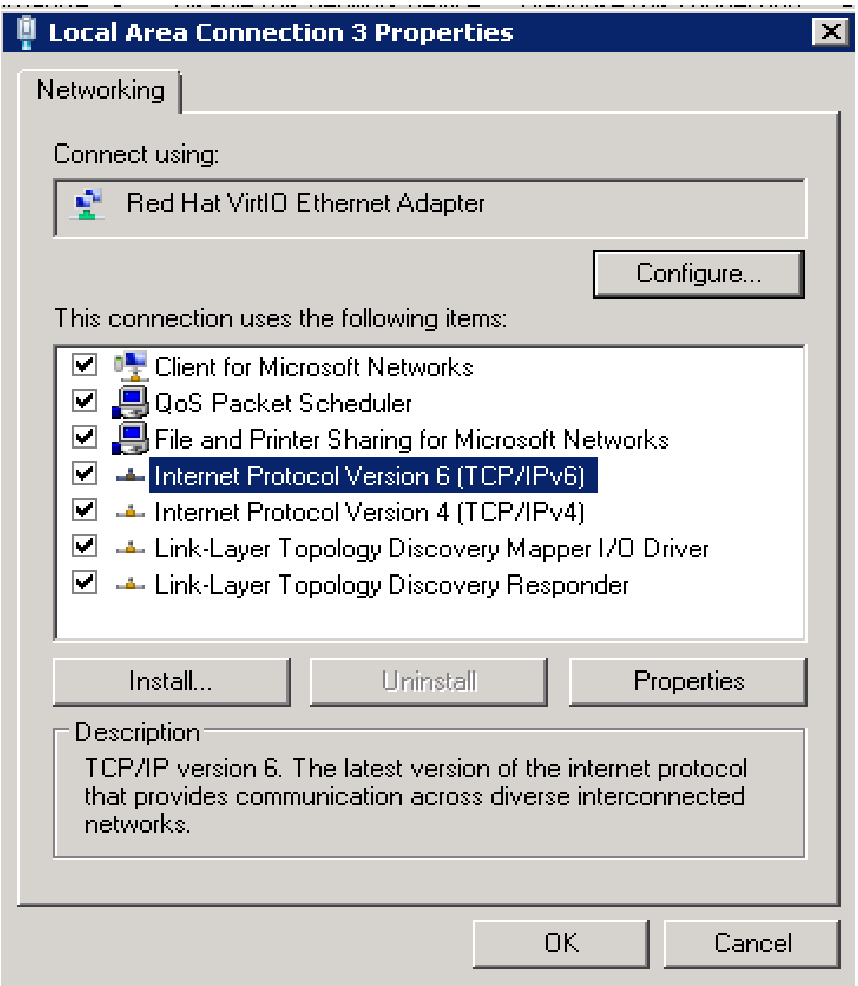

Local Area Connection 3 Properties

Networking

Connect using:

Red Hat VirtIO Ethernet Adapter

Configure...

This connection uses the following items:

☑ Client for Microsoft Networks
☑ QoS Packet Scheduler
☑ File and Printer Sharing for Microsoft Networks
☑ Internet Protocol Version 6 (TCP/IPv6)
☑ Internet Protocol Version 4 (TCP/IPv4)
☑ Link-Layer Topology Discovery Mapper I/O Driver
☑ Link-Layer Topology Discovery Responder

Install... Uninstall Properties

Description

TCP/IP version 6. The latest version of the internet protocol that provides communication across diverse interconnected networks.

OK Cancel

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 5-15 Enabling and configuring IPv6

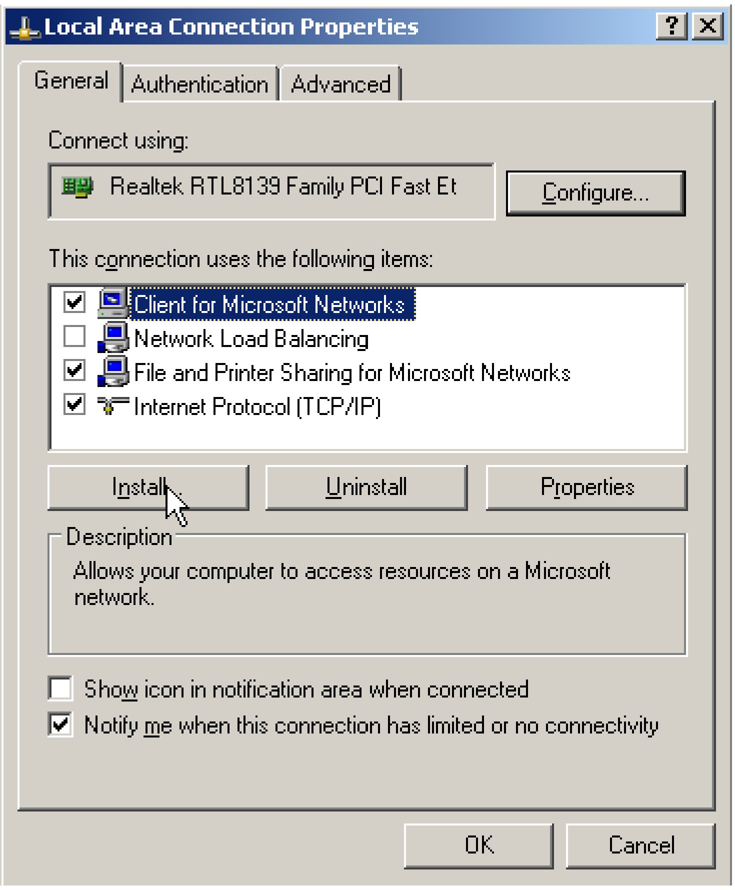

Local Area Connection Properties

General | Authentication | Advanced

Connect using:

Realtek RTL8139 Family PCI Fast Et

This connection uses the following items:

- ☑ Client for Microsoft Networks
- □ Network Load Balancing
- ☑ File and Printer Sharing for Microsoft Networks
- ☑ Internet Protocol (TCP/IP)

Description

Allows your computer to access resources on a Microsoft network.

□ Show icon in notification area when connected

☑ Notify me when this connection has limited or no connectivity

3. Select **Protocol** and click **Add**.

Figure 5-16 Adding the protocol

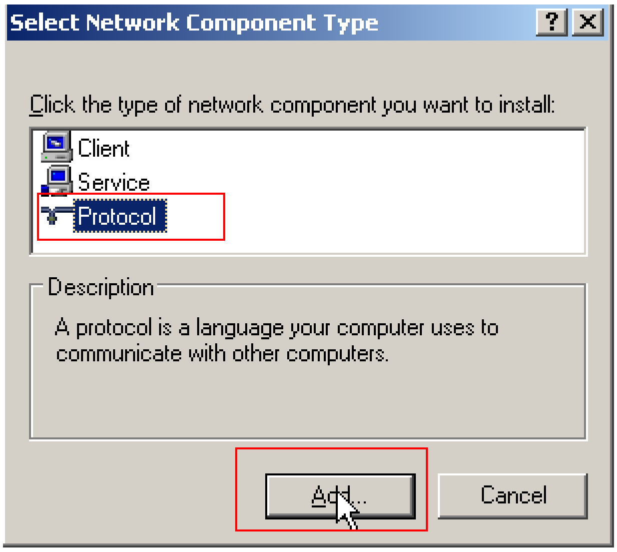

Select Network Component Type

Click the type of network component you want to install:

- Client
- Service
- Protocol

Description

A protocol is a language your computer uses to communicate with other computers.

4. Select **Microsoft TCP/IP Version 6** and click **OK**.

Figure 5-17 Network protocols

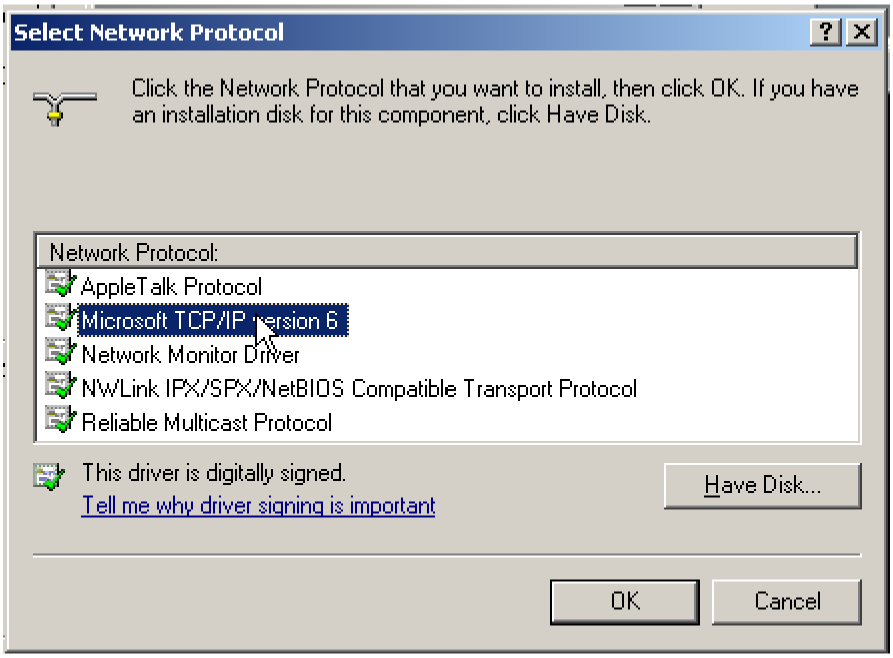

Select Network Protocol

Click the Network Protocol that you want to install, then click OK. If you have an installation disk for this component, click Have Disk.

Network Protocol:
- AppleTalk Protocol
- Microsoft TCP/IP version 6
- Network Monitor Driver
- NWLink IPX/SPX/NetBIOS Compatible Transport Protocol
- Reliable Multicast Protocol

This driver is digitally signed.
Tell me why driver signing is important

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 **Setting the Timeout Duration for IPv6 Address Assignment**.

Figure 5-18 IPv6 disabled

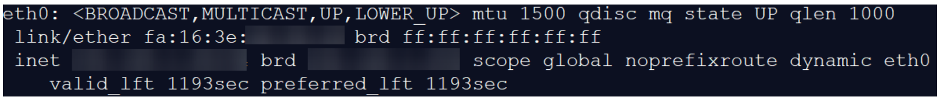

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
 link/ether fa:16:3e:[redacted] brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 1193sec preferred_lft 1193sec
```

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

Figure 5-19 IPv6 enabled

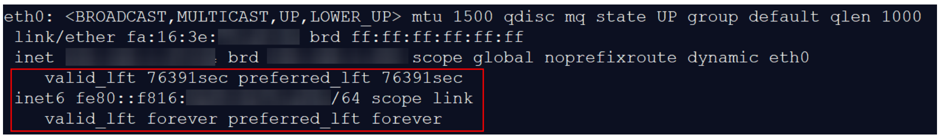

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
 link/ether fa:16:3e:[redacted] brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 76391sec preferred_lft 76391sec
 inet6 fe80::f816:[redacted]/64 scope link
    valid_lft forever preferred_lft forever
```

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

Figure 5-20 IPv6 enabled and an IPv6 address assigned

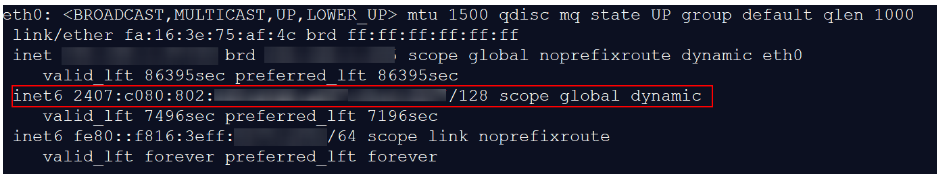

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
 link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 86395sec preferred_lft 86395sec
 inet6 2407:c080:802:[redacted]/128 scope global dynamic
    valid_lft 7496sec preferred_lft 7196sec
 inet6 fe80::f816:3eff:[redacted]/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
```

NOTE

IPv6 is enabled for Linux public images by default, as shown in **Figure 5-19**.

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 5-3** lists the paths for obtaining **ipv6-setup-rhel** and **ipv6-setup-debian**.

   Table 5-3 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 – 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 5-21 IPv6 disabled

  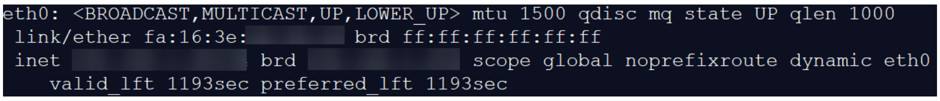

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
 link/ether fa:16:3e:[redacted] brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 1193sec preferred_lft 1193sec
```

- If a link-local address (starting with fe80) is displayed, IPv6 is enabled but dynamic assignment of IPv6 addresses is not enabled.

  Figure 5-22 IPv6 enabled

  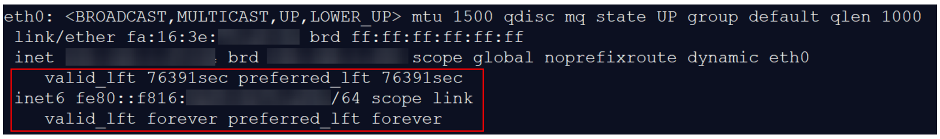

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
 link/ether fa:16:3e:[redacted] brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 76391sec preferred_lft 76391sec
 inet6 fe80::f816:[redacted]/64 scope link
    valid_lft forever preferred_lft forever
```

- If the following address is displayed, IPv6 is enabled and an IPv6 address has been assigned:

  Figure 5-23 IPv6 enabled and an IPv6 address assigned

  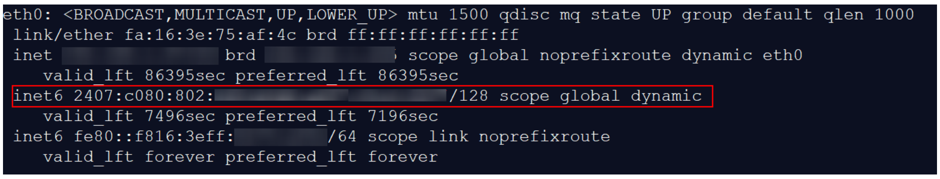

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
 link/ether fa:16:3e:75:af:4c brd ff:ff:ff:ff:ff:ff
 inet [redacted] brd [redacted] scope global noprefixroute dynamic eth0
    valid_lft 86395sec preferred_lft 86395sec
 inet6 2407:c080:802:[redacted]/128 scope global dynamic
    valid_lft 7496sec preferred_lft 7196sec
 inet6 fe80::f816:3eff:[redacted]/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
```

NOTE

IPv6 is enabled for Linux public images by default, as shown in **Figure 5-22**.

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 18.04/20.04

  a. Run the following command to access **/etc/netplan/**:

     cd /etc/netplan

  b. Run the following command to list the configuration file:

     ls

     Figure 5-24 Configuration file name

     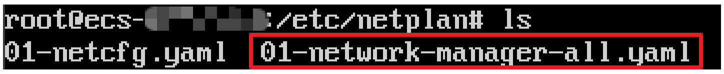

```
root@ecs-[redacted]:/etc/netplan# ls
01-netcfg.yaml  01-network-manager-all.yaml
```

  c. Run the following command to edit the configuration file **01-network-manager-all.yaml**:

     vi 01-network-manager-all.yaml

  d. Append the following content to the configuration file **01-network-manager-all.yaml** (pay attention to the YAML file format and text indentation):

```
ethernets:
 eth0:
  dhcp6: true
```

     Figure 5-25 Edited configuration file

     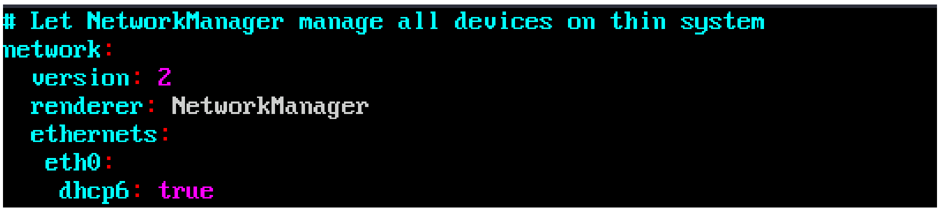

```
# Let NetworkManager manage all devices on thin system
network:
  version: 2
  renderer: NetworkManager
  ethernets:
   eth0:
    dhcp6: true
```

     Save the changes and exit.

  e. Run the following command to make the changes take effect:

     sudo netplan apply

- Ubuntu 22.04

  a. Run the following command to access **/etc/netplan/**:

     cd /etc/netplan

  b. Run the following command to list the configuration file:

     ls

Figure 5-26 Configuration file name

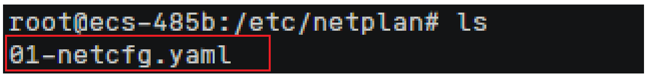

```
root@ecs-485b:/etc/netplan# ls
01-netcfg.yaml
```

c. Run the following command to edit the configuration file **01-netcfg.yaml**:

vi 01-netcfg.yaml

d. Append the following content to the configuration file **01-netcfg.yaml** (pay attention to the YAML file format and text indentation):

```
ethernets:
 eth0:
  dhcp6: true
```

Figure 5-27 Edited configuration file

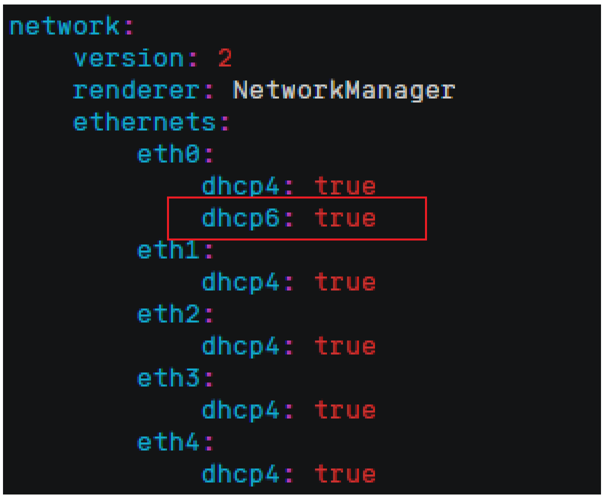

```
network:
    version: 2
    renderer: NetworkManager
    ethernets:
        eth0:
            dhcp4: true
            dhcp6: true
        eth1:
            dhcp4: true
        eth2:
            dhcp4: true
        eth3:
            dhcp4: true
        eth4:
            dhcp4: true
```

Save the changes and exit.

e. Run the following command to make the changes take effect:

sudo netplan apply

f. Run the following command to edit **/etc/NetworkManager/NetworkManager.conf**:

vi /etc/NetworkManager/NetworkManager.conf

g. Append the following content to the configuration file **NetworkManager.conf** (pay attention to the file format and indentation):

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient

[ifupdown]
managed=true

[device]
wifi.scan-rand-mac-address=no
```

Figure 5-28 Modification result

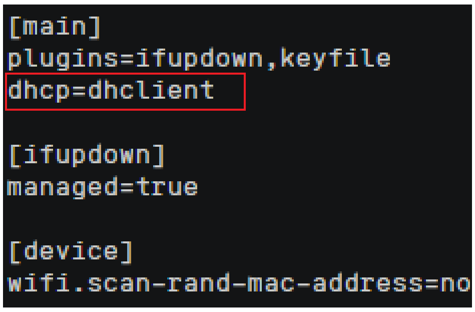

```
[main]
plugins=ifupdown,keyfile
dhcp=dhclient

[ifupdown]
managed=true

[device]
wifi.scan-rand-mac-address=no
```

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

   systemctl restart NetworkManager

- 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

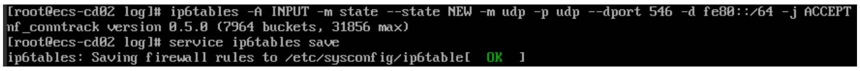

```
[root@ecs-cd02 log]# ip6tables -A INPUT -m state --state NEW -m udp -p udp --dport 546 -d fe80::/64 -j ACCEPT
nf_conntrack version 0.5.0 (7964 buckets, 31856 max)
[root@ecs-cd02 log]# service ip6tables save
ip6tables: Saving firewall rules to /etc/sysconfig/ip6table[  OK  ]
```

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

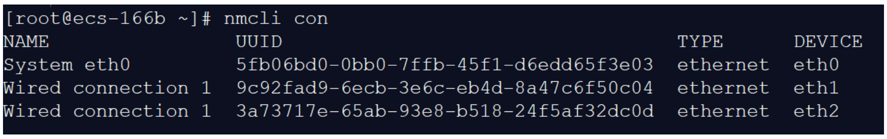

```
[root@ecs-166b ~]# nmcli con
NAME                UUID                                  TYPE      DEVICE
System eth0         5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03  ethernet  eth0
Wired connection 1  9c92fad9-6ecb-3e6c-eb4d-8a47c6f50c04  ethernet  eth1
Wired connection 1  3a73717e-65ab-93e8-b518-24f5af32dc0d  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

Checking the ECS Network Configuration

1. Run the following command to check whether the ECS network service is normal:

   systemctl status NetworkManager

   If the network service is normal, the command output shows **active (running)** and the service status is **enabled**.

   Figure 5-31 Network service status

   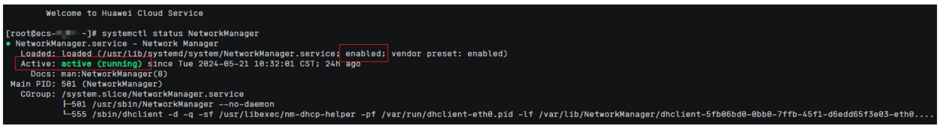

```
       Welcome to Huawei Cloud Service

[root@ecs-[redacted] ~]# systemctl status NetworkManager
● NetworkManager.service - Network Manager
   Loaded: loaded (/usr/lib/systemd/system/NetworkManager.service; enabled; vendor preset: enabled)
   Active: active (running) since Tue 2024-05-21 10:32:01 CST; 24h ago
     Docs: man:NetworkManager(8)
 Main PID: 501 (NetworkManager)
   CGroup: /system.slice/NetworkManager.service
           ├─501 /usr/sbin/NetworkManager --no-daemon
           └─555 /sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -pf /var/run/dhclient-eth0.pid -lf /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0....
```

2. Run the following command to check how the ECS NIC obtains an IP address:

   cat /etc/sysconfig/network-scripts/ifcfg-*ethx*

   NOTE

   - **ethx** needs to be replaced with a specific NIC, for example, **eth0**.
   - This command takes CentOS 7 as an example.

   Figure 5-32 Method for ECS NIC to obtain an IP address

   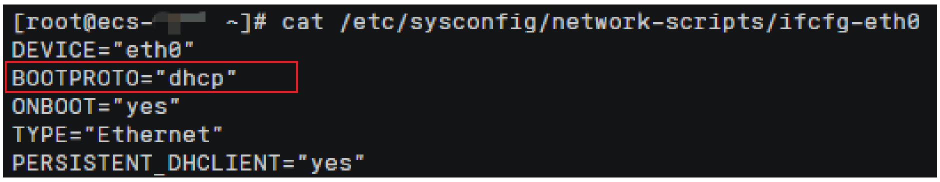

```
[root@ecs-[redacted] ~]# cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
BOOTPROTO="dhcp"
ONBOOT="yes"
TYPE="Ethernet"
PERSISTENT_DHCLIENT="yes"
```

   – If the value of **BOOTPROTO** is **dhcp**, the ECS NIC obtains an IP address using dynamic DHCP. Then, perform **3**.

   – If the value of **BOOTPROTO** is **static**, the network of the ECS NIC is statically configured. Then, perform **4**.

3. If the ECS NIC obtains an IP address using dynamic DHCP, run the following command to check whether the DHCP process is normal:

   systemctl status NetworkManager

   If the command output contains **dhclient**, the process is running properly.

   Figure 5-33 Checking the DHCP process

   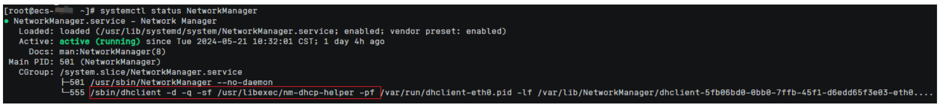

```
[root@ecs-[redacted] ~]# systemctl status NetworkManager
● NetworkManager.service - Network Manager
   Loaded: loaded (/usr/lib/systemd/system/NetworkManager.service; enabled; vendor preset: enabled)
   Active: active (running) since Tue 2024-05-21 10:32:01 CST; 1 day 4h ago
     Docs: man:NetworkManager(8)
 Main PID: 501 (NetworkManager)
   CGroup: /system.slice/NetworkManager.service
           ├─501 /usr/sbin/NetworkManager --no-daemon
           └─555 /sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -pf /var/run/dhclient-eth0.pid -lf /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0....
```

4. If the network of the ECS NIC is statically configured, run the following command to check whether the IP address configuration has taken effect:

   ip a

If the IP address configuration in the command output is **forever**, the configuration has taken effect.

