

Elastic Cloud Server

User Guide

Issue	01
Date	2025-07-25



Copyright © Huawei Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Using IAM to Grant Access to ECS..... 1

1.1 Creating a User and Granting ECS Permissions..... 1

1.2 ECS Custom Policies.....2

2 Instances..... 5

2.1 Overview..... 5

2.2 Selecting an ECS Billing Mode..... 7

2.2.1 Yearly/Monthly Billing..... 7

2.2.2 Pay-per-Use Billing..... 8

2.2.3 Spot Pricing..... 9

2.2.3.1 Spot Pricing ECSs..... 9

2.2.4 Reserved Instances..... 12

2.2.4.1 Reserved Instance Overview..... 12

2.2.4.2 Enabling and Purchasing a Reserved Instance.....18

2.2.4.3 Modifying RI Attributes.....22

2.2.5 Changing Pay-per-Use to Yearly/Monthly..... 23

2.2.6 Changing Yearly/Monthly to Pay-per-Use..... 26

2.3 Purchasing an ECS..... 27

2.3.1 Introducing ECS Purchase Options..... 27

2.3.2 Purchasing an ECS in Quick Config Mode..... 29

2.3.3 Purchasing an ECS in Custom Config Mode.....34

2.3.4 Purchasing a Spot ECS..... 65

2.3.5 Purchasing a Spot Block ECS..... 69

2.3.6 Purchasing an ECS Using a Private Image..... 72

2.3.7 Purchasing ECSs Using Auto Launch Groups..... 74

2.3.8 Purchasing an ECS in a Shared Subnet..... 80

2.3.9 Purchasing the Same ECS..... 82

2.4 Logging In to a Windows ECS..... 83

2.4.1 Login Overview (Windows)..... 83

2.4.2 Logging In to a Windows ECS Using an RDP File..... 84

2.4.3 Logging In to a Windows ECS Using VNC..... 87

2.4.4 Logging In to a Windows ECS Using MSTSC..... 88

2.4.5 Logging In to a Windows ECS from a Linux Computer.....95

2.4.6 Logging In to a Windows ECS from a macOS Server..... 97

Issue 01 (2025-07-25)

Copyright © Huawei Technologies Co., Ltd.

iii

2.4.7 Logging In to a Windows ECS from a Mobile Terminal.....	100
2.5 Logging In to a Linux ECS.....	104
2.5.1 Login Overview (Linux).....	104
2.5.2 Logging In to a Linux ECS Using CloudShell.....	106
2.5.3 Logging In to a Linux ECS Using VNC.....	109
2.5.4 Logging In to a Linux ECS Using an SSH Key Pair.....	111
2.5.5 Logging In to a Linux ECS Using an SSH Password.....	118
2.5.6 Logging In to a Linux ECS from a macOS Server.....	120
2.5.7 Logging In to a Linux ECS from a Mobile Terminal.....	121
2.6 Managing GPU Drivers of GPU-accelerated ECSs.....	134
2.6.1 GPU Driver.....	134
2.6.2 Obtaining a Tesla Driver and CUDA Toolkit.....	135
2.6.3 Manually Installing a GRID Driver on a GPU-accelerated ECS.....	136
2.6.4 Manually Installing a Tesla Driver on a GPU-accelerated ECS.....	146
2.6.5 Uninstalling a GPU Driver from a GPU-accelerated ECS.....	160
2.7 Managing ECS Configurations.....	164
2.7.1 Changing the Time Zone for an ECS.....	164
2.7.2 Enabling or Disabling Hyper-Threading.....	167
2.7.3 Obtaining Metadata and Passing User Data.....	170
2.7.3.1 Obtaining Metadata.....	171
2.7.3.2 Injecting User Data.....	180
2.7.4 Changing ECS Names.....	188
2.7.5 Migrating an ECS to a DeH.....	188
2.7.6 Migrating ECSs Across AZs.....	189
2.7.7 Managing ECS Groups.....	191
2.7.8 Configuring Mapping Between Hostnames and IP Addresses in the Same VPC.....	193
2.7.9 Starting and Stopping ECSs.....	195
2.8 Modifying ECS Specifications (vCPUs and Memory).....	197
2.8.1 Modifying Specifications of Individual ECSs.....	197
2.8.2 Changing a Xen ECS to a KVM ECS (Windows).....	201
2.8.3 Automatically Changing a Xen ECS to a KVM ECS (Linux).....	207
2.8.4 Manually Changing a Xen ECS to a KVM ECS (Linux).....	212
2.8.5 Automatically Changing Xen ECSs to KVM ECSs in a Batch (Linux).....	217
2.8.6 Changing a KVM ECS to a QingTian ECS (Windows).....	224
2.8.7 Changing a KVM ECS to a QingTian ECS (Linux).....	227
2.9 Reinstalling or Changing the OS.....	231
2.9.1 Reinstalling the OS.....	231
2.9.2 Changing the OS.....	233
2.10 Viewing ECS Information.....	237
2.10.1 Viewing the ECS Creation Status.....	237
2.10.2 Viewing Failed Tasks.....	237
2.10.3 Viewing ECS Details.....	238

2.10.4 Exporting ECS Information.....	239
2.10.5 Searching for ECSs.....	240
3 Recycle Bin (OBT).....	247
3.1 Recycle Bin Overview.....	247
3.2 Enabling Recycle Bin.....	255
3.3 Disabling Recycle Bin.....	256
3.4 Configuring a Recycle Bin Policy.....	257
3.5 Recovering an ECS from the Recycle Bin.....	258
3.6 Permanently Deleting ECSs from the Recycle Bin.....	259
3.7 Exporting ECS Information from the Recycle Bin.....	259
4 Images.....	261
4.1 Overview.....	261
4.2 Creating an Image.....	263
5 Disks.....	265
5.1 Overview.....	265
5.2 Adding a Disk to an ECS.....	267
5.3 Attaching a Disk to an ECS.....	269
5.4 Initializing Data Disks.....	270
5.5 Adding a Yearly/Monthly EVS Disk.....	272
5.6 Detaching an EVS Disk from a Running ECS.....	273
5.7 Expanding the Capacity of an EVS Disk.....	275
5.8 Expanding the Local Disks of a Disk-intensive ECS.....	276
5.9 Enabling Advanced Disk.....	277
6 Elastic Network Interfaces.....	279
6.1 Network Interface Overview.....	279
6.2 Attaching a Network Interface.....	281
6.3 Detaching a Network Interface.....	283
6.4 Changing a VPC.....	284
6.5 Modifying a Private IP Address.....	287
6.6 Managing Virtual IP Addresses.....	288
6.7 Enabling NIC Multi-Queue.....	294
6.8 Enabling IPv6 for a Network Interface.....	299
6.9 Dynamically Assigning IPv6 Addresses.....	301
6.10 Creating an Agency to Grant ECS Access to COC.....	320
7 EIPs.....	322
7.1 Overview.....	322
7.2 Binding an EIP.....	323
7.3 Unbinding an EIP.....	324
7.4 Changing an EIP.....	325
7.5 Modifying an EIP Bandwidth.....	326

7.6 Enabling Internet Connectivity for an ECS Without an EIP Bound.....	327
8 Security.....	332
8.1 Methods for Improving ECS Security.....	332
8.2 Security Groups.....	339
8.2.1 Overview.....	339
8.2.2 Default Security Groups and Rules.....	342
8.2.3 Security Group Configuration Examples.....	346
8.2.4 Configuring Security Group Rules.....	352
8.2.5 Changing a Security Group.....	359
8.3 HSS.....	360
8.4 CBH.....	363
8.5 Configuring Secure Boot for an ECS.....	366
8.6 Configuring ECS QingTian TPM.....	376
8.7 Project and Enterprise Project.....	377
8.8 Protection for Mission-Critical Operations.....	378
9 Backup Using CBR.....	382
9.1 Overview.....	382
9.2 Backing Up an ECS.....	390
10 Passwords and Key Pairs.....	394
10.1 Password Reset.....	394
10.1.1 Application Scenarios for Using Passwords.....	394
10.1.2 Resetting the Password for Logging In to an ECS on the Management Console.....	396
10.1.3 Obtaining and Deleting the Password of a Windows ECS.....	402
10.1.3.1 Obtaining the Password for Logging In to a Windows ECS.....	403
10.1.3.2 Deleting the Initial Password for Logging In to a Windows ECS.....	404
10.1.4 One-Click ECS Password Reset Plug-in.....	405
10.1.4.1 Obtaining the ECS One-Click Password Reset Plug-in.....	405
10.1.4.2 Installing the One-Click Password Reset Plug-in on an ECS.....	411
10.1.4.3 Updating the One-Click Password Reset Plug-in for an ECS.....	418
10.1.4.4 Using Scripts to Batch Update the One-Click Password Reset Plug-ins for Linux ECSs.....	422
10.1.4.5 Using Scripts to Batch Update the One-Click Password Reset Plug-ins for Windows ECSs.....	425
10.2 Key Pairs.....	429
10.2.1 Application Scenarios for Using Key Pairs.....	429
10.2.2 (Recommended) Creating a Key Pair on the Management Console.....	431
10.2.3 Creating a Key Pair Using PuTTY Key Generator.....	432
10.2.4 Importing a Key Pair.....	436
11 Launch Templates.....	439
11.1 Overview.....	439
11.2 Creating a Launch Template.....	439
11.3 Managing Launch Templates.....	441
12 Auto Launch Groups.....	442

12.1 Overview.....	442
12.2 Creating an Auto Launch Group.....	443
12.3 Managing Auto Launch Groups.....	445
13 Events.....	447
13.1 Overview.....	447
13.2 Querying an Event.....	451
13.3 Handling an Event.....	452
13.3.1 Handling an Instance Redeployment Event.....	452
13.3.2 Handling a Local Disk Replacement Event.....	458
13.3.3 Handling a System Maintenance Event.....	464
13.4 Preprocessing for Local Disk Replacement.....	467
13.4.1 Preprocessing for Disk-intensive ECSs.....	467
13.4.2 Preprocessing for Bare Metal ECSs.....	470
13.5 Preprocessing for Instance Redeployment.....	473
14 OS Dump.....	475
14.1 Configuring OS Dump.....	475
14.2 Viewing a Dump File.....	482
15 Self-Service O&M.....	484
15.1 Configuring Custom Policies for ECS Self-Service O&M.....	484
15.2 Installing UniAgent on an ECS.....	486
15.3 Setting Scheduled Tasks for ECSs.....	489
16 Resources and Tags.....	495
16.1 Tag Management.....	495
16.1.1 Overview.....	495
16.1.2 Adding Tags.....	496
16.1.3 Searching for Resources by Tag.....	499
16.1.4 Deleting Tags.....	500
16.2 Quota Adjustment.....	502
17 Cloud Eye Monitoring.....	504
17.1 Monitoring ECSs.....	504
17.2 Basic ECS Metrics.....	505
17.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed.....	513
17.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed.....	574
17.5 Creating Alarm Rules.....	581
17.6 Viewing ECS Metrics.....	584
18 Audit Using CTS.....	585
18.1 Key Operations Supported by CTS.....	585
18.2 Viewing Traces.....	586
19 QingTian Enclave Management.....	590
19.1 QingTian Enclave Overview.....	590

- 19.1.1 What Is QingTian Enclave?..... 590
- 19.1.2 QingTian Enclave Concepts..... 593
- 19.2 Getting Started with QingTian Enclave..... 595
- 19.3 Examples of Using QingTian Enclave..... 599
 - 19.3.1 Workflow..... 599
 - 19.3.2 Creating a QingTian Enclave Image..... 602
 - 19.3.3 Launching a QingTian Enclave Instance..... 603
- 19.4 Cryptographic Attestation..... 604
 - 19.4.1 PCR..... 604
 - 19.4.2 Attestation Document..... 606
 - 19.4.3 Document Signature Verification..... 607
 - 19.4.4 Integration with Huawei Cloud KMS..... 608
- 19.5 QingTian Enclave Application Development..... 609
 - 19.5.1 QingTian Enclave Application Development on Linux..... 609
 - 19.5.2 QingTian Enclave Network Proxy..... 619
 - 19.5.3 QingTian Enclave Log Forwarding Tool..... 625
- 19.6 QingTian CLI (qt CLI)..... 631
 - 19.6.1 Installation of the qt CLI..... 631
 - 19.6.2 Introduction to qt-enclave-env..... 632
 - 19.6.3 Introduction to qt enclave Subcommands..... 632
- 19.7 QingTian Error Code..... 637
- 19.8 FAQs About QingTian..... 641
 - 19.8.1 General Questions..... 642
 - 19.8.2 Development and Deployment Questions..... 644

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](#)).

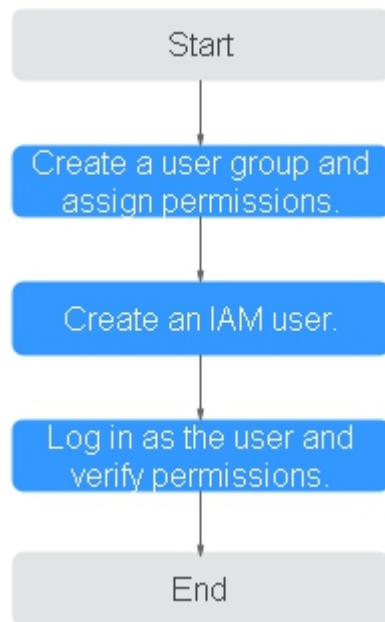
Prerequisites

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



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 **Buy 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](#) and [A Summary List of Kunpeng 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.	<ul style="list-style-type: none">• A Summary List of x86 ECS Specifications• A Summary List of Kunpeng 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. <ul style="list-style-type: none">• 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.	Network Interface 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 ($0.68/3,600 \times 30 \times 60$).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$1.02 USD ($0.68/3,600 \times 90 \times 60$).

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

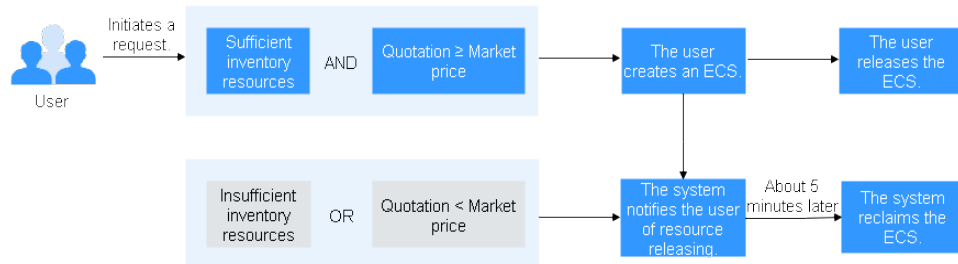
The market price for the ECSs of a certain flavor fluctuates due to supply-and-demand changes. You can purchase and use spot ECSs at a low market price to reduce computing costs.

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

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 \times 30 \times 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 \text{ USD/hour} \times 1 \text{ 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 \times 30 \times 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 \times 30 \times 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 \text{ USD/hour} \times 1 \text{ 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 \times 30 \times 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.
- RIs are in the open beta test (OBT) phase. You can [apply for OBT](#).

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	<ul style="list-style-type: none">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	<ul style="list-style-type: none">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. <p>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.</p>
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	$n \times 4$


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 a 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	<ul style="list-style-type: none">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	<ul style="list-style-type: none">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. <p>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.</p>
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

NameID	Status	Specifications	OS	Scope	AZ	Offering...	Term	Quantity	Remaini...	Operation
<input checked="" type="checkbox"/> n-7a0 01154-495291772-0	Active	2 vCPUs 4 GB c3.large.2	Linux	Regional	--	Standard	1 years	3	260 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	2 vCPUs 4 GB c3.large.2	Linux	Regional	--	Standard	1 years	1	260 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	2 vCPUs 8 GB s3.large.4	Linux	Regional	--	Standard	1 years	1	254 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 2 GB s3.medium.2	Linux	Regional	--	Standard	1 years	2	249 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	2 vCPUs 4 GB s3.large.2	Linux	Regional	--	Standard	1 years	1	254 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 2 GB s3.medium.2	Linux	Regional	--	Standard	1 years	1	254 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 2 GB s3.medium.2	Linux	Regional	--	Standard	1 years	1	254 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 2 GB s3.medium.2	Linux	Regional	--	Standard	1 years	1	234 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 1 GB s3.small.1	Linux	Regional	--	Standard	1 years	2	234 days	Buy ECS Buy Again Unsubscribe
<input type="checkbox"/>	Active	1 vCPU 1 GB s3.small.1	Linux	Regional	--	Standard	1 years	1	234 days	Buy ECS Buy Again Unsubscribe

Selected: n-7a0

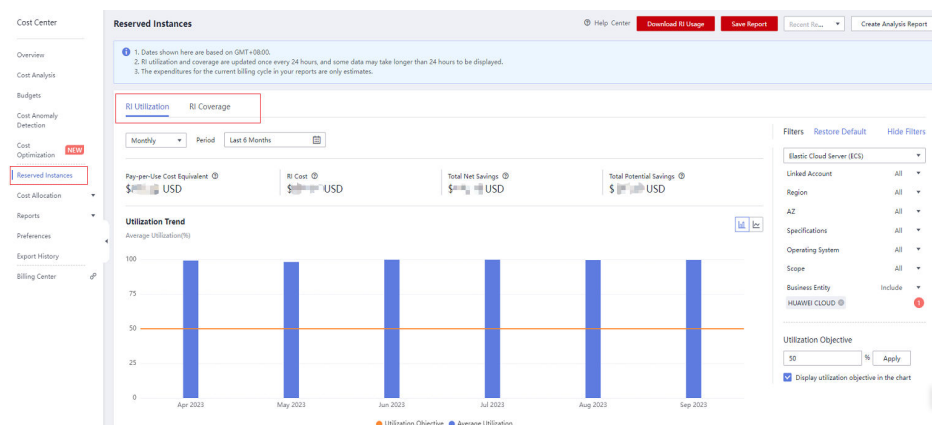
- c. Expand the RI details.

Figure 2-5 RI details

Selected: n-7a0			
Name	n-7a0	ID	01154-495291772-0
Status	Active	Order	
Specifications	2 vCPUs 4 GB 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		Obtained	Jul 17, 2023 20:00:00 GMT+08:00
Hourly Rate		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



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 instance 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 a 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. 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.

 NOTE

- For certain associated resources (such as EVS disks and EIPs), their billing modes can be changed together with the ECS to yearly/monthly.
- For associated resources whose billing modes cannot be changed together with the ECS to yearly/monthly, they will retain their original billing modes. For details, see [Billing Mode Change Rules for Associated Resources](#).
- ECSs created using ISO images do not support the billing mode change from pay-per-use to yearly/monthly.

Billing Mode Change Rules for Associated Resources

Resources associated with ECSs include disks and EIPs. [Table 2-9](#) and [Table 2-10](#) show the billing mode change rules for these associated resources.

Table 2-9 Billing mode change rules for disks attached to an ECS

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

Table 2-10 Billing mode change rules for EIPs bound to an ECS



Resource	Billing Mode	Billed By	Bandwidth Type	Changed Together with ECS to Yearly/Monthly	Measure
EIP	Pay-per-use	Bandwidth	Dedicated	Yes	None

Resource	Billing Mode	Billed By	Bandwidth Type	Changed Together with ECS to Yearly/Monthly	Measure
EIP	Pay-per-use	Traffic	Dedicated	No	On the EIP console page, change the billing mode from billing by traffic (pay-per-use) to billing by bandwidth (pay-per-use) first and then to yearly/monthly. For details, see Changing EIP Billing Mode .
EIP	Pay-per-use	Bandwidth	Shared	No	On the EIP console, change the billing mode from pay-per-use to yearly/monthly. For details, see Changing EIP Billing Mode .
EIP	Yearly/Monthly	Bandwidth	Dedicated or shared	No	The billing mode is already yearly/monthly. No actions are required.

Prerequisites

- The selected ECS is billed on a pay-per-use basis.
- 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. Choose **More > Change to Yearly/Monthly** in the **Operation** column.

 **NOTE**

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

1. Select the target ECSs.
2. Choose **More > Manage Billing > Change to Yearly/Monthly** above the ECS list.
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**


You can change the billing mode from yearly/monthly to pay-per-use in either of the following ways:

- Change to pay-per-use immediately: The pay-per-use billing mode takes effect immediately.
- Change to pay-per-use upon expiration: The pay-per-use 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.

Change to Pay-per-Use Immediately

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select one or more ECSs.
5. In the upper left corner above the ECS list, choose **More > Manage Billing > Change to Pay-per-Use Immediately**.
6. In the displayed dialog box, click **OK**. Then you are switched to Billing Center.
7. Confirm or select the ECSs for which you want to change the billing mode.
8. Confirm the refund information and click **Change to Pay-Per-Use**.
9. In the displayed dialog box, confirm the resources again and click **OK**.


Changing to Pay-per-Use upon Expiration (Console)

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select one or more ECSs.
5. In the upper left corner above the ECS list, choose **More > Manage Billing > Change to Pay-per-Use upon Expiration**.
6. In the displayed dialog box, click **OK**. Then you are switched to Billing Center.
7. Confirm or select the ECSs for which you want to change the billing mode.
8. Click **Change to Pay-Per-Use**.

Change to Pay-per-Use upon Expiration (Billing Center)

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).

Figure 2-7 Renewals



Manual Renewals (2) Auto Renewals (0) Pay-per-Use After Expiration (0) Renewals Canceled (0)						
Batch Renew Enable Auto Renewal Change to Pay-per-Use After Expiration Cancel Renewal Set Renewal Date Export						
Instance Name/ID	Product Type/Specifications	Region	Provisioned/Expires	Status	Validity Period	Operation
ecs-00000000000000000000	Elastic Cloud Server ECS	China North-01	Dec 18, 2024 17:19:10 GMT+08:00 / Jan 19, 2025 07:59:59 GMT+08:00	Provisioned	21 days until expiration (grace period)	Renew Enable Auto Renewal More

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 2-11 Purchase options of ECSs

How to Purchase	Description
Purchasing an ECS in Quick Config Mode	Learn how to purchase an ECS in quick config mode. You can specify the billing mode, region, instance specifications, image, public network access, and purchase details to quickly purchase an ECS.
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 ECSs Using Auto Launch Groups	Learn how to purchase ECSs using Auto Launch Groups. Auto Launch Groups consist of launch templates and auto launch groups. <ul style="list-style-type: none">• A launch template stores instance creation parameters (excluding the password) so that you can use it to batch create ECS instances with the same configurations.• An auto launch group lets you customize configurations and rapidly create ECSs that are of different types, billed in different modes, and distributed across multiple AZs.
Purchasing an ECS in a Shared Subnet	Learn how to purchase an ECS in a VPC subnet shared by another account. This can unify resource management, improve efficiency, and reduce O&M costs.
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 Quick 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.

The quick config mode provides basic, cost-effective, and high-performance instance specifications for you to choose from. You can specify the billing mode, region, image, public network access, and purchase details to quickly purchase an ECS.

The following describes how to purchase an ECS in quick config mode on the management console. To learn how to purchase an ECS in custom config mode, see [Purchasing an ECS in Custom Config Mode](#).

Constraints

- Only yearly/monthly and pay-per-use billing modes are supported.
- Only randomly allocated AZs are supported.
- You can only choose from the recommended instance specifications
- Only public images are supported.


Prerequisites

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](#).

Procedure

1. Log in to the management console and access the [Quick Config](#) tab.

Figure 2-8 Quick config



2. Set **Basic Configuration**.
 - a. Set **Billing Mode**.
You can select an appropriate billing mode based on the required duration and resource inventory to help you save costs.

Table 2-12 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 purchased 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 billing mode is ideal for scenarios where resource demands fluctuate and you want more flexibility on resource usage.	Pay-per-Use Billing

b. 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.

3. Set Instance.

The quick config mode provides basic, cost-effective, and high-performance instance specifications for you to choose from. Select appropriate instance specifications based on service requirements.

If the recommended specifications cannot meet your requirements, you can switch to the **Custom Config** tab. For details, see [Purchasing an ECS in Custom Config Mode](#).

Figure 2-9 Instance specifications

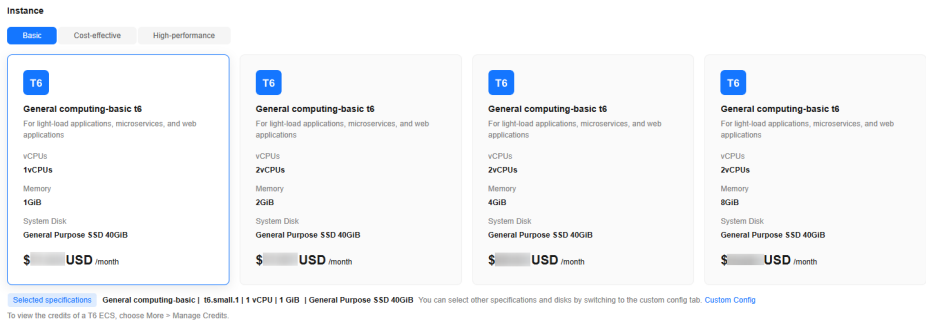



Table 2-13 Instance application scenarios

Type	Description	Recommended Flavors (Subject to Console)
Basic	Economical, suitable for medium- and light-load enterprise applications and individual users	T6: t6.small.1 and t6.large.1 S7: s7.large.2 and s7.large.4
Cost-effective	High network bandwidth and PPS, suitable for medium- and light-load enterprise applications that require high computing and network performance	S7: s7.small.1, s7.large.2, s7.xlarge.2, and s7.2xlarge.2
High-performance	Higher performance, security, and stability, suitable for medium- and heavy-load enterprise applications that require high computing and network performance	C7: c7.large.2, c7.xlarge.2, c7.2xlarge.2, and c7.3xlarge.2

4. Set **OS**.
- a. Select an image for the ECS.

Only public images are provided in quick config mode. You can select an image based on service requirements.
For more details, see [Public Image Overview](#).

Figure 2-10 OS



- b. (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. For details, see [What Should I Do When the Free Trial of HSS Basic Edition Expires?](#)

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](#).

5. Set **Public Network Access**.
 - a. You can select **Public Network Access** to buy and bind an EIP to the ECS. If you do not select it, no EIPs will be assigned.

 **NOTE**

If the default EIP route is sold out, **Public Network Access** cannot be selected.

- b. (Optional) Set **Billed By**.

This parameter is displayed only when **Public Network Access** is selected. 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.
 - c. (Optional) Set **Bandwidth Size**.

This parameter is displayed only when **Public Network Access** is selected. Select the bandwidth based on service requirements. The unit is Mbit/s.

6. Set **Other**.

Other settings use the system defaults. If you need to customize the settings, switch to the **Custom Config** tab.

For details, see [Purchasing an ECS in Custom Config Mode](#).

Table 2-14 Other settings

Parameter	Example	Description	Reference
AZ	Random	AZs are randomly allocated by default.	Region and AZ

Parameter	Example	Description	Reference
VPC	vpc-default	By default, the VPC and subnet of the latest created ECS are used. <ul style="list-style-type: none">If no ECSs were created in your account, the latest created VPC and the latest available subnet in the VPC are used by default.If no VPCs are available in your account, the system creates a VPC and subnet by default.If a VPC is available but no subnets are available, switch to the Custom Config tab to purchase ECSs. For details, see Purchasing an ECS in Custom Config Mode.The primary NIC is automatically assigned IPv4 and IPv6 addresses by default. NOTE If IPv6 is not enabled for the selected subnet, no IPv6 address is assigned.	VPC and Subnet Planning
Primary NIC	subnet-default		
Security Group	default or Sys-WebServer	The default security group is used.	Security Group Overview
ECS Name	ecs-xxxx	The name is randomly generated by the system.	N/A
Cloud Eye	Detailed monitoring enabled (Free)	Detailed monitoring is enabled by default for certain public images.	Monitoring ECSs
Enterprise Project	default	This function is provided for enterprise users. By default, the enterprise project that the latest created ECS belongs to is used. <ul style="list-style-type: none">If no ECSs were created in your account or the enterprise project that the latest created ECS belongs to is unavailable, the latest created enterprise project is used by default.If no enterprise projects are available in your account, the system creates an enterprise project by default.	Accessing the Enterprise Center

7. Set **Purchase Details**.
 - a. (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 1 year.
 - b. (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](#).
 - c. 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](#).
8. Confirm the configuration and submit the order.
 - a. In the **Configuration Summary** panel on the right side, review the ECS configuration details.