Figure 5-34 Checking the static configuration

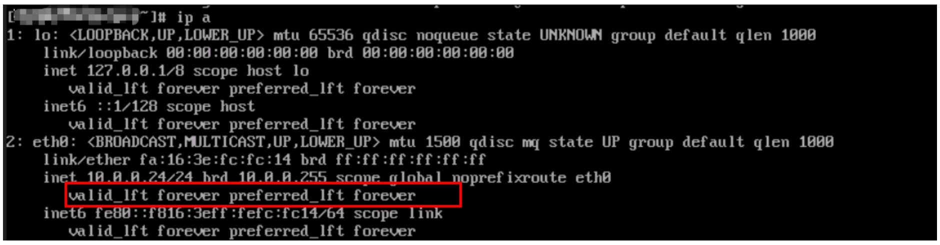

```
[[redacted]~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
       valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
       valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
    link/ether fa:16:3e:fc:fc:14 brd ff:ff:ff:ff:ff:ff
    inet 10.0.0.24/24 brd 10.0.0.255 scope global noprefixroute eth0
       valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fefc:fc14/64 scope link
       valid_lft forever preferred_lft forever
```

Setting the Timeout Duration for IPv6 Address Assignment

After automatic IPv6 address assignment is configured on an ECS running CentOS 6.x or Debian, the ECS will be created as a private image. When this image is used to create an ECS in an environment that IPv6 is unavailable, the ECS may start slow because acquiring an IPv6 address times out. Before creating the private image, you can set the timeout duration for acquiring IPv6 addresses to 30s as follows:

- 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 5-35 Modification 1

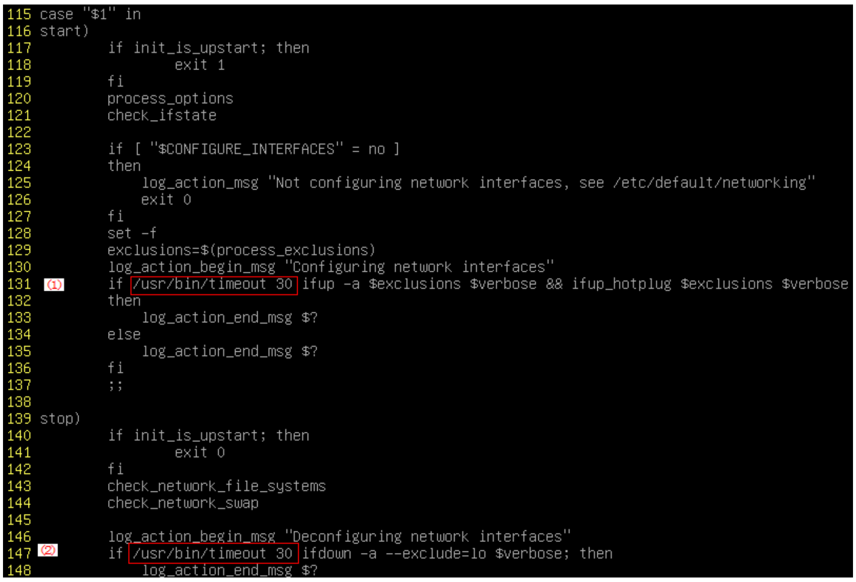

```
case "$1" in
start)
        if init_is_upstart; then
                exit 1
        fi
        process_options
        check_ifstate

        if [ "$CONFIGURE_INTERFACES" = no ]
        then
            log_action_msg "Not configuring network interfaces, see /etc/default/networking"
            exit 0
        fi
        set -f
        exclusions=$(process_exclusions)
        log_action_begin_msg "Configuring network interfaces"
        if /usr/bin/timeout 30 ifup -a $exclusions $verbose && ifup_hotplug $exclusions $verbose
        then
            log_action_end_msg $?
        else
            log_action_end_msg $?
        fi
        ;;

stop)
        if init_is_upstart; then
                exit 0
        fi
        check_network_file_systems
        check_network_swap

        log_action_begin_msg "Deconfiguring network interfaces"
        if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose; then
            log_action_end_msg $?
```

Figure 5-36 Modification 2

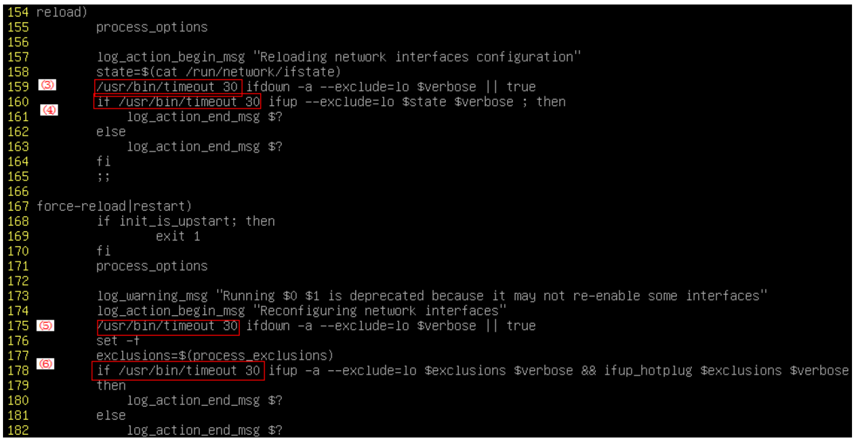

```
reload)
        process_options

        log_action_begin_msg "Reloading network interfaces configuration"
        state=$(cat /run/network/ifstate)
        /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
        if /usr/bin/timeout 30 ifup --exclude=lo $state $verbose ; then
            log_action_end_msg $?
        else
            log_action_end_msg $?
        fi
        ;;

force-reload|restart)
        if init_is_upstart; then
                exit 1
        fi
        process_options

        log_warning_msg "Running $0 $1 is deprecated because it may not re-enable some interfaces"
        log_action_begin_msg "Reconfiguring network interfaces"
        /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
        set -f
        exclusions=$(process_exclusions)
        if /usr/bin/timeout 30 ifup -a --exclude=lo $exclusions $verbose && ifup_hotplug $exclusions $verbose
        then
            log_action_end_msg $?
        else
            log_action_end_msg $?
```

- 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

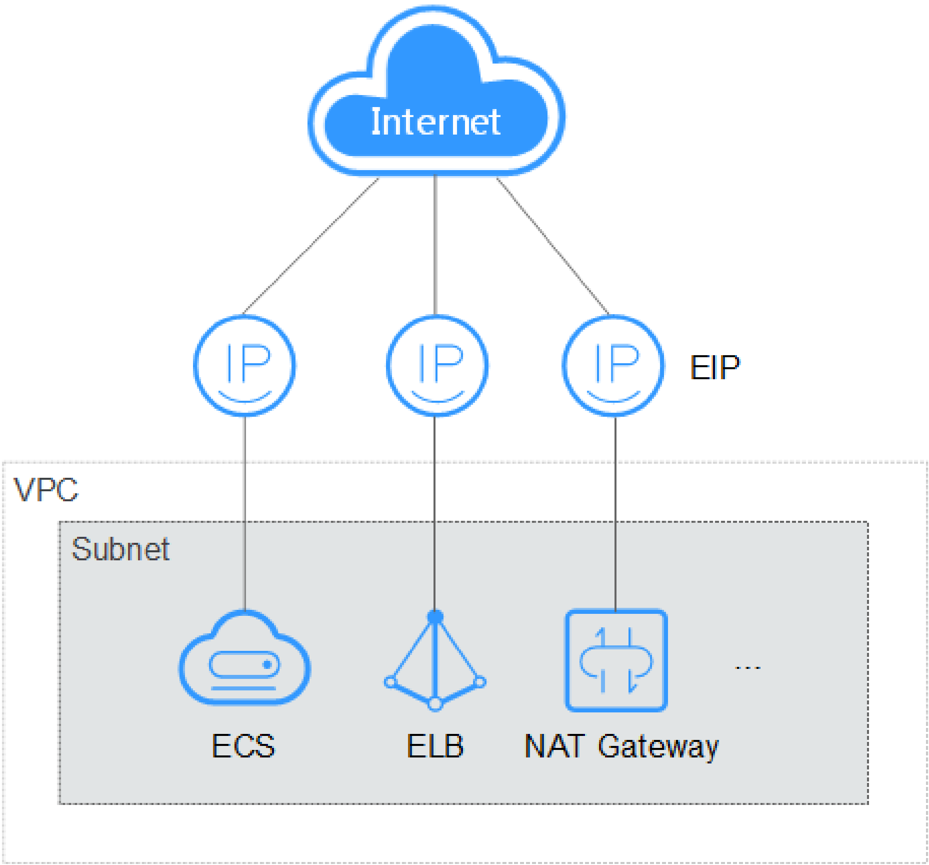

Helpful Links

- **Binding an EIP**
- **Unbinding an EIP**
- **Modifying an EIP Bandwidth**

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 that 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.

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.

⚠ CAUTION

After an EIP is unbound from an ECS, the ECS can no longer access the public network. Before unbinding an EIP, ensure that the ECS does not need to access the public network or an alternative network connection is available.

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 Network** > **Unbind EIP** in the **Operation** column.
5. Confirm the EIP to be unbound and click **OK**.

   NOTE

   Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

   For details, see **Releasing an EIP**.

6.4 Changing an EIP

Scenarios

You can change the EIP bound to your ECS as needed.

Currently, the EIP bound to the ECS cannot be directly replaced. You need to unbind the EIP first and then bind a new one to the ECS.

If there are no available EIPs, assign one first.

⚠ CAUTION

Before replacing an EIP, evaluate the impact on services to prevent network interruption.

Constraints

If an EIP is released by mistake, the system will preferentially assign you an EIP that you have released in the last 24 hours.

If you want to bind a new EIP to your ECS, you are advised to purchase one first before unbinding the original EIP.

For details, see **What Is the EIP Assignment Policy?**

Prerequisites

A new EIP has been purchased.

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 and choose **More** > **Manage Network** > **Unbind EIP** in the **Operation** column.
4. Confirm the displayed information and click **OK**.

   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 and choose **More** > **Manage Network** > **Bind EIP** in the **Operation** column.
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.

6.5 Modifying an EIP Bandwidth

Scenarios

The bandwidth of an EIP enables data transfer between the public network and an ECS. If the bandwidth of the EIP does not meet your service requirements, you can adjust the bandwidth by referring to this section.

Constraints

Reducing bandwidths may cause packet loss. Exercise caution when performing this operation.

Prerequisites

An EIP has been bound to an ECS. For details, see **Binding an EIP**.

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, locate the row containing the target ECS and choose **More** > **Manage Network** > **Modify Bandwidth** in the **Operation** column.
5. Change the bandwidth name, billing mode, or bandwidth size as prompted.

6.6 Enabling Internet Connectivity for an ECS Without an EIP Bound

Scenarios

To ensure platform security and conserve EIPs, EIPs are only assigned to specified ECSs. The ECSs that have not EIPs bound 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 that has an EIP bound to function as a proxy ECS to provide 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.

Using a Linux Proxy ECS

NOTE

Prerequisites for this solution:

- 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.

The following uses CentOS 7.9 as an example. The operations apply to CentOS 7.9 and earlier versions as well as Huawei Cloud EulerOS 2.0.

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**.

   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 (Linux)**.
8. 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 6-2 Checking connectivity

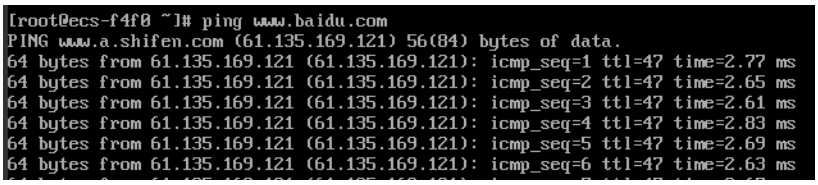

```
[root@ecs-f4f0 ~]# ping www.baidu.com
PING www.a.shifen.com (61.135.169.121) 56(84) bytes of data.
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=1 ttl=47 time=2.77 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=2 ttl=47 time=2.65 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=3 ttl=47 time=2.61 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=4 ttl=47 time=2.83 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=5 ttl=47 time=2.69 ms
64 bytes from 61.135.169.121 (61.135.169.121): icmp_seq=6 ttl=47 time=2.63 ms
```

9. (Optional) Install the iptables service and set the automatic startup of iptables.

   Perform this step only for ECSs running CentOS 7.x.

yum install iptables-services -y

systemctl start iptables

systemctl enable iptables

10. Check whether IP forwarding is enabled on the proxy ECS.

    cat /proc/sys/net/ipv4/ip_forward

    – If **0** (disabled) is displayed, go to **11**.

    – If **1** (enabled) is displayed, go to **16**.

11. Open the IP forwarding configuration file in the vi editor.

    vi /etc/sysctl.conf

12. Press **i** to enter editing mode.

13. Change the parameter value.

    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

14. Press **Esc**, type **:wq**, and press **Enter**.

    The system saves the configurations and exits the vi editor.

15. Apply the change.

    sysctl -p /etc/sysctl.conf

16. Delete the original iptables rules.

    iptables -F

17. 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 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 -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 **17** into **rc.local**, press **Esc** to exit Insert mode, and enter **:wq** to save the settings and exit.

18. Save the iptables configuration and set the automatic startup of iptables.

    service iptables save

    chkconfig iptables on

19. Check whether SNAT has been configured.

    iptables -t nat --list

    SNAT has been configured if information similar to **Figure 6-3** is displayed.

Figure 6-3 Successful SNAT configuration

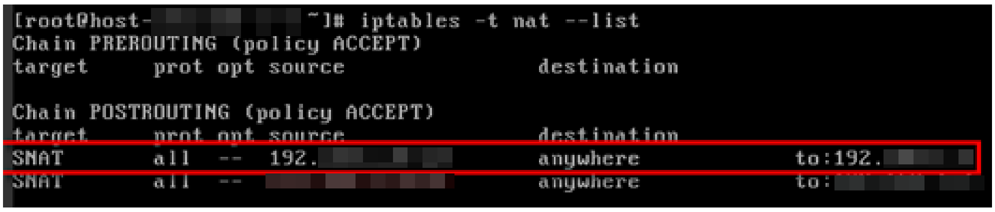

```
[root@host-[redacted] ~]# iptables -t nat --list
Chain PREROUTING (policy ACCEPT)
target     prot opt source               destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source               destination
SNAT       all  --  192.[redacted]        anywhere             to:192.[redacted]
SNAT       all  --  [redacted]            anywhere             to:[redacted]
```

20. 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 **Network**, 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.
21. Delete the added iptables rules as needed.

    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) and cloud-native anti-DDoS based on your service requirements:	• **Installing an Agent on Linux** • **Enabling the Basic/ Enterprise/Premium Edition** • **Monitoring ECSs** • **Enabling Anti-DDoS** • **Backing Up Data Periodically**

Type	Description	Protection Method
Internal security	Incorrect ports opening and weak passwords 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

Host Security Service (HSS) is designed to improve the overall security for ECSs. It helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

Before using the HSS service, install the HSS agent on your ECSs first and you will be able to check the ECS security status and risks 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.

- **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 the Agent on Huawei Cloud Servers** 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.

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.

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.

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 service interruptions. 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.

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 **Password**, ensure that the password meets complexity requirements to prevent malicious attacks. For details, see **Application Scenarios for Using Passwords**.

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.

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 7-2 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 7-2 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 Groups 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 Options**.

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 Groups and Rules**. You can also customize security group rules. For details, see **Configuring Security Group Rules**.

Constraints on Using Security Groups

- For better network performance, you are advised to associate an instance with no more than five security groups.

- A security group can have no more than 6,000 instances associated, or its performance will deteriorate.
- For inbound security group rules, the sum of the rules with **Source** set to **Security group**, of the rules with **Source** set to **IP address group**, and of the rules with inconsecutive ports, cannot exceed 120. If there are both IPv4 and IPv6 security group rules, up to 120 rules can be added for each type.

  The limits on outbound security group rules are the same as those on inbound rules.

  For example, to add inbound IPv4 rules to a security group (Sg-A), you can refer to **Table 7-3** for rules that meet the restrictions. Of these rules, rule A02 uses inconsecutive ports (TCP: 22,25,27) and security group Sg-B as the source. In this case, only one quota is occupied.

Table 7-3 Inbound security group rules

Rule No.	Action	Type	Protocol & Port	Source
Rule A01	Allow	IPv4	All	Current security group: Sg-A
Rule A02	Allow	IPv4	**TCP: 22,25,27**	**Another security group: Sg-B**
Rule A03	Allow	IPv4	TCP: 80-82	IP address group: ipGroup-A
Rule A04	Allow	IPv4	TCP: 22-24,25	IP address: 192.168.0.0/16

- Traffic from load balancers is not restricted by network ACL and security group rules if:

  Transfer Client IP Address is enabled for the listeners of a load balancer.

  The load balancer can still forward traffic to backend servers, even if there is a rule that denies traffic from the load balancer to the backend servers.

Recommendations

- Instances in a security group deny all external access requests by default, but you can add rules to allow specific requests.
- When adding a security group rule, grant the minimum permissions possible. For example, if remote login to an ECS over port 22 is allowed, only allow specific IP addresses to log in to the ECS. Do not use 0.0.0.0/0 (all IP addresses).
- Keep your configurations simple. There should be a different reason for each security group. If you use the same security group for all your different instances, the rules in the security group will likely be redundant and complex. It will make it much harder to maintain and manage.
- You can add instances to different security groups based on their functions. For example, if you want to provide website services accessible from the Internet, you can add the web servers to a security group configured for that

specific purpose and only allow external access over specific ports, such as 80 and 443. By default, other external access requests are denied. Do not run internal services, such as MySQL or Redis, on web servers that provide services accessible from the Internet. Deploy internal services on servers that do not need to connect to the Internet and associate these servers with security groups specifically configured for that purpose.

- If you have multiple IP addresses with the same security requirements, you can add them to an IP address group and select this IP address group when you configure a rule, to help you manage them in an easier way. When an IP address changes, you only need to change the IP address in the IP address group. Then, the rules in the IP address group change accordingly. You do not need to modify the rules in the security group one by one. This simplifies security group management and improves efficiency. For details, see **Using IP Address Groups to Reduce the Number of Security Group Rules**.
- Do not directly modify security group rules for active services. Before you modify security group rules used by a service, you can clone the security group and modify the security group rules in the test environment to ensure that the modified rules work. For details, see **Cloning a Security Group**.
- After you add instances to or modify rules of a security group, the security group rules are applied automatically. There is no need to restart the instances.

  If a security group rule does not take effect after being configured, see **Why Are My Security Group Rules Not Working?**

7.2.2 Default Security Groups and Rules

Note the following when using default security group rules:

- Inbound rules control incoming traffic to instances in the default security group. The instances can only communicate with each other but cannot be accessed from external networks.
- Outbound rules allow all traffic from the instances in the default security group to external networks.

Figure 7-1 shows the default security group.

Figure 7-1 Default security group

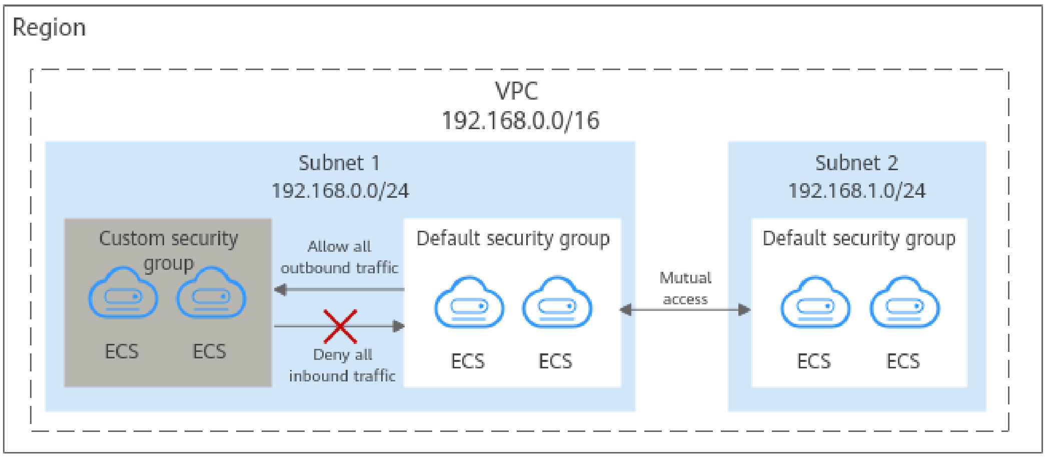

NOTE

- 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 7-4 describes the rules in the default security group.

Table 7-4 Rules in the default security group

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inbound	Allow	IPv6	All	Source: default security group (default)	Allows IPv6 instances in the security group to communicate with each other using any protocol over any port.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows all traffic from the instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows all traffic from the 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 a VPC **vpc-default** and:

- Add the **Sys-WebServer** security group.
- Add the **Sys-FullAccess** security group.
- Add security group rules to the default security group **default**.

Table 7-5 Default security group rules

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Windows ECSs through the default Windows remote desktop.
Inbound	Allow	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Linux ECSs over SSH.
Inbound	Allow	IPv4	All	Source: Default security group (default)	Allows instances in the security group to communicate with each other over IPv4 protocols.
Inbound	Allow	IPv6	All	Source: Default security group (default)	Allows instances in the security group to communicate with each other over IPv6 protocols.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows access from instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows access from instances in the security group to any IPv6 address over any port.

Table 7-6 Sys-WebServer security group rules

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	ICMP: All	Source: 0.0.0.0/0	Allows the use of the ping command to test the network connectivity over IPv4 protocols.
Inbound	Allow	IPv4	All	Source: current security group (Sys-WebServer)	Allows instances in the security group to communicate with each other over IPv4 protocols.
Inbound	Allow	IPv4	TCP: 443	Source: 0.0.0.0/0	Allows all IPv4 addresses to access websites deployed on ECSs over HTTPS.

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	TCP: 80	Source: 0.0.0.0/0	Allows all IPv4 addresses to access websites deployed on ECSs over HTTP.
Inbound	Allow	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Linux ECSs over SSH.
Inbound	Allow	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Windows ECSs through the default Windows remote desktop.
Inbound	Allow	IPv6	All	Source: current security group (Sys-WebServer)	Allows instances in the security group to communicate with each other over IPv6 protocols.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows access from instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows access from instances in the security group to any IPv6 address over any port.

Table 7-7 Sys-FullAccess security group rules

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	All	Source: current security group (Sys-FullAccess)	Allows instances in the security group to communicate with each other over IPv4 protocols.
Inbound	Allow	IPv6	All	Source: current security group (Sys-FullAccess)	Allows instances in the security group to communicate with each other over IPv6 protocols.
Inbound	Allow	IPv4	All	Source: 0.0.0.0/0	Allows all inbound data packets to pass through over IPv4 protocols.
Inbound	Allow	IPv6	All	Source address::/0	Allows all inbound data packets to pass through over IPv6 protocols.

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows access from instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows access from instances in the security group to any IPv6 address over any port.

7.2.3 Security Group Configuration Examples

When you create instances, such as cloud servers, containers, and databases, in a VPC subnet, you can use the default security group or create a security group. You can add inbound and outbound rules to the default or your security group to control traffic from and to the instances in the security group. Here are some common security group configuration examples:

- **Remotely Logging In to an ECS from a Local Server**
- **Remotely Connecting to an ECS from a Local Server to Upload or Download Files over FTP**
- **Setting Up a Website on an ECS to Provide Internet-Accessible Services**
- **Using ping Command to Verify Network Connectivity**
- **Enabling Communications Between Instances in Different Security Groups**
- **Allowing External Instances to Access the Database Deployed on an ECS**
- **Allowing ECSs to Access Only Specific External Websites**

NOTICE

If your security group rules are not applied, **submit a service ticket**.

Precautions

Note the following before configuring security group rules:

- Instances associated with different security groups are isolated from each other by default.
- Generally, a security group denies all external requests by default, while allowing instances in it to communicate with each other.

  If required, you can add inbound rules to allow specific traffic to access the instances in the security group.
- By default, outbound security group rules allow all requests from the instances in the security group to access external resources.

  If outbound rules are deleted, the instances in the security group cannot communicate with external resources. To allow outbound traffic, you need to add outbound rules by referring to **Table 7-8**.

Table 7-8 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access 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 in a security group from a local server, add an inbound rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable port 22. For details, see **Table 7-9**.
- To remotely log in to a Windows ECS using RDP, enable port 3389. 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, all external IP addresses are allowed to remotely log in to the ECS. To ensure network security and prevent service interruptions caused by network intrusions, set the source to a trusted IP address. For details, see **Table 7-11**.

Table 7-11 Remotely logging in to an ECS using a trusted 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 over FTP

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 over FTP, you need to enable FTP ports 20 and 21.

Table 7-12 Remotely connecting to an ECS from any server to upload or download files over FTP

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 0.0.0.0/0

NOTICE

- If the source is set to 0.0.0.0/0, all external IP addresses are allowed to remotely log in to the ECS to upload or download files. To ensure network security and prevent service interruptions caused by network intrusions, set the source to a trusted IP address. For details, see **Table 7-13**.
- You must first install the FTP server program on the ECSs and then check whether ports 20 and 21 are working properly.

Table 7-13 Remotely connecting to an ECS from a trusted server to upload or download files

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 192.168.0.0/24

Setting Up a Website on an ECS to Provide Internet-Accessible Services

A security group denies all external requests by default. If you set up a website on an ECS to allow access from the Internet, 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 7-14 Setting up a website on an ECS to provide internet-accessible services

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

Ping works by sending an Internet Control Message Protocol (ICMP) Echo Request. To ping an ECS from your PC to verify the network connectivity, you need to add an inbound rule to the security group of the ECS to allow ICMP traffic.