Mandatory fields that are not configured are displayed in red. You need to set them in the parameter configuration area.
 - b. Read and agree to the agreement, and click **Submit**.

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

Follow-up Operations

No login credentials are set in quick config mode. If you want to remotely log in, reset the password after the ECS is created. For details, see [Resetting the Password for Logging In to an ECS on the Management Console](#).

2.3.3 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	<ul style="list-style-type: none">• 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. For more details, see [VPC and Subnet Planning](#).
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](#).
6. Create dedicated physical resources.

To make your ECSs run on isolated physical hardware, apply for a Dedicated Computing Cluster (DCC) before creating the ECSs.


For details, see [Applying for a DCC](#).

7. Register an edge site.
- To create ECSs for CloudPond sites, register an edge site in advance.
- For details, see [Registering an Edge Site](#).

Step 1: Access the Page for Purchasing ECSs

Log in to the management console and access the [Custom Config](#) tab.

Figure 2-11 Custom config



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 2-15 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

Option	Description	Scenarios and Constraints	Reference
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
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.	<ul style="list-style-type: none">• 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** and you have applied for the open beta test (OBT) of reserved instances. If you want to associate reserved instances (RIs) with your pay-per-use ECSs, select **Associate RI** and select an RI.

NOTE

RIs are in the OBT phase. You can [apply for OBT](#).

For details, see [Reserved Instance Overview](#).

Figure 2-12 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 2-13 Spot type**Basic Configuration**

Billing Mode ?

Yearly/Monthly

Pay-per-use

Spot pricing

Spot Type

Spot

Fluctuating price, billed by usage duration

Spot block

Fixed price, billed by predefined duration

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.
- When you purchase ECSs for a CloudPond edge site, the default region is the region where the edge site is located and 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 2-16 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 .	Regions and AZs
AZ <i>N</i>	AZs supported in the selected region. <i>N</i> indicates the sequence number of an AZ.	If you want to create an ECS in a specified AZ, select that AZ. <ul style="list-style-type: none">For high availability (HA), create ECSs in different AZs.For low network latency, create ECSs in the same AZ.	
Edge AZ	Edge AZs are deployed in on-premises data centers and are dedicated to CloudPond users.	If you are purchasing ECSs for a CloudPond edge site, an edge AZ is selected by default and it cannot be changed.	What Are the Relationships Between Edge Sites, Regions, and AZs?

Instance

1. Select an instance selection mode.

Table 2-17 Instance parameters

Option	Description	Scenarios and Constraints	Reference
By Type	You can select ECS specifications based on different properties. <ul style="list-style-type: none">• 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 2-18 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

Option	Description	Scenarios and Constraints	Reference
Kunpeng	The Kunpeng CPU architecture uses Reduced Instruction Set Computer (RISC) and Huawei-developed processors, which is more cost effective. RISC executes fewer types of computer instructions at a higher speed than CISC. RISC simplifies the computer architecture and improves the execution speed.	It is suitable for the following scenarios: <ul style="list-style-type: none">Platform-independent scenarios such as e-commerce, big data, and scientific computingArm native scenarios such as mobile phone simulation	A Summary List of Kunpeng 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 2-19 Categories and sub-categories

Category	Sub-category	Description	Recommended Specification (Example)
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	C7 and C6
	Enterprise websites	A balance of compute, memory, and network resources with a baseline level of vCPU performance and high cost-effectiveness	<ul style="list-style-type: none">C7 and C6S7 and S6
	Personal application setup	A balance of compute, memory, and network resources with a baseline level of vCPU performance and high cost-effectiveness	<ul style="list-style-type: none">S7 and S6T6

Category	Sub-category	Description	Recommended Specification (Example)
	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	S7 and S6
	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.	S7 and S6
	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.	C7, C6, and C6s
Website applications/E-commerce	100,000 pageviews/ 1,000 active users	Cost-effective, flexible, elastic resources available anytime	<ul style="list-style-type: none">• S7 and S6• T6
	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	<ul style="list-style-type: none">• C7, C6, and C6s• S7 and S6
	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	C7, C6, and C6s

Category	Sub-category	Description	Recommended Specification (Example)
Gaming	Gaming	Suitable for gaming services, which require high performance, high stability, high cost-effectiveness, and low latency	C7, C6, and C6s
Databases	Compute	Stable, high-performance compute	C7, C6, and C6s
	Storage	Servers that use local disks with high storage bandwidth and IOPS to provide cost-effective mass data storage	<ul style="list-style-type: none">• M7 and M6• D7 and D6
	Network	High PPS performance, high TPS throughput, and low network latency for rapid data exchange and processing	<ul style="list-style-type: none">• E7 and E6• M7 and M6• C7, C6, and C6s
Data analytics	Management nodes	A large volume of compute resources scheduled to accelerate data processing	C7, C6, and C6s
	Compute nodes	Balanced compute with high performance and stability	<ul style="list-style-type: none">• M7 and M6• C7, C6, and C6s• I7, I7r, I3, and I3r• S7 and S6
	Storage nodes	Cost-effective, high-bandwidth storage for processing large amounts of reads and writes	D7, D6, and D3
High performance computing	High performance computing	High-performance compute clusters with large compute and high cost-effectiveness	H3
Image rendering	Animation rendering	CPU-accelerated rendering with high precision and stability	aC7, C6, and C3
	Video rendering	GPU-accelerated rendering with high processing speed	G6

Category	Sub-category	Description	Recommended Specification (Example)
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.	P2vs, P2v, and P2s
	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.	Pi2

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 2-20 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

Option	Description	Reference
KooGallery image	<p>This parameter is available only when Billing Mode is set to Yearly/Monthly or Pay-per-use.</p> <p>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.</p>	KooGallery

2. (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. For details, see [What Should I Do When the Free Trial of HSS Basic Edition Expires?](#)

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](#).

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

This parameter is displayed only when the image you select is billed. It specifies a license type for using an OS or software.

Currently, you can select **Bring your own license (BYOL)**, which allows you to use your existing OS license without the need to apply for a license again.

For more information about license types, see [License Types](#).

Storage & Backup

1. (Optional) Set **Storage Type**.

This parameter is displayed only when you have applied for a storage pool on the [Dedicated Distributed Storage](#) page.

Disks are classified as EVS disks and DSS disks based on whether they use dedicated storage resources. DSS disks provide dedicated storage resources.

- **DSS**: disks are created using resources from the dedicated storage pool.
- **EVS**: disks are created using public storage resources.

 **NOTE**

- When you create disks in a DSS storage pool, the disk type must be the same as that of the requested storage pool. For example, both are of the high I/O type.
- For more information about DSS, see [Dedicated Distributed Storage Service](#).

2. 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 2-21 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.	

Parameter	Description	Scenarios and Constraints	Reference
Throughput	<p>Amount of data read from and written into an EVS disk per second</p> <p>This parameter is displayed only when General-Purpose SSD V2 is selected for Disk Type.</p>	Configure a desired throughput based on the value range and your service requirements.	

3. (Optional) Set **Advanced Options** for the system disk.
If you want to set **SCSI** and **Encryption** for the system disk, click **Advanced Options**.

Table 2-22 Advanced options

Option	Description	Scenarios and Constraints	Reference
SCSI	<p>Specifies the SCSI device type.</p> <p>This parameter is selected by default.</p>	<p>Device types:</p> <ul style="list-style-type: none">• VBD: indicates the Virtual Block Device (VBD) mode.• SCSI: indicates the Small Computer System Interface (SCSI) mode. <p>The default device type is VBD. If SCSI is selected, the disk will support transparent SCSI command transmission.</p> <p>NOTE The disk device type is configured during the purchase process. Once disks are purchased, the device type cannot be changed.</p>	Device Types

Option	Description	Scenarios and Constraints	Reference
Encryption	<p>Encrypts the system disk.</p> <ul style="list-style-type: none">• 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.	<p>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.</p> <p>For details, see 4.</p>	Managing Encrypted EVS Disks

4. (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.

5. 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 2-23 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.	

Parameter	Description	Scenarios and Constraints	Reference
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.	

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

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

Table 2-24 Advanced options

Option	Description	Scenarios and Constraints	Reference
SCSI	Specifies the SCSI device type. This parameter is selected by default.	Device types: <ul style="list-style-type: none">• 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
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

Option	Description	Scenarios and Constraints	Reference
Encryption	Encryption is used to encrypt data disks.	Disk encryption provides strong security protection for your data. For details, see 4 .	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

7. (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?](#)
- **Cloud Backup and Recovery** is not supported for CloudPond.

8. (Optional) Set **CBRCloud 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 2-14 Network

Network

VPC

vpc-default(192.168.0.0/16) [Create VPC](#)

Primary NIC

subnet-default(192.168.0.0/24) Automatically assign IP address Available private IP addresses: 250

+ Add Extension NIC

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.
- If you want to use the shared VPC and subnet from another account, accept the resource sharing invitation first. For details, see [Responding to a Resource Sharing Invitation](#).

For more information about VPC subnet sharing, see [VPC Sharing](#) in the *Virtual Private Cloud User Guide*.

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**

- For details about how to enable IPv6 for a subnet, see [IPv6 Network](#).
- After creating an ECS, check whether the ECS has obtained an IPv6 address. If not, enable IPv6 so that the ECS dynamically obtains an IPv6 address. For details, see [Dynamically Assigning IPv6 Addresses](#).

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](#).

 NOTE

The number of extended NICs that can be attached to an ECS is determined by the ECS specifications. For details, see [A Summary List of x86 ECS Specifications](#) and [A Summary List of Kunpeng ECS Specifications](#).

- 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

- For details about how to enable IPv6 for a subnet, see [IPv6 Network](#).
- After creating an ECS, check whether the ECS has obtained an IPv6 address. If not, enable IPv6 so that the ECS dynamically obtains an IPv6 address. For details, see [Dynamically Assigning IPv6 Addresses](#).

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 2-15 Creating a security group

×

Create Security Group

Name

sg-

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

Cancel

OK

Parameters for creating a security group are as follows.

Issue 01 (2025-07-25)

Copyright © Huawei Technologies Co., Ltd.

57

Table 2-25 Creating a security group

Parameter	Description	Example value
Name	<p>This parameter is mandatory and specifies the name of a security group. The name:</p> <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.). <p>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.</p>	sg-AB
Template	<p>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.</p> <p>The following templates are provided:</p> <ul style="list-style-type: none">• 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. <p>NOTE Allowing inbound traffic on all ports poses security risks.</p> <ul style="list-style-type: none">• Fast-add rule: You can select common protocols and ports that the inbound rule will apply to.	Fast-add rule
Inbound Rules	<p>This parameter is optional. This parameter is displayed only when Fast-add rule is selected for Template.</p> <p>Currently, the following protocols and ports can be quickly added. Select protocols and ports as required.</p> <ul style="list-style-type: none">• 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 2-26 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. The initial password of user Administrator is used for authenticating Windows ECSs.	It is recommended that you set passwords with high complexity to prevent malicious attacks. The passwords must meet the requirements described in Table 2-27 . 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
Set password later	The password for logging in to the ECS is not configured during the ECS creation.	After the ECS is created, choose More > Reset Password in the Operation column, set a password for the ECS as prompted, and log in to the ECS.	Resetting the Password for Logging In to an ECS on the Management Console

Table 2-27 Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none">Consists of 8 to 26 characters.Contains at least three of the following character types:<ul style="list-style-type: none">Uppercase lettersLowercase lettersDigitsSpecial characters for Windows: !@\$%^_-=+[{()}];,./?~#*Special characters for Linux: !@\$%^_-=+[{()}];,./?~#*Cannot contain the username or the username spelled backwards.Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.)

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 2-27](#). 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 **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.

If your organization has created a tag policy for ECS, you need to add tags for ECS based on the tag policy. If a tag does not comply with the tag rules, the creation may fail. Contact the organization administrator to learn details about the tag policy.

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 1 year.

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. (Optional) Click **Save as Launch Template**.

Perform this step only when you need to create an ECS using a launch template. For details, see [Purchasing ECSs Using Auto Launch Groups](#).

The configuration cannot be saved as a launch template if the billing mode is yearly/monthly or spot block, host security is enabled, or the login mode is password.

3. Read and agree to the agreement, and click **Submit**.

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

2.3.4 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.
For details about how to back up data, see [Backing Up an ECS](#).

(Optional) Enabling Reclaim Notifications

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

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

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

Step 1 Enable CTS. For details, see [Enabling CTS](#).

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

Step 2 Configure reclaim notifications.

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


1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Under **Management & Governance**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Key Event Notifications**.
5. Click **Create Key Event Notification** in the upper right corner of the page and set parameters listed in [Table 2-28](#).

Table 2-28 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.

Type	Parameter	Configuration
Topic	Send Notification	Select Yes .
	SMN Topic	Select a topic from the drop-down list. If there are no proper SMN topics, create one. <ol style="list-style-type: none">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 & Governance**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in [Table 2-29](#) and search for traces as needed.

Table 2-29 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

Follow-up Procedure

- Resetting the password

If you set **Login Mode** to **Set password later** during the ECS creation, you can reset the password after creating the ECS.

For details, see [Resetting the Password for Logging In to an ECS on the Management Console](#).

- Connecting to an ECS

You can connect to an ECS in multiple ways. For details, see [Login Overview \(Windows\)](#) and [Login Overview \(Linux\)](#).

- Expanding storage capacity

If the system disk capacity is insufficient, you can add data disks to expand the storage capacity. For details, see [Adding a Disk to an ECS](#).

- Using an ECS

After purchasing an ECS, you can build environments, websites, or applications on it. For details, see [Best Practices Summary](#).

2.3.5 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, so back up your data.
For details about how to back up data, see [Backing Up an ECS](#).

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 a region and project.
3. Under **Management & Governance**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Key Event Notifications**.
5. Click **Create Key Event Notification** in the upper right corner of the page and set parameters listed in [Table 2-30](#).

Table 2-30 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. <ol style="list-style-type: none">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 & Governance**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in [Table 2-31](#) and search for traces as needed.

Table 2-31 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

Follow-up Procedure

- Resetting the password
If you set **Login Mode** to **Set password later** during the ECS creation, you can reset the password after creating the ECS.
For details, see [Resetting the Password for Logging In to an ECS on the Management Console](#).
- Connecting to an ECS
You can connect to an ECS in multiple ways. For details, see [Login Overview \(Windows\)](#) and [Login Overview \(Linux\)](#).
- Expanding storage capacity
If the system disk capacity is insufficient, you can add data disks to expand the storage capacity. For details, see [Adding a Disk to an ECS](#).
- Using an ECS
After purchasing an ECS, you can build environments, websites, or applications on it. For details, see [Best Practices Summary](#).

2.3.6 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

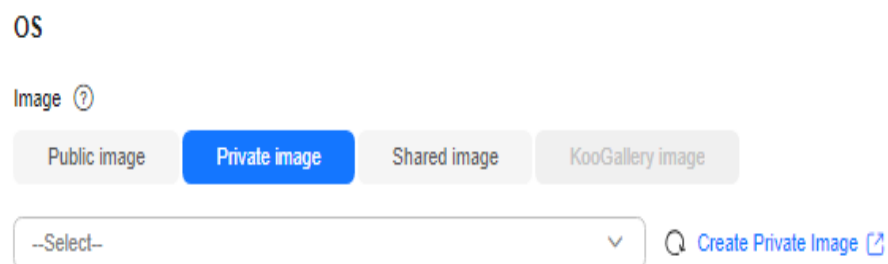
- If you use a full-ECS image to create an ECS, the EVS disks associated with the full-ECS image do not support the function of creating disks from a data disk image.
- If a full-ECS image is in **Normal** state and the system displays message "Available in AZx", the full-ECS image can be used to create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

Figure 2-16 Private 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.

Figure 2-17 Applying for a cloud server



NameID	Status	OS Type	OS	Image Type	Disk Capacity (GB)	Encrypted	Created	Enterprise Project	Operation
image-example 04630eq	Normal	Linux	CentOS 7.5 64bit	ECS system disk image(R)	40	No	May 27, 2024 18:41:27 G...	default	Apply for Server Modify More

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.

Follow-up Procedure

If an ECS is purchased using a private image, the password reset function may be unavailable because the private image may not have the password reset plug-in installed.

In this case, install the plug-in by referring to [Installing the One-Click Password Reset Plug-in on an ECS](#).

2.3.7 Purchasing ECSs Using Auto Launch Groups

Scenarios

Auto Launch Groups let you rapidly create ECSs that are billed in different modes and distributed across AZs to meet capacity targets. It consists of auto launch templates and auto launch groups.

- A launch template contains the configuration information to launch ECSs, for example, the ECS specifications, network settings, and a key pair (excluding the password). You can launch ECSs quickly without specifying the configuration parameters every time.
For details, see [Launch Template Overview](#)
- An auto launch group lets you customize configurations and rapidly create ECSs that are of different types, billed in different modes, and distributed across multiple AZs to meet capacity targets.

For details, see [Auto Launch Group Overview](#).

Notes and Constraints

- Auto launch groups are available in AP-Singapore and CN-Hong Kong.
- An auto launch group can create ECSs across AZs but cannot create ECSs across regions.
- For ECSs created with auto launch groups, the target capacity is limited as follows:
 - If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.
 - If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.
- You can specify one launch template for each auto launch group.
- Auto launch groups are free, but you will be billed for the ECSs created by the group.


Prerequisites

A launch template has been created for the ECS. For details, see [Creating a Launch Template](#).

Using a Launch Template to Create ECSs

1. In the launch template list, locate the row containing the target template and click **Buy ECS** in the **Operation** column.

Figure 2-18 Launch templates



2. On the **Buy ECS** page, modify the parameter in the template as needed.
For example, select **Password** for **Login Mode** and set a password for your ECS.
For details, see [Purchasing an ECS in Custom Config Mode](#).
3. In the **Configuration Summary** area on the right of the page, confirm the configuration details.
4. Select the checkbox before the agreement and click **Create**.
After an ECS is created, it is started by default. You can view the created ECS on the [ECS list page](#).

Using an Auto Launch Group to Create ECSs That Are Billed in Different Modes and Distributed Across AZs

1. Log in to the cloud server console, access the [Auto Launch Groups](#) page, and click **Create Group**.
2. Set the group name, target capacity, and quantity.

Figure 2-19 Configuring an auto launch group (1)

Name

alg-example

Target Capacity

ECSs

–

5

+

?

☒ Include pay-per-use ECSs. Select the quantity of ECSs or vCPUs that can be included.

Quantity

–

3

+

ECSs

Table 2-32 Parameters (1)

Parameter	Example	Description
Name	alg-example	Name of an auto launch group. The name can contain 2 to 64 characters, including letters, digits, underscores (_), and hyphens (-).
Target Capacity	ECSs: 5	The total target capacity of ECSs created using an auto launch group. You can set Target Capacity to specify the total compute delivered by the auto launch group. The target capacity can be set to the number of ECSs or the number of vCPUs. The target capacity of each auto launch group is limited. <ul style="list-style-type: none">If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.
Quantity	3	This parameter is displayed only when Include pay-per-use ECSs. Select the quantity of ECSs or vCPUs that can be included is selected. This parameter specifies the quantity of pay-per-use ECSs or vCPUs of these ECSs. The value of Quantity must be less than or equal to that of Target Capacity .

3. Set the launch template, ECS configuration, and optimization policy.

Figure 2-20 Configuring an auto launch group (2)

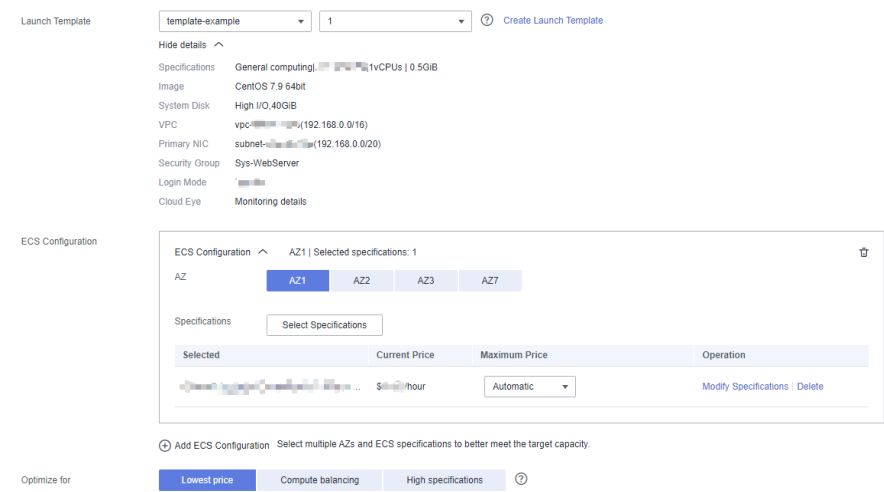


Table 2-33 Parameters (2)

Parameter	Example	Description
Launch Template	alg-example	You can select a launch template to use the configuration details it contains to launch an ECS. You can expand details to view the configuration of the selected launch template.
ECS Configuration	AZ1, s7.medium.2	This parameter specifies the AZ and specifications in the launch template. You can add and customize configurations for multiple instances. You are advised to select different AZs and instance specifications to improve the resource delivery success rate.

Parameter	Example	Description
Optimize for	Lowest price	<p>This parameter specifies the policy that instance allocation complies with.</p> <ul style="list-style-type: none">• Lowest price: The auto launch group will create the least expensive ECSs possible.• Compute balancing: The auto launch group will prioritize balancing compute loads by creating ECSs distributed across multiple AZs as evenly as possible.• High specifications: The auto launch group creates ECSs with the highest specifications possible. If you have configured a target number of ECSs, ECSs with more vCPUs will be prioritized and if the target is vCPUs, then that target will be met with as few ECSs as possible.

4. Set **Delivery Type**, **Start**, **End**, and other parameters.

Figure 2-21 Configuring an auto launch group (3)

Delivery Type

Single use

Continuous

?

Start

Immediately

Custom

?

End

Never expire

Custom

?

Global Maximum Price

☒ Maximum Price

?

Global Maximum Price

\$

/hour

ECS Deletion Settings

☒ Delete ECSs When Auto Launch Group Expires

?

☒ Delete ECSs When Target Capacity Is Exceeded

?

Table 2-34 Parameters (3)

Parameter	Example	Description
Delivery Type	Single use	<p>This parameter specifies the type of an auto launch group.</p> <ul style="list-style-type: none">• Single use: The auto launch group only attempts to create ECSs to meet the target capacity when it is started, but will not create ECSs again even if the target capacity is not reached.• Continuous: The auto launch group monitors the target and current capacity in real time and continues to create ECSs until the total target capacity is reached.
Start	Immediately	<p>Specifies the time when the auto launch group starts to launch ECSs.</p> <ul style="list-style-type: none">• Immediately: The auto launch group starts to launch ECSs immediately after the group is created.• Custom: You can specify when the auto launch group starts to launch ECSs.
End	Never expire	<p>This parameter specifies the time when the auto launch group expires. You can set both the start time and the end time to determine the validity period of the group.</p> <ul style="list-style-type: none">• Never expire: The auto launch group does not expire.• Custom: You can specify when the auto launch group expires.

Parameter	Example	Description
Global Maximum Price	0.12	<p>This parameter is displayed only when Maximum Price is selected. It specifies the allowed maximum price of a single spot ECS in the auto launch group.</p> <p>If the market price of a spot ECS in the group exceeds the global maximum price, the spot ECS will be deleted. If both the specific maximum price of a spot ECS and the global maximum price are set, the specific maximum price of the spot ECS will be used.</p>
ECS Deletion Settings	<ul style="list-style-type: none">• Delete ECSs When Auto Launch Group Expires• Delete ECSs When Target Capacity Is Exceeded	<p>This parameter is valid only when the end time of the auto launch group is specified.</p> <ul style="list-style-type: none">• Delete ECSs When Auto Launch Group Expires: ECSs in the auto launch group will be deleted when the group expires.• Delete ECSs When Target Capacity Is Exceeded: When the number of ECSs or vCPUs in the auto launch group exceeds the target capacity, the ECSs or vCPUs that exceed the target capacity will be deleted.

5. Click **Create Now**.

After an ECS is created, you can view it on the [ECS list page](#).

2.3.8 Purchasing an ECS in a Shared Subnet

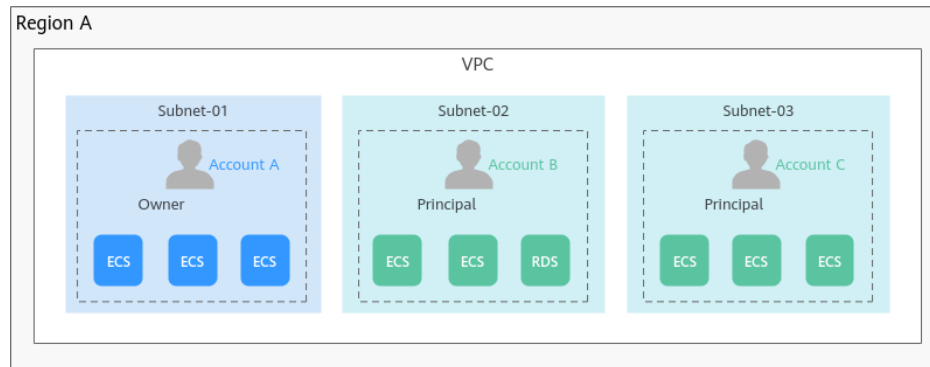
Scenarios

VPC sharing allows multiple accounts to purchase and manage cloud resources in one VPC. With VPC sharing, you can create ECSs in shared subnets for improved resource management and reduced O&M costs.

For example, to make resource management easier, you can use account A to manage basic and public IT resources, such as VPCs and subnets. In addition, account A shares subnets in the VPC with other accounts.

- Account A: the IT management account and the resource owner. It creates a VPC and subnets and shares subnets with other accounts, such as account B and C. Account A creates resources in **Subnet-01**.
- Account B: a service account and a principal of the shared subnet. Account B creates ECSs in **Subnet-02** shared by account A.

- Account C: a service account and a principal of the shared subnet. Account C creates ECSs in **Subnet-03** shared by account A.

Figure 2-22 Service planning

This section describes how to purchase an ECS in a shared subnet. For more information about VPC subnet sharing, see [VPC Sharing](#) in the *Virtual Private Cloud User Guide*.

Constraints

- A principal can receive a maximum of 100 subnet shares.
- VPC sharing is free. The principals only need to pay for the resources they create in the shared subnets.

Prerequisites

Account A, as the resource owner, has created a VPC and subnets, and specified account B as the principal. For details, see [Creating a Resource Share](#).


Procedure

- Log in to the [Resource Access Manager](#) console using account B and accept a resource sharing invitation.

For details, see [Responding to a Resource Sharing Invitation](#).

- Access the [ECS](#) console.
- Configure parameters required for purchasing an ECS.

In the **Configure Network** step, select the VPC and the subnet shared by account A.

Figure 2-23 Configuring network parameters

For other parameters, see [Purchasing an ECS](#).

2.3.9 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.

Figure 2-24 Buying the same ECS



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**.
5. Click back to ECS list. You can view the created ECS on the [ECS list page](#).

Helpful Links

- [Purchasing an ECS in Quick Config Mode](#)
- [Purchasing an ECS in Custom Config Mode](#)
- [Purchasing ECSs Using Auto Launch Groups](#)

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 the login password is forgotten, reset the password on the ECS console.
To reset a password, locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list. For details, see [Resetting the Password for Logging In to an ECS on the Management Console](#).
- 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 2-35 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.)