Table 7-15 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 Communications Between Instances in Different Security Groups

Instances in the same VPC but in different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** 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**.

Table 7-16 Enabling communications between instances in different security groups

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A

NOTICE

In the example in "Allowing Traffic from a Virtual IP Address" in **How Security Groups Are Used**, two ECSs in Subnet-A and Subnet-B are connected by a virtual IP address. If you set the source of inbound rules to the security groups associated with the ECSs, the ECSs in the two security groups cannot communicate with each other. To enable communications between them, set the source to the private IP address or subnet CIDR block of the virtual IP address.

Allowing External Instances to Access the Database Deployed on an ECS

A security group denies all external requests by default. If you have deployed a database on an ECS and want the database to be accessed from external instances on a private network, you need to add an inbound rule to the security group of the ECS to allow access over corresponding ports. Here are some common ports for databases:

- MySQL: port 3306
- Oracle: port 1521
- MS SQL: port 1433
- PostgreSQL: port 5432
- Redis: port 6379

Table 7-17 Allowing external instances to access the database deployed on an ECS

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A	Allows the ECSs in security group **sg-A** to access the MySQL database.
Inbound	1	Allow	IPv4	TCP: 1521	Security group: sg-B	Allows the ECSs in security group **sg-B** to access the Oracle database.
Inbound	1	Allow	IPv4	TCP: 1433	IP address: 172.16.3.21/32	Allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database.
Inbound	1	Allow	IPv4	TCP: 5432	IP address: 192.168.0.0/24	Allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database.

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP: 6379	IP address group: ipGroup-A	Allows ECSs whose private IP addresses are in IP address group **ipGroup-A** to access the Redis database.

NOTICE

In this example, the source IP addresses are for reference only. Replace them with actual IP addresses.

Allowing ECSs to Access Only Specific External Websites

By default, a security group allows all outbound traffic. **Table 7-19** lists the default outbound rules. If you want to allow ECSs to access only specific websites, configure the security group as follows:

1. Add outbound rules to only allow traffic over specific ports to specific IP addresses.

Table 7-18 Allowing ECSs to access only specific external websites

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	TCP: 80	IP address: 132.15.XX.XX	Allows ECSs in the security group to access the external website at http://132.15.XX.XX:80.
Outbound	1	Allow	IPv4	TCP: 443	IP address: 145.117.XX.XX	Allows ECSs in the security group to access the external website at https://145.117.XX.XX:443.

2. Delete the default outbound rules that allow all traffic.

Table 7-19 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access 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 7-2 Add Inbound Rule

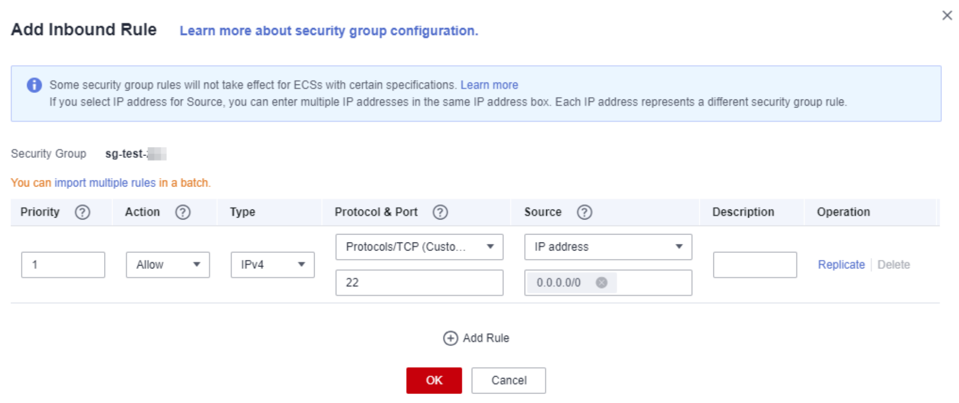

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.

Security Group sg-test-[redacted]

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1	Allow	IPv4	Protocols/TCP (Custo... 22	IP address 0.0.0.0/0		Replicate | Delete

Table 7-20 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
Action	The value can be **Allow** or **Deny**. ● If the **Action** is set to **Allow**, traffic is allowed to access the cloud servers in the security group over specified ports. ● If the **Action** is set to **Deny**, traffic is denied to access the cloud servers in the security group over specified ports. Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see **How Traffic Matches Security Group Rules**.	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. The protocol can be **All**, **TCP**, **UDP**, **GRE**, or **ICMP**.	TCP

Parameter	Description	Example Value
	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**. ● All ports: Leave it empty or enter **1-65535**.	22, 22-30
Source	Source of the security group rule. The value can be an IP address or a security group, to allow access from the IP address or the instances in the security group. ● IP address – Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) – All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) – IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) If the source is a security group, this rule will apply to all instances associated with the selected security group.	192.168.0.0/24
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 7-3 Add Outbound Rule

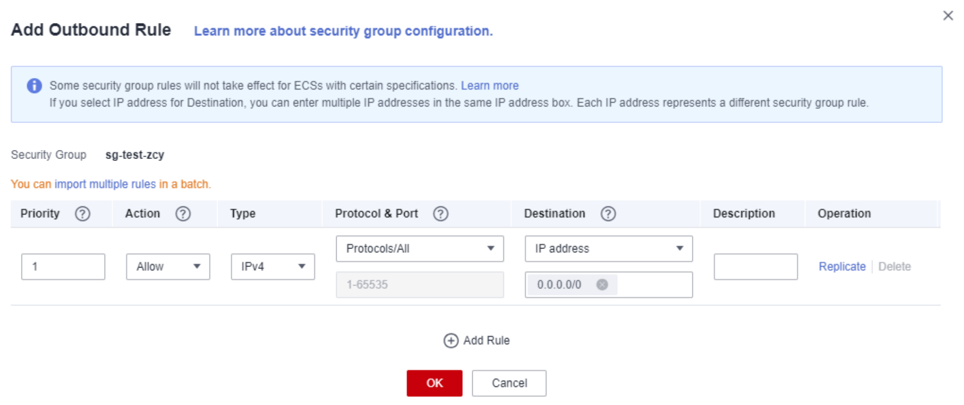

Add Outbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Destination, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.

Security Group sg-test-zcy

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Destination	Description	Operation
1	Allow	IPv4	Protocols/All 1-65535	IP address 0.0.0.0/0		Replicate | Delete

Table 7-21 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	The value can be **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. Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see **How Traffic Matches Security Group Rules**.	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. The protocol can be **All**, **TCP**, **UDP**, **GRE**, or **ICMP**.	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**. ● All ports: Leave it empty or enter **1-65535**.	22, 22-30
Destination	Destination of the security group rule. The value can be an IP address or a security group, to allow access to the IP addresses or the instances in the security group. ● IP address – Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) – All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) – IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	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

9. Click **OK** to complete the security rule configuration.

Verifying Security Group Rules

After allowing traffic over a port in a security group rule, you need to ensure that the port used by the instance is also opened.

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 7-22** shows the rule.

Table 7-22 Security group rule

Direction	Protocol & Port	Source
Inbound	TCP: 80	IP address: 0.0.0.0/0

After adding the security group rule, perform the following operations to check whether the ECS port is opened and whether the rule is applied:

1. Log in to the ECS and check whether the ECS port is opened.
   - **Checking the port of a Linux server**

     Run the following command to check whether TCP port 80 is being listened on:

     netstat -an | grep 80

     If the following figure is displayed, TCP port 80 is enabled.

     Figure 7-4 Command output for the Linux ECS

     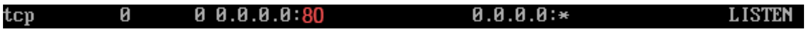

     ```
     tcp        0      0 0.0.0.0:80              0.0.0.0:*               LISTEN
     ```

   - **Checking the port of a Windows server**
     1. Choose **Start** > **Run**. Type **cmd** to open the Command Prompt.
     2. Run the following command to check whether TCP port 80 is being listened on:

        netstat -an | findstr 80

        If the following figure is displayed, TCP port 80 is enabled.

        Figure 7-5 Command output for the Windows ECS

        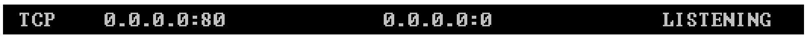

        ```
          TCP    0.0.0.0:80             0.0.0.0:0              LISTENING
        ```

2. 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, choose **More** > **Manage Network** > **Change Security Group** in the **Operation** column.

   The **Change Security Group** dialog box is displayed.
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 apply to 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) helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.

How Do I Use HSS?

Before using the HSS service, install the agent on your ECS.

- **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 the Agent on Huawei Cloud Servers** 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** > **Host Security Service**.
3. Go to the **Servers** page to view the protection status of the target servers.

Table 7-23 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.
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 step **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 more details, see **Enterprise Management User Guide**.

8 Backup Using CBR

8.1 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 is lost or inconsistent 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 initial data, and these	- 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
		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 only be rolled back to its source disk. Rolling back to another disk is not supported. ● 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 8-1 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.

8.2 Backing Up an ECS

Scenarios

Cloud Backup and Recovery (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, see **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 target ECS and choose **More** > **Manage Image/ Disk/Backup** > **Create Server Backup**.
   - 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 target ECS and choose **More** > **Manage Image/Disk/Backup** > **Create Disk Backup**.
    - 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 of the CBR console, 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**.

9 Passwords and Key Pairs

9.1 Password Reset

9.1.1 Application Scenarios for Using Passwords

Passwords are used to log in to ECSs. If you select the password login mode when purchasing an ECS, you can use the username and password to log in to your ECS. The password is very important. Keep it secure.

You can reset the password when:

- You forgot the password.
- The password has expired.
- You selected **Set password later** during the ECS purchase.

Table 9-1 provides guidance on how to reset your password in different scenarios.

Table 9-1 Resetting a password

Scenario	Prerequisites
Resetting the Password for Logging In to an ECS in the OS	N/A
Resetting the Password for Logging In to a Linux ECS	The password reset plug-in has not been installed.

Background

Table 9-2 shows the ECS password complexity requirements.

Table 9-2 Password complexity requirements

Parameter	Requirement
Password	● Consists of 8 to 26 characters. ● Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Linux ECSs: !@%-_=+[]:./^,{}? ● Cannot contain the username or the username spelled backwards.

9.1.2 Resetting the Password for Logging In to an ECS in the OS

Scenarios

This section describes how to reset the password for logging in to an ECS in the OS when the password is about to expire, the password is forgotten, or you are logging in to the ECS for the first time. It is a good practice to change the initial password upon the first login.

Prerequisites

The ECS can be logged in.

Background

Table 9-3 shows the ECS password complexity requirements.

Table 9-3 Password complexity requirements

Parameter	Requirement
Password	● Consists of 8 to 26 characters. ● Contains at least three of the following character types: – Uppercase letters – Lowercase letters – Digits – Special characters for Linux ECSs: !@%-_=+[]:./^,{}? ● Cannot contain the username or the username spelled backwards.

Linux

1. Use the existing key file to log in to the ECS as user **root** through SSH.

   For details, see **Login Using an SSH Key**.

2. Run the following command to reset the password of user **root**:

   passwd

   To reset the password of another user, replace **passwd** with **passwd username**.

3. Enter a new password that meets the requirements listed in **Table 9-3** as prompted.

   ```
   New password:
   Retype new password:
   ```

   If the following information is displayed, the password has been changed:

   ```
   passwd: all authentication tokens updates successfully
   ```

9.1.3 Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed

Scenarios

You can reset your ECS password if:

- The password is forgotten.
- The password has expired.

The method described in this section can only be used to change the password of a local Windows account, but not the password of a domain account.

Prerequisites

- A temporary Linux ECS running Ubuntu 14.04 or later is available. It is located in the same AZ and has the same CPU architecture as the target ECS.

  NOTE

  - Currently, this operation can be performed only for Ubuntu 16.04 and Ubuntu 18.04 public images.
  - You can select an existing ECS or purchase a temporary ECS.

    After the password of the purchased ECS is reset, you are advised to delete the ECS to avoid additional billing.

- You have bound an EIP to the temporary ECS and configured the apt-get source.
- You have used either of the following methods to install **ntfs-3g** and **chntpw** software packages on the temporary ECS:

  Method 1:

  Run the following command to install the **ntfs-3g** and **chntpw** software packages:

  sudo apt-get install ntfs-3g chntpw

  Method 2:

  Download the ntfs-3g and chntpw software packages of the version required by the temporary ECS OS.

Procedure

1. Stop the original ECS and detach the system disk.
   a. Log in to the management console.
   b. Click 📍 in the upper left corner and select your region and project.
   c. Click ☰ . Under **Compute**, click **Elastic Cloud Server**.
   d. Stop the original Windows ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

      > NOTE
      >
      > Do not forcibly stop the Windows ECS. Otherwise, password reset may fail.

   e. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
2. Attach the system disk to the temporary ECS.
   a. On the temporary ECS details page, click the **Disks** tab.
   b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **1.e** and attach it to the temporary ECS.
   c. Remotely log in to the temporary ECS.
   d. Run the following command to view the directory of the system disk detached from the original Windows ECS now attached to the temporary ECS:

      fdisk -l

   e. Run the following command to mount the file system of the detached system disk to the temporary ECS:

      mount -t ntfs-3g /dev/*Result obtained in step* ***2.d*** **/mnt/**

      For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

      mount -t ntfs-3g /dev/xvde2 /mnt/

      If the following error information is displayed after the preceding command is executed, the NTFS file systems may be inconsistent. In such a case, rectify the file system inconsistency.

      ```
      The disk contains an unclean file system (0, 0).
      Metadata kept in Windows cache, refused to mount.
      Failed to mount '/dev/xvde2': Operation not permitted
      The NTFS partition is in an unsafe state. Please resume and shutdown
      Windows fully (no hibernation or fast restarting), or mount the volume
      read-only with the 'ro' mount option.
      ```

      Back up the disk data, run the following command to rectify the NTFS file system inconsistency, and attach the system disk:

      ntfsfix /dev/*Result obtained in step* ***2.d***

      For example, if the result obtained in step **2.d** is **xvde2**, run the following command:

      ntfsfix /dev/xvde2
3. Change the password of the specified user and clear the original password.
   a. Run the following command to back up the SAM file:

cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/config/SAM.bak

b. Run the following command to change the password of the specified user:

chntpw -u Administrator /mnt/Windows/System32/config/SAM

c. Enter **1**, **q**, and **y** as prompted, and press **Enter**.

The password has been reset if the following information is displayed:

```
Select: [q] > 1
Password cleared!
Select: [q] > q
Hives that have changed:
#Name
0<SAM>
Write hive files? (y/n) [n] : y
0<SAM> - OK
```

4. Stop the temporary ECS, detach the system disk, and attach the system disk to the original Windows ECS.

   a. Stop the temporary ECS, go to the ECS details page, and click the **Disks** tab.

   b. Click **Detach** to detach the data disk temporarily attached in step **2.b**.

   c. On the original Windows ECS details page, click the **Disks** tab.

   d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in step **4.b** and attach it to the original ECS as the system disk.

5. Start the original Windows ECS and set a new login password.

   a. Click **Start** to start the original Windows ECS. After the status becomes **Running**, click **Remote Login** in the **Operation** column.

   b. Click **Start**. Enter **CMD** in the search box and press **Enter**.

   c. Run the following command to set a new password. The new password must meet the password complexity requirements described in **Application Scenarios for Using Passwords**.

   net user Administrator *New password*

9.1.4 Resetting the Password for Logging In to a Linux ECS

Scenarios

Keep your password secure. Reset the password if:

- The password is forgotten.
- The password has expired.

This section describes how to reset the password of user **root**. After resetting the password, you can log in to the ECS, and change the private key or reset the password of a non-**root** user.

Prerequisites

- A temporary Linux ECS is available. It is located in the same AZ and has the same CPU architecture as the target ECS.

NOTE

You can select an existing ECS or purchase a temporary ECS.

After the password of the purchased ECS is reset, you are advised to delete the ECS to avoid additional billing.

- You have bound an EIP to the temporary ECS.

Procedure

1. Download the script for resetting the password and upload the script to the temporary ECS.

   Download and decompress the password reset script. Use a connection tool, such as WinSCP, to upload the obtained **changepasswd.sh** script to the temporary ECS.

   To download WinSCP, log in at **https://winscp.net/**.

2. Stop the original Linux ECS, detach the system disk from it, and attach the system disk to the temporary ECS.

   a. Log in to the management console.

   b. Click 📍 in the upper left corner and select your region and project.

   c. Click ☰ . Under **Compute**, click **Elastic Cloud Server**.

   d. Stop the original ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

      NOTE

      Do not forcibly stop the original ECS. Otherwise, password reset may fail.

   e. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.

3. Attach the system disk to the temporary ECS.

   a. On the page providing details about the temporary ECS, click the **Disks** tab.

   b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step **2.e** and attach it to the temporary ECS.

4. Log in to the temporary ECS remotely and reset the password.

   a. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.

   b. Run the following command to view the directory of the system disk detached from the original Linux ECS now attached to the temporary ECS:

      fdisk -l

Figure 9-1 Viewing the directory of the system disk

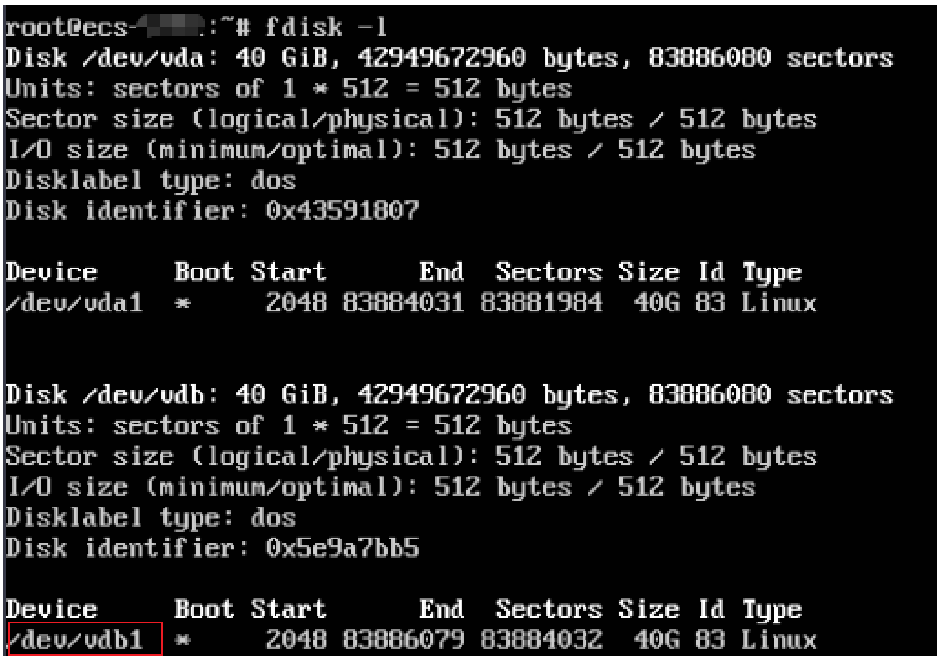

```
root@ecs-[redacted]:~# fdisk -l
Disk /dev/vda: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x43591807

Device     Boot Start      End  Sectors Size Id Type
/dev/vda1  *     2048 83884031 83881984  40G 83 Linux


Disk /dev/vdb: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x5e9a7bb5

Device     Boot Start      End  Sectors Size Id Type
/dev/vdb1  *     2048 83886079 83884032  40G 83 Linux
```

c. Run the following commands in the directory where the **changepasswd.sh** script is stored to run the script for resetting the password:

chmod +x changepasswd.sh

./changepasswd.sh

When you run the password reset script, if the system displays a message indicating that there is no command related to logical volume manager (LVM), such as the message "no lvs command", install an LVM tool on the temporary ECS. The LVM2 tool is recommended, which can be installed by running the **yum install lvm2** command.

NOTE

If the original ECS and the temporary ECS both run CentOS 7, a mount failure may occur during script execution. To resolve this issue, replace **mount $dev $mountPath** with **mount -o nouuid $dev $mountPath** in the script.

d. Enter the new password and the directory obtained in step **4.b** as prompted.

If the following information is displayed, the password has been changed:

```
set password success.
```

5. (Optional) Enable remote root login for non-root users.

vi /etc/ssh/sshd_config

Modify the following settings:

- Change **PasswordAuthentication no** to **PasswordAuthentication yes**.

  Alternatively, uncomment **PasswordAuthentication yes**.

- Change **PermitRootLogin no** to **PermitRootLogin yes**.

  Alternatively, uncomment **PermitRootLogin yes**.

- Change the value of **AllowUsers** to **root**.

Search for **AllowUsers** in the file. If **AllowUsers** is missing, add **AllowUsers root** at the end of the file.

6. Stop the temporary ECS, detach the system disk, attach the system disk to the original Linux ECS, and restart the original Linux ECS.
   a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
   b. Click **Detach** to detach the data disk temporarily attached in step **3**.
   c. On the page providing details about the original Linux ECS, click the **Disks** tab.
   d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in **6.b**.
7. Restart the original Linux ECS.

9.2 Key Pairs

9.2.1 Application Scenarios for Using Key Pairs

Key Pairs

Key pairs (SSH 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, rather than a password, to log in to the ECS. 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.

Scenarios

When purchasing an ECS, you are advised to select the key pair login mode.

- Logging in to a Linux ECS

  You can directly use a key pair to log in a Linux ECS.
  - During the ECS creation, select the key pair login mode. For details, see "Set **Login Mode**" in **Step 3: Configure Advanced Settings**.

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 methods:
  - 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 PuTTY Key Generator**. 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 your public key.

  > NOTE
  >
  > If the public key of the existing key pair is stored by clicking **Save public key** on puttygen.exe, 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?**

Notes and 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.

9.2.2 (Recommended) Creating a Key Pair on the Management Console

Scenarios

You can create a key pair on the management console. After the key pair is created, the public key is automatically stored in the system, and the private key is stored in your local computer. After a key pair is created for an ECS on the management console, ensure that you store your private key in a secure place. Without a private key, you will not be able to log in to 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 navigation pane on the left, choose **Key Pair**.

5. On the displayed page, click **Create Key Pair**.
6. Enter a key pair name.

   A key pair name consists of two parts: KeyPair and four random digits (KeyPair-xxxx).
7. Click **OK**.
8. Manually or automatically download a .pem private key file with the name that you specify as the key name. Store it in a secure place and click **OK**.

   NOTE

   This is the only chance for you to save the private key file. Keep it secure. You'll need to provide the key pair name when you create an ECS, and the corresponding private key each time you connect to the ECS through SSH.

9.2.3 Creating a Key Pair Using PuTTY Key Generator

Scenarios

You can use the third-party tool puttygen.exe to create a key pair. After the key pair is created, both the public key and private key are stored 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 9-2 PuTTY Key Generator

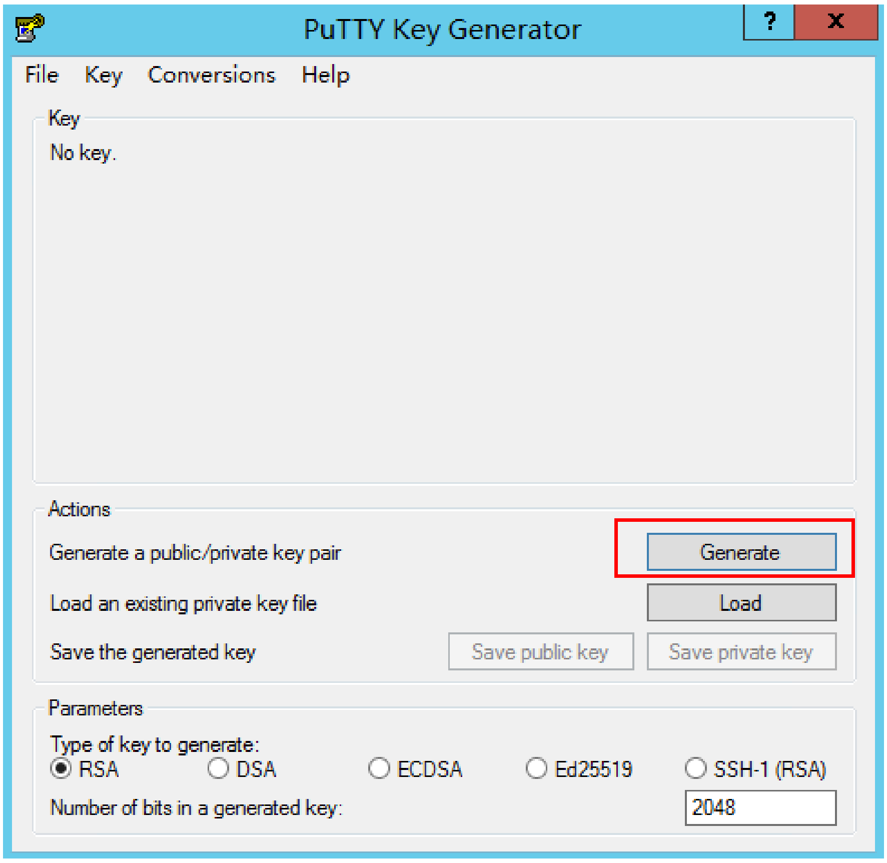

No key.

Type of key to generate:
● RSA ○ DSA ○ ECDSA ○ Ed25519 ○ SSH-1 (RSA)

Number of bits in a generated key: 2048

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 9-3** is the public key.

Figure 9-3 Generating the public and private keys

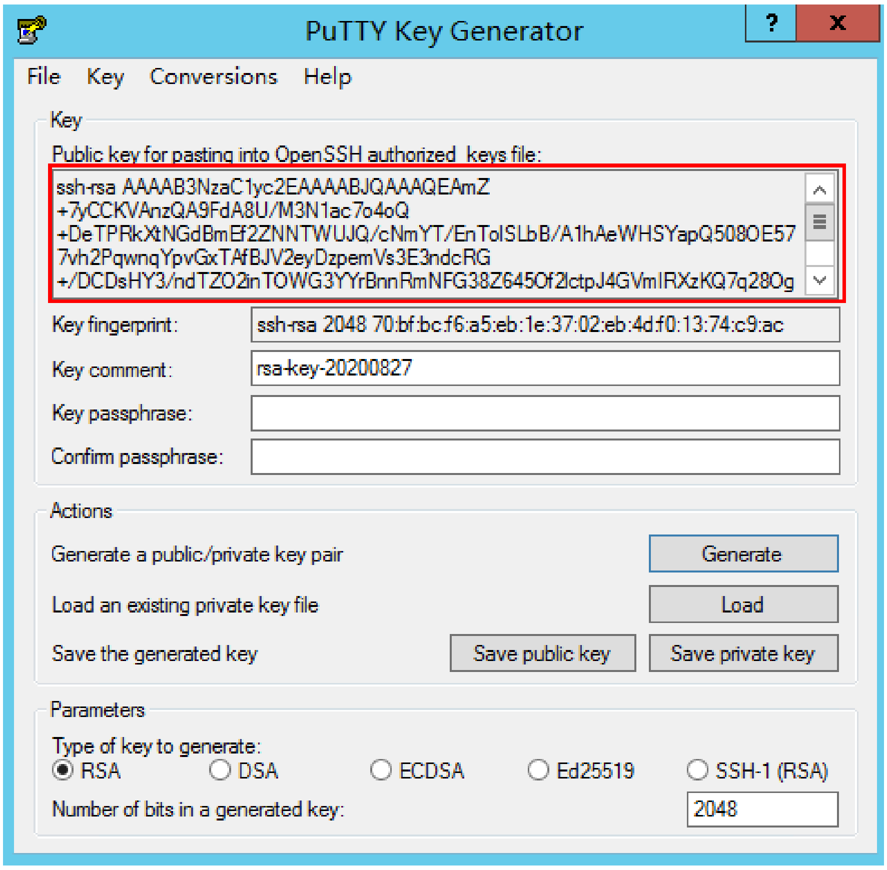

Public key for pasting into OpenSSH authorized_keys file:

```
ssh-rsa AAAAB3NzaC1yc2EAAAABJQAAAQEAmZ
+7yCCKVAnzQA9FdA8U/M3N1ac7o4oQ
+DeTPRkXtNGdBmEf2ZNNTWUJQ/cNmYT/EnToISLbB/A1hAeWHSYapQ508OE57
7vh2PqwnqYpvGxTAfBJV2eyDzpemVs3E3ndcRG
+/DCDsHY3/ndTZO2inTOWG3YYrBnnRmNFG38Z645Of2lctpJ4GVmIRXzKQ7q28Og
```

Key fingerprint: ssh-rsa 2048 70:bf:bc:f6:a5:eb:1e:37:02:eb:4d:f0:13:74:c9:ac

Key comment: rsa-key-20200827

Key passphrase:

Confirm passphrase:

Generate a public/private key pair

Load an existing private key file

Save the generated key

Type of key to generate:
● RSA ○ DSA ○ ECDSA ○ Ed25519 ○ SSH-1 (RSA)

Number of bits in a generated key: 2048

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 Key Generator 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 9-4 Saving a private key

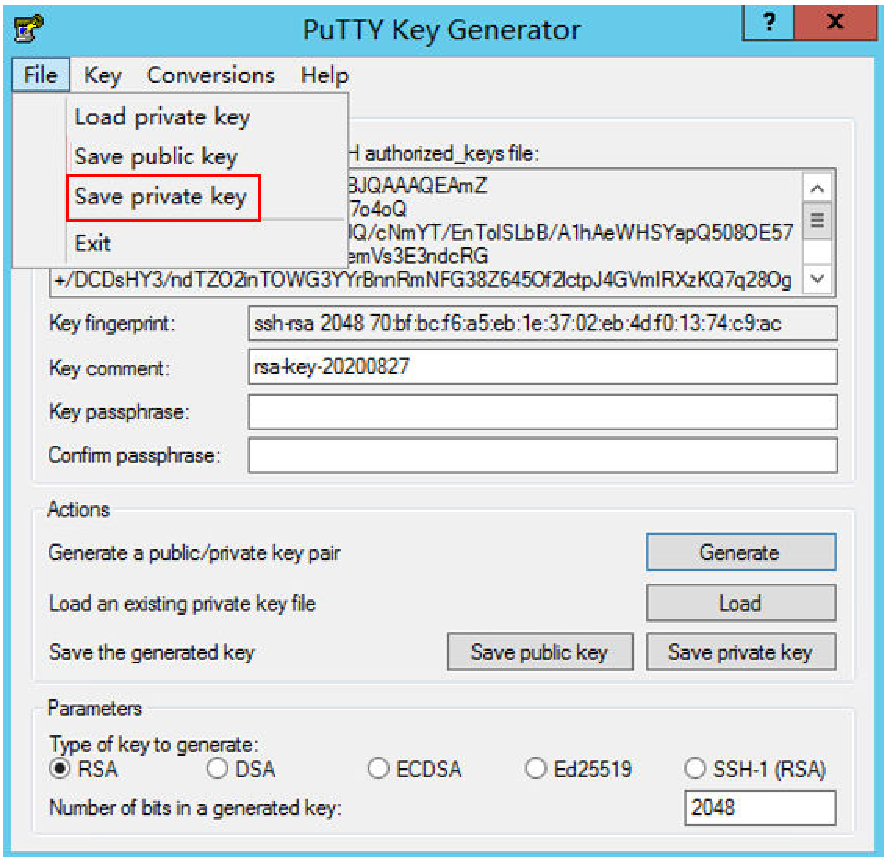

H authorized_keys file:

```
BJQAAAQEAmZ
7o4oQ
IQ/cNmYT/EnTolSLbB/A1hAeWHSYapQ508OE57
emVs3E3ndcRG
+/DCDsHY3/ndTZO2inTOWG3YYrBnnRmNFG38Z645Of2lctpJ4GVmIRXzKQ7q28Og
```

Key fingerprint: ssh-rsa 2048 70:bf:bc:f6:a5:eb:1e:37:02:eb:4d:f0:13:74:c9:ac

Key comment: rsa-key-20200827

Key passphrase:

Confirm passphrase:

Generate a public/private key pair

Load an existing private key file

Save the generated key

Type of key to generate:
● RSA ○ DSA ○ ECDSA ○ Ed25519 ○ SSH-1 (RSA)

Number of bits in a generated key: 2048

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:

5. After you have saved the key pair, import your public key to the ECS by referring to **Importing a Key Pair**.

9.2.4 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
     1. 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 PuTTY Key Generator**).

        NOTE

        Make sure that the file to be imported is a public key file.
     2. Click **OK**.

        After the public key is imported, you can change its name.
   - Copying the public key content
     1. Copy the public key content from the locally stored .txt file into the **Public Key Content** text box.
     2. 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?**

9.2.5 Obtaining and Deleting the Password of a Windows ECS

9.2.5.1 Obtaining the Password for Logging In to a Windows ECS

Scenarios

Password authentication is required to log in to a Windows ECS. You must use the private key bound to the ECS when the ECS was created to obtain the administrator password generated during the 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.

Prerequisites

You have obtained the private key file (.pem file) which was generated during the ECS creation.

When you selected the key pair login mode during the ECS creation, a private key file (.pem file) was generated during the creation of the key pair. For details about how to create and use a key pair, see **Application Scenarios for Using Key Pairs**.

If the private key file is lost, you can **reset the key pair** and bind it to the ECS. If you select **I agree to host the private key of the key pair**, you can export the managed private key as required. For details, see **Exporting a Private Key**.

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**.

   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 .
7. Use either of the following methods to obtain the password through the private key:
   - Click **Select File** and upload the private key from a local directory.
   - Copy the content of the private key file and paste it into the text box.
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 the *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/{*project_id*}/servers/{*server_id*}/os-server-password".

   NOTE

   For details, see "Obtaining the Password for Logging In to an ECS" in the *Elastic Cloud Server 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.

9.2.5.2 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.

10 Resources and Tags

10.1 Tag Management

10.1.1 Overview

Scenarios

A tag identifies an ECS. Adding tags to an ECS facilitates ECS identification and management.

You can add tags to an ECS either during or after the ECS creation. A maximum of 10 tags can be added to an ECS.

NOTE

Tags added during the ECS creation will also be added to the EIP and EVS disks (including the system disk and data disks) of the ECS. If the ECS uses an existing EIP, the tags will not be added to the EIP.

After creating the ECS, you can view the tags on the pages providing details about the ECS, EIP, and EVS disks.

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, usage, owner, or environment).

Figure 10-1 Example tags

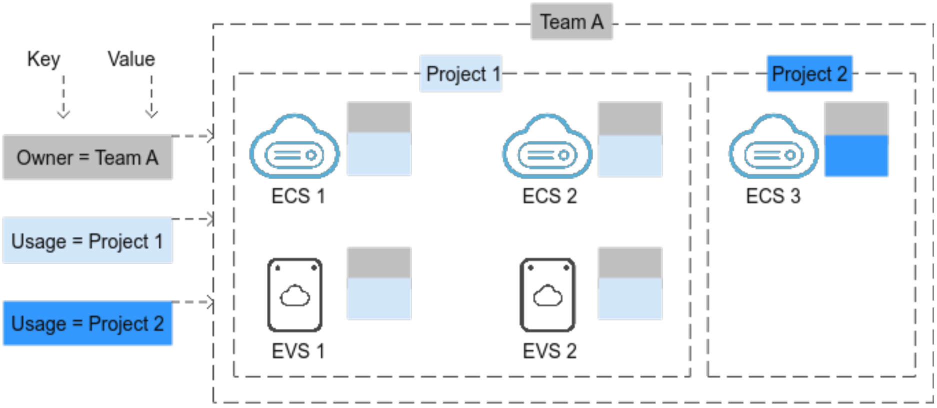

Figure 10-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Usage**. Each tag has a value.

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 10-1** lists the tag key and value requirements.

Table 10-1 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

10.1.2 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 10-1**.

   NOTE

   For details about other parameters, see **Purchasing an ECS**.

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 10-1**.
6. Click **OK**.

   After tags are added, you can click **Edit** in the **Operation** column to edit them.

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.
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.
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.

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 **Compute** > **Elastic Cloud Server** from the service list and select the predefined tag keys and values.

10.1.3 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.

10.1.4 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 tag to be deleted, and click **Delete** in the **Operation** column.
6. Click **OK**.

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.
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.

10.2 Quota Adjustment

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECS 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.
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.
3. Click **Increase Quota** in the upper right corner of the page.
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**.

11 Monitoring Using Cloud Eye

11.1 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 is classified into 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**.

11.2 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 11-1**.

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 Linux ECS, see "Installing and Configuring the Agent (Linux)" in *Cloud Eye User Guide*.

Table 11-1 Basic ECS metrics

Metric	LinuxECS	
None	Xen	KVM
CPU Usage	Supported	Supported
Memory Usage	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Disk Usage	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Disk Read Bandwidth	Supported	Supported
Disk Write Bandwidth	Supported	Supported

Metric	LinuxECS	
Disk Read IOPS	Supported	Supported
Disk Write IOPS	Supported	Supported
Inband Incoming Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Inband Outgoing Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Outband Incoming Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband incoming rate.)	Supported
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
Inbound Bandwidth	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
Outbound Bandwidth	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
Inbound PPS	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported

Metric	LinuxECS	
Outbound PPS	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
New Connections	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported

Table 11-2 describes these basic ECS metrics.

The monitoring intervals for the following ECSs with raw monitoring metrics are as follows:

- Xen ECSs: 4 minutes
- KVM ECSs: 5 minutes

Table 11-2 Basic metric description

Metric ID	Parameter	Description	Value Range	Unit	Conversion Rule	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?** Formula: CPU usage of an ECS/Number of vCPUs in the ECS	0-100	%	N/A	ECS	5 minutes
mem_util	Memory Usage	Memory usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Formula: Used memory of an ECS/Total memory of the ECS **NOTE** The memory usage of QingTian ECSs cannot be monitored.	0-100	%	N/A	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
disk_util_inband	Disk Usage	Disk usage of an ECS This metric is unavailable if the image has no UVP VMTools installed. Formula: Used capacity of an ECS-attached disk/Total capacity of the ECS-attached disk	0-100	%	N/A	ECS	5 minutes
disk_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from an ECS-attached disk per second Formula: Total number of bytes read from an ECS-attached disk/ Monitoring interval byte_out = (rd_bytes - last_rd_bytes)/ Time difference	≥ 0	byte/s	1024(IEC)	ECS	5 minutes
disk_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to an ECS-attached disk per second Formula: Total number of bytes written to an ECS-attached disk/ Monitoring interval	≥ 0	byte/s	1024(IEC)	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
disk_read_requests_rate	Disk Read IOPS	Number of read requests sent to an ECS-attached disk per second Formula: Total number of read requests sent to an ECS-attached disk/ Monitoring interval req_out = (rd_req - last_rd_req)/Time difference	≥ 0	Request/s	N/A	ECS	5 minutes
disk_write_requests_rate	Disk Write IOPS	Number of write requests sent to an ECS-attached disk per second Formula: Total number of write requests sent to an ECS-attached disk/ Monitoring interval req_in = (wr_req - last_wr_req)/Time difference	≥ 0	Request/s	N/A	ECS	5 minutes
network_incoming_bytes_rate_inband	Inband Incoming Rate	Number of incoming bytes on an ECS per second Formula: Total number of inband incoming bytes on an ECS/Monitoring interval	≥ 0	byte/s	1024(IEC)	ECS	5 minutes
network_outgoing_bytes_rate_inband	Inband Outgoing Rate	Number of outgoing bytes on an ECS per second Formula: Total number of inband outgoing bytes on an ECS/Monitoring interval	≥ 0	byte/s	1024(IEC)	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_incoming_bytes_aggregate_rate	Outband Incoming Rate	Number of incoming bytes on an ECS per second on the hypervisor Formula: Total number of outband incoming bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	byte/s	1024(IEC)	ECS	5 minutes
network_outgoing_bytes_aggregate_rate	Outband Outgoing Rate	Number of outgoing bytes on an ECS per second on the hypervisor Formula: Total number of outband outgoing bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	byte/s	1024(IEC)	ECS	5 minutes
network_vm_connections	Network Connections	Total number of TCP and UDP connections to an ECS **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	Count	N/A	ECS	5 minutes
network_vm_bandwidth_in	Inbound Bandwidth	Number of public and private bits received by the ECS per second	≥ 0	Byte/s	1024(IEC)	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_vm_bandwidth_out	Outbound Bandwidth	Number of public and private bits sent by the ECS per second	≥ 0	Byte/s	1024(IEC)	ECS	5 minutes
network_vm_pps_in	Inbound PPS	Number of public and private packets received by the ECS per second	≥ 0	Packet/s	N/A	ECS	5 minutes
network_vm_pps_out	Outbound PPS	Number of public and private packets sent by the ECS per second	≥ 0	Packet/s	N/A	ECS	5 minutes
network_vm_newconnections	New Connections	Number of new connections (including TCP, UDP, and ICMP) created on the ECS	≥ 0	connect/s	N/A	ECS	5 minutes

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

11.3 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 the 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.

Namespace

AGT.ECS

OS Metrics: CPU

Table 11-3 CPU metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object ● 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	%	N/A	ECS	1 minute
cpu_usage_idle	(Agent) Idle CPU Usage	Percentage of time that CPU is idle ● 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	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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	%	N/A	ECS	1 minute
cpu_usage_system	(Agent) Kernel Space CPU Usage	Percentage of time that the CPU is used by kernel space ● 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	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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	%	N/A	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 ● 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	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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	%	N/A	ECS	1 minute
cpu_usage_irq	(Agent) CPU Interrupt Time	Percentage of time that the CPU is servicing interrupts ● 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	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_softirq	(Agent) CPU Software Interrupt Time	Percentage of time that the CPU is servicing software interrupts ● 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	%	N/A	ECS	1 minute

OS Metric: CPU Load

Table 11-4 CPU load metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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	N/A	N/A	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	N/A	N/A	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	N/A	N/A	ECS	1 minute

OS Metric: Memory

Table 11-5 Memory metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 - 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	GB	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object ● 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	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_free	(Agent) Idle Memory	Amount of memory that is not being used ● Linux: Obtain the metric value from **/proc/meminfo**. ● Windows is not supported currently.	≥0	GB	N/A	ECS	1 minute
mem_buffers	(Agent) Buffer	Amount of memory that is being used for buffers ● 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	GB	N/A	ECS	1 minute
mem_cached	(Agent) Cache	Amount of memory that is being used for file caches ● 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	GB	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
total_open_files	(Agent) Total File Handles	Total handles used by all processes ● Linux: Use the **/proc/{pid}/fd** file to summarize the handles used by all processes. ● Windows is not supported currently.	≥ 0	Count	N/A	ECS	1 minute

OS Metric: Disk

NOTE

- Currently, only physical disks are monitored. The NFS-attached disks cannot be monitored.
- By default, Docker-related mount points are shielded. The prefix of the mount point is as follows:

```
/var/lib/docker;/mnt/paas/kubernetes;/var/lib/mesos
```

Table 11-6 Disk metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks ● 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 (-), periods (.), 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 (-), periods (.), and swung dashes (~).	≥0	GB	N/A	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	Total space on the disks, including used and free ● 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 (-), periods (.), 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 (-), periods (.), and swung dashes (~).	≥0	GB	N/A	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	Used space on the disks ● 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 (-), periods (.), 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 (-), periods (.), and swung dashes (~).	≥0	GB	N/A	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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** ● 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 (-), periods (.), 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 (-), periods (.), and swung dashes (~).	0-100	%	N/A	ECS - Mount point	1 minute

OS Metric: Disk I/O

Table 11-7 Disk I/O metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows: – The disk I/O data is obtained through the Win32_PerfFormattedData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The	≥ 0	byte/s	1024(IEC)	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
		instantaneous value returned by the object indicates the metric value 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 (-), periods (.), and swung dashes (~). – 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	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows: – The disk I/O data is obtained through the Win32_PerfFormattedData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous	≥ 0	Requests/s	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
		value returned by the object indicates the metric value 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 (-), periods (.), and swung dashes (~). – 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	Unit	Conversion Rule	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 - 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 (-), periods (.), and swung dashes (~). - Windows: – The disk I/O data is obtained through the Win32_PerfFormattedData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous value returned	≥ 0	byte/s	1024(IEC)	- ECS - Disk - ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
		by the object indicates the metric value 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 (-), periods (.), and swung dashes (~). – 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	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows: – The disk I/O data is obtained through the Win32_PerfFormattedData_PerfDisk_LogicalDisk object in WMI. The object is obtained once in each collection period. The instantaneous	≥ 0	Requests/s	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
		value returned by the object indicates the metric value 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 (-), periods (.), and swung dashes (~). – 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	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_readTime	(Agent) Average Read Request Time	Average amount of time that read requests have waited on the disks ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	ms/Count	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	ms/Count	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	0-100	%	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_queue_length	(Agent) Disk Queue Length	This metric reflects the disk usage in a specified period and can be used to evaluate the disk I/O performance. A larger value indicates a busier disk and poorer I/O performance. ● Linux: The metric value is calculated by dividing the data changes in the fourteenth column of the corresponding device in **/proc/diskstats** in a collection period by the metric collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters (case-sensitive), hyphens (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	count	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	Byte/op	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	Byte/op	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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 (-), periods (.), and swung dashes (~). ● Windows is not supported currently.	≥ 0	ms/op	N/A	● ECS - Disk ● ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_device_used_percent	Block Device Usage	Percentage of the physical disk usage of the monitored object. Calculation formula: Used storage space of all mounted disk partitions/Total disk storage space ● Collection method for Linux ECSs: Obtain the disk usage of each mount point, calculate the total disk storage space based on the disk sector size and the number of sectors, and then you can calculate the used storage space in total. ● Windows does not support this metric.	0-100	%	N/A	ECS - Disk	1 minute

OS Metric: File System

Table 11-8 File system metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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. Value: **0** (read and write) or **1** (read only) Linux: Check file system information in the fourth column in file **/proc/mounts**.	- **0**: readable and writable - **1**: read-only	N/A	N/A	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 (-), periods (.), and swung dashes (~).	≥ 0	Count	N/A	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
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 (-), periods (.), and swung dashes (~).	≥ 0	Count	N/A	ECS - Mount point	1 minute
disk_inodesUsedPercent	(Agent) Percentage of Total inode Used	Number of used index nodes on the disk 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 (-), periods (.), and swung dashes (~).	0-100	%	N/A	ECS - Mount point	1 minute

NOTE

The Windows OS does not support the file system metrics.

OS Metric: NIC

Table 11-9 NIC metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_bitRecv	(Agent) Outbound Bandwidth	Number of bits sent by this NIC per second ● 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	1024(IEC)	ECS	1 minute
net_bitSent	(Agent) Inbound Bandwidth	Number of bits received by this NIC per second ● 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	1024(IEC)	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second ● 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	N/A	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second ● 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	N/A	ECS	1 minute
net_errin	(Agent) Receive Error Rate	Percentage of receive errors detected by this NIC per second ● Linux: Check metric value changes in file **/proc/net/dev** in a collection period. ● Windows is not supported currently.	0-100	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_errout	(Agent) Transmit Error Rate	Percentage of transmit errors detected by this NIC per second ● Linux: Check metric value changes in file **/proc/net/dev** in a collection period. ● Windows is not supported currently.	0-100	%	N/A	ECS	1 minute
net_dropin	(Agent) Received Packet Drop Rate	Percentage of packets received by this NIC which were dropped per second ● Linux: Check metric value changes in file **/proc/net/dev** in a collection period. ● Windows is not supported currently.	0-100	%	N/A	ECS	1 minute
net_dropout	(Agent) Transmitted Packet Drop Rate	Percentage of packets transmitted by this NIC which were dropped per second ● Linux: Check metric value changes in file **/proc/net/dev** in a collection period. ● Windows is not supported currently.	0-100	%	N/A	ECS	1 minute

OS Metric: NTP

Table 11-10 NTP metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
ntp_offset	(Agent) NTP Offset	NTP offset of the monitored object ● Collection method for Linux ECSs: Run **chronyc sources -v** to obtain the offset. ● Windows does not support this metric.	≥ 0	ms	N/A	ECS	1 minute

OS Metric: TCP

Table 11-11 TCP metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_total	(Agent) Total TCP Connections	Total number of TCP connections in all states ● 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	Count	N/A	ECS	1 minute
net_tcp_established	(Agent) TCP ESTABLISHED Connection	Number of TCP connections in ESTABLISHED state ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_sys_sent	(Agent) TCP SYS_SENT Connections	Number of TCP connections that are being requested by the client ● 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	Count	N/A	ECS	1 minute
net_tcp_sys_recv	(Agent) TCP SYS_RECV Connections	Number of pending TCP connections received by the server ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_fin_wait1	(Agent) TCP FIN_WAIT1 Connections	Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client ● 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	Count	N/A	ECS	1 minute
net_tcp_fin_wait2	(Agent) TCP FIN_WAIT2 Connections	Number of TCP connections in the FIN_WAIT2 state ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_time_wait	(Agent) TCP TIME_ WAIT Connections	Number of TCP connections in TIME_WAIT state ● 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	Count	N/A	ECS	1 minute
net_tcp_close	(Agent) TCP CLOSE Connections	Number of closed TCP connections ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_close_wait	(Agent) TCP CLOSE_WAIT Connections	Number of TCP connections in CLOSE_WAIT TCP state ● 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	Count	N/A	ECS	1 minute
net_tcp_last_ack	(Agent) TCP LAST_ACK Connections	Number of TCP connections waiting for ACK packets when the connections are being passively closed by the client ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_listen	(Agent) TCP LISTEN Connections	Number of TCP connections in the LISTEN state ● 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	Count	N/A	ECS	1 minute
net_tcp_closing	(Agent) TCP CLOSING Connections	Number of TCP connections to be automatically closed by the server and the client at the same time ● 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	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_retrans	(Agent) TCP Retransmission Rate	Percentage of packets that are resent ● 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	%	N/A	ECS	1 minute

OS Metric: GPU

Table 11-12 GPU metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_status	GPU Health Status	Overall measurement of the GPU health - 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.	N/A	N/A	- ECS - ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_usage_encoder	Encoding Usage	Encoding capability usage of the GPU - 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	%	N/A	- ECS - ECS - GPU	1 minute
gpu_usage_decoder	Decoding Usage	Decoding capability usage of the GPU - 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	%	N/A	- ECS - ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
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. ● 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	Count	N/A	● ECS ● ECS - GPU	1 minute
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. ● 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	Count	N/A	● ECS ● ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_aggregate_correctable	Aggregate Correctable ECC Errors	Aggregate correctable ECC errors on the GPU ● 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	Count	N/A	● ECS ● ECS-GPU	1 minute
gpu_aggregate_uncorrectable	Aggregate Uncorrectable ECC Errors	Aggregate uncorrectable ECC Errors on the GPU ● 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	Count	N/A	● ECS ● ECS-GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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 ● 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	Count	N/A	● 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 ● 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	Count	N/A	● ECS ● ECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_performance_state	(Agent) Performance Status	GPU performance of the monitored object - 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**: indica	N/A	N/A	- ECS - ECS-GPU	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
			tes the unknown performance status.				
gpu_usage_mem	(Agent) GPU Memory Usage	GPU memory usage of the monitored object ● 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	%	N/A	● ECS ● ECS-GPU	1 minute
gpu_usage_gpu	(Agent) GPU Usage	GPU usage of the monitored object ● 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	%	N/A	● ECS ● ECS-GPU	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

11.4 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.