ECS OS	Local OS	Connection Method	Requirement
	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 .	
	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 an RDP File

Scenarios

Remote Desktop Protocol (RDP) is a multi-channel remote login protocol provided by Microsoft. This section describes how to use an RDP file to remotely log in to a Windows ECS.

NOTE

Each RDP file downloaded from the management console can be used to log in to only one ECS, and the file is named in the format of "ECS name-EIP".

Prerequisites

- The target ECS is running.
- If your ECS uses key pair authentication, you have obtained the password for logging in to the Windows ECS. For details, see [Obtaining the Password for Logging In to a Windows ECS](#).

- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group to which the ECS belongs. For details, see [Configuring Security Group Rules](#).
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- 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](#).

Login from Windows

If the local computer runs Windows, you can use the RDP file to log in to the target Windows ECS.



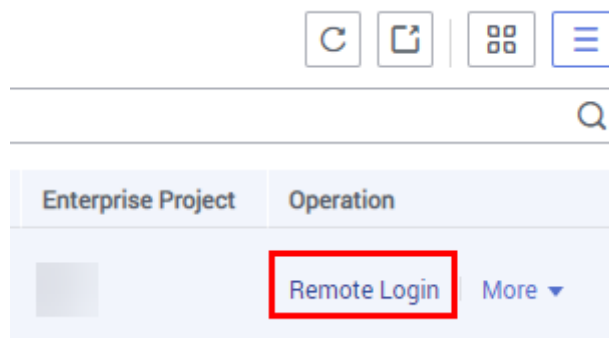
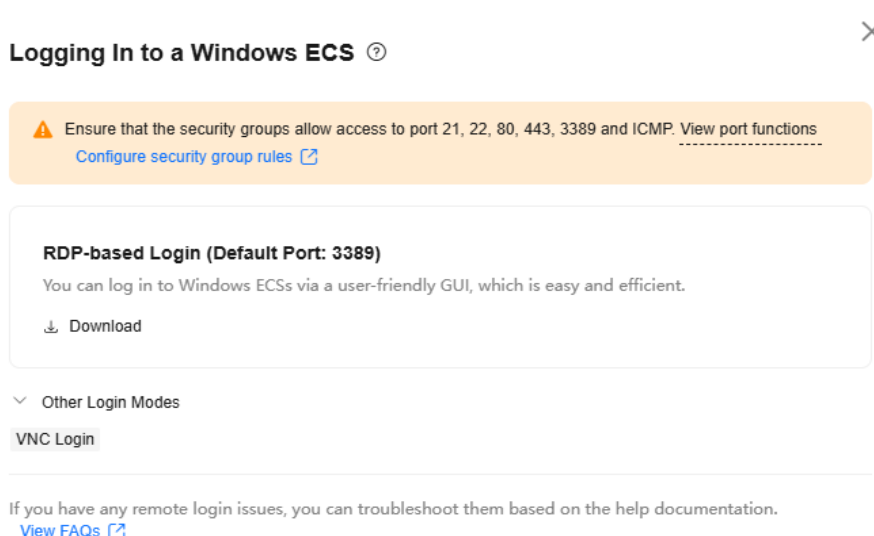
1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click  and choose **Compute > Elastic Cloud Server**.
4. In the **Operation** column of the target ECS, click **Remote Login**.

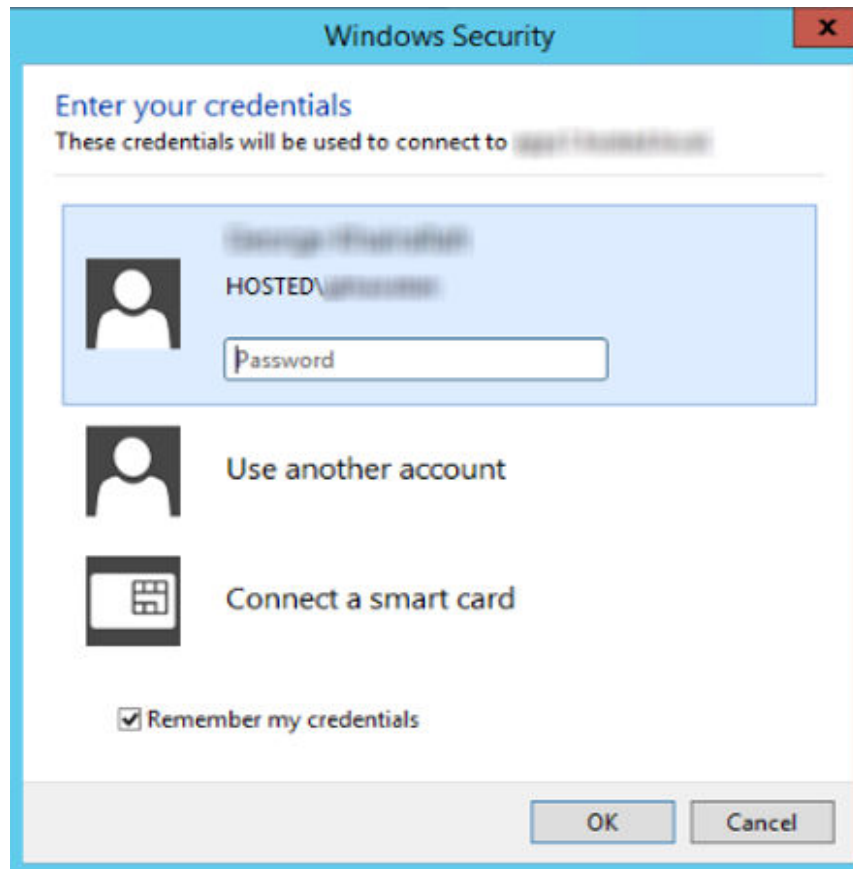
Figure 2-25 Remote Login



5. In the **Logging In to a Windows ECS** dialog box, select **RDP-based Login** and click **Download RDP File** to download the RDP file to the local computer.

Figure 2-26 Download RDP File

6. Double-click the downloaded RDP file to remotely access the Windows ECS.
 - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
 - If your ECS uses key pair authentication, obtain the password by following the instructions provided in [Obtaining the Password for Logging In to a Windows ECS](#).
 - If your password has been forgotten, see [Resetting the Password for Logging In to an ECS on the Management Console](#).

Figure 2-27 Logging in to a Windows ECS using an RDP file

2.4.3 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.

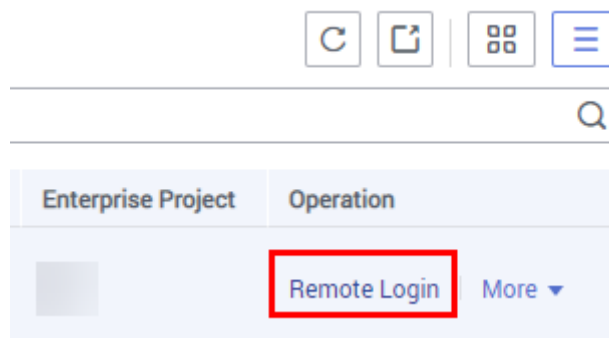
Prerequisites

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. For details, see [Obtaining the Password for Logging In to a Windows ECS](#).


Logging In to a Windows ECS

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

4. Obtain the password for logging in to the ECS.
Before logging in to the ECS, you must have the login password.
 - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
 - If your ECS uses key pair authentication, obtain the password by following the instructions provided in [Obtaining the Password for Logging In to a Windows ECS](#).
5. In the **Operation** column of the target ECS, click **Remote Login**.

Figure 2-28 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 2-29 Ctrl+Alt+Del

8. Enter the ECS password as prompted.

Helpful Links

- [Login Password Resetting](#)
- [Multi-User Login Issues](#)
- [Remote Logins](#)

2.4.4 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.

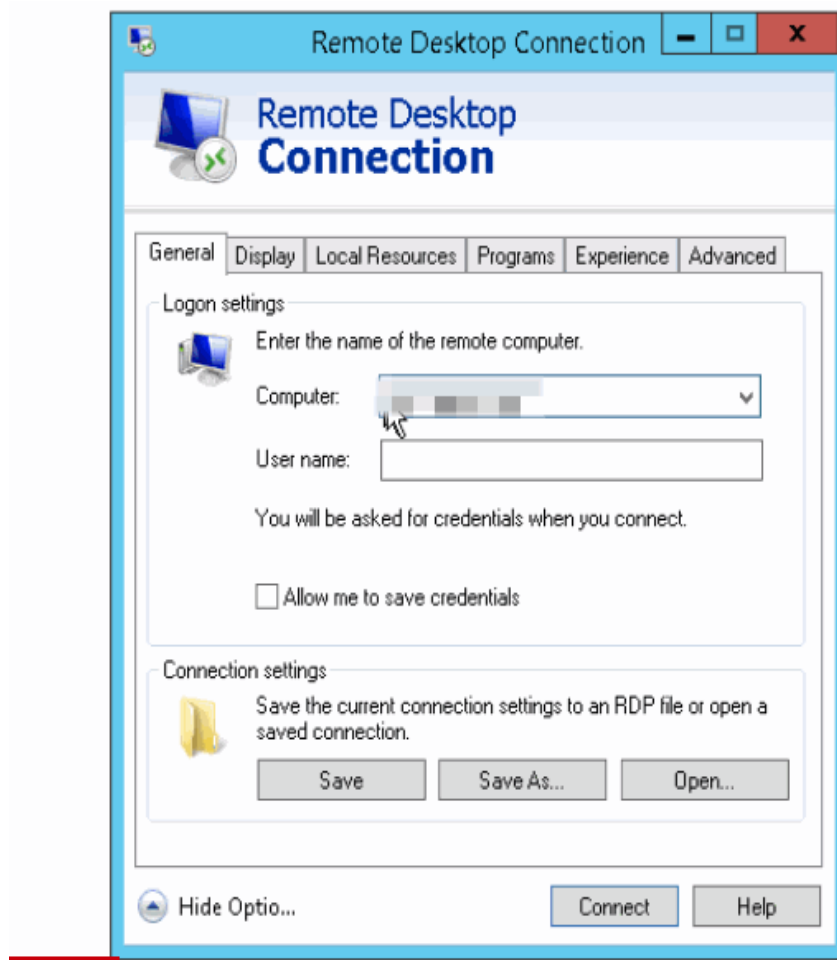
- If your ECS uses key pair authentication, you have obtained the password for logging in to the Windows ECS. For details, see [Obtaining the Password for Logging In to a Windows ECS](#).
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.
- Access to port 3389 is allowed in the inbound direction of the security group 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.

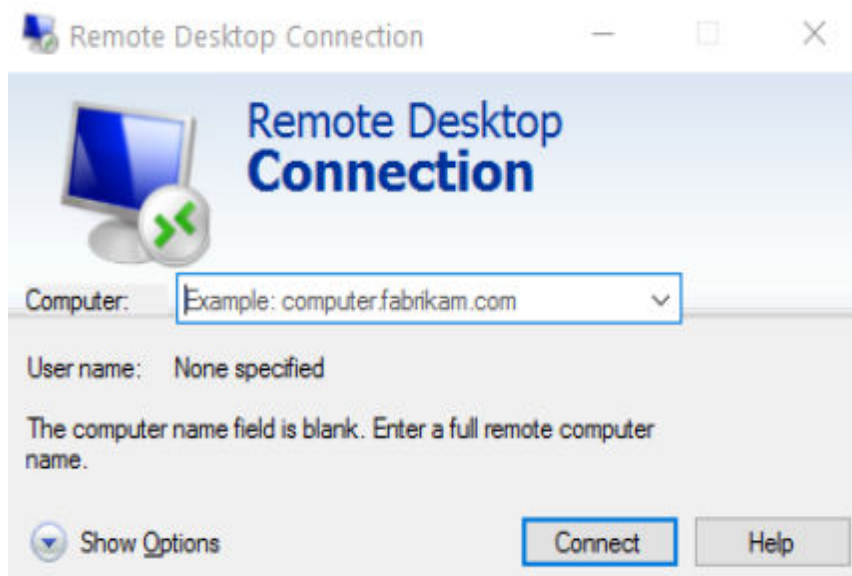
Figure 2-30 Logging in to an ECS using MSTSC



For details, see the following procedure:

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

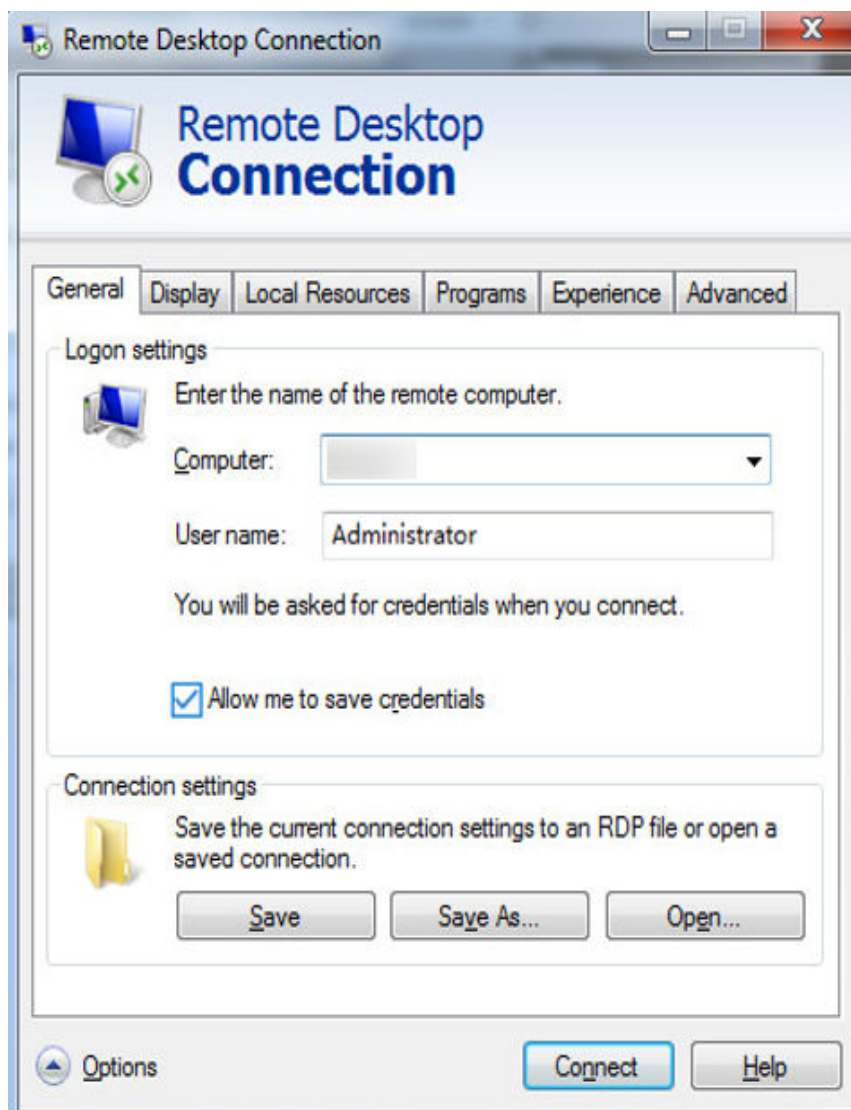
Figure 2-31 Showing options



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

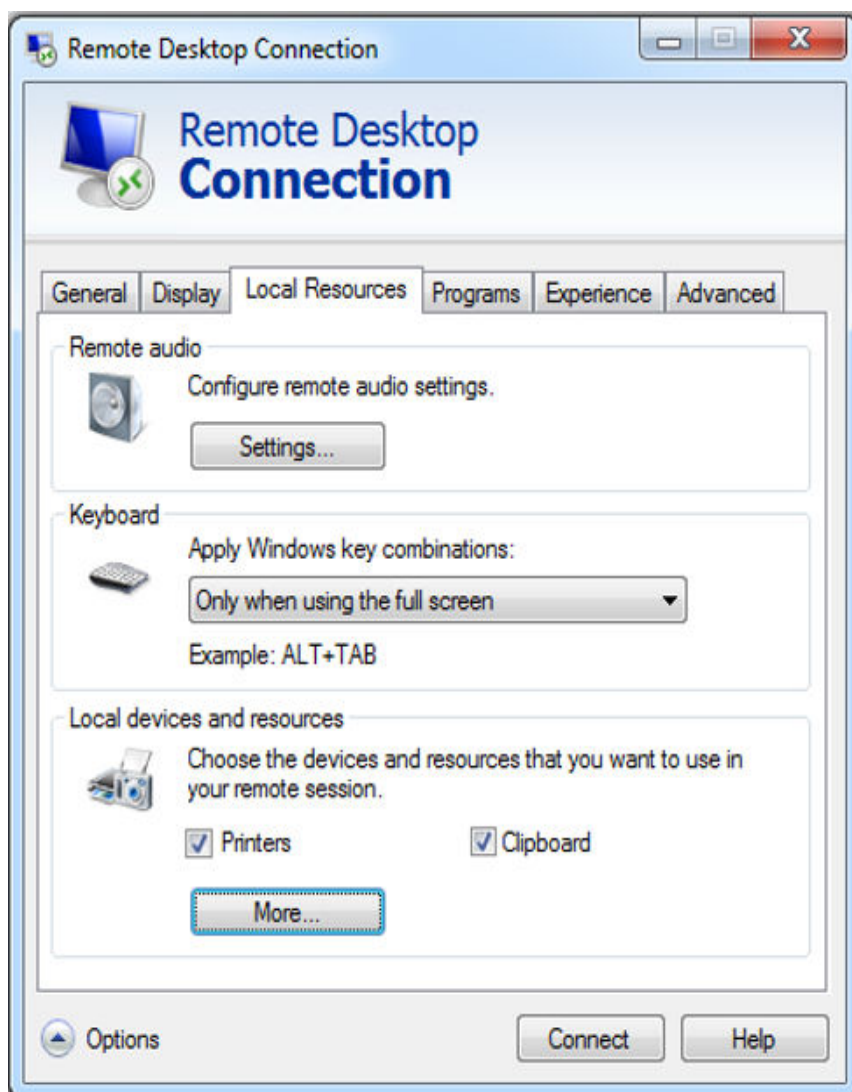
NOTE

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

Figure 2-32 Remote Desktop Connection

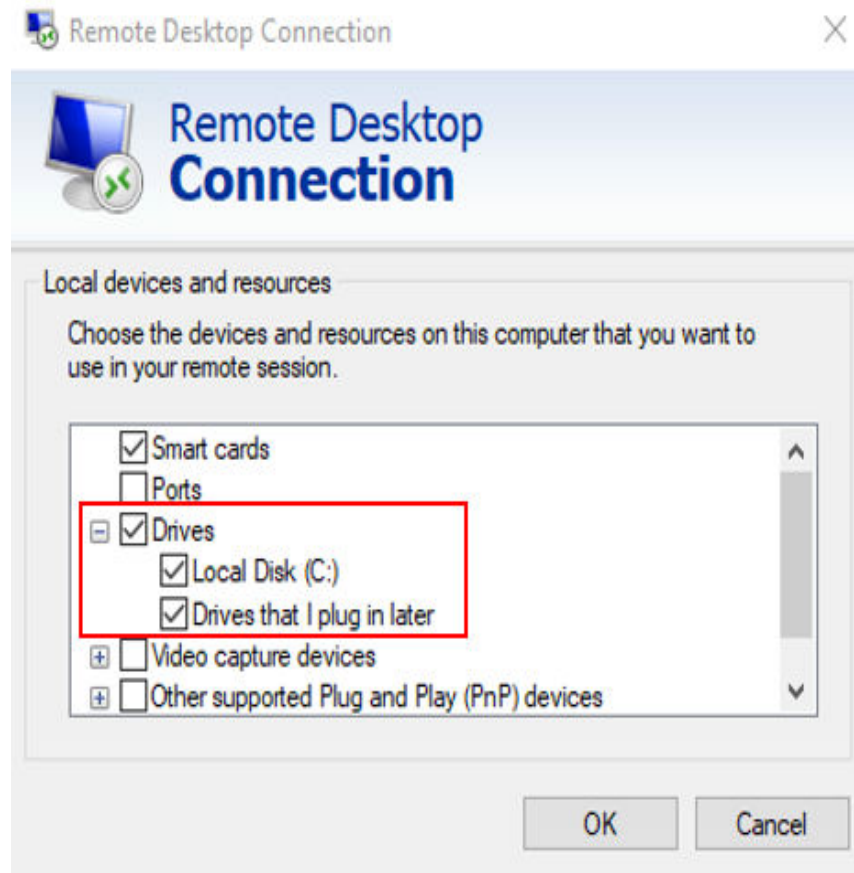
5. (Optional) To use local server resources in a remote session, configure parameters on the **Local Resources** tab.
To copy data from the local server to your ECS, select **Clipboard**.

Figure 2-33 Clipboard




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

Figure 2-34 Drives



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

Figure 2-35 Adjusting the size of the desktop

7. Click **Connect** and enter the login password as prompted to log in to the ECS. To ensure system security, change the login password after you log in to the ECS for the first time.
8. (Optional) Copy local files to the Windows ECS using clipboard. If the file size is greater than 2 GB, an error will occur. To resolve this issue, see [troubleshooting cases](#).

Enabling RDP

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


NOTE

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

1. Log in to the Windows ECS using VNC. For details, see [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 2-36 System Properties



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

Helpful Links

- [Login Password Resetting](#)
- [Multi-User Login Issues](#)
- [Remote Logins](#)

2.4.5 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 2-36 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.

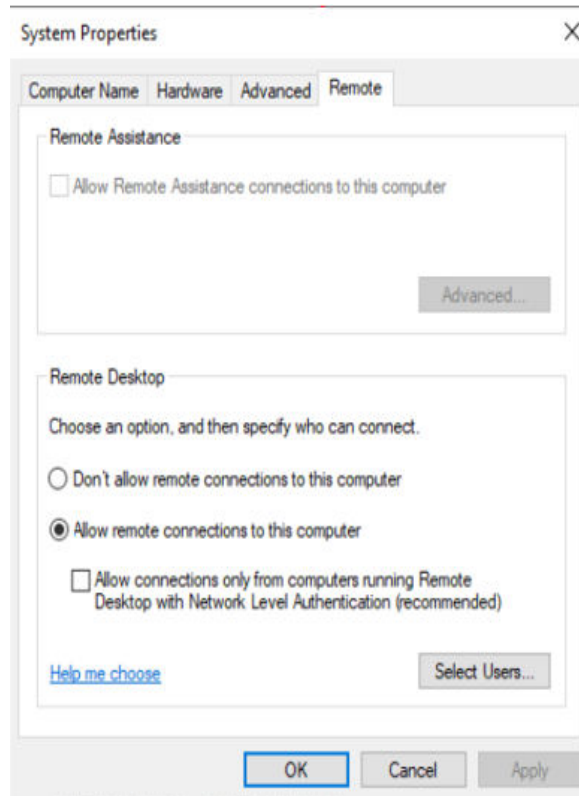
Enabling RDP

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 2-37 System Properties

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

2.4.6 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. If you have forgotten the password, reset the password by referring to [Resetting the Password for Logging In to an ECS on the Management Console](#).


- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group 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 2-38 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 2-39 Add user account



The 'Add a User Account' dialog box contains the following fields and controls:

- Username:** A text input field.
- Password:** A password input field with masked characters (dots). Below it is a checkbox labeled 'Show password'.
- Friendly name:** A text input field with the placeholder text 'Optional'.
- Buttons:** 'Cancel' and 'Add' buttons at the bottom right.

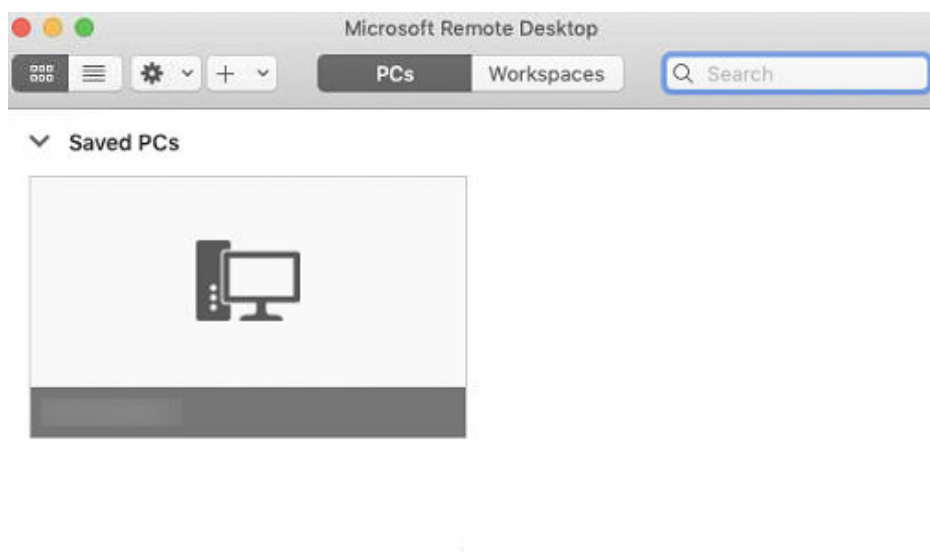
Figure 2-40 Add PC




The 'Add PC' dialog box contains the following fields and controls:

- PC name:** A text input field.
- User account:** A dropdown menu.
- Tabs:** 'General' (selected), 'Display', 'Devices & Audio', and 'Folders' tabs.
- Friendly name:** A text input field with the placeholder text 'Optional'.
- Group:** A dropdown menu with 'Saved PCs' selected.
- Gateway:** A dropdown menu with 'No gateway' selected.
- Checkboxes:**
 - ☒ Bypass for local addresses
 - ☒ Reconnect if the connection is dropped
 - ☐ Connect to an admin session
 - ☐ Swap mouse buttons
- Buttons:** 'Cancel' and 'Add' buttons at the bottom right. The 'Add' button is highlighted with a red rectangle.

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

Figure 2-41 Double-click for login

5. Confirm the information and click **Continue**.
You have logged in to the Windows ECS.

Figure 2-42 Successful login

2.4.7 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.

If your mobile device does not support remote desktop applications, you can use other login modes. For details, see [Login Overview \(Windows\)](#).

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If you have forgotten the password, reset the password by referring to [Resetting the Password for Logging In to an ECS on the Management Console](#).
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group 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 2-43 Remote desktop



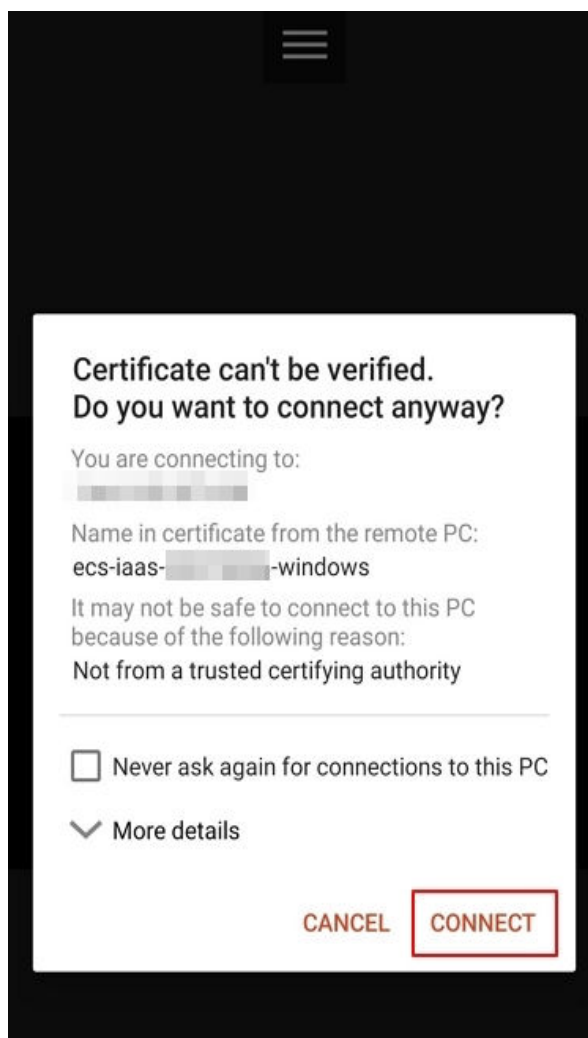
3. In the **Add desktop** dialog box, enter the Windows ECS hostname or EIP for **PC name**.
4. Click **SAVE**.
5. On the **Remote Desktop** page, click the icon of the Windows ECS to be logged in to.

Figure 2-44 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 2-45 Confirmation



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 2-46 Successful login

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**.
- If the login password is forgotten, reset the password on the ECS console.
To reset a password, locate the row containing the target ECS, click **More** in the **Operation** column, and select **Reset Password** from the drop-down list. For details, see [Resetting the Password for Logging In to an ECS on the Management Console](#).

Login Modes

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

Table 2-37 Linux ECS login modes

ECS OS	Local OS	Connection Method	Requirement
Linux	Windows	(Recommended) Use CloudShell available on the management console to log in to the ECS. Logging In to a Linux ECS Using CloudShell	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.)
	Windows	Use a remote login tool, such as PuTTY or Xshell. <ul style="list-style-type: none">• Password-authenticated: Logging In to a Linux ECS from a Local Windows Server• Key-pair-authenticated: Logging In to a Linux ECS from a Local Windows Server	
	Linux	Run commands. <ul style="list-style-type: none">• 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 CloudShell

Scenarios

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

For instructions about how to copy and paste data on CloudShell pages after the ECS login, see [Common CloudShell Operations](#).

Constraints

For details about the supported regions, see "Linux ECS Login" in the [Function Overview](#).

Prerequisites

- The target ECS is running.
- The login port (port 22 by default) was opened. If a different port is required, use the default port to log in to the ECS and then change the port number.

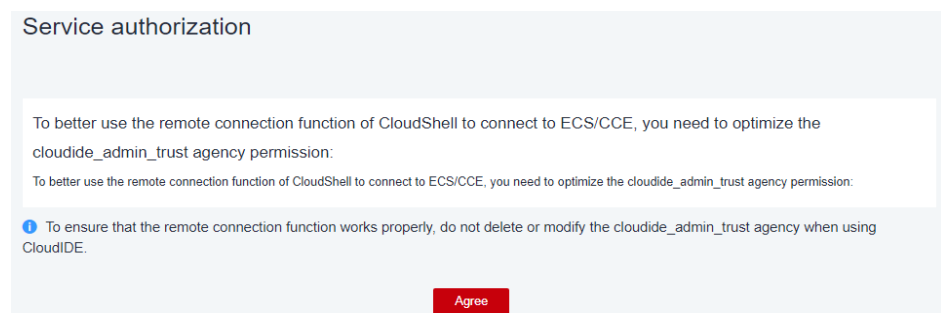
For details about how to change the remote login port, see [How Can I Change a Remote Login Port?](#) For details about how to configure security group rules, see [Configuring Security Group Rules](#).

- The password for logging in to the target ECS has been set. If you did not set a password when creating the ECS, reset the password before logging in to the ECS.
- You can use CloudShell to connect to an ECS through a public or private network. When you choose to connect through a private network, service authorization needs to be performed by a user with the Security Administrator permissions.

- If the **Service authorization** page is displayed, it means you have the Security Administrator permissions. Click **Agree**.

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

Figure 2-47 Service authorization



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



Perform the following steps to assign permissions:

- i. Create a user group and assign the Security Administrator permissions to the user group. For details, see [Creating a User Group and Assigning Permissions](#).
- ii. Add the user to the user group. For details, see [Adding Users to a User Group](#).

NOTE

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

Procedure

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

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


NOTE

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

- If you select the EIP bound to the ECS:
 - i. In the CloudShell configuration wizard, set the port (22 by default), username, authentication type, and password (or key) of the ECS.
 - ii. Click **Connect**.

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

Figure 2-48 CloudShell configuration wizard (EIP)



Region : Refresh

ECS :

☒ (EIP) ☐ (Private IP)

Port : 22

User : root

Auth-Type : Password-based

Password :

Session Name : root@

☒ Open Remote Host Filesystem


Note:

- To ensure the security of the connection, the system will automatically disconnect sessions that have not been active for more than 20 minutes.
- Please make sure to add inbound rules to allow external network traffic from CloudShell Proxy Server (SSH default port 22) to be sent to the ECSs in the security group.
- When operations get stuck after remote login, please check the CPU and memory of the machine. Please configure Cloud Eye to send alarm notifications when abnormal ECS events occur.
- Huawei CloudShell will not save your password, please keep it properly.

Connect Cancel


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

Figure 2-49 Successful login



- If you select the private IP address of the ECS:
 - i. Click **Go** to open the CloudShell configuration wizard.

Figure 2-50 CloudShell configuration wizard (private IP) 1




- ii. In the CloudShell configuration wizard, set the port (22 by default), username, authentication type, and password (or key) of the ECS.
 - iii. Click **Connect**.

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

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

Figure 2-51 Successful login



Common CloudShell Operations

- **New session**

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

- **Keyboard shortcuts**

Use keyboard shortcuts to edit commands.

Table 2-38 Keyboard shortcuts for CloudShell

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

- **Copy & Paste**

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

- **Historical records**

Scroll up or down the terminal to view historical records. By default, only the latest 1,000 lines of historical records are retained for terminals.

- **Customized layout for multiple terminals**

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

2.5.3 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



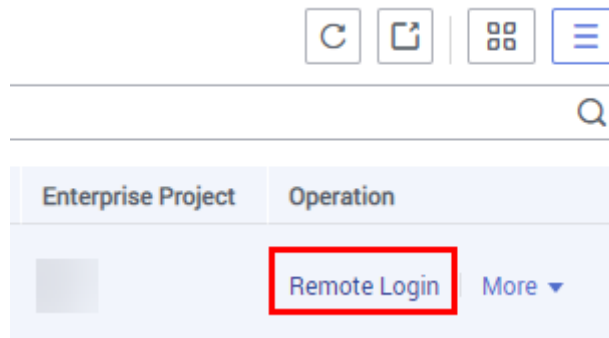
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the **Operation** column of the target ECS, click **Remote Login**.

Figure 2-52 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 2-53 Username (root as an example) and password




Follow-up Procedure

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

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

Figure 2-54 Paste & Send



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

 **NOTE**

There is a low probability that data is lost when you use 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?](#)

2.5.4 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:

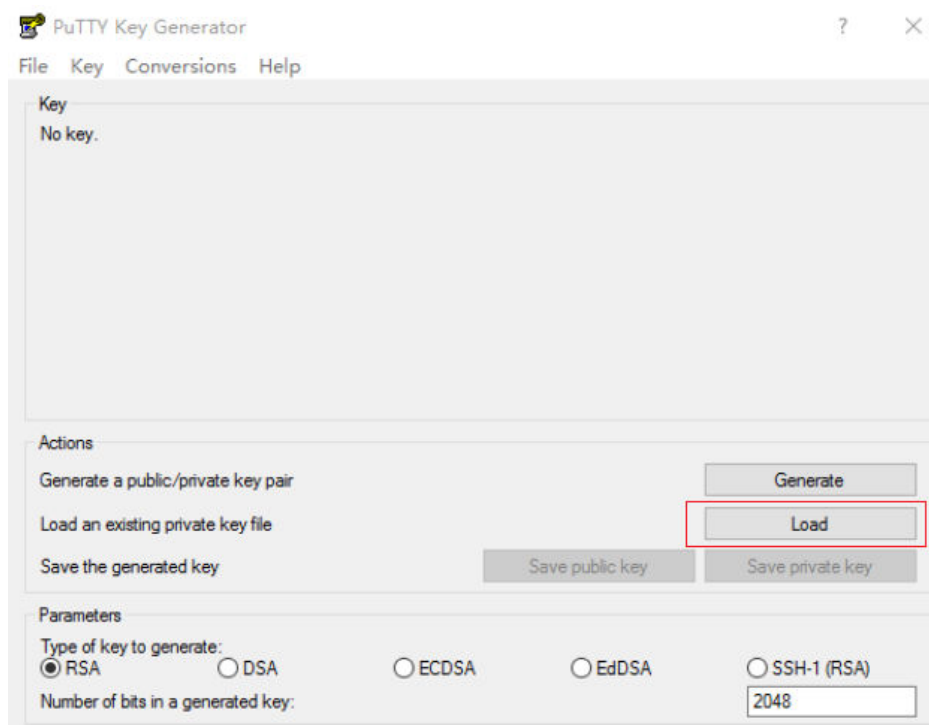
<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>

NOTE

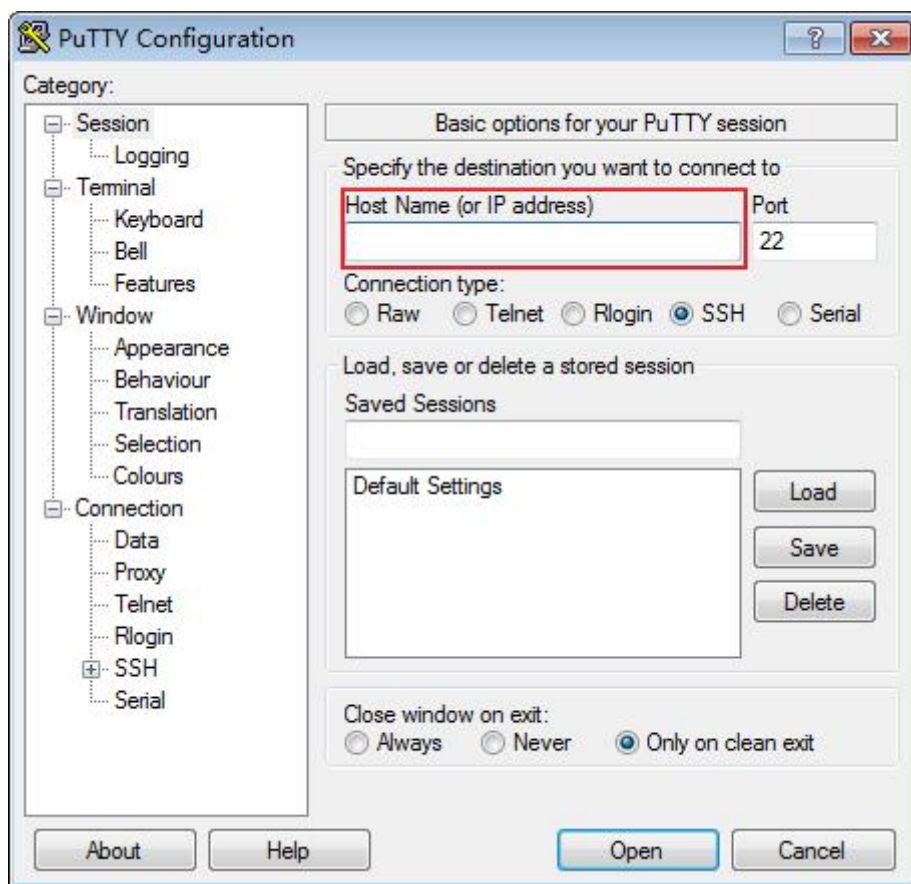
PuTTYgen is a key generator, which is used to create a key pair that consists of a public key and a private key for PuTTY.

3. Run PuTTYgen.
4. In the **Actions** pane, click **Load** and import the private key file that you stored during ECS creation.

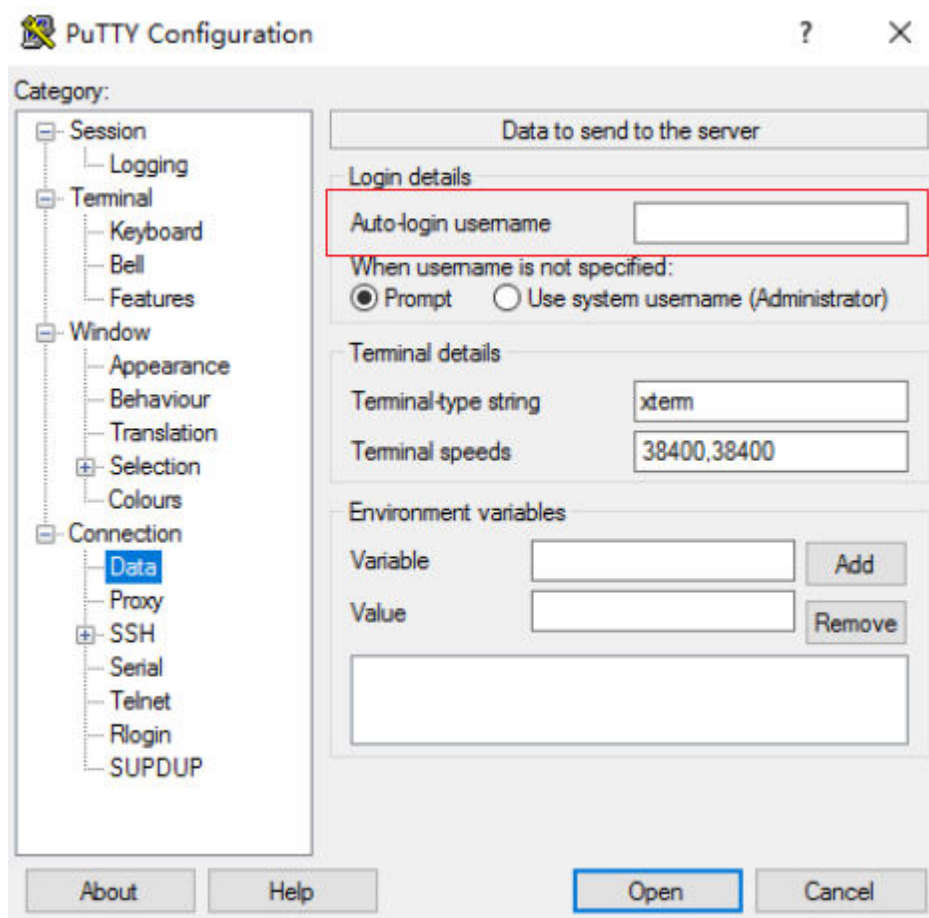
Ensure that the format of **All files (*.*)** is selected.

Figure 2-55 Importing the private key file

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 2-56 Configuring the EIP


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

Figure 2-57 Entering the username**NOTE**

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

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

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

Figure 2-58 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 2-59 SSH Security Warning

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

Logging In to a Linux ECS from a Local Linux Server

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

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

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

NOTE

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

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

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

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

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

NOTE

In the preceding command:

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

Follow-up Procedure

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

Helpful Links

- [What Can I Do If I Forget My Password for Remote Login?](#)
- [Why Can't I Log In to My Linux ECS?](#)

2.5.5 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 2-60 Session



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

If you log in to the ECS for the first time, PuTTY displays a security warning dialog box, asking you whether to accept the ECS security certificate. Click **Yes** to save the certificate to your local registry.

6. After the SSH connection to the ECS is set up, enter the username and password as prompted to log in to the ECS.

NOTE

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

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

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

Logging In to a Linux ECS from a Local Linux Server

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

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

```
ssh xx.xx.xx.xx
```

NOTE

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

2. Verify the SSH fingerprint of the ECS and enter **yes**.
The authenticity of host '`xx.xx.xx.xx (xx.xx.xx.xx)`' can't be established.
ECDSA key fingerprint is SHA256:rnKuzrUSYS03MCoaXXXXXXXXXXXXXXXXXXXXXXXXXXXX.
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?](#)

2.5.6 Logging In to a Linux ECS from a macOS Server

Scenarios

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 password-based SSH login, the **Login Mode** has been set to **Password** during the ECS purchase process and the username and password for logging in to the Linux ECS have been obtained.
If the password is forgotten, reset the password by referring to [Resetting the Password for Logging In to an ECS on the Management Console](#).
- 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.

2.5.7 Logging In to a Linux ECS from a Mobile Terminal

Scenarios

If you want to manage Linux ECSs anytime, anywhere, you can log in to them from an iOS or Android terminal.

This section describes how to remotely log in to a Linux ECS from a mobile terminal using the username and password.

- 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

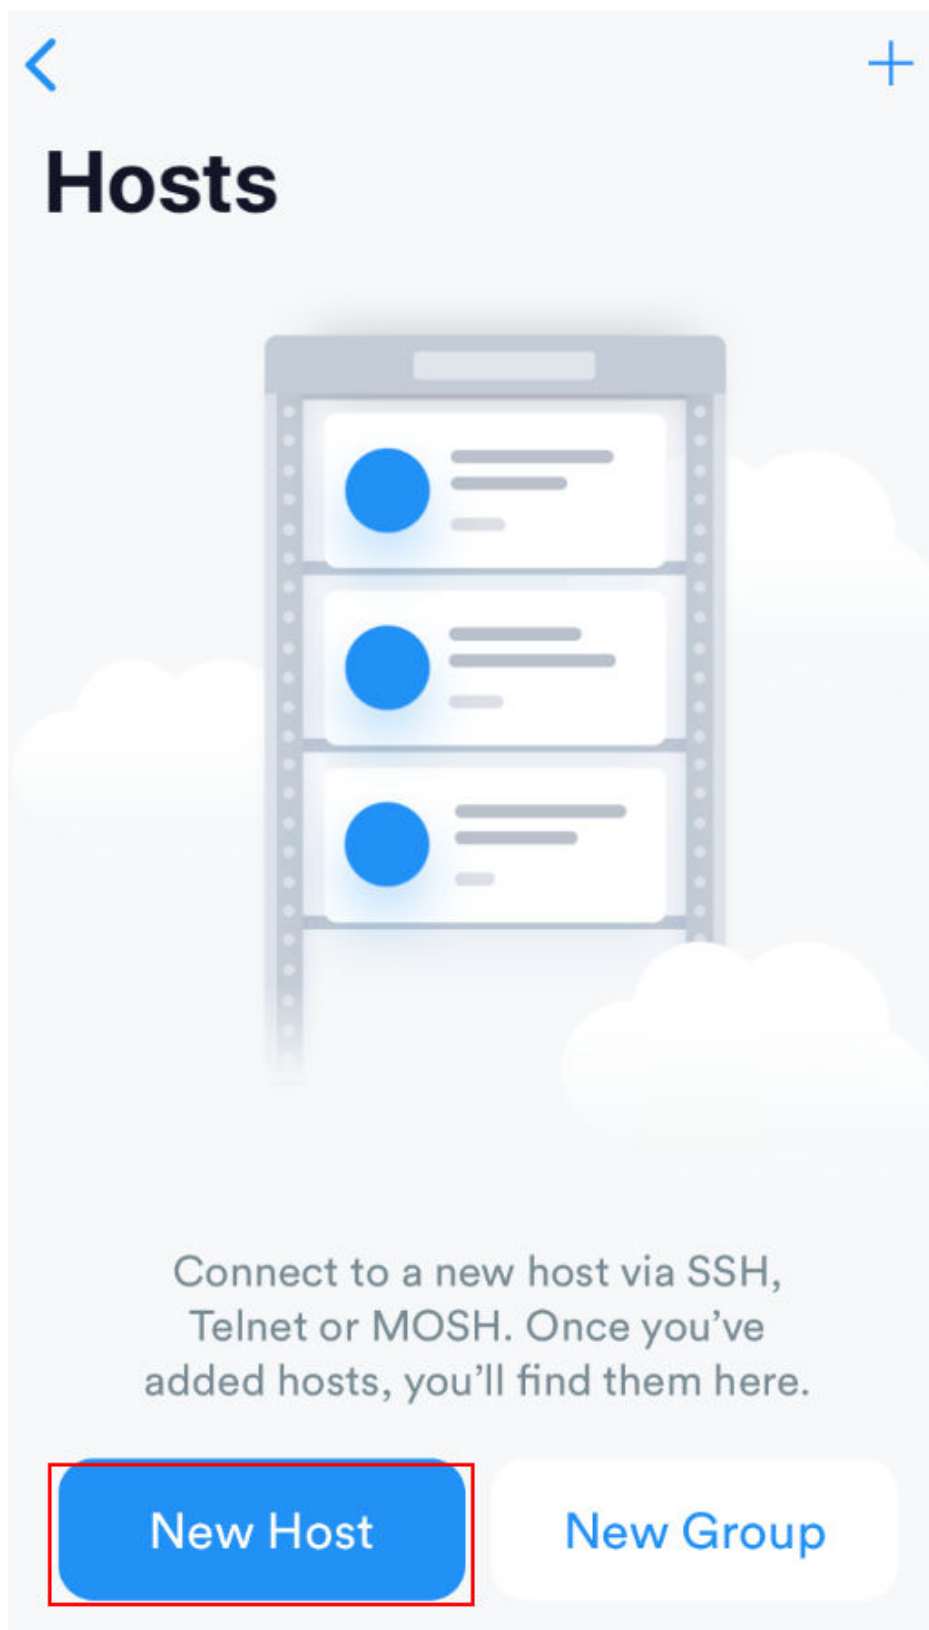
- The target ECS is running.
- You have obtained the username and password for logging in to the ECS. If the password is forgotten, reset the password by referring to [Resetting the Password for Logging In to an ECS on the Management Console](#).
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group 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 2-61 New Host



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

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

Figure 2-62 Setting parameters

Cancel New Host Save

1 Alias

2 Hostname

Group >

Tags >

Backspace as CTRL+H

SSH / MOSH

3 Use SSH

Use Mosh (Beta)

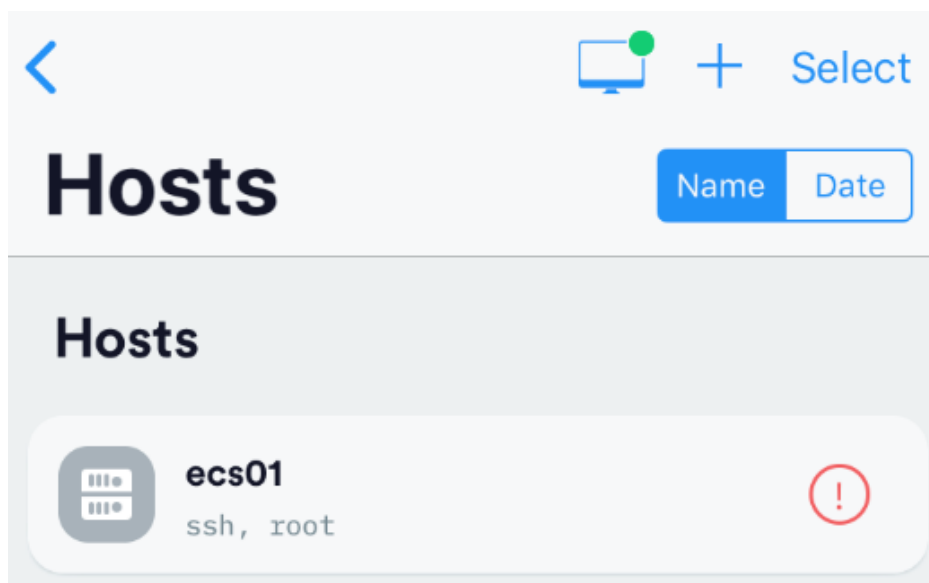
4 Port 22 Default

5 Username root

6 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 2-63 Login information



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

Figure 2-64 Connected

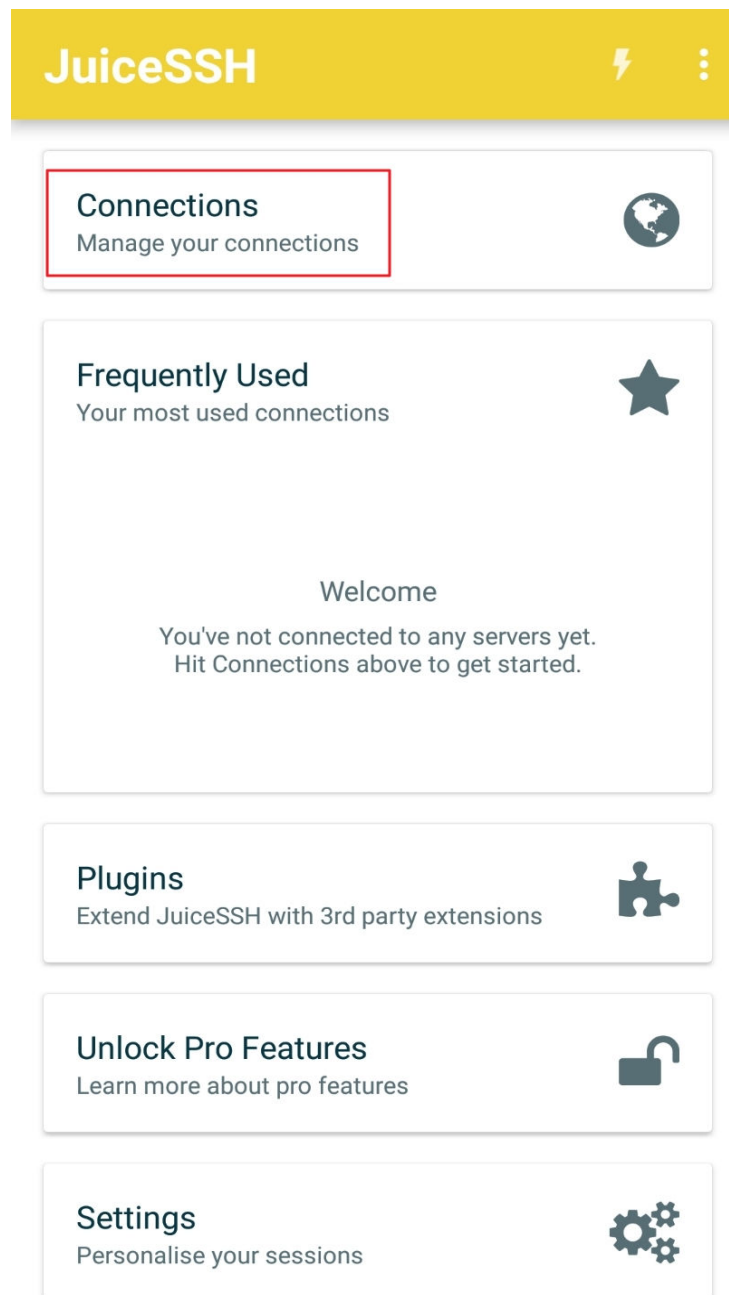


Logging In to a Linux ECS from an Android Terminal


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

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

Figure 2-65 Starting JuiceSSH




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

Figure 2-66 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 2-67 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.

GENERATE SNIPPET

- **Port:** Enter port number **22**.