Namespace

AGT.ECS

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 11-13 Process metrics

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_cpu	(Agent) Process CPU Usage	CPU consumed by a process. **pHashId** (process name and process ID) is the value of **md5**. ● Linux: Check metric value changes in file **/proc/pid/stat**.	0–1 x Number of vCPUs.	%	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_mem	(Agent) Process Memory Usage	Memory consumed by a process. **pHashId** (process name and process ID) is the value of **md5**. ● 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**.	0-100	%	N/A	ECS	1 minute
proc_pHashId_file	(Agent) Process 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.	≥0	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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.	≥0	Count	N/A	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.	≥0	Count	N/A	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.	≥0	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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.	≥0	Count	N/A	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.	≥0	Count	N/A	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.	≥0	Count	N/A	ECS	1 minute

Metric	Parameter	Description	Value Range	Unit	Conversion Rule	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.	≥0	Count	N/A	● ECS ● ECS - Process	1 minute
specified_process_file	(Agent)Files Opened by a Process	Number of files opened by a specific process of the monitored object ● Linux: Run the **ls -l /proc/pid/fd** command to view the number of opened files.	≥0	Count	N/A	● ECS ● ECS - Process	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

11.5 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 region, the agent of the latest version is used with simplified monitoring metrics:

EU-Dublin

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 11-14 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.	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.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	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**	0-100	ECS	1 minute
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 (~).	≥ 0	ECS	1 minute
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 (~).	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 (~).	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 (~).	0-100	ECS	1 minute
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.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
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.	≥ 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.	≥ 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.	≥ 0	ECS	1 minute
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.

11.6 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 & Deployment**, 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**.

11.7 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 to search for the target ECS.
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.

12 Audit Using CTS

12.1 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 enabled.

Key ECS Operations Recorded by CTS

Table 12-1 ECS operations recorded by CTS

Operation	Resource Type	Trace
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	Trace
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

12.2 Viewing Traces

Scenarios

After you enable CTS and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, the system starts recording operations on data in Object Storage Service (OBS) buckets. Cloud Trace Service (CTS) stores operation records (traces) generated in the last seven days.

This section describes how to query or export operation records of the last seven days on the CTS console.

- **Viewing Real-Time Traces in the Trace List of the New Edition**
- **Viewing Real-Time Traces in the Trace List of the Old Edition**

Constraints

- You can only query operation records of the last seven days on the CTS console. To store operation records for longer than seven days, you must configure transfer to OBS or Log Tank Service (LTS) so that you can view them in OBS buckets or LTS log groups.
- After performing operations on the cloud, you can query management traces on the CTS console one minute later and query data traces five minutes later.
- Data traces are not displayed in the trace list of the new version. To view them, you need to go to the old version.
- These operation records are retained for seven days on the CTS console and are automatically deleted upon expiration. Manual deletion is not supported.

Viewing Real-Time Traces in the Trace List of the New Edition

1. Log in to the management console.
2. Click ☰ in the upper left corner and choose **Management & Governance** > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. On the **Trace List** page, use advanced search to query traces. You can combine one or more filters.
   - **Trace Name**: Enter a trace name.
   - **Trace ID**: Enter a trace ID.
   - **Resource Name**: Enter a resource name. If the cloud resource involved in the trace does not have a resource name or the corresponding API operation does not involve the resource name parameter, leave this field empty.
   - **Resource ID**: Enter a resource ID. Leave this field empty if the resource has no resource ID or if resource creation failed.
   - **Trace Source**: Select a cloud service name from the drop-down list.
   - **Resource Type**: Select a resource type from the drop-down list.
   - **Operator**: Select one or more operators from the drop-down list.
   - **Trace Status**: Select **normal**, **warning**, or **incident**.
     - **normal**: The operation succeeded.
     - **warning**: The operation failed.
     - **incident**: The operation caused a fault that is more serious than the operation failure, for example, causing other faults.
   - Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range within the last seven days.
5. On the **Trace List** page, you can also export and refresh the trace list, and customize columns to display.
   - Enter any keyword in the search box and press **Enter** to filter desired traces.
   - Click **Export** to export all traces in the query result as an .xlsx file. The file can contain up to 5,000 records.
   - Click ↻ to view the latest information about traces.
   - Click ⚙ to customize the information to be displayed in the trace list. If **Auto wrapping** is enabled (●), excess text will move down to the next line; otherwise, the text will be truncated. By default, this function is disabled.
6. For details about key fields in the trace structure, see **Trace Structure** and **Example Traces**.
7. (Optional) On the **Trace List** page of the new edition, click **Go to Old Edition** in the upper right corner to switch to the **Trace List** page of the old edition.

Viewing Real-Time Traces in the Trace List of the Old Edition

1. Log in to the management console.
2. Click ☰ in the upper left corner and choose **Management & Governance** > **Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. Each time you log in to the CTS console, the new edition is displayed by default. Click **Go to Old Edition** in the upper right corner to switch to the trace list of the old edition.
5. Set filters to search for your desired traces. The following filters are available.
   - **Trace Type**, **Trace Source**, **Resource Type**, and **Search By**: Select a filter from the drop-down list.
     - If you select **Resource ID** for **Search By**, specify a resource ID.
     - If you select **Trace name** for **Search By**, specify a trace name.
     - If you select **Resource name** for **Search By**, specify a resource name.
   - **Operator**: Select a user.
   - **Trace Status**: Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
   - Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range within the last seven days.
6. Click **Query**.
7. On the **Trace List** page, you can also export and refresh the trace list.
   - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5,000 records.
   - Click ↻ to view the latest information about traces.
8. Click ⌄ on the left of a trace to expand its details.

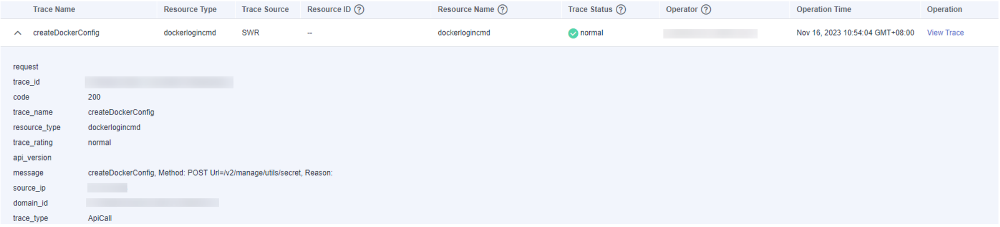

	Trace Name	Resource Type	Trace Source	Resource ID ?	Resource Name ?	Trace Status ?	Operator ?	Operation Time	Operation
^	createDockerConfig	dockerlogincmd	SWR	--	dockerlogincmd	✓ normal	[redacted]	Nov 16, 2023 10:54:04 GMT+08:00	View Trace

request	
trace_id	[redacted]
code	200
trace_name	createDockerConfig
resource_type	dockerlogincmd
trace_rating	normal
api_version	
message	createDockerConfig, Method: POST Url=/v2/manage/utils/secret, Reason:
source_ip	[redacted]
domain_id	[redacted]
trace_type	ApiCall

9. Click **View Trace** in the **Operation** column. The trace details are displayed.

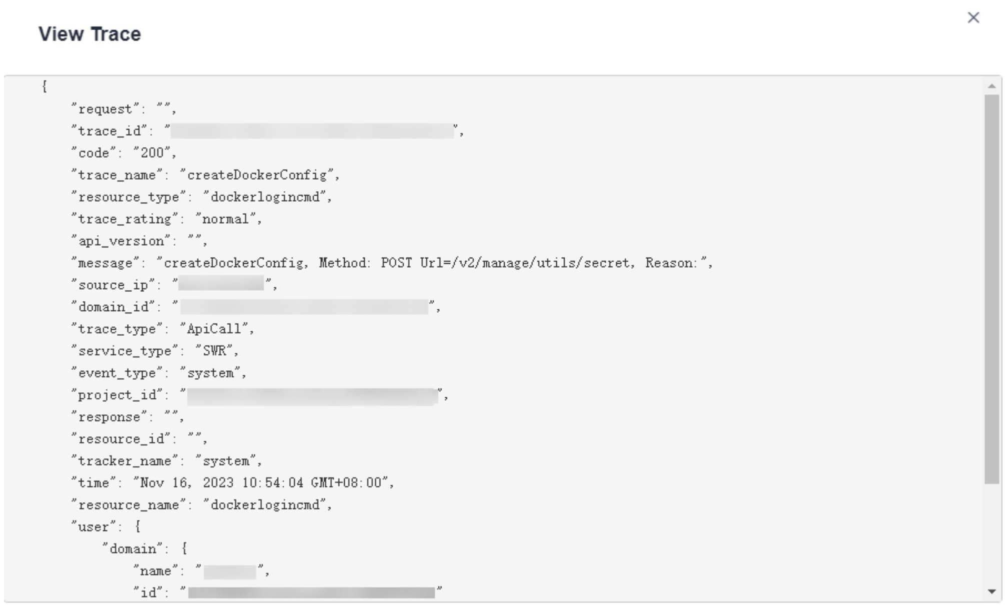

```
{
    "request": "",
    "trace_id": "[redacted]",
    "code": "200",
    "trace_name": "createDockerConfig",
    "resource_type": "dockerlogincmd",
    "trace_rating": "normal",
    "api_version": "",
    "message": "createDockerConfig, Method: POST Url=/v2/manage/utils/secret, Reason:",
    "source_ip": "[redacted]",
    "domain_id": "[redacted]",
    "trace_type": "ApiCall",
    "service_type": "SWR",
    "event_type": "system",
    "project_id": "[redacted]",
    "response": "",
    "resource_id": "",
    "tracker_name": "system",
    "time": "Nov 16, 2023 10:54:04 GMT+08:00",
    "resource_name": "dockerlogincmd",
    "user": {
        "domain": {
            "name": "[redacted]",
            "id": "[redacted]"
```

10. For details about key fields in the trace structure, see **Trace Structure** and **Example Traces** in the *CTS User Guide*.

11. (Optional) On the **Trace List** page of the old edition, click **New Edition** in the upper right corner to switch to the **Trace List** page of the new edition.

13 QingTian Enclave Management

13.1 QingTian Enclave Overview

13.1.1 What Is QingTian Enclave?

- QingTian Enclave instances are secure and isolated virtual machines (VMs) using the QingTian architecture. The instance that has the ownership of QingTian Enclave instances is called the parent instance. QingTian Enclave instances are completely independent VMs and have no persistent storage, interactive access, or external networking. They communicate with the parent instance through a secure local channel, which is called vsock. Even the **root** user of the parent instance cannot access or SSH into QingTian Enclave instances.
- The QingTian Hypervisor isolates the vCPUs and memory of QingTian Enclave instances from the parent instance to provide an isolated environment and greatly reduce the attack surface area. QingTian Enclave helps you protect sensitive core data and applications and enhance the security of your services running in QingTian Enclave instances.
- QingTian Enclave also supports attestation that allows you to verify the trusted measurements of QingTian Enclave instances. Huawei Cloud **Key Management Service (KMS)** provides built-in support for attestation to only allow applications in specific QingTian Enclave instances to be able to call KMS APIs for sensitive data processing.

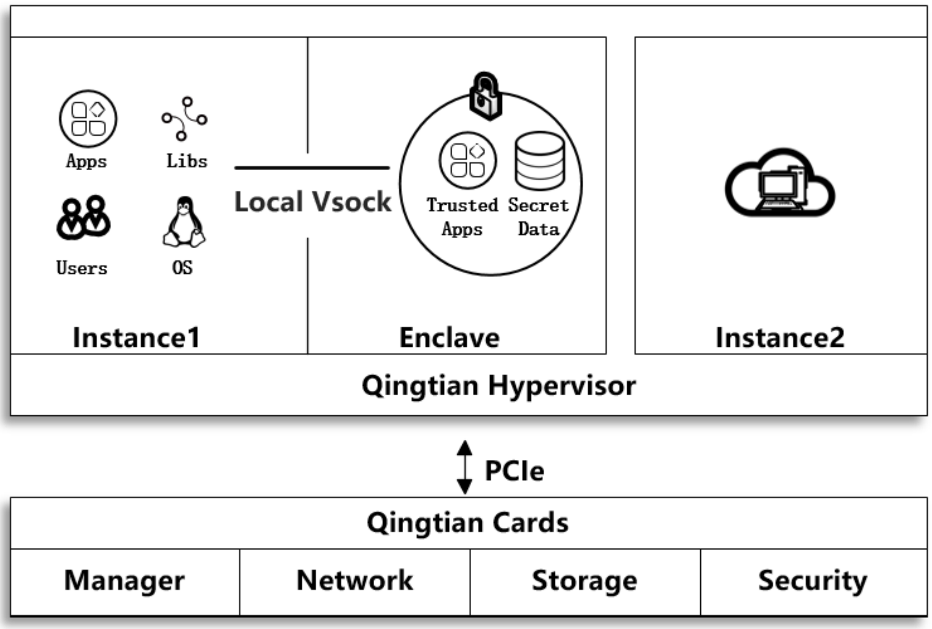

Constraints

QingTian Enclave instances have the following constraints.

Name	Constraints
Parent instance (primary VM)	1. At least two vCPUs and 512 MiB of memory are required. 2. Only the Linux OS is supported.

Name	Constraints
QingTian Enclave instances (secondary VMs)	1. BMSs do not support QingTian Enclave. 2. Only the Linux OS is supported. 3. The memory must be at least 128 MiB and cannot be less than four times the size of the QingTian Enclave Image File (EIF) to launch a QingTian Enclave instance. 4. If 2 MiB hugepages are configured in the configuration file to launch a QingTian Enclave instance, the maximum memory allowed is 512 MiB. 5. If 1 GiB hugepages are configured in the configuration file to launch a QingTian Enclave instance, the maximum memory allowed is 256 GiB. 6. All vCPUs and memory allocated to QingTian Enclave instances must come from the same NUMA node. 7. The number of the vCPUs must be an even number and cannot exceed the number of vCPUs per NUMA node on the parent instance minus 2. The total number of vCPUs cannot exceed 62. 8. Applications running in QingTian Enclave instances need to be packaged with the OS (kernel, ramdisk, and init) into a QingTian Enclave Image File (EIF).

NOTE

For details about isolating vCPUs and memory, see **Resource Isolation**.

The relationship between QingTian Enclave instances and their parent instance are as follows:

1. A maximum of two QingTian Enclave instances can be created from a parent instance.
2. QingTian Enclave instances cannot share the same physical core with their parent instance.
3. QingTian Enclave instances are running only when the parent instance is running. If the parent instance is stopped or terminated, QingTian Enclave instances are also stopped or terminated.
4. Resources (vCPUs and memory) of QingTian Enclave instances come from the parent instance. The memory range must be a continuous physical range aligned by 2 MiB/1 GiB.

You also need to note the following:

1. The parent instance must be a C7t and kC2 instance.
2. QingTian Enclave is available in the following regions: CN North-Beijing4, CN East-Shanghai1, CN South-Guangzhou, AP-Singapore, and TR-Istanbul.
3. If your services running in the QingTian Enclave instances are terminated unexpectedly, you need to manually run the services again.

4. By default, 1 GiB hugepages are configured for QingTian Enclave instances, with 1 GiB of memory and 2 vCPUs.

Billing

QingTian Enclave is free during the open beta test (OBT). You only need to pay for the ECSs you purchase.

Related Services

QingTian Enclave integrates with the following Huawei Cloud services:

1. KMS

   Key Management Service (KMS) is a core service provided by Huawei Cloud Data Encryption Workshop (DEW). KMS is a highly available cloud service that helps users create, store, manage, and audit keys. KMS uses Hardware Security Modules (HSMs) to protect keys and can be integrated with multiple Huawei Cloud services. Additionally, you can develop customized encryption applications using KMS.

2. IAM

   The Identity and Access Management (IAM) provides permissions management to securely manage access to your Huawei Cloud services and resources.

13.1.2 QingTian Enclave Concepts

- QingTian Enclave instances

  QingTian Enclave instances are completely independent VMs whose vCPUs and memory all come from the parent instance. QingTian Enclave instances have no external networking or persistent storage. Resources in the QingTian Enclave instances cannot be accessed by the processes, applications, kernel, or users of the parent instance.

- Parent instance

  The parent instance is an ECS instance that is used to allocate its vCPUs and memory to QingTian Enclave instances. These resources can be used during the lifetime of the QingTian Enclave instances. QingTian Enclave instances can only communicate with the parent instance after they are successfully launched.

- QingTian Enclave image file

  A QingTian Enclave image file (.eif) provides system information required for launching a QingTian Enclave instance and running QingTian Enclave applications in the instance, including a Linux operating system, other third-party libraries, and QingTian Enclave applications. For details about image creation, see **QingTian Enclave Application Development on Linux**.

- QingTian CLI

  QingTian CLI (qt CLI) is a command line tool that can be used to create, terminate, and query QingTian Enclave instances. The qt CLI must be installed and used on the parent instance. For details, see **QingTian CLI (qt CLI)**.

- Enclave SDK

Enclave SDK consists of a series of open-source libraries to develop your QingTian Enclave applications. Enclave SDK integrates APIs for interacting with Huawei Cloud KMS, such as encryption, decryption, and random number generation, and provides built-in support for remote attestation.

- QingTian cryptographic attestation

  QingTian cryptographic attestation is a process during which a QingTian Enclave instance proves its identity when interacting with the KMS service. Attestation is completed using a signed attestation document generated by the QingTian Hypervisor. Information contained in a QingTian Enclave attestation document can be used as a condition for third-party service authentication. You can use kms:RecipientAttestation-related condition keys in IAM to control access to specific KMS APIs, such as APIs for random number generation, encryption, and decryption.

- Attestation document

  An attestation document is generated and signed by the QingTian Hypervisor. The document contains QingTian Enclave information, including platform configuration registers (PCRs), cryptographic digest, and user statement. External services can use attestation documents to verify the identity of QingTian Enclave instances to establish trust. You can use attestation documents to build your own trustworthy system and interact with KMS. For details, see **Attestation Document**.

- qt-proxy

  The qt-proxy is a network proxy service running on the parent instance. The qt-proxy enables the parent instance to forward network packets from the QingTian Enclave instances so that the instances can communicate with external networks. This is the only way for QingTian Enclave instances to communicate with external services.

- PCR

  Platform configuration registers (PCRs) are cryptographic measurements that are unique to QingTian Enclave instances. Some PCRs are automatically generated when QingTian Enclave instances are created, and they can be used to verify the QingTian Enclave instance integrity since it was created. You can also manually create other PCRs that can be used to ensure that the QingTian Enclave instance is running on your expected platform. In addition, several PCRs included in attestation documents can be used to create condition keys of IAM access control policies for stronger access control. For details, see **PCR**.

- Local Vsock Connection

  The Local Vsock Connection is the only secure local channel between QingTian Enclave instances and the parent instance.

- QingTian Security Module

  The QingTian Security Module (QTSM) consists of the qtsm-lib function library and qtsm-server service. You can call the qtsm-lib user-mode API in your QingTian Enclave applications, and the qtsm-server will process specific QTSM requests and return messages. The qtsm-lib user-mode APIs can be used to query the PCR value of a specified index (qtsm_describe_pcr), extend the PCR value of a specified index (qtsm_extend_pcr), lock the PCR value of a specified index (qtsm_lock_pcr), lock the PCR values of specified indexes in batches (qtsm_lock_pcrs), obtain the QTSM information (qtsm_get_describe), and obtain the attestation document (qtsm_get_attestation).

13.2 Getting Started with QingTian Enclave

This section guides you through the process of understanding and using QingTian Enclave features. It shows you how to launch a parent instance, build a QingTian Enclave image file, query a running QingTian Enclave instance, and stop a QingTian Enclave instance.

1. Prepare an instance that supports the QingTian Enclave feature.
   - When purchasing an instance, select QingTian Enclave in the advanced settings and select a Linux OS image. Huawei Cloud EulerOS 2.0 is recommended.
   - Connect to the parent instance. For details, see the *Huawei Cloud Elastic Cloud Server User Guide*.
   - Install the qt CLI on the parent instance, set parameters for resource isolation in the configuration file as required, and enable resource isolation. To use qt CLI, you must install Python libraries. For details, see **QingTian CLI (qt CLI)**.
     - Install the qt CLI and other necessary RPM packages.
       ```
       yum install qt-enclave-bootstrap
       yum install virtio-qtbox
       yum install qingtian-tool
       ```
     - Set parameters for resource isolation in the configuration file and enable resource isolation. In this tutorial, 1 GiB of memory and 2 vCPUs are used by default.
       ```
       systemctl start qt-enclave-env
       ```
2. Use the binary mode to install Docker on the parent instance. You can download the required Docker version from the official website of Docker. The following uses docker-27.0.1 as an example.
   ```
   wget https://download.docker.com/linux/static/stable/x86_64/docker-27.0.1.tgz
   ```
   Decompress the downloaded package.
   ```
   tar zxf docker-27.0.1.tgz
   ```
   After the decompression is complete, copy all the files in the Docker directory to the **/usr/bin** directory.
   ```
   cp docker/* /usr/bin
   ```
   Start the Docker service and set the log level to error.
   ```
   dockerd -l error &
   ```
   Verify the Docker version.
   ```
   docker version
   ```
   Run the **hello-world** container to check whether Docker is installed.
   ```
   docker run hello-world
   ```
3. Build a QingTian Enclave image file.

   Here we use the following **hello_enclave.sh** script as the QingTian Enclave application.
   ```
   #!/bin/bash
   while true
   do
       echo "hello enclave!"
       sleep 2
   done
   ```

The Dockerfile content is as follows:

```
FROM ubuntu:latest
COPY hello_enclave.sh /root/hello_enclave.sh
CMD ["/root/hello_enclave.sh"]
```

Check that the script has execution permissions.

```
chmod +x hello_enclave.sh
```

Build a Docker image named **hello-enclave**.

```
docker build -f Dockerfile -t hello-enclave .
```

Run the **qt enclave make-img** command to convert the Docker image to a QingTian Enclave image file named **hello-enclave.eif**.

```
qt enclave make-img --docker-uri hello-enclave --eif hello-enclave.eif
```

The output is as follows:

```
# qt enclave make-img --docker-uri hello-enclave --eif hello-enclave.eif
{
   "digest":   "SHA384",
   "PCR0":
"63bf78ece7d2388ff773d0cad2ebc9a3070359db46d567ba271ff8adfb8b0b091be4ff4d5dda3f1c8310909
6e3656f3b",
   "PCR8":
"00000000000000000000000000000000000000000000000000000000000000000000000000000000000
0000000000000"
}
```

The QingTian Enclave image file named **hello-enclave.eif** has now been built. The command contains a set of PCR values, including PCR0 and PCR8. (In this example, no certificates or keys are specified during image creation, so PCR8 is 0.) These hash values are measurements of QingTian Enclave images and they are generally used as expected measurements (compared with the measurements in the attestation document during the boot-up process).

– Run a QingTian Enclave instance.

  You can now run the QingTian Enclave image using the following command:

```
qt enclave start --mem 1024 --cpus 2 --eif hello-enclave.eif --cid 4 --debug-mode
```

The QingTian Enclave instance runs in debug mode. For details about debug mode, see **Introduction to qt enclave Subcommands**.

The output is as follows:

```
# qt enclave start --cpus 2 --mem 1024 --eif hello-enclave.eif --cid 4 --debug-mode
Started enclave with EnclaveID : 0, EnclaveCID : 4, NumberOfCPUs : 2, MemoryMiB : 1024
{
   "EnclaveID":    0,
   "EnclaveCID":   4,
   "NumberOfCPUs": 2,
   "MemoryMiB":    1024,
   "LaunchMode":   "debug"
}
```

In this tutorial, 2 vCPUs and 1024 MiB of memory are allocated to the QingTian Enclave instance, and the EnclaveCID is set to 4. The EnclaveCID can be used as the IP address of the local socket between the QingTian Enclave instance and the parent instance.