Figure 2-68 Port

 **New Connection** 

BASIC SETTINGS

Nickname: linux_test

Type: SSH ▼

Address: 

Identity: linux_test ▼

ADVANCED SETTINGS

Port:

22

Connect Via: (Optional) ▼

Run Snippet: (Optional) ▼

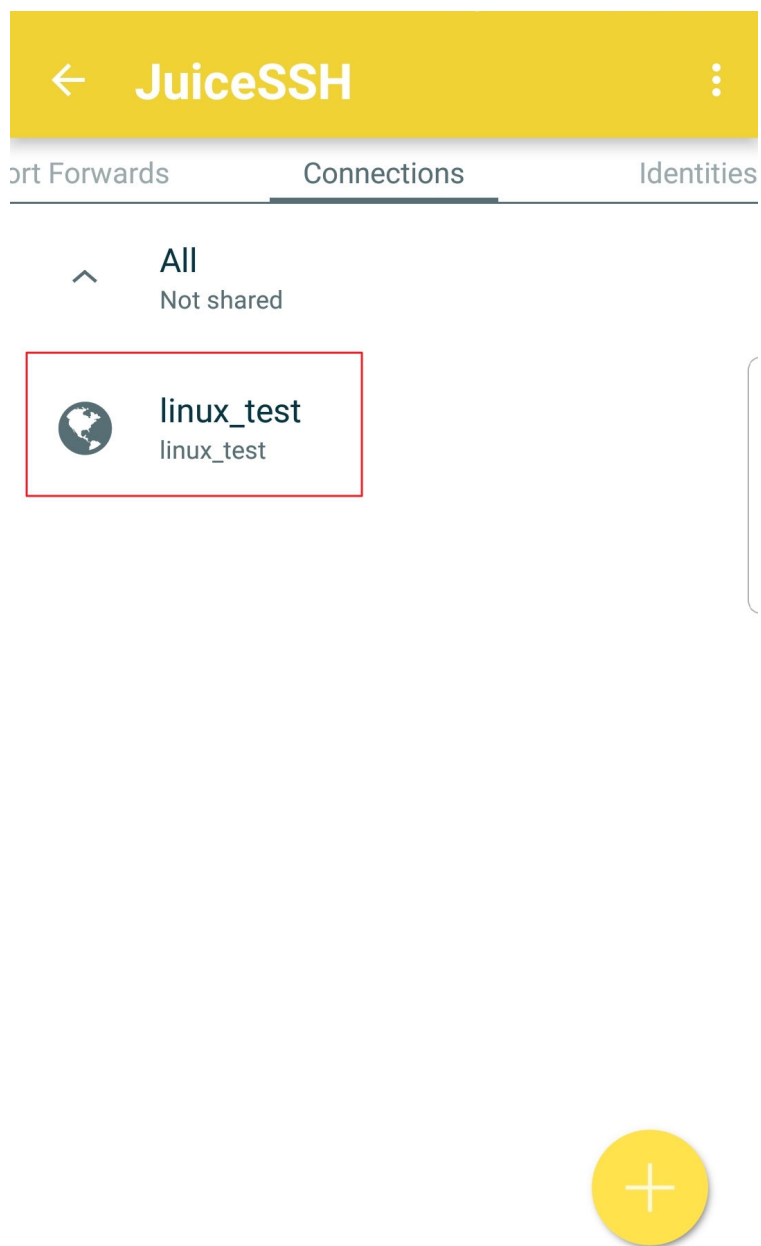
Backspace: Default (sends DEL) ▼

GROUPS

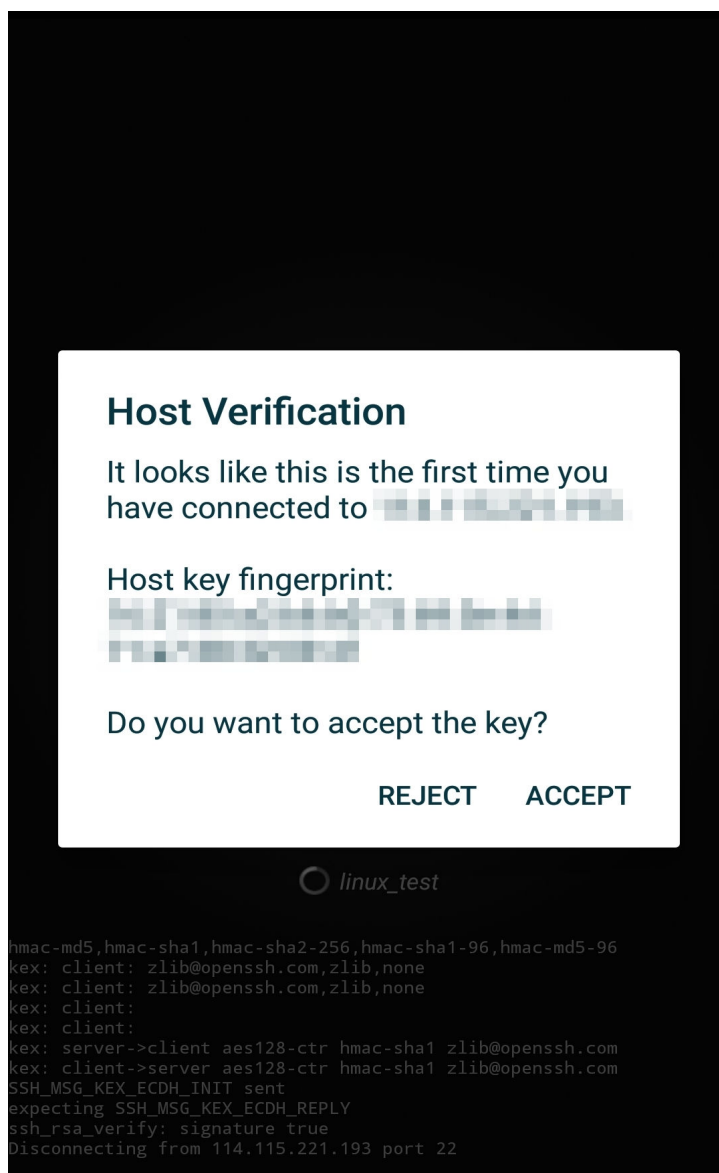
ADD TO GROUP

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

Figure 2-69 Connections

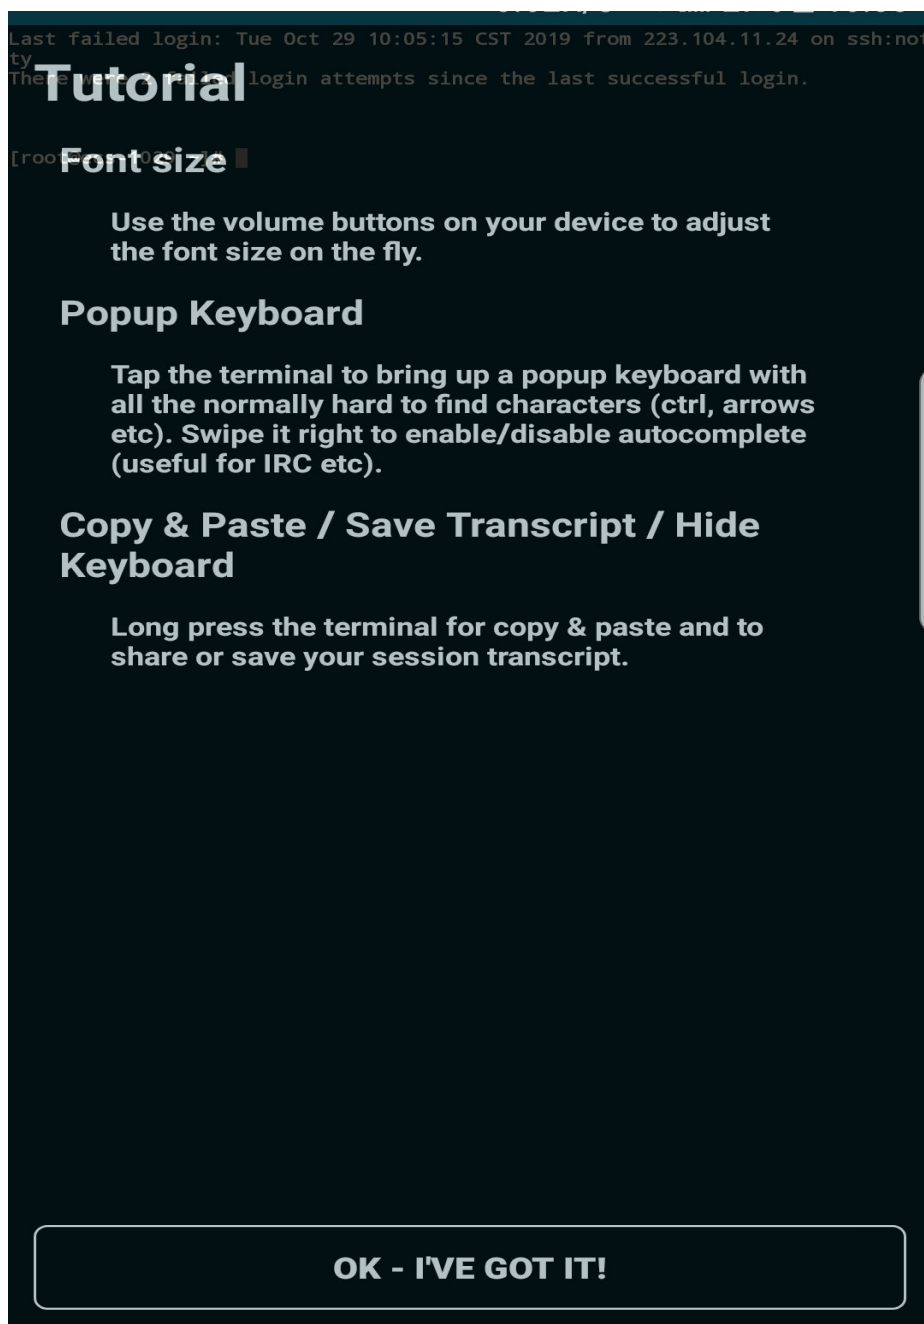


5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 2-70 Confirming the information


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

Figure 2-71 Tutorial



You have logged in to the Linux ECS.

Figure 2-72 Successful login



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 2-39 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 2-40 Mapping between Tesla drivers and ECS types

ECS Type	Driver	Product Series	Product
P2s	Tesla	V	V100
P2v	Tesla	V	V100
Pi2	Tesla	T	T4
PI1	Tesla	P	P4

Downloading a CUDA Toolkit

Download the [CUDA software package](#) and select the corresponding CUDA Toolkit software package based on the instance type and driver version.

NOTE

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 2-73 Selecting the CUDA Toolkit version

Manual Driver Search

Data Center / Tesla

T-Series

Tesla T4

Linux 64-bit

11.6

English (US)

Find

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](#)
4. [Installing the GRID Driver and Configuring the License](#)

NOTE

- NVIDIA allows you to apply for a 90-day trial license.
- For details about GPU-accelerated ECSs with different specifications and application scenarios, see [GPU-accelerated ECSs](#).

Purchasing a GRID License

- Purchase a license.
To obtain an official license, contact NVIDIA or their NVIDIA agent in your local country or region.
- Apply for a trial license.
Log in at the [official NVIDIA website](#) and enter desired information.
For details about how to 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 2-74 Applying for a trial license

Environment

* Certified Server: Other	▼	* NVIDIA GPUs: V100	▼
* VDI Hypervisor: RedHat Virtualization	▼	* VDI Remoting Client: Select an Option	▼
* VDI Seats: Select an Option	▼	Other Applications:	
* Primary Application: Select (1-3) Options	▼		

Downloading GRID Driver and Software License Packages

1. Obtain the driver installation package required for an OS. For details, see [Table 2-41](#).

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.

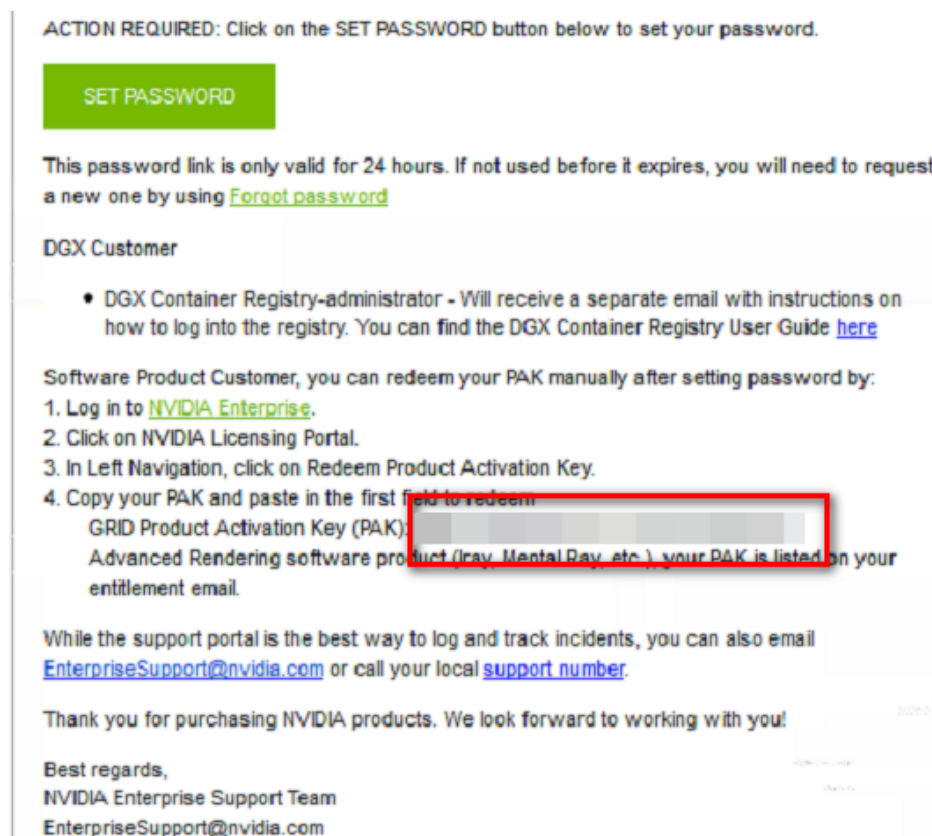
Table 2-41 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
G5.xlarge.4	GPU passthrough	<ul style="list-style-type: none">CentOS 7.6 64bitCentOS 7.5 64bitUbuntu 20.04 64bitUbuntu 18.04 64bit	Select a version as needed.	x86_64
P2s	GPU passthrough	<ul style="list-style-type: none">Huawei Cloud EulerOS 2.0 64bitCentOS 8.2 64bitCentOS 7.9 64bitCentOS 7.8 64bitCentOS 7.7 64bitCentOS 7.6 64bitCentOS 7.5 64bitUbuntu 22.04 server 64bitUbuntu 20.04 server 64bitUbuntu 18.04 server 64bitUbuntu 16.04 server 64bit	Select a version as needed.	x86_64
P2v	GPU passthrough	<ul style="list-style-type: none">CentOS 7.4 64bitEulerOS 2.2 64bitUbuntu 20.04 server 64bitUbuntu 18.04 server 64bitUbuntu 16.04 server 64bit	Select a version as needed.	x86_64

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
PI2	GPU passthrough	<ul style="list-style-type: none">• Huawei Cloud EulerOS 2.0 64bit• CentOS 8.2 64bit• CentOS 8.1 64bit• CentOS 8.0 64bit• CentOS 7.9 64bit• CentOS 7.8 64bit• CentOS 7.7 64bit• CentOS 7.6 64bit• CentOS 7.5 64bit• Ubuntu 22.04 server 64bit• Ubuntu 20.04 server 64bit• Ubuntu 18.04 server 64bit• Ubuntu 16.04 server 64bit	Select a version as needed.	x86_64
PI1	GPU passthrough	<ul style="list-style-type: none">• CentOS 7.3 64bit• Ubuntu 20.04 server 64bit• Ubuntu 16.04 server 64bit• Ubuntu 14.04 server 64bit	Select a version as needed.	x86_64

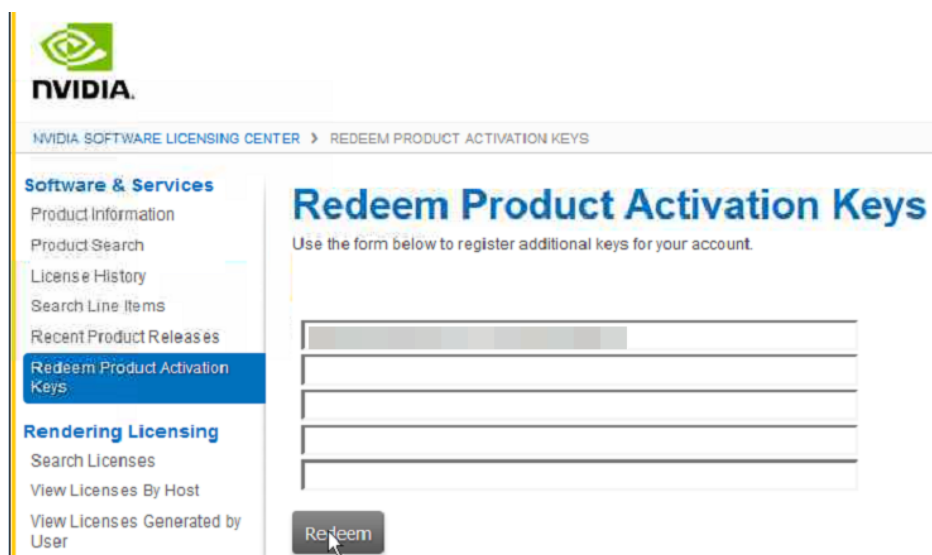
2. After the registration, log in at the [official NVIDIA website](#) and enter the account.
3. Check whether NVIDIA is used for the first time.
 - a. If yes, go to step [4](#).
 - b. If no, go to step [6](#).
4. Refer to [Figure 2-75](#) to obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 2-75 PAK




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

Figure 2-76 Redeem Product Activation Keys




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

Figure 2-77 Logging in to the official NVIDIA website

The image shows the NVIDIA Enterprise login page. At the top is the NVIDIA logo and the word "Enterprise". Below this is a "LOG IN" section. It contains two input fields: "Username:" and "Password:". The password field is filled with dots. To the right of the password field is a green "LOGIN" button. Below the login button are two links: "Forgot password?" and "Need help?".

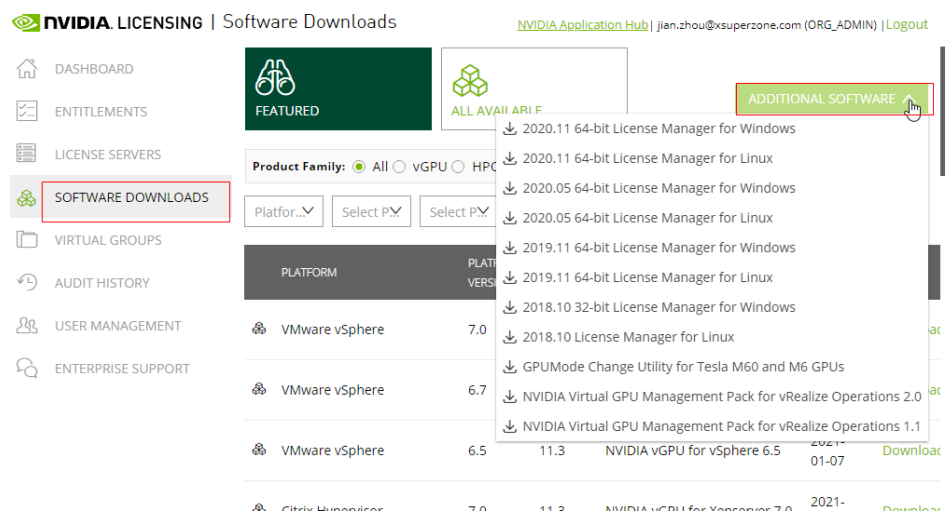
7. Log in at the official NVIDIA website as prompted and select **SOFTWARE DOWNLOADS**.

Figure 2-78 SOFTWARE DOWNLOADS page

The image shows the NVIDIA Software Downloads page. On the left is a sidebar with navigation links: DASHBOARD, ENTITLEMENTS, LICENSE SERVERS, SOFTWARE DOWNLOADS (highlighted), VIRTUAL GROUPS, AUDIT HISTORY, USER MANAGEMENT, and ENTERPRISE SUPPORT. The main content area has a header with "FEATURED" and "ALL AVAILABLE" tabs, and a "Product Family:" dropdown menu with options: All (selected), vGPU, HPC, and PGI. Below this is a search bar with "Platform...", "Select P...", "Select P...", and "Search" fields, and "CLEAR" and "COLUMNS" buttons. The main content is a table with columns: PLATFORM, PLATFORM VERSION, PRODUCT VERSION, DESCRIPTION, and RELEASE DATE. The table lists four entries for VMware vSphere and Citrix Hypervisor. Each entry has a "Download" link. At the bottom, there is a "page size: 25" dropdown and pagination controls showing "1 of 8".

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

8. Download the GRID driver of the required version. For details, see [Table 2-41](#).
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 2-79 ADDITIONAL SOFTWARE

Deploying and Configuring the License Server

The following uses an ECS running CentOS 7.5 as an example to describe how to deploy and configure the license server on the ECS.

NOTE

- The target ECS must have at least 2 vCPUs and 4 GiB of memory.
- Ensure that the MAC address of the target ECS has been recorded.
- If the license server is used in the production environment, deploy it in high availability mode. For details, see [official NVIDIA documentation for license server high availability](#).

1. Configure the network.
 - If the license server is to be accessed using the VPC, ensure that the license server and the GPU-accelerated ECS with the GRID driver installed are in the same VPC subnet.
 - If the license server is to be accessed using a public IP address, configure the security group 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
-----

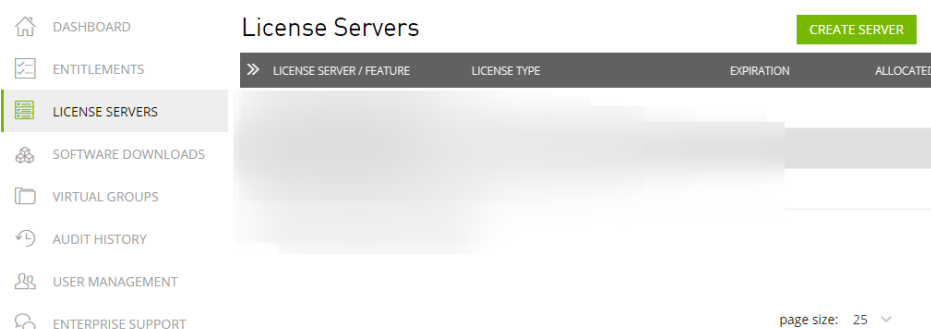
License Server has been successfully installed to:

    /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 2-80 LICENSE SERVERS



- b. Click **CREATE SERVER**.
- c. On the displayed **Create License Server** page, configure parameters.

Figure 2-81 Create License Server

Create License Server

Server Name

Name this license server

Description

Provide a short description

MAC Address

MAC Address (XX:XX:XX:XX:XX:XX or XX-XX-XX-XX-XX-XX)

Failover License Server

Failover License Server

Failover MAC Address

Failover MAC Address

Feature

Select a feature

Licenses

1

ADD

Added Features

FEATURE	COUNT
No features have been added yet	

CANCEL

RESET

CREATE LICENSE SERVER

Table 2-42 Parameters for creating a license server

Parameter	Description
Server Name	License server name, which can be customized.
Description	License description information.
MAC Address	MAC address of the ECS where the license server is deployed. You can log in to the ECS and run ipconfig -a to query the MAC address.
Feature	Select a feature, enter the number of required licenses in the Licenses text box, and click ADD . In active/standby deployment, enter the name of the standby server in Failover License Server and enter the MAC address in Failover MAC Address .

- d. Click **CREATE LICENSE SERVER**.
- e. Download the license file.

Figure 2-82 Downloading the license file

DASHBOARD

ENTITLEMENTS

LICENSE SERVERS

SOFTWARE DOWNLOADS

VIRTUAL GROUPS


License Servers

CREATE SERVER

LICENSE SERVER / FEATURE	LICENSE TYPE	EXPIRATION	ALLOCATED
test			
<div><div>DOWNLOAD LICENSE FILE</div><div>MANAGE LICENSES</div><div>REMOVE FEATURES</div><div>ADD FEAT</div><div>DELETE SERVER</div></div>			
Quadro-Virtual-DWS-5.0	CONCURRENT_COUNTED_SINGLE		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 2-83 Uploading a license file



Installing the GRID Driver and Configuring the License

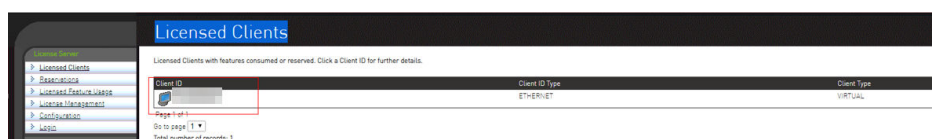
1. Install the GRID driver of a desired version, for example, on a GPU-accelerated Windows ECS.

 NOTE

Microsoft remote login protocols do not support GPU 3D hardware acceleration. To use this function, install third-party desktop protocol-compliant software, such as VNC, PCoIP, or NICE DCV, and access the ECS through the client.

2. Open the NVIDIA control panel on the Windows control panel.
3. Enter the IP address and port number of the deployed license server in the level-1 license server, and then click **Apply**. If the message indicating that you have obtained a GRID license is displayed, the installation is successful. Additionally, the MAC address of the GPU-accelerated ECS with the GRID driver installed is displayed on the **Licensed Clients** page of the license server management console.