– Query a running QingTian Enclave instance.

After the QingTian Enclave instance is created, run the following commands to check whether the instance is running:

```
qt enclave query --enclave-id 0
# qt enclave query --enclave-id 0
[{
        "EnclaveID":    0,
        "ProcessID":    29990,
        "EnclaveCID":   4,
        "NumberOfCPUs": 2,
        "MemoryMiB":    1024,
        "LaunchMode":   "debug"
    }]
```

The command can query information about the QingTian Enclave instance, including EnclaveID, ProcessID, EnclaveCID, number of vCPUs, memory size, and its running mode. You can run the **qt enclave console** command to view the read-only console output of the QingTian Enclave instance because the instance is launched in debug mode.

```
hello enclave!
hello enclave!
hello enclave!
hello enclave!
```

You can see that **hello enclave!** is printed to the console every two seconds.

– Stop a QingTian Enclave instance.

If you want to stop a QingTian Enclave instance, run the following commands:

```
# qt enclave stop --enclave-id 0
stop enclave 0 successfully
{
    "EnclaveID":    0
}
```

13.3 Examples of Using QingTian Enclave

In this chapter, we will show how to use QingTian Enclave instances together with KMS (sub-service of DEW), IAM, and OBS.

13.3.1 Workflow

Roles

The typical usage of a QingTian Enclave instance involves the following roles:

- Data security administrator: has control permissions for the confidential data and Huawei Cloud KMS keys. A data security administrator owns a Huawei Cloud account and has the highest permissions. For example, a data security administrator can create IAM users and grant them the minimum permissions, such as creating encryption keys and encrypting sensitive data. In this example, we suppose that the data security administrator is also responsible for building the QingTian Enclave image file. This party obtains the expected measurements PCR0 and PCR8 and uses these values as condition keys in IAM policies.
- Parent instance administrator: is authorized by the data security administrator and has permission to access the parent instance and manage the lifecycle of QingTian Enclave instances. This party launches a QingTian Enclave instance using the QingTian Enclave image file built by the data security administrator.
- QingTian Enclave application developer: develops applications running in the QingTian Enclave instances. In this example, the application needs to obtain the ciphertext object from OBS bucket **Bucket1**, call the kms-decrypt API to decrypt the ciphertext, process the data, and generate the results to **Bucket2**.

Data and Environment Preparation

The following gives an overview of the data encryption process, attestation settings, and QingTian Enclave instance creation.

1. The data security administrator creates keys in KMS (a sub-service of DEW).
2. The data security administrator uses KMS keys to encrypt a piece of sensitive data, for example, bank card information. For details, see **Example 1: Encrypting or Decrypting Small Volumes of Data**.
3. The data security administrator uses the command line tool obsutil to upload the encrypted ciphertext to a Huawei Cloud OBS bucket. For details, see **Uploading an Object**.
4. The data security administrator compiles and packages the QingTian Enclave application by creating a Docker image and using the qt CLI to convert the Docker image into a QingTian Enclave image file. For details, see descriptions about how to build a QingTian Enclave image file. The data security administrator records PCR0 and PCR8 generated when the QingTian Enclave image file is built.
5. The data security administrator sets PCR0 and PCR8 as condition keys of the IAM access control policies (controlling the kms-decrypt API).
6. The parent instance administrator boots the parent instance, starts the qt-proxy service, and boots the QingTian Enclave instance using the QingTian Enclave image file.

Remote Attestation and Data Decryption

The following describes the execution process of a QingTian Enclave application.

1. With the qt-proxy service, the QingTian Enclave application downloads the ciphertext from the Huawei OBS bucket to the QingTian Enclave instance.
2. The QingTian Enclave application generates a pair of RSA public and private keys (pubKey and priKey) for end-to-end data encryption with the KMS service. The encryption does not depend on HTTPS. Then, the QingTian Enclave SDK is used to call the KMS-provided kms-decrypt API that supports the attestation document as the input parameter. The attestation document includes the QingTian Enclave instance's PCRs and the encrypted pubKey generated by the application.
3. Huawei Cloud KMS receives the request and verifies whether the attestation document is signed by the QingTian Attestation PKI. During the access control check of the kms-decrypt API, PCRs in the attestation document will be compared with those in the IAM policies. If they match, the API can be called. If they do not match, the access will be denied.
4. KMS decrypts the data first, encrypts the data using the pubKey provided by the attestation document, and sends the encrypted data to the QingTian Enclave application. The QingTian Enclave application uses the priKey to decrypt the received ciphertext data.

13.3.2 Building a QingTian Enclave Image

After a QingTian Enclave application is developed, you need to build a QingTian Enclave image file (.eif) in a trusted environment. The QingTian Enclave image file contains everything required to launch a QingTian Enclave instance, including the application code, runtime dependencies, operating system, and file system. This section describes how to build a QingTian Enclave image file.

1. Create a Docker source image.

   Package the QingTian Enclave application and its execution environment into a Docker image. For details, see **QingTian Enclave Application Development on Linux**.

2. Obtain the image from the image library.

   The following uses the Ubuntu image provided in the Docker repository as an example. Obtain the image source from Docker (networking needs to be configured for query). Run the following command to query the image source:

   ```
   docker search ubuntu
   ```

   Pull the Ubuntu image locally:

   ```
   docker pull ubuntu
   ```

   After the Ubuntu image is pulled locally, run the following command:

   ```
   docker image ls
   ```

   If you use a Docker image locally, directly perform step 3 to convert the image.

3. Convert the Docker image to a QingTian Enclave image file.

   First, use OpenSSL or other tools to create a private key (private-key.pem) and a certificate (server.pem). This operation is optional. The mandatory parameters in the ***qt make-img*** command are the Docker source image and the generated target QingTian Enclave image.

   ```
   # qt enclave make-img --docker-uri ubuntu --eif  /home/docker/ubuntu.eif --
   private-key  /home/docker/private-key.pem--signing-certificate
   /home/docker/server.pem
   {
       "digest":       "SHA384",
       "PCR0":
   "b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c0364a9b
   88dd58017a62"
       "PCR8":
   "dbf4a7f9fab7f18619b5899c407081981ad6762fb9a809da78548821b5021965423181584acd7b2017033
   76f1133a546"
   }
   ```

   Then you have obtained a QingTian Enclave image file. You will get a set of PCR0 and PCR8. These hashes are measurements of the instance and can be used as condition keys in IAM authorization policies to implement conditional access control over KMS APIs. For details, see **PCR**.

13.3.3 Launching a QingTian Enclave Instance

Resource Isolation

Before launching a QingTian Enclave instance, you need to isolate resources in the parent instance for QingTian Enclave instances first. Isolated resources include vCPUs and memory. You can specify the isolated resources by modifying the ***/etc/qingtian/enclave/qt-enclave-env.conf*** configuration file in the instance.

```
1 GiB hugepage:
hugepage_size:1024
Memory: 1 GiB
memory_mib:1024
vCPUs
```

```
cpu_count:2
# vCPU list
# cpu_list:2,3
```

Do not repeatedly enable resource isolation, or the hugepage memory may become insufficient and the launch of QingTian Enclave instances or resource isolation may fail. In this example, retain the default settings of 1 GiB hugepage, 2 vCPUs, and 1 GiB of memory. After confirming the parameters in the configuration file, run the following command:

```
systemctl restart qt-enclave-env.service
```

There are constraints between configuration items in the ***/etc/qingtian/enclave/qt-enclave-env.conf*** configuration file. For details, see **Introduction to qt-enclave-env**.

Launching a QingTian Enclave Instance

On the parent instance, run the ***qt enclave start*** command with the QingTian Enclave instance image file specified to create a QingTian Enclave instance. After the QingTian Enclave instance is launched, the QingTian Enclave application and its dependencies will be booted from the QingTian Enclave image file to the QingTian Enclave instance. For example, if you want to create a QingTian Enclave instance with 2 vCPUs, 1 GiB of memory, and an EnclaveCID of 4, run the following commands:

```
[root@localhost ~]# qt enclave start --cpus 2 --mem 1024 --eif /home/docker/ubuntu.eif --cid 4
Started enclave with EnclaveID : 0, EnclaveCID : 4, NumberOfCPUs : 2, MemoryMiB : 1024
{
    "EnclaveID":    0,
    "EnclaveCID":   4,
    "NumberOfCPUs": 2,
    "MemoryMiB":    1024,
    "LaunchMode":   "debug"
}
```

In this instance, the CMD statement in the original Ubuntu image is **/bin/bash**, so the QingTian Enclave instance executes the statement after being launched. After the statement is executed, the QingTian Enclave application exits, and the QingTian enclave instance is closed.

13.4 Cryptographic Attestation

QingTian Enclave instances support cryptographic attestation. The instances use cryptographic attestation to prove their identities and build trust with external services. The attestation process uses an attestation document that includes the measurements of the QingTian Enclave runtime environment. These measurements can be used to create access control policies in external services to control access to specific operations for specific QingTian Enclave instances.

You can use the QingTian Enclave SDK to obtain an attestation document from the QingTian Hypervisor. The attestation document includes unique measurements and digital signature. This document can be attached to requests from the QingTian Enclave instance to an external service. The external service can validate whether the measurements included in the attestation document match the values in the access control policies to determine whether to grant the QingTian Enclave instance access or establish trust.

13.4.1 PCR

A QingTian Enclave instance's measurements include a series of hashes calculated using standard trusted measurements and are stored in the platform configuration registers (PCRs) of the QingTian Security Module (QTSM).

NOTE

A QingTian Enclave instance's measurements can support a maximum of 32 PCRs. The QingTian Enclave system occupies the PCRs with indexes 0 to 15 (PCR0-PCR15), and the QingTian Enclave application can use the PCRs with indexes 16 to 31 (PCR16-PCR31).

Image verification is not performed for QingTian Enclave instances that are launched in debug mode. PCR0 to PCR15 used by QingTian Enclave are made up entirely of zeros to prevent data leaks. Your QingTian Enclave application can continue to use PCR16 to PCR31.

System PCRs

PCR	Measurement	Remarks
PCR0	QingTian Enclave image file	A measurement of the content of the QingTian Enclave image file, excluding the certificate and signature information
PCR3	IAM Agency	A contiguous measurement of the IAM agency assigned to the parent instance. This ensures that the attestation process succeeds only when the parent instance has the correct IAM agency. It is delivered only once when the parent instance is launched. After it is reset, the instance needs to be restarted to apply the change.
PCR4	Instance UUID of the parent instance	A contiguous measurement of the UUID of the parent instance. This ensures that the attestation process succeeds only when the parent instance has a specific instance UUID.

PCR	Measurement	Remarks
PCR8	QingTian Enclave image file signing certificate	A measurement of the signing certificate for the QingTian Enclave image file

Currently, QingTian Enclave provides the measurements for PCR0 and PCR8 and it will have more measurements for future use.

1. PCR0 is the measurement of the QingTian Enclave image file and is a determined value since the QingTian Enclave image file is built. Example PCR0:

   ```
   EXTEND_PCR: index: 0
   EXTEND_PCR: data:
   0d1ae7330f437ee563178df30a7c7b7634125d31cac14f6784933db5e90080008438b38fdbb39c886ffe058
   6ab099b56
   EXTEND_PCR res: data:
   b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c0364a9b8
   8dd58017a62
   ```

2. To further enhance the security policy of QingTian Enclave, you can create an IAM agency and attach it to the parent instance. In the condition keys of KMS key policies, you can use the SHA384 hash value of IAM agency as PCR3. This ensures that only QingTian Enclaves running on instances with the correct IAM agency can perform specific KMS actions on KMS keys. You can generate the hash using any tool that can convert a string to an SHA384 hash. Example PCR3:

   ```
   $IAM_AGENCY="iam::6c031a4leefc480bb60f20c003891fcd:agency:cddd"; \
    python -c"import hashlib, sys; \
    h=hashlib.sha384(); h.update(b'\0'*48); \
    h.update(\"$IAM_AGENCY\".encode('utf-8')); \
    print(h.hexdigest())"
   ```

3. PCR4 is based on SHA384 of the parent instance's UUID, so you can generate the PCR after launching the parent instance. You can generate the hash using any tool that can convert a string to an SHA384 hash. Example PCR4:

   ```
   $INSTANCE_ID="ecb23eec- 51d4-462f-8dbd-63bfbae7869b"; \
    python -c"import hashlib, sys; \
    h=hashlib.sha384(); h.update(b'\0'*48); \
    h.update(\"$INSTANCE_ID\".encode('utf-8')); \
    print(h.hexdigest())"
   ```

4. PCR8 is a measurement of the signing certificate of the QingTian Enclave image file. You can sign the QingTian Enclave image file using your signing certificate and private key. PCR8 is available only when the QingTian Enclave image file is signed with the signing certificate and private key. PCR8 can be used to verify that the image is signed by using a specific signing certificate. As long as the specified signing certificate is not changed, PCR8 remains unchanged, even if the image file is changed. Details of PCR8 are as follows:

   ```
   EXTEND_PCR: index: 8
   EXTEND_PCR: data:
   c5b3e075e00c261e7fc364f1541067b2a42d4b793225ab10e5cfb8eaca31b3d598af9dd2e491828c2569a9
   953401abcb
   EXTEND_PCR res: data:
   4f8b066ce5ac24150612ba9a55bbb9211f626152ada40ede160f4d7ecbfa214c2a549181f6611a3d16a12e
   c88a577a01
   ```

13.4.2 Attestation Document

An attestation document is used to verify the reliability measurement results of QingTian Enclave instances. An attestation document is generated by the QingTian Hypervisor. It includes the PCR list, the QingTian Public Key Infrastructure (PKI) certificate chain, cryptographic algorithm declaration, and user-defined data for the QingTian Enclave application. The attestation document is signed by the Huawei Cloud QingTian Attestation PKI.

The attestation document generated by the QingTian Hypervisor is encoded in Concise Binary Object Representation (CBOR) and signed in Object Signing and Encryption (COSE). For details, see **RFC 8949: Concise Binary Object Representation (CBOR)**.

The structure of the QingTian Enclave attestation document complies with the Concise Data Definition Language (CDDL) (RFC 8610):

```
AttestationDocument = {
    module_id: text,                ; Security module ID
    timestamp: uint .size 8,        ; Timestamp
    digest: digest,                 ; Digest algorithm
    pcrs: { + index => pcr },       ; PCRs
    certificate: cert,              ; Signing certificate of the QingTian Enclave's attestation document
    cabundle: [* cert],             ; QingTian PKI certificate chain
    ? user_data: user_data,         ; (Optional) Application data
    ? nonce: user_data,             ; (Optional) Data not repeatedly used
    ? pubkey: user_data,        ; (Optional) Application public key
}

cert = bytes .size (1..4096)        ; DER encoding certificate
user_data = bytes .size (0..4096)
pcr = bytes .size (48)        ; PCR content
index = 0..31
digest = "SHA384"
```

The optional parameters (**pubkey**, **user_data**, and **nonce**) in the attestation document allow for custom (or auto-negotiated) application-level security protocols between the QingTian Enclave instances and external entities. For example, the QingTian Enclave application can create an asymmetric key pair (PriKey and PubKey) and provide trusted attestation for the PubKey using the QingTian Enclave attestation document. Then, the external entity can use some custom application-level security protocols such as trusted key distribution and trusted key agreement based on the PubKey attestation.

13.4.3 Document Signature Verification

This section introduces the verification process of the attestation document. When you request an attestation document from the QingTian Hypervisor, you will receive a binary blob containing the signed attestation document. The signed attestation document is encoded in CBOR and signed in COSE. The verification process is as follows:

1. Decode the CBOR object and map it to the COSE_Sign1 structure.

2. Extract the attestation document from the COSE_Sign1 structure.

3. Verify the validity of the CA certificate chain in the attestation document.

4. Verify the validity of the digital signature of the attestation document.

The attestation document is signed by the Huawei Cloud QingTian Attestation PKI. The QingTian Enclave's root certificate can be downloaded at **https://qingtian-enclave.obs.myhuaweicloud.com/huawei_qingtian-enclaves_root-G1.zip**. The SHA-256 hash value of the compressed file is as follows:

```
99e9203a64cfb0c6495afd815051e97bea8a37895dc083d715674af64adeadfe
```

The root certificate of the QingTian Attestation PKI can be valid for up to 30 years. The subject of the root certificate is in the following format:

```
CN=huaweicloud.qingtian-enclaves, C=CN, O=Huawei Technologies, OU=Huawei Cloud
```

COSE and CBOR

The COSE_Sign1 signature structure is usually used to sign a single signature for a message. The content and signature parameters are placed in the protected header. The COSE_Sign1 data structure is a CBOR array that includes the following fields:

```
[
  protected header;       // Protected header information
  unprotected header;  // Unprotected header information
  payload;             // Signed data and attestation document's CBOR encapsulation data
  signature;           // Signature
]
```

In the context of the attestation document, an example array is as follows:

```
18(                      // COSE_Sign1 CBOR tag
 {1: -35},               // Algorithm: ECDS 384
 {},                     // Empty
 attestation doc,        // Attestation document
 signature,                // Signature
)
```

Certificate Verification

Verifying the certificate chain is an indispensable phase of the certificate verification. The CA bundle in the attestation document contains a list of root and intermediate certificates which are sequenced as follows:

```
[ ROOT_CERT - INTERM_1 - INTERM_2 ... -INTERM_N ]
   0          1         2           N
```

To verify the validity of the target certificate (certificate in the attestation document) using certain certificate verification tools, you may need to verify the certificates in the following sequence:

```
[ TARGET_CERT - INTERM_N - INTERM_N-1 ... - ROOT_CERT]
```

13.4.4 Integration with Huawei Cloud KMS

Huawei Cloud Key Management Service (KMS) has built-in attestation support for QingTian Enclave instances. You can use the Huawei Cloud KMS APIs included in the QingTian Enclave SDK to perform Huawei Cloud KMS actions, such as decryption, random number generation, and encryption in QingTian Enclave instances based on the attestation documents. KMS can ingest attestation documents from QingTian Enclave instances and validates the measurements in the attestation documents against these specified in the IAM policies to determine whether QingTian Enclave instances can access KMS APIs.

The following is an example IAM authorization policy. This policy allows you to call KMS APIs for decrypting data or data keys only in the QingTian Enclave environment, and the measurements for PCR0 and PCR8 of QingTian Enclave must be the same as the specified measurements.

After the authorization is successful, the exported JSON information is as follows:

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "kms:cmk:decrypt",
                "kms:dek:decrypt"
            ],
            "Resource": "*",
            "Condition": {
                "StringEqualsIgnoreCase": {
                    "kms:RecipientAttestation/PCR0": [

"c5158cb6ee9dbb0ead648c3dc80e472c85e0d67f19fb53fbd3fb94c3371aec63cdb93b80d727a7084248873b1d
8e8b41"

                    ],
                    "kms:RecipientAttestation/PCR8": [

"705afb1012d27f4e07a25e674e6a17dec57305e29cd412184b7bcb78d9e67f16a0cc26d8706a4fab418a5da578
8bc949"
                    ]
                }
            }
        }
    ]
}
```

JSON information description:

- Action: actions allowed by the identity policy.
- Resource: resources that can be obtained by the identity policy.
- *: all resources can be obtained.
- Condition: request condition. In QingTian Enclave, any combinations of PCR0, PCR3, PCR4, and PCR8 can be used as request conditions. PCR0 and PCR8 are used in this example.

For details about how to create a user group, see **Creating a User Group and Assigning Permissions**.

For details about how to create a custom policy, see **Creating a Custom Policy**.

13.5 QingTian Enclave Application Development

A fully featured QingTian Enclave application consists of at least two components:

1. An application with low security requirements running on the parent instance
2. An application with high security requirements running inside a QingTian Enclave instance

Due to the isolated environment of the QingTian Enclave instance, the only channel for the applications running on the parent instance to communicate with those inside the QingTian Enclave instance is the vsock socket.

13.5.1 QingTian Enclave Application Development on Linux

QingTian Enclave SDK

The QingTian Enclave SDK consists of a series of open-source libraries for you to develop your own QingTian Enclave applications. It includes the qtsm-lib function library provided by QingTian Security Module (QTSM). In addition, the QingTian Enclave SDK integrates with KMS APIs that provide built-in support for obtaining attestation documents and other KMS-related services. The typical examples show how to call KMS APIs for decryption in QingTian Enclave.

Table 13-1 API description

Type	API	Description
libqtsm APIs	qtsm_describe_pcr	Queries the PCR value of a specified index.
	qtsm_extend_pcr	Extends the PCR value of a specified index.
	qtsm_lock_pcr	Locks the PCR value of a specified index.
	qtsm_lock_pcrs	Locks the PCR values of specified indexes in batches.
	qtsm_get_describe	Obtains the QTSM information.
	qtsm_get_attestation	Obtains the attestation document.
	qtsm_get_random	Obtains a random hardware number.
KMS APIs	kms_generate_datakey_blocking	Generates a new key pair and obtains the public key and private key.
	kms_generate_datakey_blocking_with_proxy	Integrates the qtproxy and obtains the key pair.
	kms_gen_random_blocking	Obtains a random number.
	kms_gen_random_blocking_with_proxy	Integrates the qtproxy and obtains a random number.
	kms_decrypt_data_blocking	Decrypts data.
	kms_decrypt_data_blocking_with_proxy	Integrates the qtproxy and decrypts data.

You can obtain the source code for free from the open-source repository at **https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian/tree/master/enclave** and develop your own QingTian Enclave application based on the test example.

Vsock Communication

The following uses vsock as an example to describe how to develop QingTian Enclave applications on Linux. The vsock application in this example can only run on Linux instances.

The vsock application helps developers know how information is exchanged between the parent instance and the QingTian Enclave instance. The vsock application includes two parameters: **Server** and **Client**. You can specify the two parameters to define the roles (client or server application) for the parent instance and the QingTian Enclave instance. In the vsock application, the client application sends a simple text message over the vsock to the server application and the server application listens to the vsock and prints the message to the terminal once it receives the message.

The following describes how a QingTian Enclave instance functioning as the server application receives the **hello world** message from the parent instance functioning as the client application.

1. Compile a SocketCommunication.py program.

```
#!/usr/local/env python3
import argparse
import socket
import sys

CID_DEFAULT = 3
PORT_DEFAULT = 9999
TIMEOUT = 5
BLACKLOG_DEFAULT = 5

class Client:
    def __init__(self, cid, port):
        self.clientAddr = (cid, port)
        self.connect()

    def connect(self):
        self.socket = socket.socket(socket.AF_VSOCK, socket.SOCK_STREAM)
        self.socket.settimeout(TIMEOUT)
        print("connecting to the server")
        try:
            self.socket.connect(self.clientAddr)
        except socket.error:
            print("client's socket connection err")
            sys.exit(1)

    def send(self, msg):
        print("client sends hello to the server")
        self.socket.sendall(msg)

    def disconnect(self):
        self.socket.close()

    def receiveData(self):
        while True:
            try:
                message = self.socket.recv().decode()
            except (socket.error, UnicodeDecodeError):
                break
            if message:
                print(message, end = " ", flush = True)
        print()

def clientHandler(args):
    client = Client(args.cid, args.port)
    message = "Hello world"
    client.send(message.encode())
    client.disconnect()

class Server:
    def __init__(self, port):
        self.socket = socket.socket(socket.AF_VSOCK, socket.SOCK_STREAM)
        self.serverAddr = (socket.VMADDR_CID_ANY, port)
        self.socket.bind(self.serverAddr)
        self.socket.listen(BLACKLOG_DEFAULT)

    def receiveData(self):
        while True:
            print("waiting for a connection")
            (conn, clientAddr) = self.socket.accept()
            try:
                print("connection from ", clientAddr)
                while True:
                    try:
                        data = conn.recv(256).decode()
                    except (socket.error, UnicodeDecodeError):
                        break
                    if data:
                        print("data: ", data)
                    else:
                        print("connection close")
                        break
            finally:
                conn.close()

def serverHandler(args):
    server = Server(args.port)
    server.receiveData()


def main():
    parser = argparse.ArgumentParser(description = "Hello world demo", prog='SocketCommunication')
    subparsers = parser.add_subparsers(description = "Communication roles")
    parserClient = subparsers.add_parser("Client", description = "Client",
                                    help = "Communicate with server using a given cid and port.")
    parserClient.add_argument("-c", "--cid", default = CID_DEFAULT, type = int, help = "Client's Cid")
    parserClient.add_argument("-p", "--port", default = PORT_DEFAULT, type = int, help = "Client's
port")
    parserClient.set_defaults(func = clientHandler)
    parserServer = subparsers.add_parser("Server", description = "Server", help = "Listen on a given
port")
    parserServer.add_argument("-p", "--port", default = PORT_DEFAULT, type = int, help = "Server's
Port")
    parserServer.set_defaults(func = serverHandler)
    if len(sys.argv) < 2:
        parser.print_usage()
        sys.exit(1)
    args = parser.parse_args()
    args.func(args)


if __name__ == "__main__":
    main()
```

2. Create a file named **Dockerfile**.