Figure 2-84 License server management console



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 Tesla Driver on a Windows ECS](#)
- Installing a CUDA Toolkit
 - [Installing the CUDA Toolkit on a Linux ECS](#)
 - [Installing the CUDA Toolkit on a Windows ECS](#)

Installing a Tesla Driver on a Linux ECS

The following uses Ubuntu 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 2-85 Selecting a NVIDIA driver version

Manual Driver Search

Search by product, product type or series

Data Center / Tesla

i

T-Series

Tesla T4

Linux 64-bit

11.6

English (US)

Find

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

Figure 2-86 Selecting a driver version

Tesla Driver for Linux x64					
Driver Version: 418.67	CUDA Toolkit: 10.1	Release Date: Tue May 07, 2019	File Size: 107.23 MB	Info: 	View

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`

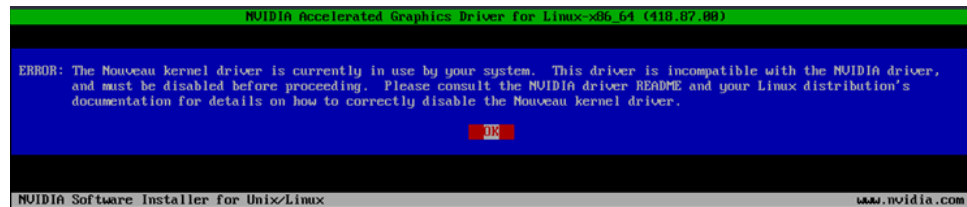
Figure 2-87 Obtaining the installation package

```
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 2-88 Disabling the Nouveau driver

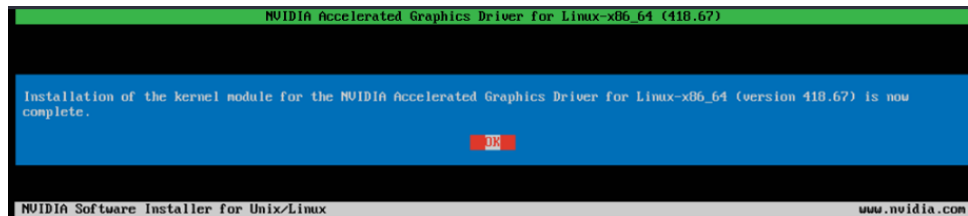


- a. Run the following command to check whether the Nouveau driver has been installed:
lsmod | grep nouveau
 - If the command output contains information about the Nouveau driver, the Nouveau driver has been installed and must be disabled. Then, go to step 9.b.
 - If the command output does not contain information about the Nouveau driver, the Nouveau driver has been disabled. Then, go to step 10.
- b. Edit the **blacklist.conf** file.
If the **/etc/modprobe.d/blacklist.conf** file is unavailable, create it.
vi /etc/modprobe.d/blacklist.conf
Add the following statement to the end of the file:

```
blacklist nouveau
options nouveau modeset=0
```
- c. Run the following command to back up and create an initramfs application:
 - Ubuntu
sudo update-initramfs -u
 - CentOS:
mv /boot/initramfs-\$(uname -r).img /boot/initramfs-\$(uname -r).img.bak
dracut -v /boot/initramfs-\$(uname -r).img \$(uname -r)

- d. Restart the ECS:
reboot
10. Select **OK** for three consecutive times as prompted to complete the driver installation.

Figure 2-89 Completing the NVIDIA driver installation



11. Run the following command to set systemd:
systemctl set-default multi-user.target
12. Run the **reboot** command to restart the ECS.
13. Log in to the ECS and run the **nvidia-smi** command. If the command output contains the installed driver version, the driver has been installed.

Figure 2-90 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 T4                Off        | 00000000:21:01.0 Off |                    |
| 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 a Tesla Driver on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install a Tesla driver on a GPU-accelerated ECS.

1. Log in to the ECS.
2. Download the NVIDIA driver package.
Select a driver version at [NVIDIA Driver Downloads](#) based on the ECS type.

Figure 2-91 Selecting a driver type (Windows)

Manual Driver Search

Search by product, product type or series

Data Center / Tesla

i

T-Series

Tesla T4

Windows Server 2016

10.1

English (US)

Find

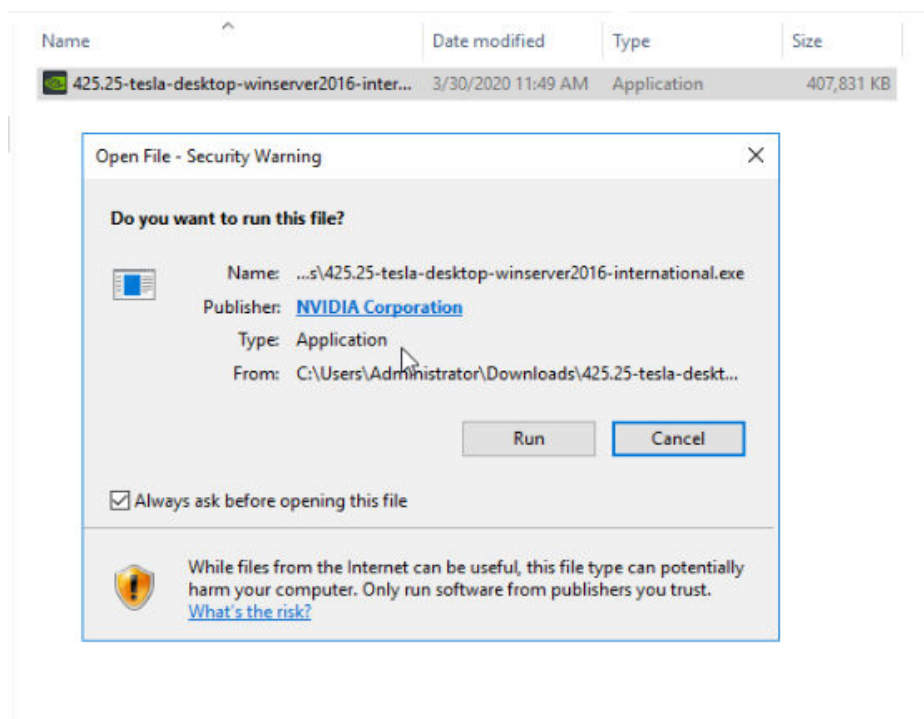
3. Select a driver version as required. The following uses Tesla 425.25 as an example.

Figure 2-92 Selecting a driver version (Windows)


Tesla Driver for Windows

Driver Version:	CUDA Toolkit:	Release Date:	File Size:	Info:	View
425.25	10.1	Tue May 07, 2019	417.62 MB	Production Branch WHQL	

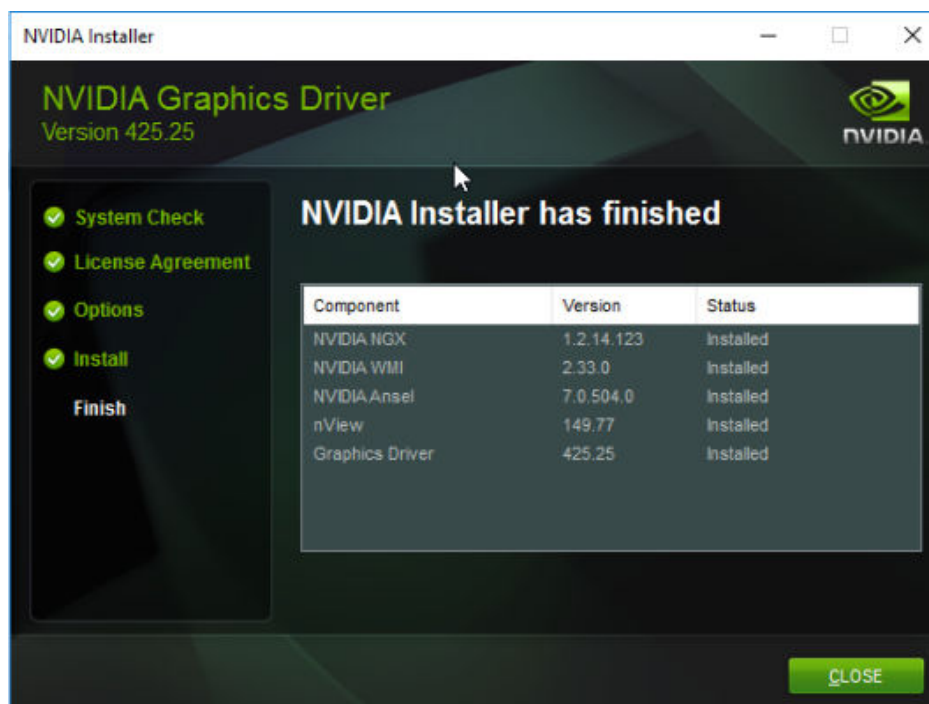
4. Click **View** in the row containing the driver to be downloaded.
5. Click **Download** to download the installation package.
6. Double-click the driver and click **Run**.

Figure 2-93 Running the NVIDIA driver installation program

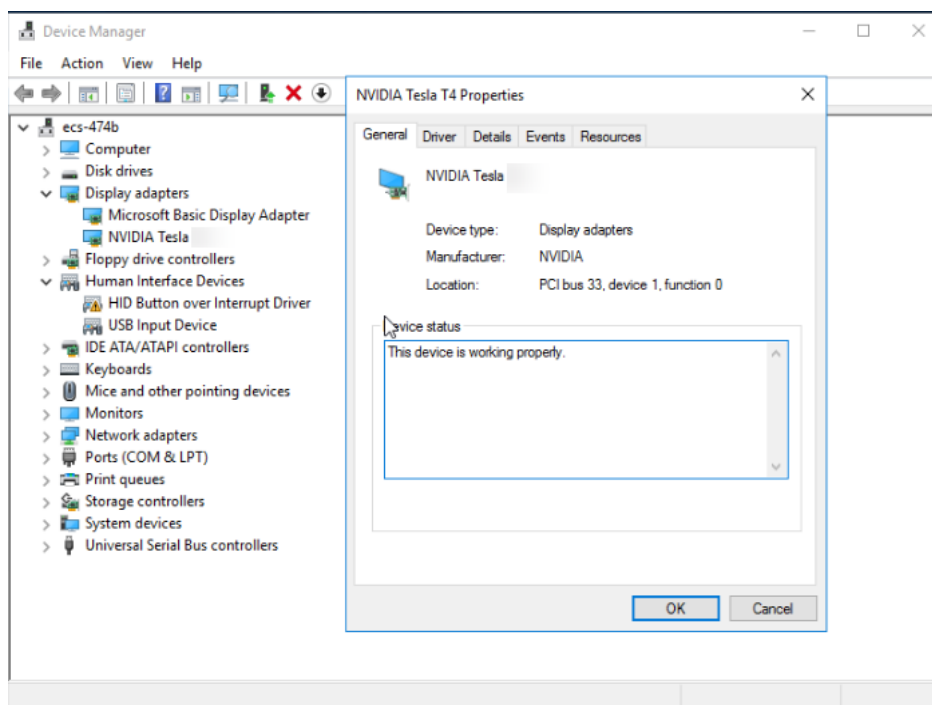
7. Select an installation path and click **OK**.

Figure 2-94 Selecting an installation path

8. Install the NVIDIA program as prompted.

Figure 2-95 Completing the driver installation

9. Restart the ECS.
10. Check whether the NVIDIA driver has been installed.
 - a. Switch to **Device Manager** and click **Display adapters**.

Figure 2-96 Display adapters

- b. Open the **cmd** window on the ECS and run the following commands:
cd C:\Program Files\NVIDIA Corporation\NVSMI

nvidia-smi

If the command output contains the installed driver version, the driver has been installed.

Figure 2-97 Viewing the NVIDIA driver version

```
C:\Program Files\NVIDIA Corporation\NVSMI>nvidia-smi
2020
+-----+
| NVIDIA-SMI 425.25      | Driver Version: 425.25      | CUDA Version: 10.1      |
+-----+-----+
| GPU  Name      | TCC/WDDM | Bus-Id      | Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf | Pwr:Usage/Cap |      Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
| 0   Tesla      | TCC      | 00000000:21:01.0 Off |   0x   |      0%      Default |
| N/A   33C    P8   | 11W / 70W | 0MiB / 15205MiB |           |              |
+-----+-----+
+-----+
| Processes:                 | GPU Memory Usage |
| GPU      PID  Type  Process name      | Usage |
+-----+
| No running processes found |
+-----+
C:\Program Files\NVIDIA Corporation\NVSMI>
```

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 2-98 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
Version	20.04	22.04	24.04					
Installer Type	deb (local)	deb (network)	runfile (local)					

- Find the link for downloading CUDA 10.1 corresponding to Ubuntu 20.04 64bit and copy the link.
- 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 2-99 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
--2022-11-10 12:27:59-- 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?FFJqELqU0U0EH-9_YUp4zpYc3t1j8q9J5clzHL2kX5Uo-f721Gy8tP3hm2JdJ2QntuMutpR-7nmp4Eb6Qjd7YB4BhuU56Di4U3prth0y22SX6Xuyw950k94i6Sb9he5nISJHhbPXUgVLzUhhijXhuBUCkMBMsuvoxn4X_uk_k2AY1oWbtU7D2sMaF8oU8 [following]
--2022-11-10 12:27:59-- https://developer.download.nvidia.com/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.39_linux.run?FFJqELqU0U0EH-9_YUp4zpYc3t1j8q9J5clzHL2kX5Uo-f721Gy8tP3hm2JdJ2QntuMutpR-7nmp4Eb6Qjd7YB4BhuU56Di4U3prth0y22SX6Xuyw950k94i6Sb9he5nISJHhbPXUgVLzUhhijXhuBUCkMBMsuvoxn4X_uk_k2AY1oWbtU7D2sMaF8oU8
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?FFJqELqU0U0EH-9_YUp4zpYc3t1j8q9J5clzHL2kX5Uo-f721Gy8tP3hm2JdJ2QntuMutpR-7nmp4Eb6Qjd7YB4BhuU56Di4U3prth0y22SX6Xuyw950k94i6Sb9he5nISJHhbPXUgVLzUhhijXhuBUCkMBMsuvoxn4X_uk_k2AY1oWbtU7D2sMaF8oU8 [following]
--2022-11-10 12:27:59-- https://developer.download.nvidia.cn/compute/cuda/10.1/secure/Prod/local_installers/cuda_10.1.105_418.39_linux.run?FFJqELqU0U0EH-9_YUp4zpYc3t1j8q9J5clzHL2kX5Uo-f721Gy8tP3hm2JdJ2QntuMutpR-7nmp4Eb6Qjd7YB4BhuU56Di4U3prth0y22SX6Xuyw950k94i6Sb9he5nISJHhbPXUgVLzUhhijXhuBUCkMBMsuvoxn4X_uk_k2AY1oWbtU7D2sMaF8oU8
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 (2.36) [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
```

- Install CUDA.
Follow the instructions provided on the official NVIDIA website.
- Run the following command to install CUDA:
sh cuda_10.1.243_418.87.00_linux.run
- Select **accept** on the installation page and press **Enter**.

Figure 2-100 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
```

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

Figure 2-101 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 2-102 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_Uilities/deviceQuery**:

cd /usr/local/cuda-10.1/samples/1_Uilities/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 2-103 deviceQuery common output

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

CUDA Device Query (Runtime API) version (CUDA RT static linking)
Detected 1 CUDA Capable device(s)

Device 0: "Tesla K80"
  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 2-104 Checking the CUDA version

```
root@ecs-474b:/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:/usr/local/cuda/bin#
```

12. Run the following command to enable the persistent mode:

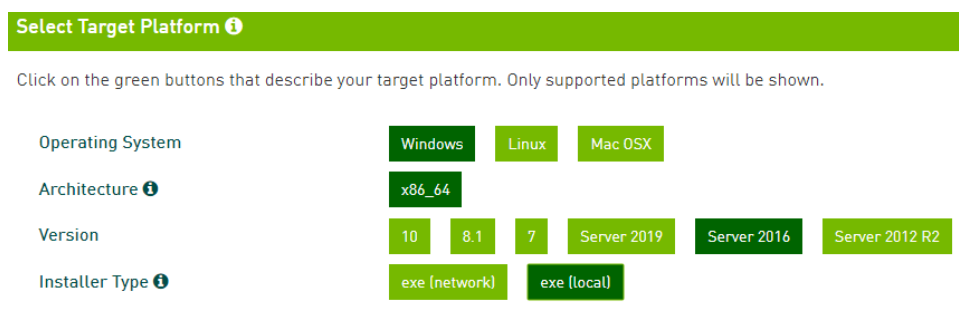
```
sudo nvidia-smi -pm 1
```

Enabling the persistent mode optimizes the GPU performance on Linux ECSs.


Installing the CUDA Toolkit on a Windows ECS

The following uses Windows Server 2016 Standard 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.

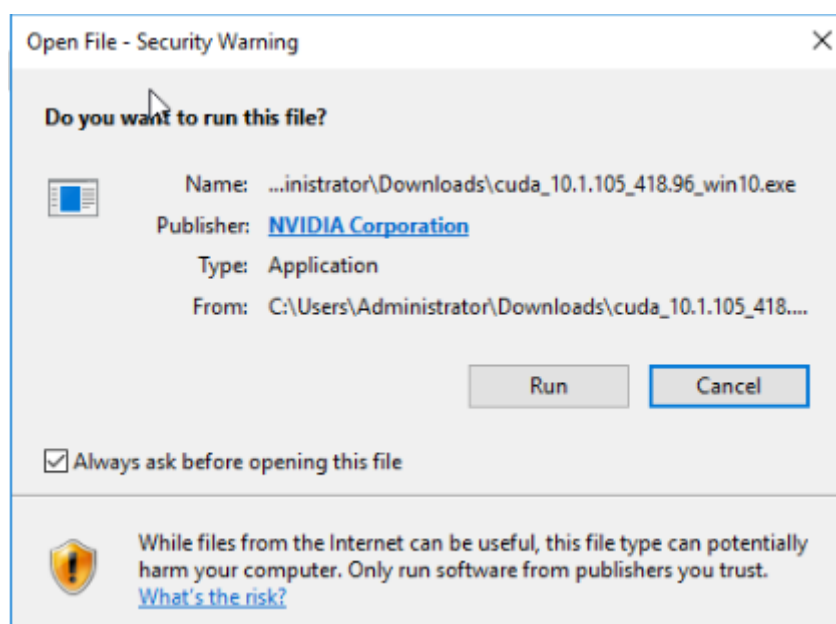
1. Log in to the ECS.
2. On the CUDA download page, set parameters according to the information shown in [Downloading a CUDA Toolkit](#).

Figure 2-105 Selecting a CUDA version


- Find the link for downloading CUDA 10.1.

Figure 2-106 Finding the link for downloading CUDA


- Click **Download** to download the CUDA toolkit.
- Double-click the installation file and click **Run** to install the CUDA toolkit.

Figure 2-107 Installing CUDA


- On the **CUDA Setup Package** page, select an installation path and click **OK**.

Figure 2-108 Selecting an installation path

7. Install the CUDA toolkit as prompted.

Figure 2-109 Completing the installation

8. Check whether CUDA has been installed
Open the **cmd** window and run the following command:
nvcc -V
If the command output contains the CUDA version, CUDA has been installed.

Figure 2-110 Successful installation

2.6.5 Uninstalling a GPU Driver from a GPU-accelerated ECS

Scenarios

You can manually uninstall a GPU driver from a GPU-accelerated ECS.

This section describes how to uninstall a GPU driver from a Windows ECS and a Linux ECS.

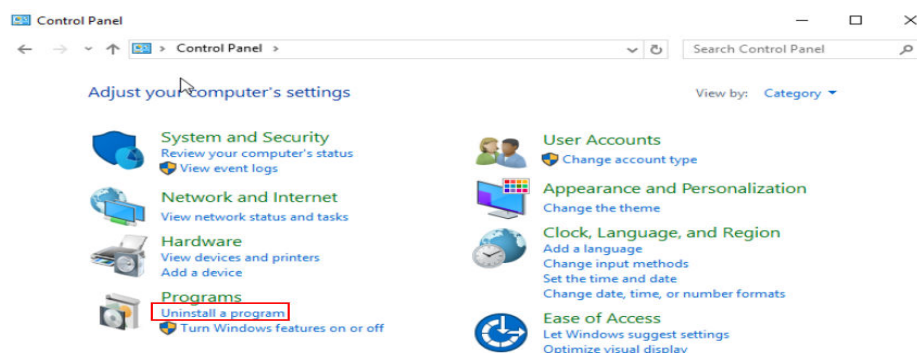
- [Uninstalling a GPU Driver from a Windows ECS](#)
- [Uninstalling a GPU Driver from a Linux ECS](#)

Uninstalling a GPU Driver from a Windows ECS

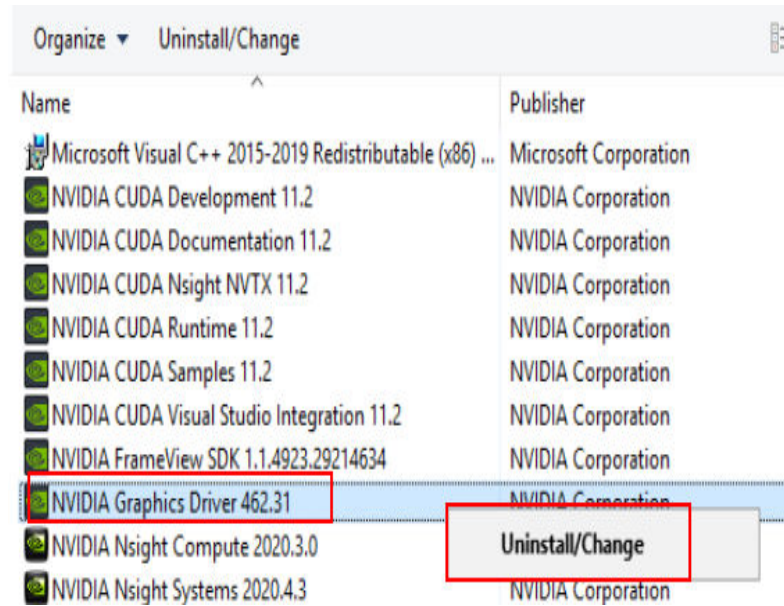
This section uses Windows Server 2016 Datacenter Edition 64-bit as an example to describe how to uninstall the NVIDIA driver (driver version: 462.31) from a GPU-accelerated ECS.

1. Log in to the ECS.
2. Click **Start** in the task bar and choose **Control Panel**.
3. In Control Panel, click **Uninstall a program** under **Programs**.


Figure 2-111 Uninstalling a program.



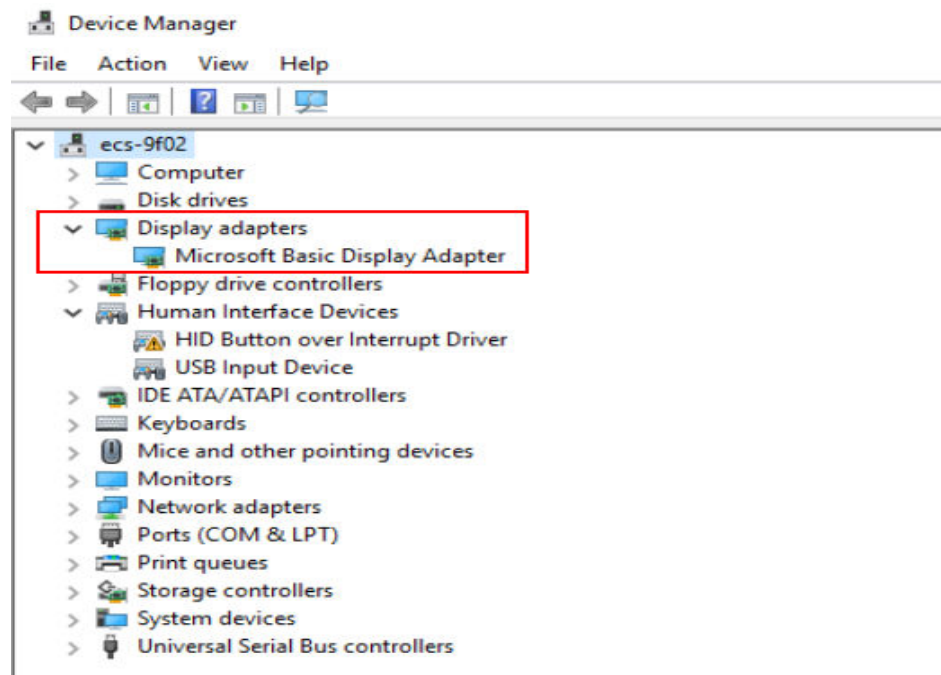
4. Right-click the NVIDIA driver to be uninstalled and choose **Uninstall/Change** from the shortcut menu.

Figure 2-112 Uninstalling a NVIDIA driver


5. In the displayed **NVIDIA Uninstaller** window, click **UNINSTALL**.

Figure 2-113 Confirming the uninstallation

6. After the uninstallation is complete, click **RESTART LATER**.
7. Check whether the NVIDIA driver has been uninstalled.
 - a. In Control Panel, click **Device Manager**.
If no NVIDIA graphics cards are not displayed under **Display adapters**, the driver is uninstalled successfully.

Figure 2-114 Viewing Display adapters

- b. Open the cmd window of the ECS and run the following commands:
- ```
cd C:\Program Files\NVIDIA Corporation\NVSMI
nvidia-smi.exe
```

**Figure 2-115** Command output

If the command output indicates that the file does not exist, the driver is uninstalled successfully.

After the NVIDIA driver is uninstalled, you can install a new NVIDIA driver without restarting the ECS.

## Uninstalling a GPU Driver from a Linux ECS

For NVIDIA Tesla drivers installed using .run Packages, you are advised to perform the following steps to uninstall it.

### NOTE

If you use .run Packages to install the NVIDIA Grid driver, you only need to perform [step 1](#) to uninstall the NVIDIA driver.

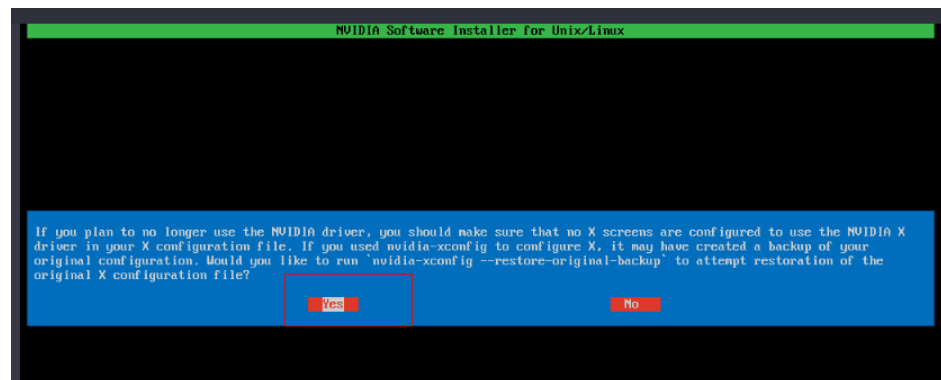
The following uses 64-bit Ubuntu Server 20.04 as an example to describe how to uninstall Tesla 460.73.01 and CUDA 11.2.

1. Uninstall the NVIDIA driver.
  - a. Query the path where **nvidia-uninstall** is stored.  
**whereis nvidia-uninstall**  
Generally, **nvidia-uninstall** is stored in the **/usr/bin/** directory.


**Figure 2-116** Querying the nvidia-uninstall path

```
root@ecs-# whereis nvidia-uninstall
nvidia-uninstall: /usr/bin/nvidia-uninstall
```


- b. Uninstall the driver from the path where **nvidia-uninstall** is stored.  
**/usr/bin/nvidia-uninstall**
  - c. Select **Yes** and press **Enter**.

**Figure 2-117** NVIDIA driver uninstallation (1)

- d. Select **OK** and press **Enter**.

**Figure 2-118** NVIDIA driver uninstallation (2)

- e. After the driver is uninstalled, press **Enter**.

**Figure 2-119** NVIDIA driver uninstallation (3)

2. Uninstall the CUDA and CUDA Deep Neural Network (cuDNN) libraries.

To upgrade the CUDA driver version, uninstall the corresponding CUDA library and then install a new one with the target version.

- a. Uninstall the CUDA library.

**/usr/local/cuda/bin/cuda-uninstaller**

Generally, **cuda-uninstaller** is stored in the **/usr/local/cuda/bin** directory.


 **NOTE**

The uninstallation command varies depending on CUDA versions. If the **cuda-uninstaller** file is not found, check whether a file starting with **uninstall\_cuda** exists in the **/usr/local/cuda/bin/** directory.

If such a file exists, replace **cuda-uninstaller** in the preceding command with the file name.

- b. On the uninstallation page, select all options, move the cursor to **Done**, and press **Enter**.

**Figure 2-120** Uninstalling a CUDA driver



If the CUDA library is uninstalled, the message "Successfully uninstalled" is displayed.

- c. Remove the CUDA and cuDNN libraries.  
**rm -rf /usr/local/cuda-11.2**

## 2.7 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 **/usr/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 **/usr/share/zoneinfo** includes both time zones and directories. The directories contain time zone files for specific cities. Locate the time zone for the city in which the ECS is located.

For example:

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

4. Set the target time zone.
  - a. Run the following command to open the **/etc/sysconfig/clock** file:  
**vim /etc/sysconfig/clock**
  - b. Locate the **ZONE** entry and change its value to the name of the desired time zone file.

For example:

- If the target time zone is for Shanghai, China, change the **ZONE** entry value as follows:  
**ZONE="Asia/Shanghai"**
- If the target time zone is for Paris, France, change the **ZONE** entry value as follows:  
**ZONE="Europe/Paris"**

5. Press **Esc**. Then, run the following command to save and exit the **/etc/sysconfig/clock** file:  
**:wq**
6. Run the following command to check whether the **/etc/localtime** file is available on the ECS:  
**ls /etc/localtime**

- If the file is available, go to step 7.
  - If the file is not available, go to step 8.
7. Run the following command to delete the existing **/etc/localtime** file:  
**rm /etc/localtime**
  8. Run the following command to create a symbolic link between **/etc/localtime** and your time zone file so that the ECS can find this time zone file when it references the local time:  
**ln -sf /usr/share/zoneinfo/Asia/city1 /etc/localtime**
  9. Run the following command to restart the ECS so that all services and applications running on the ECS use the new time zone:  
**reboot**
  10. Log in to the ECS again and run the following command as user **root** to check whether the time zone has been changed:

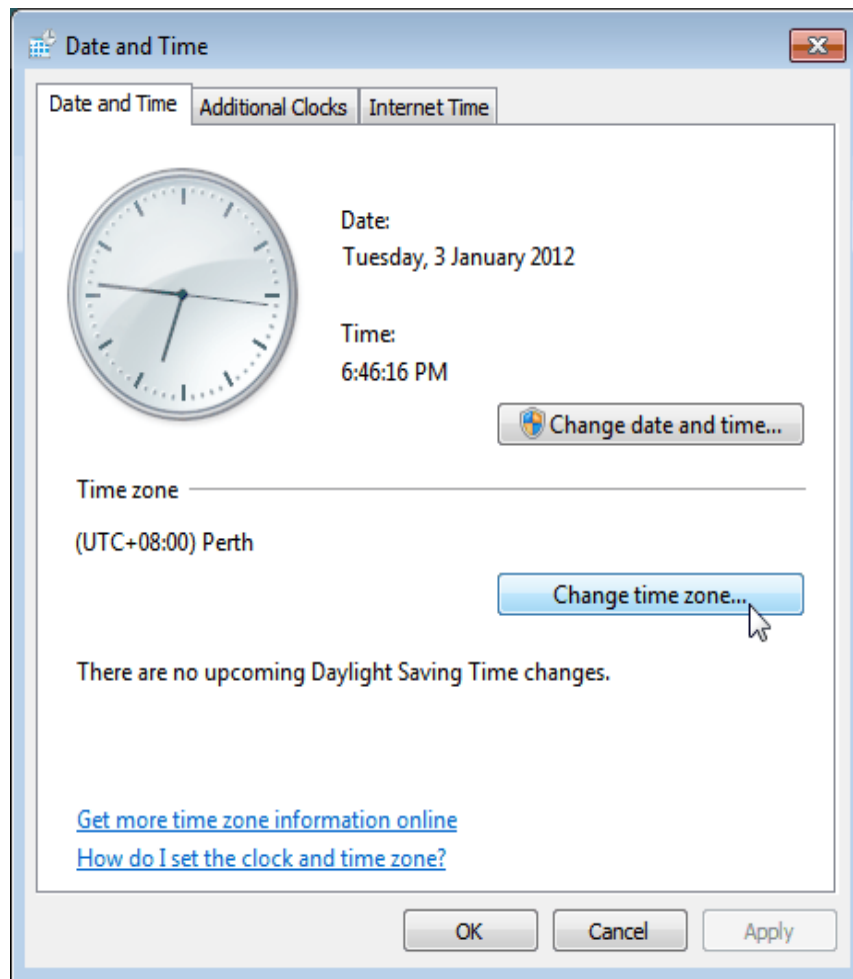
**ls -lh /etc/localtime**

The following information is displayed:

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

## For Windows ECSs

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

**Figure 2-121** Date and Time

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

## 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.

For example, for an ECS with the c7.xlarge.2 flavor (hyper-threading enabled by default), it has four vCPUs, which means four hyperthreads. After hyper-threading is disabled, it has two vCPUs, which means two CPU cores.


## Notes and Constraints

- After an x86 ECS is created, 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 in Custom Config Mode](#).
2. Select **Configure now** to configure advanced options.
3. Select **Specify CPU options**.



**Figure 2-122** Specifying CPU options



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.

Before modifying the specifications, stop the ECS or select **Authorize ECS auto-stop**.

Figure 2-123 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: ecs-... ID: 405a23e1-...  
Image: CentOS 7.5 64bit Billing Mode: Pay-per-use  
Specifications: General computing-plus | c7n.large.2 | 2 vCPUs | 4 GiB Region: CN-Hong Kong

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

**General computing-plus** | General computing | Memory-optimized

| ECS Type                  | Flavor Name   | vCPUs    | Memory | CPU                    | Assured / Maximum Bandwidth | Packets Per Second | IPv6 |
|---------------------------|---------------|----------|--------|------------------------|-----------------------------|--------------------|------|
| General computing-plus... | c7n.large.2   | 2 vCPUs  | 4 GiB  | Intel Ice Lake 2.6G... | 0.8 / 4 Gbit/s              | 400,000 PPS        | Yes  |
| General computing-plus... | c7n.large.4   | 2 vCPUs  | 8 GiB  | Intel Ice Lake 2.6G... | 0.8 / 4 Gbit/s              | 400,000 PPS        | Yes  |
| General computing-plus... | c7n.xlarge.2  | 4 vCPUs  | 8 GiB  | Intel Ice Lake 2.6G... | 1.6 / 8 Gbit/s              | 800,000 PPS        | Yes  |
| General computing-plus... | c7n.xlarge.4  | 4 vCPUs  | 16 GiB | Intel Ice Lake 2.6G... | 1.6 / 8 Gbit/s              | 800,000 PPS        | Yes  |
| General computing-plus... | c7n.2xlarge.2 | 8 vCPUs  | 16 GiB | Intel Ice Lake 2.6G... | 3 / 16 Gbit/s               | 1,500,000 PPS      | Yes  |
| General computing-plus... | c7n.2xlarge.4 | 8 vCPUs  | 32 GiB | Intel Ice Lake 2.6G... | 3 / 16 Gbit/s               | 1,500,000 PPS      | Yes  |
| General computing-plus... | c7n.3xlarge.2 | 12 vCPUs | 24 GiB | Intel Ice Lake 2.6G... | 5 / 17 Gbit/s               | 2,000,000 PPS      | Yes  |

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

**CPU options** ☒ **Specify CPU options**

Threads per core: 2

Price: ... Next

6. Click **Next**.
7. Confirm the settings, read and agree to 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 2-124](#), **CPU(s)** indicates the number of logical cores.

Figure 2-124 Viewing the result