```
#start the Docker image from ubuntu
FROM ubuntu AS base-img
WORKDIR /home/builder
# COPY vsocket example
COPY . vsocket
# install relative dependencies
RUN apt-get update && \
    apt-get install python3 -y && \
    apt-get install gcc -y && \
    apt-get install gawk -y
```

```
# Launch a client
CMD ["python3", "/home/builder/vsocket/SocketCommunication.py","Server","-p 9999"]
```

3. Build a Docker image.
   ```
   sudo docker build -t vsock-sample-client -f Dockerfile .
   ```
4. Convert the Docker image to a QingTian Enclave image file.
   ```
   qt enclave make-img --docker-uri vsock-sample-client --eif vsock_sample.eif
   ```
5. Boot the QingTian Enclave instance in debug mode using the QingTian Enclave image file **vsock_sample.eif**.
   ```
   qt enclave start --cpus 2 --mem 4096 --eif vsock_sample.eif --debug-mode --cid 4
   ```
   Run the **qt enclave console** command to view the read-only terminal output in the QingTian Enclave instance.
   ```
   qt enclave console --enclave-id 0
   waiting for a connection
   ```
6. Boot a parent instance terminal and start the client program.
   ```
   python3 SocketCommunication.py Client -c 4 -p 9999
   ```
7. Check that the following information is displayed on the terminal after the server application receives the message over the vsock.
   ```
   connection from  (3, 4180219645)
   data:  Hello world
   connection close
   waiting for a connection
   ```

Example of Using libqtsm and SDKs

The following describes how to use libqtsm and SDK APIs in QingTian Enclave applications based on open-source sample code. The sample application can only run on Linux instances.

1. Install the libqtsm development package.

   yum install libqtsm-devel

2. Obtain the open-source sample code from the following URL and copy it to the environment where the QingTian Enclave image is built:

   https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian/tree/master/enclave/qtsm

3. Create a file named **Dockerfile**.
   ```
   # start the Docker image from ubuntu
   FROM ubuntu AS base-img
   WORKDIR /home/builder
   # COPY libqtsm example
   COPY ./qtsm qtsm_tests/
   # install relative dependencies
   RUN apt-get update && \
       apt-get install gcc -y && \
       apt-get install make -y && \
       apt-get install libssl-dev -y && \
       apt-get install libglib2.0-dev -y && \
       apt-get install curl -y && \
       apt-get install libcurl4-openssl-dev –y  && \
       apt-get install -y libcbor-dev && \
       apt-get install -y libjson-c-dev
   # build a test demo
   RUN cd qtsm_tests/tests/ && \
           make
   RUN cp /home/builder/qtsm_tests/tests/gtest_libqtsm /root/
   # Launch a client
   CMD "/root/gtest_libqtsm"
   ```

4. Build a Docker image, convert it to a QingTian Enclave image file, and boot the QingTian Enclave instance.
5. Obtain the open-source sample code of SDK APIs from the following URL:

   https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian/tree/master/enclave/qtsm-sdk-c/samples

13.5.2 QingTian Enclave Network Proxy

Overview

QingTian Enclave proxy (qproxy) is a network proxy tool of QingTian Enclave. With this tool, you can smoothly migrate network services that are deployed on QingTian-based VMs to QingTian Enclave instances without any modifications.

The qproxy tool is an executable binary file. It needs to be executed using different commands on the parent instance and QingTian Enclave instance.

- On the parent instance, run **/path/to/qproxy host --config=/path/to/config_qproxy.toml <cid>** to enable qproxy.
- On a QingTian Enclave instance, run **/path/to/qproxy host --config=/path/to/config_qproxy.toml <cid>** to enable qproxy.

For example, a network service is directly deployed on a VM to process requests from end users. End users initiate network requests through the port exposed by the service and wait for the service to respond. After the network service is migrated to QingTian Enclave using qproxy, the same network service is provided.

You can run **qlog host** and **qlog enclave** on the parent instance and QingTian Enclave instance, respectively, to execute a given qproxy binary file. A local vsock-based communication link is established between them. The qproxy component in the parent instance listens to port 5050, receives user requests, and forwards the received requests to the qproxy component in the QingTian Enclave instance through the local vsock. The qproxy sends requests to the listened port 5050 in the QingTian Enclave instance. After the network service processing is complete, the response is returned along the original path.

Figure 13-1 Seamless migration of a network service to QingTian Enclave

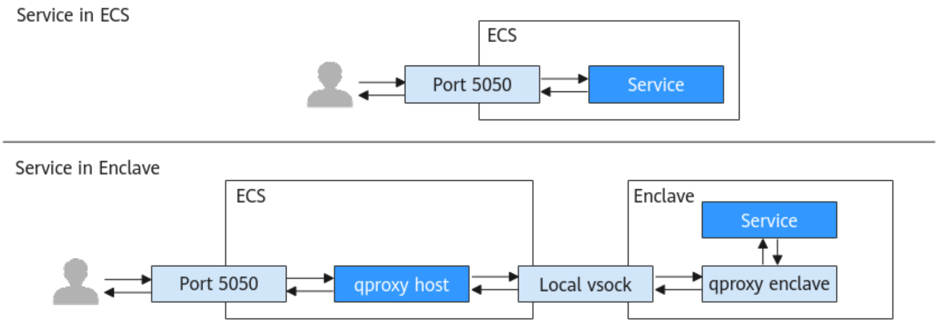

The following describes how to use qproxy.

Prerequisites

1. You have obtained the qproxy source code by performing the following:

   Clone the QingTian Enclave code repository.

   ```
   git clone https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian.git
   ```

2. You have obtained the cargo tool chain by performing the following:

   a. Install rustup.

   ```
   curl --proto '=https' --tlsv1.2 -sSf https://sh.rustup.rs | sh
   ```

   b. After the installation is complete, load rustup.

   ```
   source $HOME/.cargo/env
   ```

   c. Check that rustc and cargo are installed.

   ```
   rustc -V
   cargo -V
   ```

3. You have learned about the pre-dependency of qproxy.

Table 13-2 Pre-dependency

Dependency Item	Earliest Test Version
cargo	1.77.0
libcbor	0.10.2
libssl (libcrypto)	3.0.0
libcurl	4.0.0

4. You have prepared the QingTian Enclave environment by performing the following:

   a. Install the qt CLI tool and required RPM packages.

   b. Install Docker.

   c. Install Python 3 and required Python modules (docker and knack).

   For details, see **Getting Started with QingTian Enclave** and **Installation of the qt CLI**.

Procedure

Step 1 Build qproxy.

The generated qproxy binary file is compiled in **qingtian-tools/qproxy/target/release**.

Step 2 Create a workspace.

Create a workspace and copy the qproxy binary file to the workspace. Store the files generated subsequently in the workspace.

Step 3 Configure the **config_qproxy.toml** file.

In the workspace directory, create the **config_qproxy.toml** file with the following content:

```
[[inbound_connections]]
host_port = 5050
```

```
enclave_port = 5050
vsock_port = 9995

[[inbound_connections]]
host_port = 5443
enclave_port = 5443
vsock_port = 9994

[log_location]
host_log = "host.log" # qproxy host log name, e.g./var/log/qproxy/host.log
enclave_log = "enclave.log" # qproxy enclave log name, e.g./var/log/qproxy/enclave.log
log_level = "info" # qproxy logger level, e.g. "off", "info", "warn", "error", "debug", "trace"
host_log_dir = "/var/log/qproxy" # qproxy host log dir, and its default value is "/var/log/qproxy"
enclave_log_dir = "/var/log/qproxy" # qproxy enclave log dir, and its default value is "/var/log/qproxy"
```

Step 4 Create a QingTian Enclave image that contains qproxy.

1. In the workspace directory, create the HTTP request test script **app.sh**.

   ```
   #!/bin/bash

   PORT=5050

   while true; do
        echo -e "HTTP/1.1 200 OK\nContent-Type: text/plain\n\nHello world!" | nc -l 127.0.0.1  $PORT
   done
   ```

2. Run the following command in the workspace to grant **app.sh** execute permissions:

   ```
   chmod +x app.sh
   ```

3. In the workspace directory, create the start.sh script.

   ```
   #!/bin/bash
   ip link set lo up
   /root/qproxy enclave --config=/root/config_qproxy.toml &
   /root/app.sh
   ```

4. Run the following command in the workspace to grant **start.sh** execute permissions:

   ```
   chmod +x start.sh
   ```

5. In the workspace directory, create a Dockerfile.

   ```
   FROM ubuntu:latest
   COPY ./qproxy /root/qproxy
   COPY ./config_qproxy.toml /root/config_qproxy.toml
   COPY ./start.sh /root/start.sh
   COPY ./app.sh /root/app.sh
   RUN apt-get update && \
       apt-get install -y netcat-openbsd && \
       apt-get install -y iproute2
   CMD "/root/start.sh"
   ```

6. Run the following command in the workspace to create a Docker image:

   ```
   docker build -f Dockerfile -t test_qproxy_enclave .
   ```

7. Run the following command in the workspace to create a QingTian Enclave image:

   ```
   qt enclave make-img --docker-uri test_qproxy_enclave --eif test_qproxy_enclave.eif
   ```

Step 5 Start qproxy.

1. Run the following command in the workspace to start a QingTian Enclave instance:

   ```
   qt enclave start --cpus 2 --mem 1024 --cid 4 --eif test_qproxy_enclave.eif
   ```

2. Run the following command in the workspace to start qproxy in the parent instance:

   ```
   ./qproxy host --config=./config_qproxy.toml 4 &
   ```

3. Run the following curl command on the parent instance:

   ```
   curl localhost:5050
   ```

"Hello world!" is displayed.

Step 6 Set the qproxy environment variables.

The qproxy contains two sub-commands, one is executed on the QingTian Enclave instance (qproxy enclave), and the other is executed on the parent instance (qproxy host).

You can set the RUST_LOG environment variable to control the output of different levels of logs of a binary file.

- RUST_LOG=OFF: All logs are not displayed.
- RUST_LOG=info: Logs of the "info", "warn", and "error" levels are displayed.
- By default, logs of the "warn" and "error" levels are displayed.
- For more information, see the **EnvFilter documentation**.

----End

qproxy Help Information

qproxy help

```
$ qproxy --help
Usage: qproxy <COMMAND>

Commands:
  host          The part of qproxy that runs outside the enclave
  enclave       The part of qproxy that runs inside the enclave
  check-config  Check the qproxy configuration file
  help          Print this message or the help of the given subcommand(s)

Options:
  -h, --help     Print help
  -V, --version  Print version
```

qproxy enclave help

```
$ qproxy enclave --help
The part of qproxy that runs inside the enclave

Usage: qproxy enclave [OPTIONS]

Options:
      --parent-cid <PARENT_CID>
          The CID of the parent VM of this enclave

          [env: QPROXY_PARENT_CID=]
          [default: 3]

      --config <CONFIG>
          Path to the configuration file

  -t, --threads <THREADS>
          Number of threads the async runtime is allowed to use

          [env: TOKIO_WORKER_THREADS=]

      --control-port <CONTROL_PORT>
          The port where to listen for control messages from the enclave

          Leave at default value unless you know what you are doing

          [env: QPROXY_CONTROL_PORT=]
          [default: 6666]
          [0..=65535]
  -h, --help
          Print help (see a summary with '-h')
```

qproxy host help

```
$ qproxy host --help
The part of qproxy that runs outside the enclave

Usage: qproxy host [OPTIONS] <CID>

Arguments:
  <CID>
          The CID of the enclave

          [env: QPROXY_LISTEN_CID=]

Options:
      --config <CONFIG>
          Path to the configuration file

  -t, --threads <THREADS>
          Number of threads the async runtime is allowed to use

          [env: TOKIO_WORKER_THREADS=]

      --ipv4
          Only resolve IPv4 addresses

      --ipv6
          Only resolve IPv6 addresses

      --control-port <CONTROL_PORT>
          The port where to listen for control messages from the enclave

          Leave at default value unless you know what you are doing

          [env: QPROXY_CONTROL_PORT=]
          [default: 6666]
          [0..=65535]

  -h, --help
          Print help (see a summary with '-h')
```

Configuration Information

Configuration Parameters

- **outbound_connections**: used to forward traffic from QingTian Enclave instances to specific external services (hostnames/IP addresses) and ports.
- **inbound_connections**: used to forward traffic received by the specified port of the parent instance to QingTian Enclave instances.

Table 13-3 Configuration parameters

Bound	Variable	Type	Description
outbound_connections	hostname	String	Hostname used to forward traffic, for example, **api.myservice.com**. It also can be an IP address.

Bound	Variable	Type	Description
	vsock_port	u32	Port used by the vsock in the QingTian Enclave instance. It must be unique and cannot conflict with the qproxy port number.
	tcp_port	u32	Port used by the instance to forward traffic (connections will be made to hostname:port). It must be unique and cannot conflict with qproxy ports (8080 by default).
inbound_connections	host_port	u32	Port that qproxy on the host listens to. For example, if **host_port** is set to **80**, qproxy on the host listens to 0.0.0.0:80.
	enclave_port	u32	Port that the qproxy on the QingTian Enclave instance listens to. For example, if **enclave_port** is set to **80**, the qproxy on the QingTian Enclave instance listens to 127.0.0.1:80.
	vsock_port	u32	Port used by the vsock in the QingTian Enclave instance. It must be unique and cannot conflict with the qproxy port number.

Log File Configuration

Table 13-4 Log file configuration

Variable	Type	Description
host_log	String	Log file name of qproxy on the host. For example, if **host_log** is set to **host.log**, the log file name of qproxy on the host is **host.log**.
enclave_log	String	Log file name of qproxy on the QingTian Enclave instance. For example, if **enclave_log** is set to **enclave.log**, the log file name of qproxy on the QingTian Enclave instance is **enclave.log**.
log_level	String	qproxy log level. For example, if **log_level** is set to **info**, logs of the "info", "warn", and "error" levels are displayed. All log levels are "off", "info", "warn", "error", "debug", and "trace".
host_log_dir	String	Log directory of qproxy on the host. The default value is **/var/log/qproxy**.
enclave_log_dir	String	Log directory of qproxy on the QingTian Enclave instance. The default value is **/var/log/qproxy**.

Configuration File Reference

```
[[outbound_connections]]
hostname = "api.myservice.com"
vsock_port = 7777
tcp_port = 443

[[outbound_connections]]
hostname = "another.api.com"
vsock_port = 7778
tcp_port = 5555

[[inbound_connections]]
host_port = 80
enclave_port = 80
vsock_port = 9000

[[inbound_connections]]
host_port = 443
enclave_port = 443
vsock_port = 9001

[log_location]
host_log = "host.log"
enclave_log = "enclave.log"
log_level = "info"
host_log_dir = "/var/log/qproxy"
enclave_log_dir = "/var/log/qproxy"
```

13.5.3 QingTian Enclave Log Forwarding Tool

Overview

QingTian Enclave log (qlog) is an O&M tool for QingTian Enclave. A QingTian Enclave instance is a completely isolated sub-VM running in a QingTian VM. Even the **root** user cannot log in to the QingTian Enclave instance via SSH. To help O&M personnel monitor services running in QingTian Enclave and locate faults, the qlog tool is provided. qlog can collect specified log files and resource usage (CPU and memory usages) of QingTian Enclave instances and send the collected data to the parent instance.

The qlog tool is an executable binary file. It needs to be executed using different commands on the parent instance and QingTian Enclave instance.

- On the parent instance, run **/path/to/qlog receive-file <cid>/path/to/config_qlog.toml** to enable qlog.
- On a QingTian Enclave instance, run **/path/to/qlog monitor /path/to/config_qlog.toml** to enable qlog.

Service logs of traditional VMs are stored in a directory similar to **/var/log/service.log**.

After services are migrated to QingTian Enclave, you can use qlog to export the log files (stored in **/var/log/service.log**) to the parent instances.

You can run **qlog receive-file** and **qlog monitor** on the parent instance and QingTian Enclave instance, respectively, to execute a given qlog binary file. A local vsock-based communication link is established between them. The qlog component running in a QingTian Enclave instance collects specified service logs or resource usage of the QingTian Enclave instance, and sends the collected data to the qlog component of the parent instance. The qlog component of the parent instance stores the received data in a specified directory, for example, **/var/log/qlog/service.log**.

Figure 13-2 Seamless migration of a network service to QingTian Enclave

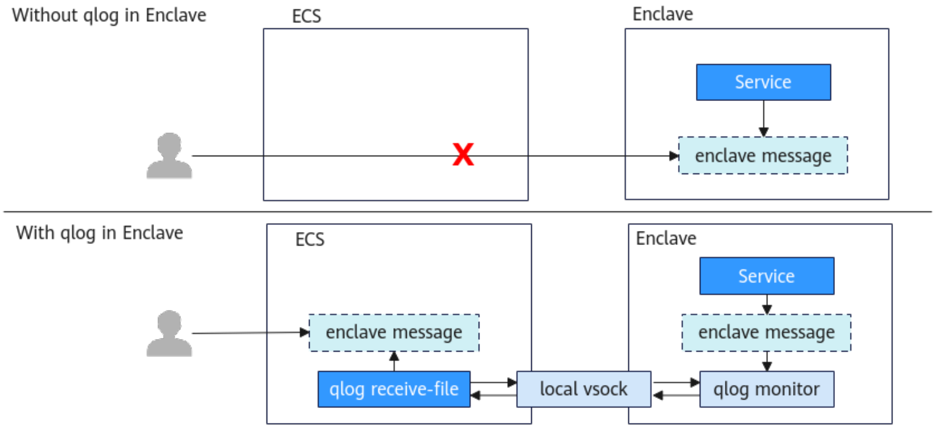

The following describes how to use qlog.

Prerequisites

1. You have obtained qlog by performing the following:

   Clone the QingTian Enclave code repository.

   ```
   git clone https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian.git
   ```

2. You have obtained the cargo tool chain by performing the following:

   a. Install rustup.

   ```
   curl --proto '=https' --tlsv1.2 -sSf https://sh.rustup.rs | sh
   ```

   b. After the installation is complete, load rustup.

   ```
   source $HOME/.cargo/env
   ```

   c. Check that rustc and cargo are installed.

   ```
   rustc -V
   cargo -V
   ```

3. You have learned about the pre-dependency of qlog.

Table 13-5 Pre-dependency

Dependency Item	Earliest Test Version
glibc	2.34
cargo	1.77.0

4. You have prepared the QingTian Enclave environment by performing the following:

   a. Install the qt CLI tool and required RPM packages.

   b. Install Docker.

   c. Install Python 3 and required Python modules (docker and knack).

   For details, see **Getting Started with QingTian Enclave** and **Installation of the qt CLI**.

Procedure

Step 1 Build qlog.

Go to the **qingtian-tools/qlog** directory and run the following cargo command:

```
cargo build --release
```

The generated qlog binary file is compiled in **qingtian-tools/qlog/target/release**.

Step 2 Create a working directory.

Create a working directory named workspace and copy the qlog binary file to the workspace. Store the files generated subsequently in the workspace.

Step 3 Configure **config_qlog.toml**.

In the workspace directory, create the **config_qlog.toml** file with the following content:

```
port: 6000
workspace: /var/log/qlog
server_logfile: server.log
client_logfile: client.log
server_threads: 4
client_threads: 1
log_level: info
rotate_size: 65536
rotate_num: 10
monitor_items:
  - name: service  #Service name
    monitor_type: file
    monitor_path: /var/log/service.log  #QingTian Enclave service log path
    outputfile: service.log  #Name of the log file synchronized to the parent instance
  - name: resource
    monitor_type: resource
    monitor_internel: 15
    outputfile: resource.log
```

Step 4 Create a QingTian Enclave image that contains qlog.

1. In the workspace directory, create the **start.sh** script.

```
#/bin/bash
/root/qlog monitor /root/config_qlog.toml &

LOG_FILE="/var/log/service.log"
LOG_MESSAGE="Hello, service."

while true; do
    TIMESTAMP=$(date '+%Y-%m-%d %H:%M:%S')
    echo "$TIMESTAMP - $LOG_MESSAGE" >> "$LOG_FILE"
    sleep 3
done
```

2. Run the following command in the workspace to grant **start.sh** execute permissions:

```
chmod +x start.sh
```

3. In the workspace directory, create the **Dockerfile** file with the following content:

```
FROM ubuntu:latest
COPY ./qlog /root/qlog
COPY ./config_qlog.toml /root/config_qlog.toml
COPY ./start.sh /root/start.sh
CMD /root/start.sh
```

4. Run the following command in the workspace to grant **start.sh** execute permissions:

```
chmod +x start.sh
```

5. Run the following command in the workspace to create a Docker image:

```
docker build -f Dockerfile -t test_qlog_enclave .
```

6. Run the following command in the workspace to create a QingTian Enclave image:

```
qt enclave make-img --docker-uri test_qlog_enclave --eif test_qlog_enclave.eif
```

Step 5 Start qlog.

1. Run the following command in the workspace to start a QingTian Enclave instance:

```
qt enclave start --cpus 2 --mem 1024 --cid 4 --eif test_qlog_enclave.eif
```

2. Run the following command in the workspace to start qlog in the parent instance:

```
./qlog receive-file 4 ./config_qlog.toml &
```

3. Run the following command on the parent instance to view the service log:

```
tail -F /var/log/qlog/service.log
```

A line of "Hello, service." is printed every 3 seconds.

4. Run the following command on the parent instance to check the resource usage of the QingTian Enclave instance:

```
tail -F /var/log/qlog/resource.log
```

The CPU usage and memory usage are printed every 15 seconds.

----End

qlog Help Information

qlog help

```
$ qlog --help
A tool to monitor logs and resource usage over a Vsock connection

Usage: qlog <COMMAND>

Commands:
  monitor       Monitor resource usage
  receive-file  Receive data from qlog monitor
  help          Print this message or the help of the given subcommand(s)

Options:
  -h, --help  Print help
```

qlog monitor help

```
$ qlog help monitor
Monitor logs and/or resource usage

Usage: qlog monitor [OPTIONS] <CONFIG>

Arguments:
  <CONFIG>  A configuration file in yaml format, which defines items to be monitored

Options:
  -c, --cid <CID>  CID to listen on (defaults to VMADDR_CID_ANY) [default: 4294967295]
  -h, --help       Print help
```

qlog receive-file help

```
$ qlog help receive-file
Receive data from qlog monitor

Usage: qlog receive-file [OPTIONS] <CID> <CONFIG>
```

```
Arguments:
  <CID>      Enclave VM's CID
  <CONFIG>  A configuration file in yaml format, which defines items to be monitored

Options:
  -w, --workspace <WORKSPACE>  Set to workspace if specified, prior to configuration file
  -h, --help                   Print help
```

Configuration Information

Table 13-6 Configuration parameters

Variable	Type	Description
port	u32	Vsock port
workspace	String	Workspace for writing process logs and monitoring data
server_logfile	String	Writing process logs of the server (qlog monitor) to <workspace>/<server_logfile>
client_logfile	String	Writing process logs of the client (qlog receive-file) to <workspace>/<client_logfile>
server_threads	u32	Number of threads on the server (qlog monitor)
client_threads	u32	Number of threads on the client (qlog receive-file)
log_level	String	Log levels (TRACE, DEBUG, INFO, WARN, ERROR, and OFF)
rotate_size	u32	(Optional) Log file dump size. The default value is 2 MB.
rotate_num	u32	(Optional) Number of old log files to be retained during log file dump. The default value is 10.
monitor_items	Vector	Monitored item list

Table 13-7 Monitoring item parameters

Variable	Type	Description
name	String	Name of the monitoring item.
monitor_type	String	Monitoring type. The value **file** indicates log files, and the value **resource** indicates resource usage.
monitor_path	String	(Optional) Path of the log file to be monitored in the QingTian Enclave instance. This parameter can be specified only when **monitor_type** is **file**.
monitor_internel	u32	(Optional) Resource monitoring interval, in seconds. The default value is 15. This parameter can be specified only when **monitor_type** is **resource**.
outputfile	String	Name of a monitoring data file. The path for writing monitoring data is <workspace>/<outputfile>.

Configuration File Reference

```
port: 6000
workspace: /var/log/qlog
server_logfile: server.log
client_logfile: client.log
server_threads: 4
client_threads: 1
log_level: info
rotate_size: 65536
rotate_num: 10
monitor_items:
  - name: item1
    monitor_type: file
    monitor_path: /var/log/item1.log
    outputfile: output1.log
  - name: item2
    monitor_type: file
    monitor_path: /var/log/item2.log
    outputfile: output2.log
  - name: item3
    monitor_type: file
    monitor_path: /var/log/item3.log
    outputfile: output3.log
  - name: item4
    monitor_type: file
    monitor_path: /var/log/item4.log
    outputfile: output4.log
  - name: item5
    monitor_type: resource
    monitor_internel: 15
    outputfile: output5.log
```

13.6 QingTian CLI (qt CLI)

13.6.1 Installation of the qt CLI

If you intend to install the qt CLI on a parent instance using Linux images other than Huawei Cloud EulerOS, compile and install it in the Huawei Cloud QingTian open-source repository. If you intend to install the qt CLI on a parent instance using the Huawei Cloud EulerOS image, run the following command:

```
yum install qingtian-tool
```

The RPM package contains the following parts:

qt-enclave-env: supports resource isolation. Before the QingTian Enclave instances are created, the vCPUs and memory of a parent instance need to be allocated to QingTian Enclave instances to build an isolated, secure runtime environment.

qt CLI: a QingTian command line tool to build QingTian Enclave image files required for booting QingTian Enclave instances and manage the lifecycle of the instances.

Note that before using the qt CLI, you need to install python3 and the necessary python modules, such as docker and knack. Run the following command to install the necessary modules:

```
pip3 install docker knack
```

13.6.2 Introduction to qt-enclave-env

The qt-enclave-env is a service. After the service is started, it reads information to be isolated from the **qt-enclave-env.conf** configuration file and isolates resources. Resources must be isolated before the QingTian Enclave instances are created.

The following describes the content in the configuration file **/etc/qingtian/enclave/qt-enclave-env.conf**.

```
#Configure the size of hugepages to be isolated for the QingTian Enclave instance. The value can be 2 or 1024, indicating 2 MiB or 1 GiB hugepages, respectively.
hugepage_size:1024
# Configure the size of the memory to be isolated. The value must be an integer multiple of the hugepage size.
memory_mib:1024
# Configure the number of vCPUs to be isolated. This configuration item and cpu_list are mutually exclusive, or the service will fail to be started.
cpu_count:2
# Configure a list of vCPUs to be isolated. A CPU ID other than 0 can be entered. This configuration item and cpu_count are mutually exclusive, or the service will fail to be started.
# cpu_list:2,3
```

Note that whether the hugepage memory is successfully reserved by the qt-enclave-env service is affected by memory fragmentation of the parent instance. If the system has been running for a long time or the qt-enclave-env service is restarted repeatedly, the hugepage memory may fail to be reserved. To avoid this issue, you are advised to start the qt-enclave-env once after the system is started, which helps to reserve sufficient memory.

13.6.3 Introduction to qt enclave Subcommands

qt is a level-1 command. It contains a level-2 subcommand **enclave**.

```
[root@localhost ~]# qt
  ____  _          _______ _
 / __ \(_)        |__   __(_)
| |  | |_ _ __   __ _| |   _  __ _ _ __
| |  | | | '_ \ / _` | |  | |/ _` | '_ \
| |__| | | | | | (_| | |  | | (_| | | | |
 \___\_\_|_| |_|\__, |_|  |_|\__,_|_| |_|
                 __/ |
                |___/

Welcome to the cool QingTian new CLI!

    enclave : Enclave life-circle management.
```

qt enclave contains subcommands for building QingTian Enclave image files, and starting, stopping, and querying QingTian Enclave instances.

```
[root@localhost ~]# qt enclave
usage: qt enclave [-h] {make-img,start,stop,query,console} ...
qt enclave: error: the following arguments are required: _subcommand
enclave command line interface
[root@localhost ~]# qt enclave -h

Group
  qt enclave : Enclave life-circle management.

Commands:
  console  : Console an enclave via the enclave-id while debugging.
  make-img : Make an eif image from a docker image.
  query    : Query an enclave via the enclave-id or query all enclaves.
  start    : Start an enclave via an eif image.
  stop     : Stop an enclave via the enclave-id.
```

qt enclave make-img

This command is used to convert a Docker image to a QingTian Enclave image file. The command format is as follows:

```
[root@localhost ~]# qt enclave make-img -h

Command
    qt enclave make-img : Make an eif image from a docker image.

Arguments
    --docker-uri [Required]
    --eif        [Required]
    --private-key
    --signing-certificate

Global Arguments
    --debug              : Increase logging verbosity to show all debug logs.
    --help -h            : Show this help message and exit.
    --only-show-errors   : Only show errors, suppressing warnings.
    --output -o          : Output format.  Allowed values: json, jsonc, none, table, tsv, yaml,
                      yamlc.  Default: json.
    --query             : JMESPath query string. See http://jmespath.org/ for more information
                      and examples.
    --verbose           : Increase logging verbosity. Use --debug for full debug logs.

Examples
    Given docker-uri and eif to make an eif image
      qt enclave make-img --docker-uri [DOCKER-URI] --eif [EIF]

    Make an eif image with private-key and signing-certificate
      qt enclave make-img --docker-uri [DOCKER-URI] --eif [EIF] --private-key [PRIVATE-KEY]
      --signing-certificate [SIGNING-CERTIFICATE]
```

Mandatory: **--docker-uri**, which specifies the Uniform Resource Identifier (URI) of the Docker image in a Docker repository. You can run the **docker image ls** command to query the URI of the current local image.

Mandatory: **--eif**, which specifies the path used to store the generated EIF.

Optional: **--private-key**, which specifies the absolute path of the private key used to sign the QingTian Enclave image. If you specify **PRIVATE-KEY**, you must also specify **SIGNING-CERTIFICATE**.

Optional: **--signing-certificate**, which specifies the absolute path of the certificate used to sign the QingTian Enclave image. If you specify **SIGNING-CERTIFICATE**, you must also specify **PRIVATE-KEY**.

Returned value: If the preceding two optional parameters are configured, ensure that the certificates are valid. If the certificates are valid, the command output contains additional PCR0 and PCR8, which is used for measuring the QingTian Enclave image and signature certificate. If the certificates are invalid, the QingTian Enclave image fails to be built.

Example command of building an image:

```
[root@localhost docker]# qt enclave make-img --docker-uri ubuntu --eif /home/docker/ubuntu.eif --private-
key  /home/docker/private-key.pem --signing-certificate /home/docker/server.pem
{
   "digest":       "SHA384",
   "PCR0":
"b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c0364a9b88dd5
8017a62"
   "PCR8":
"dbf4a7f9fab7f18619b5899c407081981ad6762fb9a809da78548821b5021965423181584acd7b201703376f11
33a546"
}
```

qt enclave start

This command is used to launch a QingTian Enclave instance. The command format is as follows:

```
[root@localhost ~]# qt enclave start -h

Command
  qt enclave start : Start an enclave via an eif image.

Arguments
  --cid        : Default: 4.
  --eif        [Required]
  --cpus       : Default: 2.
  --debug-mode
  --mem        : Default: 1024.

Global Arguments
```

```
 --debug       : Increase logging verbosity to show all debug logs.
 --help -h      : Show this help message and exit.
 --only-show-errors : Only show errors, suppressing warnings.
 --output -o     : Output format.  Allowed values: json, jsonc, none, table, tsv, yaml, yamlc.
           Default: json.
 --query       : JMESPath query string. See http://jmespath.org/ for more information and
           examples.
 --verbose      : Increase logging verbosity. Use --debug for full debug logs.

Examples
 Given an eif image, an unused cid, the number of cpus and memory needed
   qt enclave start  [--cpus CPUS] [--mem MEM] --eif EIF [--cid CID]
```

Optional: **--cpus**, which specifies the number of vCPUs to be allocated to the QingTian Enclave instance. The value cannot be greater than the number of isolated vCPUs. If this parameter is not specified, the default value is 2.

Optional: **--mem**, which specifies the memory size (MiB) allocated to the QingTian Enclave instance. The value cannot be greater than the isolated memory size and must be greater than the QingTian Enclave image size. If this parameter is not specified, the default value is 1024 MiB.

Mandatory: **--eif**, which specifies the path of the EIF.

Optional: **--cid**, which specifies the context identifier (CID) of the QingTian Enclave instance. The CID is the socket IP address for communication between the parent instance and the QingTian Enclave instance using vsock. The available CID range is from 4 to 4294967294. If this parameter is not specified, the default value is 4.

Optional: **--debug-mode**, which specifies whether to start the QingTian Enclave instance in debug mode. If you enable debug mode, PCRs that are made up entirely of zeros can be used to collect and print internal logs of QingTian Enclave instances.

Returned value: Details of the created QingTian Enclave instance

Example command of launching a QingTian Enclave instance:

```
qt enclave start --cpus 2 --mem 1024 --eif /home/docker/ubuntu.eif --cid 4
```

qt enclave query

This command is used to query information about the QingTian Enclave instance on a parent instance. The command format is as follows:

```
[root@localhost ~]# qt enclave query -h

Command
 qt enclave query : Query an enclave via the enclave-id or query all enclaves.

Arguments
 --enclave-id

Global Arguments
 --debug       : Increase logging verbosity to show all debug logs.
 --help -h      : Show this help message and exit.
 --only-show-errors : Only show errors, suppressing warnings.
 --output -o     : Output format.  Allowed values: json, jsonc, none, table, tsv, yaml, yamlc.
           Default: json.
 --query       : JMESPath query string. See http://jmespath.org/ for more information and
           examples.
 --verbose      : Increase logging verbosity. Use --debug for full debug logs.
```

```
Examples
  Given an enclave-id to query an enclave
      qt enclave query --enclave-id [ENCLAVE-ID]

  Query all enclaves without enclave-id
    qt enclave query
```

Optional: **--enclave-id**. If this parameter is specified, information about the specified QingTian Enclave instance is queried. If this parameter is not specified, information about all existing QingTian Enclave instances is queried.

The returned value is the information about the queried QingTian Enclave instance.

- **EnclaveID**: specifies the ID of the QingTian Enclave instance.
- **ProcessID**: specifies the process identifier (PID) of the process holding the QingTian Enclave instance's resources in the parent instance.
- **EnclaveCID**: specifies the vsock socket ID used for communication between the QingTian Enclave instance and the parent instance.
- **NumberOfCPUs**: specifies the number of vCPUs allocated from the parent instance to the QingTian Enclave instance.
- **MemoryMiB**: specifies the memory size (MiB) allocated from the parent instance to the QingTian Enclave instance.

Example command of querying a QingTian Enclave instance:

```
[root@localhost ~]#qt enclave query
[{
        "EnclaveID":    0,
        "ProcessID":    29990,
        "EnclaveCID":   4,
        "NumberOfCPUs": 2,
        "MemoryMiB":    1024,
        "LaunchMode":   "debug"
   }]
```

If there are no QingTian Enclave instances available, the command output is empty.

If the **--enclave-id** parameter is specified but the QingTian Enclave instance identified by the specified **--enclave-id** is not found, the command output is empty.

qt enclave stop

This command is used to stop a QingTian Enclave instance. The command format is as follows:

```
[root@localhost ~]# qt enclave stop -h

Command
  qt enclave stop : Stop an enclave via the enclave-id.

Arguments
  --enclave-id [Required]

Global Arguments
  --debug          : Increase logging verbosity to show all debug logs.
  --help -h        : Show this help message and exit.
  --only-show-errors    : Only show errors, suppressing warnings.
  --output -o       : Output format.  Allowed values: json, jsonc, none, table, tsv, yaml,
               yamlc.  Default: json.
  --query        : JMESPath query string. See http://jmespath.org/ for more information
              and examples.
  --verbose       : Increase logging verbosity. Use --debug for full debug logs.

Examples
  Given an enclave-id to stop an enclave
   qt enclave stop --enclave-id [ENCLAVE-ID]
```

Mandatory: **--enclave-id**, which identifies the QingTian Enclave instance to be stopped.

Returned value: If a successful message is returned, the instance is stopped. If no message is returned, the instance failed to be stopped.

Example command of stopping a QingTian Enclave instance:

```
[root@localhost ~]# qt enclave stop --enclave-id 1
stop 1 success
```

qt enclave console

This command is used to view the read-only console output of the QingTian Enclave instance in the parent instance when the instance is started in debug mode. The command format is as follows:

```
[root@localhost ~]# qt enclave console -h

Command
    qt enclave console : Console an enclave via the enclave-id while debugging.

Arguments
    --enclave-id [Required]

Global Arguments
    --debug              : Increase logging verbosity to show all debug logs.
    --help -h            : Show this help message and exit.
    --only-show-errors   : Only show errors, suppressing warnings.
    --output -o          : Output format.  Allowed values: json, jsonc, none, table, tsv, yaml,
                          yamlc.  Default: json.
    --query              : JMESPath query string. See http://jmespath.org/ for more information
                          and examples.
    --verbose            : Increase logging verbosity. Use --debug for full debug logs.

Examples
    Given an enclave-id to console an enclave
        qt enclave console --enclave-id [ENCLAVE-ID]
```

Mandatory: **--enclave-id**, which specifies the enclave-id of the QingTian Enclave instance whose read-only console output is to be obtained.

After the command is executed successfully, the read-only console output of the QingTian Enclave instance is displayed as follows:

```
hello enclave!
hello enclave!
hello enclave!
hello enclave!
```

You can press **Ctrl+C** to exit the command. Note that the **qt enclave console** command can be executed on only one specified QingTian Enclave instance at a time.

13.7 QingTian Error Code

Error Code	Error Message	Description	Solution
01	Missing necessary argument.	Mandatory parameters are missing.	Check the command parameters.
02	Invalid argument provided.	Invalid parameter.	Check the command parameters.
03	File operation failure.	File operation error.	Check whether the target file or directory exists.
04	Ioctl get sandbox capacity failure.	Failed to get sandbox capacity using Ioctl.	Contact Huawei Cloud technical support.
05	Ioctl define sandbox failure.	Failed to define sandbox using Ioctl.	Contact Huawei Cloud technical support.
06	Invalid parameters provided in configuration file.	Invalid parameters exist in the configuration file.	Check the corresponding configuration file.
07	Missing necessary parameters in configuration file.	Mandatory parameters are missing in the configuration file.	Check the corresponding configuration file.
08	Mmap memory failure.	Mmap memory error.	Contact Huawei Cloud technical support.
09	Ioctl add memory failure.	Failed to increase memory using Ioctl.	Contact Huawei Cloud technical support.
10	Load image failure because provided memory is too small.	Failed to load the image. The possible cause is that the memory size is insufficient.	Add the memory setting when launching QingTian Enclave instances.

Error Code	Error Message	Description	Solution
11	Ioctl add cpu failure.	Failed to add vCPUs using Ioctl.	Contact Huawei Cloud technical support.
12	Lock acquire failure.	Failed to acquire locks.	View the qt CLI logs and check whether permissions are correctly configured for lock files.
13	Socket initialization failure.	Failed to initialize the socket.	Contact Huawei Cloud technical support.
14	Socket binding failure.	Failed to bind the socket.	Contact Huawei Cloud technical support.
15	Socket listen failure.	Failed to listen to the socket.	Contact Huawei Cloud technical support.
16	Socket accept failure.	Failed to receive the socket execution.	Contact Huawei Cloud technical support.
17	Write heartbeat to the enclave failure.	An error occurred when the heartbeat message is written to the QingTian Enclave instance.	Contact Huawei Cloud technical support.
18	Read heartbeat from the enclave failure.	An error occurred when the heartbeat messages is read from the QingTian Enclave instance.	Contact Huawei Cloud technical support.
19	Ioctl start an enclave failure.	Failed to start the QingTian Enclave instance using Ioctl.	Contact Huawei Cloud technical support.
20	Wait heartbeat timeout.	Waiting for the heartbeat messages timed out.	Add the memory setting when launching QingTian Enclave instances. If the fault persists, contact Huawei Cloud technical support.

Error Code	Error Message	Description	Solution
21	Get json print object failure.	Failed to obtain the JSON printing object.	Check whether the cjson library is normal.
22	Write enclave's configuration file failure.	Failed to generate the QingTian Enclave configuration file.	View the qt CLI logs and check whether permissions are correctly configured for the lock files. If the fault persists, contact Huawei Cloud technical support.
23	Socket connection failure	Incorrect Socket connection	Contact Huawei Cloud technical support.
24	Write cmd to the enclave server failure.	Failed to write commands to the QingTian Enclave server.	Contact Huawei Cloud technical support.
25	Read message from the enclave server failure.	Failed to obtain the QingTian Enclave server information.	Contact Huawei Cloud technical support.
26	Create cjson object failure.	Failed to create the cjson object.	Check whether the cjson library is normal.
27	Create cjson array failure.	Failed to create the cjson array.	Check whether the cjson library is normal.
28	The required enclave is not running.	The requested QingTian Enclave instance is not running.	Run the **qt enclave query** command to query the running QingTian Enclave instance.
29	Invalid enclave pid.	Invalid QingTian Enclave PID.	Contact Huawei Cloud technical support.
30	Add number into cjson object failure.	Failed to add a number to the cjson printing object.	Check whether the cjson library is normal.
31	Add string into cjson object failure.	Failed to add a character string to the cjson printing object.	Check whether the cjson library is normal.

Error Code	Error Message	Description	Solution
32	The required enclave is not running in the debug mode.	The requested QingTian Enclave instance is not running in debug mode.	Run the **qt enclave query** command to query the running mode of the QingTian Enclave instance.
33	Enclave console read failure.	Failed to read commands on the QingTian Enclave instance console.	Contact Huawei Cloud technical support.
34	Write img header failure.	Failed to write the image header file during image creation.	Contact Huawei Cloud technical support.
35	Write cmdline failure.	Failed to write cmdline during image creation.	Contact Huawei Cloud technical support.
36	Write kernel failure.	Failed to write the kernel during image creation.	Contact Huawei Cloud technical support.
37	Write initrd failure.	Failed to write initrd during image creation.	Contact Huawei Cloud technical support.
38	Write certificate failure.	Failed to write the certificate during image creation.	Contact Huawei Cloud technical support.
39	Get pcr failure.	Failed to obtain the PCR value.	Contact Huawei Cloud technical support.
40	Add signature failure.	An error occurred for image signature during image creation.	Contact Huawei Cloud technical support.

Error Code	Error Message	Description	Solution
41	Check enclave image info failure while building an eif file.	The enclave image information is abnormal during the building of the QingTian Enclave image file.	Contact Huawei Cloud technical support.
42	The required enclave is in maintenance state.	The requested QingTian Enclave instance is in the maintenance phase.	Contact Huawei Cloud technical support.
43	The cid has already been used.	The CID is in use.	Specify an idle CID.

13.8 FAQs About QingTian

13.8.1 General Questions

What Is QingTian Enclave?

QingTian Enclave provides an isolated and highly-constrained environment where you can deploy your security-sensitive applications to reduce the attack surface area.

What Are the Advantages of QingTian Enclave?

QingTian Enclave allows you to create isolated compute environments from general ECSs to process your highly sensitive data.

QingTian Enclave instances are completely independent VMs and have no persistent storage, interactive access, or external networking. They communicate with your ECSs through a secure local channel.

When Should I Use QingTian Enclave?

When you process security-sensitive data and want the data to be isolated from users, applications, or third-party libraries, you can use QingTian Enclave to provide an independent, isolated environment for your data processing.

You can develop and run various applications in QingTian Enclave, such as personal privacy information processing, proprietary code and algorithm operation, and multi-party computation.

How Do I Get Started with QingTian Enclave?

You can refer to **Getting Started with QingTian Enclave** to start your journey with QingTian Enclave.

What Is vCPU and Memory Isolation?

vCPU and memory isolation prevents users, applications, and third-party libraries on the parent instance from directly accessing the vCPUs and memory of QingTian Enclave instances. You can use the QingTian CLI (qt CLI) to boot a QingTian Enclave instance with isolated vCPUs and memory. For details, see **QingTian CLI (qt CLI)**.

How Are vCPUs and Memory of QingTian Enclave Instances Isolated from Their Parent Instance?

QingTian Enclave uses the verified vCPU-based technology for isolation, combined with the unique design of the QingTian architecture and a root of trust based on Huawei-developed iNIC. The QingTian Hypervisor, which is developed and designed by Huawei Cloud, can divide physical resources on a server into partitions. It discards all unnecessary functions compared with other virtualization technologies. QingTian Enclave extends the isolation capabilities of the QingTian Hypervisor to protect and isolate the vCPUs and memory allocated to QingTian Enclave instances from those of the parent instance, creating isolated execution environments.

Which Instance Types Support QingTian Enclave?

Currently, C7t and kC2 ECSs support QingTian Enclave.

What Is an Attestation Document?

An attestation document is used to verify the reliability measurement results of QingTian Enclave instances. An attestation document is generated by the QingTian Hypervisor. It includes the platform Configuration Register (PCR) list, the QingTian Public Key Infrastructure (PKI) certificate chain, cryptographic algorithm declaration, and user-defined data for the QingTian Enclave application.

The attestation document is signed by the Huawei Cloud QingTian Attestation PKI. Huawei Cloud Key Management Service (KMS) has built-in attestation support for QingTian Enclave instances. You can use the Huawei Cloud KMS APIs included in the QingTian Enclave SDK to perform KMS options, such as decryption, random number generation, and encryption in QingTian Enclave instances based on the attestation document. KMS can ingest attestation documents from QingTian Enclave instances and validates the measurements in the attestation documents against these specified in the IAM policies to determine whether QingTian Enclave instances can access KMS APIs.

What Is the Root of Trust of QingTian Enclave's Attestation Document and How Can I Verify It?

The attestation document is signed by the Huawei Cloud QingTian Attestation PKI. You can download the QingTian Enclave's root certificate at **https://qingtian-**

enclave.obs.myhuaweicloud.com/huawei_qingtian-enclaves_root-G1.zip. For details about how to verify the document signature, see **Document Signature Verification**.

How Is QingTian Enclave Billed?

Currently, QingTian Enclave is free, and you only need to pay for the ECSs you purchase.

Why Does the Isolation Command (systemctl start qt-enclave-env) Fail?

During system runtime, fragmented memory is inevitably generated. As a result, continuous huge pages cannot be obtained during service isolation.

In this case, run **systemctl status qt-enclave-env** to check whether the error log contains **allocating hugepages error**.

If the error log contains **allocating hugepages error**, the number of available continuous huge pages provided by the system is less than the expected number (**$wanted_mem_num**).

You can use either of the following solutions to handle this issue:

- Solution 1:
  1. Check the maximum number of available continuous huge pages (**$free_mem_num**) provided by the system.

     ```
     cat /sys/devices/system/node/node0/hugepages/hugepages-1048576kB/free_hugepages
     ```
  2. Modify the **/etc/qingtian/enclave/qt-enclave-env.conf** configuration file to ensure that the value of **memory_mib** is less than the value of **$free_mem_num** multiplied by 1024.

     ```
     vim /etc/qingtian/enclave/qt-enclave-env.conf
     ```

     To prevent memory fragmentation caused by repeated executions of the isolation command, you are advised to execute the resource isolation command immediately after the system is started.
- Solution 2:
  1. Modify the **/etc/default/grub** file.

     ```
     vim /etc/default/grub
     ```

     Add **default_hugepagesz=1G hugepagesz=1G hugepages=$wanted_mem_num** to the Linux command line parameter **GRUB_CMDLINE_LINUX** in the GRUB file.
  2. Apply the modification.

     ```
     grub2-mkconfig -o /boot/efi/EFI/hce/grub.cfg
     ```
  3. Restart the VM and check the number of available huge pages again.

     ```
     reboot
     ```

13.8.2 Development and Deployment Questions

How Many QingTian Enclave Instances Can I Create from an ECS?

You can create a maximum of two QingTian Enclave instances from an ECS.

What Is Vsock and How Can I Use it to Communicate With a QingTian Enclave Instance?

Vsock is a type of socket interface defined by a context identifier (CID) and port number. The CID is the same as the IP address in a TCP/IP connection.

Vsock communicates with a QingTian Enclave instance using standard and well-defined POSIX Socket APIs (for example, connect, listen, and accept). You can develop your own QingTian Enclave applications using vsock. For details, see **QingTian Enclave Application Development on Linux**. Applications can also send HTTP requests using vsock through a proxy.

Why Does the Vsock Performance Deteriorate After QingTian Enclave Instances Are Launched from kC2 Instances?

For kC2 instances, if all isolated vCPUs are used to launch QingTian Enclave instances, the vsock performance will deteriorate.

When enabling the qt-enclave-env service, you are advised to isolate two more vCPUs for QingTian Enclave. Some vCPUs are used to launch QingTian Enclave instances and some are used to forward data through the vsock channel. This helps prevent the vsock performance deterioration.

1. Modify the **/etc/qingtian/enclave/qt-enclave-env.conf** configuration file of the qt-enclave-env service.

   vim /etc/qingtian/enclave/qt-enclave-env.conf

   - Method 1: Change the value of **cpu_count** to the number of the QingTian Enclave's vCPUs plus 2.
   - Method 2: Change the number of vCPUs in **cpu_list** to the number of the QingTian Enclave's vCPUs plus 2.
2. Restart the qt-enclave-env service.

   systemctl restart qt-enclave-env
3. Restart the QingTian Enclave instance.

   qt enclave start --cpus ${isolated_cpus_count-2} --mem ${wanted_mem} --eif ${eif_file_location} --cid ${wanted_cid}

   Where:

   - **isolated_cpus_count-2** indicates the number of isolated vCPUs minus 2.
   - **wanted_mem** indicates the expected memory size.
   - **eif_file_location** indicates the eif file location.
   - **wanted_cid** indicates the expected CID value.

   Restart the QingTian Enclave instance and check the vsock performance.

Why Does the qt-enclave-env Service Fail to Be Started After SELinux Is Enabled on an ECS?

Symptom: After SELinux is enabled on an ECS, the qt-enclave-env service fails to be started. The message "insmod virtio-qtbox.ko Permission denied" is displayed in the qt-enclave-env service logs.

Possible Cause: SELinux provides powerful security mechanisms including mandatory access control, fine-grained access control, policy enforcement, type enforcement, security context, and auditing to protect the Linux system from malicious attacks and data leakage threats. As a result, the qt-enclave-env service cannot directly use the **insmod virtio-qtbox.ko** command to insert the kernel module.

Solution: Run the **insmod /opt/qingtian/enclave/virtio-qtbox.ko** command or disable SELinux first and then restart the qt-enclave-env service.