```
[root@ecs-~]# 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 2-125 Viewing the result



## 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 2-43](#). The following describes the URI and methods of using the supported ECS metadata.

#### Notes

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

Perform the following configuration on the firewall:

- Windows

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

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

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

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

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

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

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

- Linux

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

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

#### ECS Metadata Types

[Table 2-43](#) 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 2-43** ECS metadata types

| Metadata Type  | Metadata Item             | Description                                                                                                                                                                                                                                                                                                                                                                 |
|----------------|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| OpenStack      | /meta_data.json           | Displays ECS metadata.<br>For the key fields in the ECS metadata, see <a href="#">Table 2-44</a> .                                                                                                                                                                                                                                                                          |
| OpenStack      | /password                 | Displays the password for logging in to an ECS.<br>This metadata is used by Cloudbase-Init to store ciphertext passwords during initialization of key-pair-authenticated Windows ECSs.                                                                                                                                                                                      |
| OpenStack      | /user_data                | Displays ECS user data.<br>This metadata allows you to specify scripts and configuration files for initializing ECSs. For details, see <a href="#">Injecting User Data</a> .<br>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.<br>Before enabling an ECS to obtain a temporary AK and SK, authorize agency permissions to the <b>op_svc_ecs</b> account and ECSs in IAM.<br><b>NOTE</b><br>You can determine what permissions are granted to the agency based on the principal of least privilege (PoLP).<br>ECSs will not use agencies to perform operations on resources. |
| OpenStack      | /spot/instance-action     | Queries the prompt of stopping a spot ECS.                                                                                                                                                                                                                                                                                                                                  |
| EC2-compatible | /meta-data/hostname       | Displays the name of the host accommodating an ECS.<br>To remove the suffix <b>.novalocal</b> from an ECS, see:<br><a href="#">Is an ECS Hostname with Suffix .novalocal Normal?</a>                                                                                                                                                                                        |
| EC2-compatible | /meta-data/local-hostname | The meaning of this field is the same as that of hostname.                                                                                                                                                                                                                                                                                                                  |

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

**Table 2-44** Metadata key fields

| Parameter             | Type   | Description                                                                                                                                                                           |
|-----------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uuid                  | String | Specifies an ECS ID.                                                                                                                                                                  |
| availability_zone     | String | Specifies the AZ where an ECS locates.                                                                                                                                                |
| meta                  | Dict   | Specifies the metadata information, including the image name, image ID, and VPC ID.                                                                                                   |
| hostname              | String | Specifies the name of the host accommodating an ECS.<br>To remove the suffix <b>.novalocal</b> from an ECS, see:<br><a href="#">Is an ECS Hostname with Suffix .novalocal Normal?</a> |
| 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**

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

**Invoke-RestMethod http://169.254.169.254/openstack/latest/meta\_data.json | ConvertTo-Json**

```
{
 "random_seed": "rEocCViRS+dNwlYdGlxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRny4kKGoNPEVBCc05Hg1TcDbLAPfJwgJS1okqEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAJjqsahnU+g/tawABTVySLBKlAT8fMGax1mTGgArucn/
WzDcy19DGioKPE7F8lLtSQ4Ww3VCLK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbo3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmxth7IGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNIHA/NvlEsxDxqBCoss/Jfe+yCmUFyxov/
+L8oNkTzkmtCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA
+qyoM6CVyftEoz1zIRryo09bJ65Eg6Jd8dj1UCVsDqRY1pljgzE/
Mzsw6AaaCVhaMJL7u7YmVdyKzA6z65Xtvujz0Vo=",
 "uuid": "ca9e8b7c-f2be-4b6d-a639-f10b4d994d04",
 "availability_zone": "lt-test-1c",
 "enterprise_project_id": "0",
 "hostname": "ecs-ddd4.novalocal",
 "launch_index": 0,
 "instance_type": "s3.medium.2",
 "meta": {
 "metering.image_id": "3a64bd37-955e-40cd-ab9e-129db56bc05d",
 "metering.imagetype": "gold",
 "metering.resourcespeccode": "s3.medium.2.linux",
 "metering.cloudServiceType": "hws.service.type.ec2",
 "image_name": "CentOS 7.6 64bit",
 "metering.resourcetype": "1",
```

```
"vpc_id": "3b6c201f-aeb3-4bce-b841-64756e66cb49",
"os_bit": "64",
"cascaded.instance_extrainfo": "pcibridge:1",
"os_type": "Linux",
"charging_mode": "0"
},
"region_id": "xxx",
"project_id": "6e8b0c94265645f39c5abbe63c4113c6",
"name": "ecs-ddd4"
}
```

## User Data (OpenStack Metadata API)

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

- URI  
/169.254.169.254/openstack/latest/user\_data

- Usage method  
Supports GET requests.

- Example  
Linux:

**curl http://169.254.169.254/openstack/latest/user\_data**

Windows:

**Invoke-RestMethod http://169.254.169.254/openstack/latest/user\_data**

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkl0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhllH
BsYWNlIHRRvIGdvlG5vdy4gQnV0IHRob2ZSBza3kga25vd3MgdGhllHJlYXNvbnMgYW5kIHRob2ZSBwYXR0ZXJl
cyBiZWphbmQgYWxsIGNsb3VkcycgYW5kIHlvdSB3aWxsIGtub3csIHRobywgY2h1biB5b3UgbGlmdCB5b3
Vyc2VsZiBoaWd0IGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=
```

### NOTE

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

Figure 2-126 404 Not Found

```
[root@python3dktmest--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.

</body>
</html>
```

## Network Data (OpenStack Metadata API)

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

- URI  
/openstack/latest/network\_data.json
- Usage method  
Supports GET requests.

- Example

 NOTE

**instance\_max\_bandwidth** and **instance\_min\_bandwidth** are in the unit of Mbit/s. If the value is -1, the bandwidth is not limited.

Linux:

**curl http://169.254.169.254/openstack/latest/network\_data.json**

Windows:

**Invoke-RestMethod http://169.254.169.254/openstack/latest/network\_data.json | ConvertTo-Json**

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

## Security Key (OpenStack Metadata API)

This API is used to obtain 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-Tokensecuritytoken**} in the message header. **securitytoken** is the value returned when a call is made to the API.
- URI  
/openstack/latest/securitykey
- Usage method

Supports GET requests.

- Examples

Linux:

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

Windows:

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

## Instance Action (OpenStack Metadata API)

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

### NOTE

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

- URI

/openstack/latest/spot/instance-action

- Usage method

Supports GET requests.

- Example

Linux:

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

Windows:

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

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

## User Data (EC2 Compatible API)

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

- URI

/169.254.169.254/latest/user-data

- Usage method

Supports GET requests.

- Example

Linux:

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

Windows:

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

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkI0IGZlZWxzIGFuIGltcHVsc2lubi4uLnRoXMGaXMgdGhllHBsYWNUIHHRvIGdviG5vdy4gQnV0IHROZSBza3kga25vd3MgdGhllHJlYXNvbnMgYW5kIHROZSBwYXR0ZXJuYyBiZWphbmQgYWxsIGNsbnVkb3VkcWwYw5kiHlvdSB3aWxsIGtub3csiHRvbywgd2hIbiB5b3UgbGlmCB5b3Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA==
```

## Hostname (EC2 Compatible API)

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

- URI  
/169.254.169.254/latest/meta-data/hostname
- Usage method  
Supports GET requests.
- Example  
Linux:  
**curl http://169.254.169.254/latest/meta-data/hostname**  
Windows:  
**Invoke-RestMethod http://169.254.169.254/latest/meta-data/hostname**  
vm-test.novalocal

## Instance Type (EC2 Compatible API)

This API is used to query an ECS flavor.

- URI  
/169.254.169.254/latest/meta-data/instance-type
- Usage method  
Supports GET requests.
- Example  
Linux:  
**curl http://169.254.169.254/latest/meta-data/instance-type**  
Windows:  
**Invoke-RestMethod http://169.254.169.254/latest/meta-data/instance-type**  
s3.medium.2

## Local IPv4 (EC2 Compatible API)

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

- URI  
/169.254.169.254/latest/meta-data/local-ipv4
- Usage method  
Supports GET requests.
- Example  
Linux:  
**curl http://169.254.169.254/latest/meta-data/local-ipv4**  
Windows:  
**Invoke-RestMethod http://169.254.169.254/latest/meta-data/local-ipv4**  
192.1.1.2

## Availability Zone (EC2 Compatible API)

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

- URI  
/169.254.169.254/latest/meta-data/placement/availability-zone
- Usage method  
Supports GET requests.
- Example  
Linux:  
**curl http://169.254.169.254/latest/meta-data/placement/availability-zone**  
Windows:  
**Invoke-RestMethod http://169.254.169.254/latest/meta-data/placement/availability-zone**  
az1.dc1

## Public IPv4 (EC2 Compatible API)

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

- URI  
/169.254.169.254/latest/meta-data/public-ipv4
- Usage method  
Supports GET requests.
- Example  
Linux:  
**curl http://169.254.169.254/latest/meta-data/public-ipv4**  
Windows:  
**Invoke-RestMethod http://169.254.169.254/latest/meta-data/public-ipv4**  
46.1.1.2

## Public Keys (EC2 Compatible API)

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

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

```
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADII5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/hioggSnYu/AD0Yqm8vVO0kWLun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/WRenxlwR00KkcH5Jro763+wYcwKieb4eKRxaQoQvoFgVjLBULXAJH4eKoKTVNtMXAvPP9aMy2SLgsJNtMb9ArfziAibIQynq7UIfLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvbLe8sUteBsJZwL6K4i+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBiklOBfuUENIJUhaBGenerated-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 from a public image, or a private image with Cloud-Init installed.
  - The format of the user data 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 injected.
- Windows
  - The image that is used to create ECSs must have Cloudbase-Init installed.
  - The user data to be specified must be less than or equal to 32 KB.
  - If user data is uploaded as text, the data can contain only ASCII characters. If user data is uploaded using a file, the file can contain any characters and the file size cannot exceed 32 KB.
  - The image that is used to create ECSs must be a public image, a private image created based on a public image, or a private image with Cloudbase-Init installed.
  - DHCP must be enabled on the VPC network, and port 80 must be enabled for the security group in the outbound direction.

## Injecting User Data

1. Create a user data script that 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.

**Figure 2-127** User data injection



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

## User Data Scripts of Linux ECSs

User data scripts (referred to as scripts) of Linux ECSs are based on the open-source Cloud-Init architecture. This architecture uses ECS metadata as the data source for configuring the ECSs. User data scripts are compatible with the open-source Cloud-Init. For details about Cloud-Init, see <http://cloudinit.readthedocs.io/en/latest/topics/format.html>.

- Script execution time: A 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 types: user-data and Cloud-Config data scripts

**Table 2-45** 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.

Item	User-Data Script	Cloud-Config Data Script
Format	The first line must start with <b>#!</b> (for example, <b>#!/bin/bash</b> or <b>#!/usr/bin/env python</b> ), 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 <b>#cloud-config</b> , 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 user data injected into a Linux ECS?
  - a. Log in to the ECS.
  - b. Run the following command to view the user data as user **root**:  
**curl http://169.254.169.254/openstack/latest/user\_data**
- Script usage examples  
The following 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 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 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 2-128** Viewing operating results

```
localhost localhost.localdomain localhost4 localhost4.localdomain4
::1 localhost localhost.localdomain localhost6 localhost6.localdomain6
192.168.1.130 us.archive.ubuntu.com
```

## User Data Scripts of Windows ECSs

User data scripts of Windows ECSs are based on the open-source Cloudbase-Init architecture. This architecture uses ECS metadata as the data source for initializing and configuring the ECSs. User data scripts are compatible with the open-source Cloudbase-Init. For details about Cloudbase-Init, see <https://cloudbase-init.readthedocs.io/en/latest/userdata.html>.

- Script types: batch-processing program and PowerShell scripts

**Table 2-46** Windows ECS script types

Item	Batch-Processing Program Script	PowerShell Script
Format	The script must start with <b>rem cmd</b> and use it as the first line. No space is allowed at the beginning of the first line.	The script must start with <b>#ps1</b> and use it as the first line. No space is allowed at the beginning of the first line.
Constraint	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.	Before Base64 encoding, the size of the script, including the first line, cannot exceed 32 KB.

- How can I view the user data injected into a Windows ECS?
  - Log in to the ECS.
  - In the address bar of a browser, access the following URL and view the user data:

**`http://169.254.169.254/openstack/latest/user_data`**

- Script usage examples

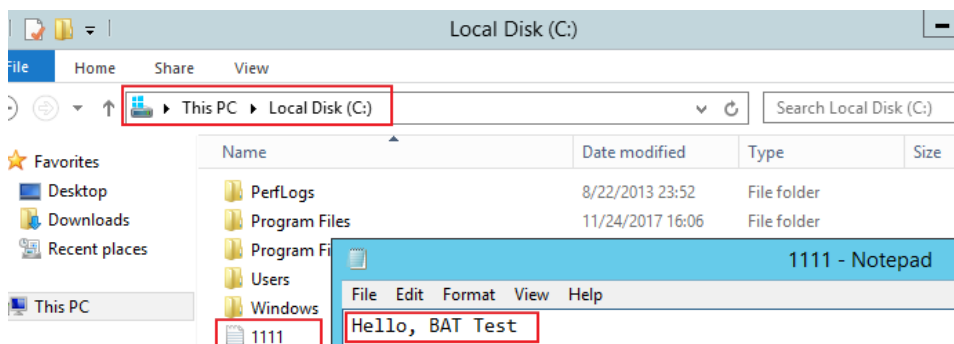
The following describes how to inject scripts in different formats into Windows ECSs and view script execution results.

### Example 1: Inject a batch-processing program script.


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

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

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

**Figure 2-129** Creating text file (Batch)

To view the user data injected into the Windows ECS, log in at [http://169.254.169.254/openstack/latest/user\\_data](http://169.254.169.254/openstack/latest/user_data).

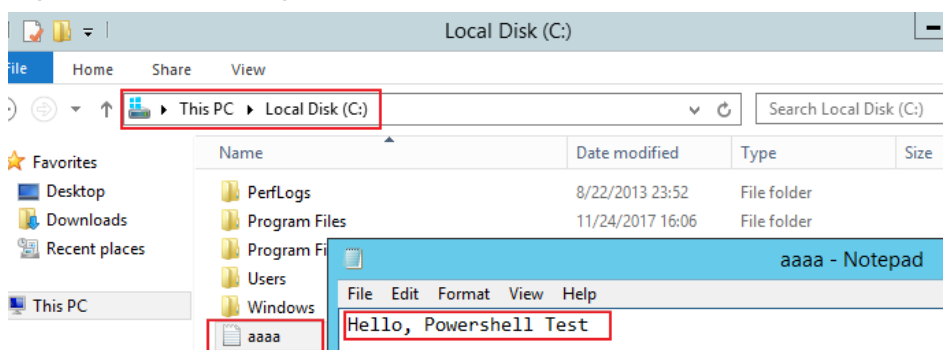
**Figure 2-130** Viewing user data (Batch)

### Example 2: Inject a PowerShell script.

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

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

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

**Figure 2-131** Creating text file (PowerShell)

To view the user data injected into the Windows ECS, log in at [http://169.254.169.254/openstack/latest/user\\_data](http://169.254.169.254/openstack/latest/user_data).

**Figure 2-132** Viewing user data (PowerShell)

## Case 1

This case illustrates how to inject user data to simplify Linux ECS configurations.

To enable syntax highlighting, display line numbers, and set the tab stop to 4 for Vim, create a `.vimrc` configuration file and inject it 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 2-47](#).

**Table 2-47** Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none"><li>Consists of 8 to 26 characters.</li><li>Contains at least three of the following character types:<ul style="list-style-type: none"><li>Uppercase letters</li><li>Lowercase letters</li><li>Digits</li><li>Special characters for Linux: <code>!@\$%^_-=+[{ }],./?~#*</code></li></ul></li><li>Cannot contain the username or the username spelled backwards.</li></ul>

User data example:

Using a ciphertext password (recommended)

```
#!/bin/bash
echo 'root:6V6azyelWcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig' | chpasswd -e;
```

In this command, `$6$V6azyelWcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlig` is the ciphertext password, which can be generated by performing the following steps:

- Run the following command to generate an encrypted ciphertext value:  
**`python -c "import crypt, getpass, pwd;print crypt.mksalt()"`**

The following information is displayed:

```
6V6azyeLwcD3CHlpY
```

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:

```
6V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlq
```

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

#### NOTE

When you specify the **adminPass** field during Linux ECS creation, you can refer to this example to set the password for the ECS through user data injection.

## 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 2-48](#).

**Table 2-48** Password complexity requirements

Parameter	Requirement
Password	<ul style="list-style-type: none"><li>• Consists of 8 to 26 characters.</li><li>• Contains at least three of the following character types:<ul style="list-style-type: none"><li>– Uppercase letters</li><li>– Lowercase letters</li><li>– Digits</li><li>– Special characters for Linux: !@\$%^&amp;_+=+[{ }],./?~#*</li></ul></li><li>• Cannot contain the username or the username spelled backwards.</li></ul>

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

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

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

## Case 4

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

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

### NOTE

The new password must meet the password complexity requirements listed in [Table 2-48](#).

User data example:

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

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

## Case 5

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

User data example:

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

## Case 6

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

User data example:

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

## Helpful Links

For more information about user data injection cases, visit the official Cloud-init or 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 a region and project.
3. Click . Under **Compute**, click **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 .

### Batch Changing ECS Names

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Select the target ECSs.
5. Above the ECS list, choose **More > Change ECS Name**.
6. Enter a new name.
7. Click **OK**.  
If you batch change ECS names, the new ECS names are the same, for example, all are **ecs-test**.

## 2.7.5 Migrating an ECS to a DeH

### Scenarios

ECSs can be migrated:

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

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



 NOTE

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

## Constraints

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

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, locate the ECS to be migrated, choose **More > View O&M and Monitoring > Reallocate ECS** in the **Operation** column.
5. In the displayed dialog box, select the target DeH.

 NOTE

If no DeHs are available, create a DeH first. For details, see [Buying DeHs](#).

6. Click **OK**.

## 2.7.6 Migrating ECSs Across AZs

### Scenarios

Migration Center (MgC) provides a single place for you to easily migrate, modernize, and optimize your applications using tools built based on Huawei Cloud migration methodologies and best practices. With MgC, you can create a migration workflow quickly and easily based on different migration scenarios.

You can use the cross-AZ migration function of MgC to efficiently and visually migrate resources and switch services between AZs. This section describes how to migrate ECSs across AZs on the ECS console.

### Preparations


- Prepare a HUAWEI ID or an IAM user that has required permissions to use MgC. For details, see [Preparations](#).
- On the MgC console, create an independent project for the migration and set **Project Type** to **Application migration**. For details, see [Managing Projects](#).
- On the MgC console, create an application for the migration. Set **Business Scenario** to **Cross-AZ migration-Cross-AZ migration**, **Target Region** to the destination region, and **Target AZ** to the AZ where ECSs will be migrated. For details, see [Creating an Application](#).

## Constraints

- Servers whose system disk is greater than 1 TB cannot be migrated.
- Frozen servers in the retention period cannot be migrated.
- After the migration is complete, the passwords of Linux servers remain unchanged, but those of Windows servers will be changed. For details, see [Are There Any Precautions I Need to Take When Performing a Cross-AZ Migration?](#)

For more constraints about cross-AZ migration, see [Cross-AZ Server Migration](#).

## Procedure

1. Log in to the management console and access the [Elastic Cloud Server](#) page.
2. Click  in the upper left corner and select a region and project.
3. In the **Operation** column of the ECS to be migrated, choose **More > View O&M and Monitoring > Migrate Across AZs**.  
The **Migration Center (MgC)** page is displayed.
4. In the left navigation pane, select the created project for application migration from the drop-down list.
5. Configure cross-AZ migration for ECSs.

The table below lists the steps for configuring cross-AZ migration. For details, see [Migrating Servers Across AZs on Huawei Cloud](#).

**Table 2-49** Configuring cross-AZ migration

Step	Description
Step 1: Discover servers in the source AZ	On the <b>Cloud Discovery</b> page, if you select an application from the <b>Application (Optional)</b> drop-down list in the <b>Basic Information</b> area, all discovered servers will be grouped into the application.  If you only want to migrate a specified ECS, go to the next step. You do not need to specify <b>Application (Optional)</b> here.
Step 2: Group servers as an application	Select the servers to be grouped as an application.
Step 3: Get target recommendations	Assess the servers grouped into an application. If the assessment fails, you can click <b>Modify Target Configuration</b> to manually update the specifications, disk type, and disk storage for a target server.
Step 4: Create a cross-AZ migration workflow	Create a cross-AZ migration workflow and migrate the servers as instructed.

### Helpful Links

- [Where Is MgC Available?](#)
- [What Are the Known Errors Related to Cross-AZ Migration Workflows and How Can I Fix Them?](#)

## 2.7.7 Managing ECS Groups

### Scenarios

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

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.

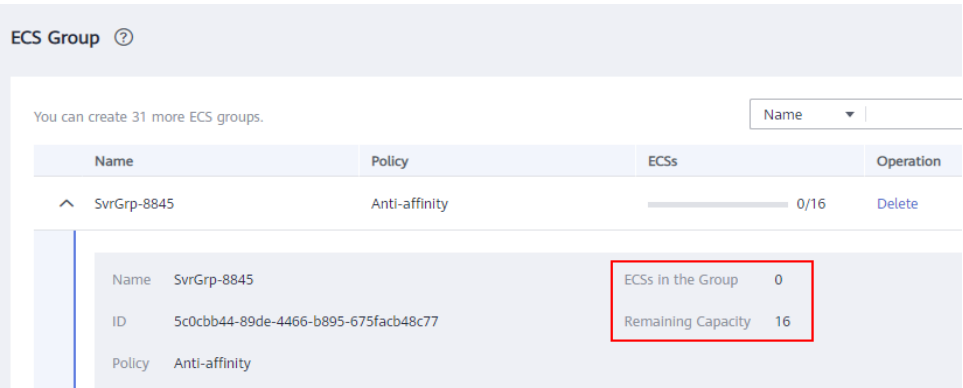
You can perform the following operations on an ECS group:

- [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 different hosts.
- If the maximum number of ECS groups is reached, you need to contact customer service to increase the quota.
- The maximum number of ECSs that can be added to an ECS group varies depending on the region. You can view the quota on the **ECS Group** page, as shown in [Figure 2-133](#).

**Figure 2-133** Maximum number of ECSs that can be added to an ECS group



## Creating an ECS Group

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



1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **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 the ECS group.
7. Select a policy for the ECS group.
8. Click **OK**.

## Adding an ECS to an ECS Group

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

### NOTE

- After an ECS is added to an ECS group, the system reallocates a host to run this ECS to ensure that ECSs in this group run on different hosts. When you attempt to restart the ECS, 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.
- Existing ECSs cannot be added to any ECS group if they have local disks attached (such as disk-intensive or ultra-high I/O ECSs), GPU cards attached (GPU-accelerated ECSs), FPGA cards attached (FPGA-accelerated ECSs), or AI cards attached (AI-accelerated ECSs). They can only be added to an ECS group during the creation process.

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
6. On the **Add ECS** page, select an ECS to be added.
7. (Optional) Stop the ECS.



If an ECS in the **Running** state fails to be added to an ECS group, you can perform this step to stop the ECS.

- a. Click **Stop** in the **Operation** column.
- b. Select a stop option.
  - **Stop**: The ECS will be stopped normally.
  - **Force stop**: This operation will cause loss of unsaved data. Exercise caution when performing this operation.

- c. Click **OK**.
8. Click **OK**. The ECS is added to the ECS group.



## Removing an ECS from an ECS Group

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

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

## 2.7.8 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

## 2.7.9 Starting and Stopping ECSs


You can start, stop, restart, unsubscribe, or delete the ECS.

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

### NOTE

For bare metal ECSs (with physical flavors), do not run commands such as **shutdown**, **poweroff**, or **halt** in the OS because the commands may be invalid or the ECS may fail to be started after being stopped.


### Starting ECSs

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select the target ECSs.
5. Click **Start** in the upper left corner of the list.
6. In the displayed window, click **OK** to start the selected ECSs.

### NOTE

Contact the administrator if the ECS has been in the **Starting** state for more than 30 minutes.

### Stopping ECSs

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select the target ECSs.
5. Click **Stop** in the upper left corner of the list.
6. In the displayed dialog box, select a stop option based on your service requirements.

---

### NOTICE

After an ECS is forcibly stopped, unsaved data on the ECS will be lost.


---

7. Click **OK** to stop the selected ECSs.

### NOTE

Contact the administrator if the ECS has been in the **Stopping** state for more than 30 minutes.

## Restarting ECSs

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select the target ECSs.
5. Click **Restart** in the upper left corner of the list.

---

### NOTICE

After an ECS is forcibly restarted, unsaved data on the ECS will be lost.


---

6. Click **OK** to restart the selected ECSs.

### NOTE

Contact the administrator if the ECS has been in the **Restarting** state for more than 30 minutes.

## Deleting or Unsubscribing from ECSs

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, select **Elastic Cloud Server**.
4. In the ECS list, select the target ECSs.
5. Choose **More > Delete** or **More > Unsubscribe** in the upper left corner of the list.
  - Deleting a pay-per-use ECS
    - i. Select **Delete the EIPs bound to the ECSs** and **Delete all data disks attached to the ECSs** to delete the associated resources.
    - ii. Click **Next** to confirm the deletion.
    - iii. Click **OK** to complete the deletion.
  - Unsubscribing from a yearly/monthly ECS
    - i. Click **OK** to switch to the **Unsubscribe** page of the Billing Center.
    - ii. Select the resources to be unsubscribed from and the reason for unsubscription.
    - iii. Confirm the information and select **After being unsubscribed from, the resource not in the recycle bin will be deleted immediately and cannot be restored. I have backed up data or no longer need the data.**
    - iv. Click **Unsubscribe** and confirm the resources to be unsubscribed from.
    - v. Click **Unsubscribe** again to unsubscribe from the yearly/monthly resources.

 NOTE

Contact the administrator if the ECS has been in the **Deleting** state for more than 30 minutes.

## 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.

- Before changing a Xen ECS to a KVM ECS, you need to manually install the required drivers on the ECS first, or the ECS will be unavailable (such as OS startup failure) after the modification is complete. The following section describes how to change a Xen ECS to a KVM ECS. For Linux, you are advised to use a script to automatically change a Xen ECS to a KVM ECS.
  - [Changing a Xen ECS to a KVM ECS \(Windows\)](#)
  - [Automatically Changing a Xen ECS to a KVM ECS \(Linux\)](#)
  - [Manually Changing a Xen ECS to a KVM ECS \(Linux\)](#)

 NOTE

- ECSs can be classified as the following based on the virtualization types:
  - Xen ECS flavors: S1, C1, C2, and M1.
  - KVM ECS flavors: See the **Virtualization** column in [ECS 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 the specifications of a Windows ECS, modify the SAN policy by following the instructions provided in [Why Does a Disk Attached to a Windows ECS Go Offline?](#) to prevent disks from going offline after the specifications are modified.

- 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.

## Pricing

Modifying specifications will change how much you need to pay for the ECS. For details, see [Pricing of a Changed Specification](#).

## Preparations

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

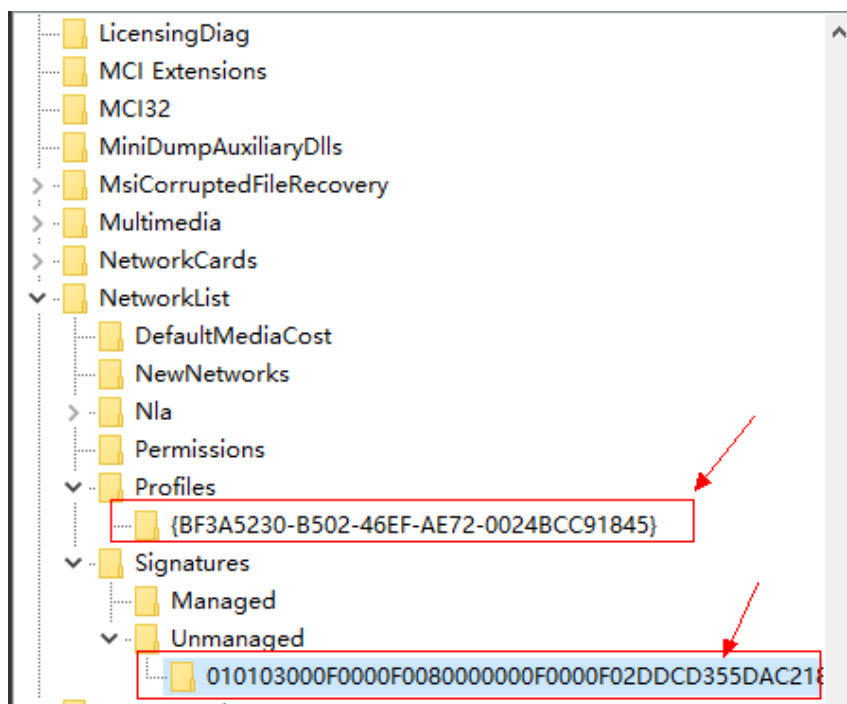
### NOTE

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**
- Windows  
Delete the following directories in the registry on the ECS:  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion  
\NetworkList\Profiles  
HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows NT\CurrentVersion  
\NetworkList\Signatures\Unmanaged

Figure 2-134 Registry



## Step 1: Modify Specifications

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Choose **More > Modify Specifications** in the **Operation** column.  
The **Modify ECS Specifications** page is displayed.
5. Select the new ECS type, vCPUs, and memory.  
Before modifying the specifications, stop the ECS or select **Authorize ECS auto-stop**.

**Figure 2-135** Modifying 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:  ID:

Image: CentOS 7.6 64bit Billing Mode: Pay-per-use

Specifications: General computing | s7n.small.05 | 1 vCPU | 0.5 GiB Region:

vCPUs: --Select vCPUs-- Memory: --Select Memory-- Flavor Name:

General computing-plus **General computing** Memory-optimized General computing-basic

ECS Type	Flavor Name	vCPUs	Memory	CPU	Assured / Maximum Bandwidth	Packets Per Second	IPv6
<input type="radio"/> General computing s7n	s7n.small.05	1 vCPU	0.5 GiB	Intel Ice Lake 2.6G...	0.1 / 0.8 Gbit/s	100,000 PPS	No
<input checked="" type="radio"/> General computing s7n	s7n.small.1	1 vCPU	1 GiB	Intel Ice Lake 2.6G...	0.1 / 0.8 Gbit/s	100,000 PPS	No
<input type="radio"/> General computing s7n	s7n.medium.2	1 vCPU	2 GiB	Intel Ice Lake 2.6G...	0.1 / 0.8 Gbit/s	100,000 PPS	No
<input type="radio"/> General computing s7n	s7n.medium.4	1 vCPU	4 GiB	Intel Ice Lake 2.6G...	0.1 / 0.8 Gbit/s	100,000 PPS	No
<input type="radio"/> General computing s7n	s7n.large.025	2 vCPUs	0.5 GiB	Intel Ice Lake 2.6G...	0.2 / 1.5 Gbit/s	150,000 PPS	No
<input type="radio"/> General computing s7n	s7n.large.2	2 vCPUs	4 GiB	Intel Ice Lake 2.6G...	0.2 / 1.5 Gbit/s	150,000 PPS	No

6. Click **Next**.
7. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.
8. Check whether the specifications have been modified.
  - a. On the console, check whether **Failures** is displayed by referring to [Viewing Failed Tasks](#).
    - If yes, go to step **8.b**.
    - If no, the specifications have been modified.
  - b. Click **Failures**. Then, in the **Failures** dialog box, click **Operation Failures** and check whether the task is contained in the list by **Name/ID**, **Operated At**, or **Task**.
    - If yes, the specifications failed to be modified. See [Follow-up Procedure](#) for failure causes.
    - If no, the specifications have been modified.

## Step 2: Check Disk Attachment

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

- Windows ECS  
For details, see [Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?](#)
- Linux ECS  
For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

## Follow-up Procedure

Perform the following operations if the specifications fail to be modified:

1. Log in to the management console.
2. Under **Management & Governance**, choose **Cloud Trace Service**.
3. In the navigation pane on the left, choose **Trace List**.
4. In the **Trace Name** column, locate the **resizeServer** event by resource ID.  
The resource ID is the ID of the ECS whose specifications failed to be modified.
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.

## 2.8.2 Changing a Xen ECS to a KVM ECS (Windows)

### Scenarios

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

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

#### NOTE

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

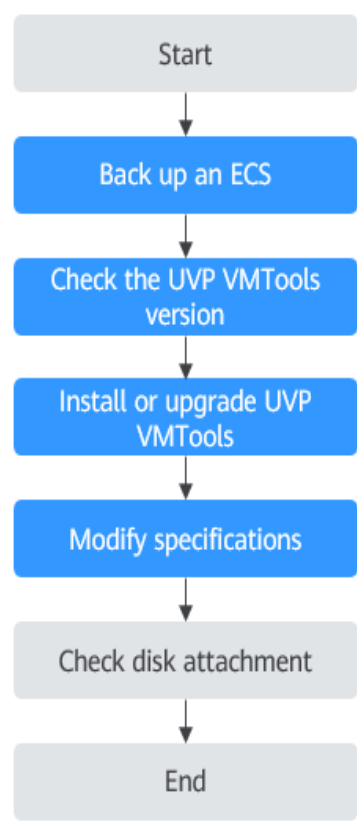
### Constraints

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- If a Windows ECS is attached with a cross-region disk, the ECS specifications cannot be modified. Otherwise, ECS data may be lost.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

### Procedure

[Figure 2-136](#) shows the flowchart for changing a Xen ECS to a KVM ECS.

**Figure 2-136** Flowchart for changing a Xen ECS to a KVM ECS



**Table 2-50** describes the operations for changing a Xen ECS to a KVM ECS.

**Table 2-50** Procedure for changing a Xen ECS to a KVM ECS

Step	Description
1	<b>Step 1: Back Up an ECS</b>
2	<b>Step 2: Check the UVP VMTools Version</b>
3	<b>Step 3: Install or Upgrade UVP VMTools</b>
4	<b>Step 4: Modify Specifications</b>
5	<b>(Optional) Step 5: Check Disk Attachment</b>

**Step 1: Back Up an ECS**

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

1. Check the ECS.

Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.

For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

 NOTE

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Check the UVP VMTools Version

Before modifying specifications, check the UVP VMTools version.

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

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

URL for downloading the script: [https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/windows/server\\_resize/check\\_kvm\\_drivers.vbs](https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/windows/server_resize/check_kvm_drivers.vbs)

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

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

## Step 3: Install or Upgrade UVP VMTools

When you install or upgrade UVP VMTools, if the PV driver has been installed on the ECS, the system will check the PV driver version. Ensure that the PV driver

version meets service requirements. Otherwise, installing UVP VMTools will fail on the ECS. This section describes how to check the installation of the PV driver and UVP VMTools.

---

**⚠ CAUTION**

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

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

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

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

```
pvdriVer=5.0.104.010
```


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

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

**📖 NOTE**

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



- a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
- b. Record the **HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\EnableLUA** value.

**Figure 2-137 EnableLUA**

3. Uninstall the PV driver of the old version.
  - a. On the ECS OS, choose **Start > Control Panel**.
  - b. Click **Uninstall a program**.
  - c. Uninstall **GPL PV Drivers for Windows x.x.x.xx** as prompted.
  - d. Restart the ECS on the management console.
4. Install the PV driver of the new version.
  - a. Download the PV driver installation package.  
Download the PV driver at <https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/pvdriver-windows.zip>.
  - b. Decompress the PV driver software package.
  - c. Double-click **pvdriver-windows.iso**.
  - d. Run **Setup.exe** and install the PV driver as prompted.  
Wait until the driver installation is complete. Do not click **Setup.exe** during the installation.
  - e. Restart the ECS as prompted for the PV driver to take effect.
5. Check and restore the UAC configuration.
  - a. In the **Run** dialog box, enter **regedit** and click **OK** to open the registry editor.
  - b. Check the **HKEY\_LOCAL\_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\System\EnableLUA** value and compare it with the value you recorded. If they are different, change the value to the one recorded in step 2.
6. Install or upgrade UVP VMTools.
  - a. Download the UVP VMTools installation package.  
Download UVP VMTools at <https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/vmtools-windows.zip>.
  - b. Decompress the UVP VMTools installation package.
  - c. Double-click **vmtools-windows.iso**.
  - d. Run **Setup.exe** and install UVP VMTools as prompted.  
The installation program will automatically adapt to the OS version and identify whether UVP VMTools is newly installed or upgraded.  
Wait until the installation is complete. Do not click **Setup.exe** during the installation.
  - e. Restart the ECS as prompted for UVP VMTools to take effect.

- f. Check whether UVP VMTools has been installed. For details, see [Step 2: Check the UVP VMTools Version](#).

## Step 4: Modify Specifications

1. Log in to 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.
7. (Optional) Set **DeH**.  
If the ECS is created on a DeH, you can change the DeH where the ECS resides.  
To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with new specifications.
8. Select the checkbox to confirm that operations in [Step 3: Install or Upgrade UVP VMTools](#) has been performed.
9. Click **Next**.
10. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.

### NOTE

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

## (Optional) Step 5: Check Disk Attachment

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

- Windows ECS  
For details, see [Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?](#)

## Follow-up Procedure

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

 **NOTE**

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

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

## 2.8.3 Automatically Changing a Xen ECS to a KVM ECS (Linux)

### Scenarios

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

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

 **NOTE**

- Xen ECS flavors: S1, C1, C2, and M1.
- KVM ECS flavors: See the **Virtualization** column in [ECS Specifications](#).
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

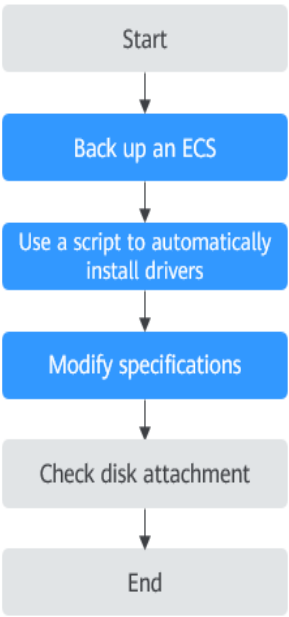
### Constraints

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

### Procedure

[Figure 2-138](#) shows the flowchart for automatically changing a Xen ECS to a KVM ECS.

**Figure 2-138** Flowchart for automatically changing a Xen ECS to a KVM ECS



**Table 2-51** describes the operations for automatically changing a Xen ECS to a KVM ECS using a script.

**Table 2-51** Procedure for automatically changing a Xen ECS to a KVM ECS using a script

Step	Description
1	<b>Step 1: Back Up an ECS</b>
2	<b>Step 2: Use a Script to Automatically Install Drivers</b>
3	<b>Step 3: Modify Specifications</b>
4	<b>(Optional) Step 4: Check Disk Attachment</b>

**Step 1: Back Up an ECS**

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

1. Check the ECS.
- Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

## 2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.

For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

### NOTE

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Use a Script to Automatically Install Drivers

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

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

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

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

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

URL for downloading the script: [https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/linux/server\\_resize/resize\\_ecs\\_modify\\_linux.sh](https://latin-server-resize.obs.na-mexico-1.myhuaweicloud.com/linux/server_resize/resize_ecs_modify_linux.sh)

3. Run the following command to execute the script which automatically checks and installs the native Xen PV driver and virtio driver:

```
bash resize_ecs_modify_linux.sh
```

Figure 2-139 Executing the script

```
susellisp3:/home # bash resize_ecs_modify_linux.sh
2018-08-21 11:04:23 Info:*****BEGIN Modify*****
2018-08-21 11:04:23 Info:get linux system type and version...
2018-08-21 11:04:23 Info:system type: susell
2018-08-21 11:04:23 Info:search grub file...
2018-08-21 11:04:23 Info:find grub file: /boot/grub/menu.lst
2018-08-21 11:04:23 Info:search initrd file list...
2018-08-21 11:04:23 Info:find initrd file: /boot/initrd-3.0.76-0.11-default
2018-08-21 11:04:23 Info:begin to modify grub file...
2018-08-21 11:04:23 Info:modify grub file: /boot/grub/menu.lst
2018-08-21 11:04:23 Info:backup file: /boot/grub/menu.lst
2018-08-21 11:04:23 Info:modify grub file success!
2018-08-21 11:04:23 Info:backup file: /boot/grub/menu.lst
2018-08-21 11:04:23 Info:add xen_platform_pci.dev_unplug=all in /boot/grub/menu.lst
2018-08-21 11:04:23 Info:begin to modify fstab file...
2018-08-21 11:04:23 Info:modify fstab file: /etc/fstab
2018-08-21 11:04:23 Info:backup file: /etc/fstab
2018-08-21 11:04:23 Info:modify fstab file success!
2018-08-21 11:04:23 Info:check xen/ide driver is already exist in initrd** or not
```

4. Wait until the script is executed.

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

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

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

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



Figure 2-140 Successful script execution

```
161.548762] device-mapper: uevent: version 1.0.3
161.551753] device-mapper: ioctl: 4.37.1-ioctl (2018-04-03) initialised: dm-devel@redhat.com
generating grub configuration file ...
found linux image: /boot/vmlinuz-3.10.0-1062.12.1.el7.x86_64
found initrd image: /boot/initramfs-3.10.0-1062.12.1.el7.x86_64.img
found linux image: /boot/vmlinuz-3.10.0-957.el7.x86_64
found initrd image: /boot/initramfs-3.10.0-957.el7.x86_64.img
found linux image: /boot/vmlinuz-0-rescue-8f5b018f6eb344909f6cfec5ad0839ef
found initrd image: /boot/initramfs-0-rescue-8f5b018f6eb344909f6cfec5ad0839ef.img
162.148361] SELinux: SELinux with ACLs, security attributes, no debug enabled
162.189514] xor: automatically using best checksumming function:
162.202066] avx : 22448.000 MB/s
162.233066] raid6: sse2x1 gen() 7382 MB/s
162.277075] raid6: sse2x2 gen() 8589 MB/s
162.297084] raid6: sse2x4 gen() 10273 MB/s
162.318073] raid6: avx2x1 gen() 13410 MB/s
162.337070] raid6: avx2x2 gen() 16503 MB/s
162.356066] raid6: avx2x4 gen() 18976 MB/s
162.358393] raid6: using algorithm avx2x4 gen() (18976 MB/s)
162.361600] raid6: using avx2x2 recovery algorithm
162.431572] Btrfs loaded, crc32c=crc32c-intel
162.446525] fuse init (API version 7.23)
done
2020-09-24 15:12:13 Info:check xen/ide driver is already exist in /boot/initramfs-0-rescue-8f5b018f6eb344909f6cfec5ad0839ef.img or not
2020-09-24 15:12:23 Info:xen driver:yes
2020-09-24 15:12:23 Info:ide driver:no
2020-09-24 15:12:23 Info:check virtio driver is already exist in /boot/initramfs-0-rescue-8f5b018f6eb344909f6cfec5ad0839ef.img or not
2020-09-24 15:12:41 Info:virtio driver:yes
2020-09-24 15:12:41 Info:check xen/ide driver is already exist in /boot/initramfs-3.10.0-1062.12.1.el7.x86_64.img or not
2020-09-24 15:12:45 Info:xen driver:yes
2020-09-24 15:12:45 Info:ide driver:no
2020-09-24 15:12:45 Info:check virtio driver is already exist in /boot/initramfs-3.10.0-1062.12.1.el7.x86_64.img or not
2020-09-24 15:12:52 Info:virtio driver:yes
2020-09-24 15:12:52 Info:check xen/ide driver is already exist in /boot/initramfs-3.10.0-957.el7.x86_64.img or not
2020-09-24 15:12:55 Info:xen driver:yes
2020-09-24 15:12:55 Info:ide driver:no
2020-09-24 15:12:55 Info:check virtio driver is already exist in /boot/initramfs-3.10.0-957.el7.x86_64.img or not
2020-09-24 15:13:03 Info:virtio driver:yes
2020-09-24 15:13:03 Info:centos7 already contain xen and virtio driver
% Total % Received % Xferd Average Speed Time Time Time Current
 % 0 0 0 0 0 0 0 0 100 30 0 0 0 0 0 0 0 0 0
2020-09-24 15:13:03 Info:Success to set kvm meta!
```

 NOTE

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

### Step 3: Modify Specifications

1. Log in to the management console.
2. Click  in the upper left corner and select a 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.
7. (Optional) Set **DeH**.  
If the ECS is created on a DeH, you can change the DeH where the ECS resides.  
To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with new specifications.
8. Select the checkbox to confirm that the configuration is complete.
9. Click **Next**.
10. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.

 NOTE

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

### (Optional) Step 4: Check Disk Attachment

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

- Linux ECS  
For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

## Follow-up Procedure

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

### NOTE

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

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

## 2.8.4 Manually Changing a Xen ECS to a KVM ECS (Linux)

### Scenarios

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

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

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

### NOTE

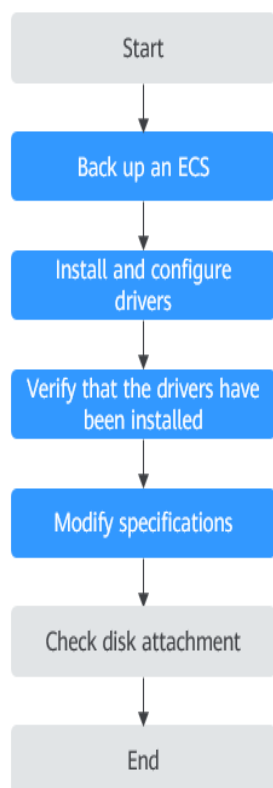
- Xen ECS flavors: S1, C1, C2, and M1.
- KVM ECS flavors: See the **Virtualization** column in [ECS Specifications](#).
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

### Constraints

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- To prevent data loss, the specifications of Linux ECSs that use LVM or RAID arrays cannot be modified.
- A Xen ECS with more than 24 VBD disks attached cannot be changed to a KVM ECS.
- A Xen ECS can be changed to a KVM ECS, but a KVM ECS cannot be changed to a Xen ECS.

### Procedure

[Figure 2-141](#) shows the flowchart for manually changing a Xen ECS to a KVM ECS.

**Figure 2-141** Flowchart for manually changing a Xen ECS to a KVM ECS**Table 2-52** Procedure for manually changing a Xen ECS to a KVM ECS

Step	Description
1	<b>Step 1: Back Up an ECS</b>
2	<b>Step 2: Install and Configure Drivers</b>
3	<b>Step 3: Verify that the Drivers Have Been Installed</b>
4	<b>Step 4: Modify Specifications</b>
5	<b>(Optional) Step 5: Check Disk Attachment</b>

## Step 1: Back Up an ECS

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

1. Check the ECS.

Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.

For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

 **NOTE**

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Install and Configure Drivers

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

1. Log in to the ECS.

2. Uninstall tools from the ECS.

For details, see [Uninstalling PV Drivers](#).

3. Change the GRUB disk ID to UUID.

For details, see [Changing the Disk Identifier in the GRUB Configuration File to UUID](#).

4. Change the fstab disk ID to UUID.

For details, see [Changing the Disk Identifier in the fstab File to UUID](#).

5. Install native Xen and KVM drivers.

For details, see [Installing Native Xen and KVM Drivers](#).

## Step 3: Verify that the Drivers Have Been Installed

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

 **NOTE**

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

1. Log in to the ECS.

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

**cat /boot/grub/grub.cfg**

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

```
.....menuentry 'Ubuntu Linux, with Linux 3.13.0-24-generic' --class ubuntu --class gnu-linux --class
gnu --class os --unrestricted $menuentry_id_option 'gnulinux-3.13.0-24-generic-advanced-
ec51d860-34bf-4374-ad46-a0c3e337fd34' {
 recordfail
 load_video
 gfxmode $linux_gfx_mode
 insmod gzio
 insmod part_msdos
 insmod ext2
 if [x$feature_platform_search_hint = xy]; then
 search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34
 else
 search --no-floppy --fs-uuid --set=root ec51d860-34bf-4374-ad46-a0c3e337fd34
 fi
 echo 'Loading Linux 3.13.0-24-generic ...'
 linux /boot/vmlinuz-3.13.0-24-generic root=UUID=ec51d860-34bf-4374-ad46-a0c3e337fd34 ro
 echo 'Loading initial ramdisk ...'
 initrd /boot/initrd.img-3.13.0-24-generic
}
```

 **NOTE**

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

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

**cat /etc/fstab**

- If yes, the disk ID has been changed to UUID.
- If no, the modification failed. In such a case, change the fstab disk ID to UUID again by referring to [Step 2: Install and Configure Drivers](#).

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

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

- If the boot virtual file system is initramfs, run the following commands:  
**lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r` | grep xen**  
**lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r` | grep virtio**
- If the boot virtual file system is initrd, run the following commands:  
**lsinitrd /boot/initrd-`uname -r` | grep `uname -r` | grep xen**  
**lsinitrd /boot/initrd-`uname -r` | grep `uname -r` | grep virtio**

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

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



```
[root@CTU10000xxxxx home]# lsinitrd /boot/initramfs-`uname -r`.img | grep `uname -r` | grep virtio
-rwxr--r-- 1 root root 23448 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
```

```
block/virtio_blk.ko
-rwxr--r-- 1 root root 50704 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/
drivers/net/virtio_net.ko
-rwxr--r-- 1 root root 28424 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
scsi/virtio_scsi.ko
drwxr-xr-x 2 root root 0 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio
-rwxr--r-- 1 root root 14544 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio.ko
-rwxr--r-- 1 root root 21040 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio_pci.ko
-rwxr--r-- 1 root root 18016 Jul 16 17:53 lib/modules/2.6.32-573.8.1.el6.x86_64/kernel/drivers/
virtio/virtio_ring.ko
```

 **NOTE**

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

## Step 4: Modify Specifications

1. Log in to the management console.
2. Click  in the upper left corner and select a 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.
7. (Optional) Set **DeH**.  
If the ECS is created on a DeH, you can change the DeH where the ECS resides.  
To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that DeH resources are insufficient and cannot be used to create the ECS with new specifications.
8. Select the checkbox to confirm that the configuration is complete.
9. Click **Next**.
10. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.

 **NOTE**

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

## (Optional) Step 5: Check Disk Attachment

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

- Linux ECS

For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

## Follow-up Procedure

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

### NOTE

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

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

## 2.8.5 Automatically Changing Xen ECSs to KVM ECSs in a Batch (Linux)

### Scenarios

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

If a large number of Xen ECSs running Linux need to be changed to KVM ECSs, follow the instructions provided in this section to automatically install desired drivers and configure automatic disk attachment for these ECSs in a batch using a script. After installing the drivers, use the **Modify Specifications** function available on the management console to change the Xen ECSs to KVM ECSs.

### NOTE

- Xen ECS flavors: S1, C1, C2, and M1.
- KVM ECS flavors: See the **Virtualization** column in [ECS Specifications](#).
- To support both Xen and KVM, Linux ECSs require the xen-pv and virtio drivers. Before changing a Xen ECS to a KVM ECS, make sure that the Linux ECS has been configured, including driver installation and automatic disk attachment.

### Constraints

- Batch driver installation only applies to Linux ECSs.
- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- The executor ECS must run CentOS 7, have an EIP bound, and be able to communicate with the ECSs where drivers are to be batch installed.

 **NOTE**

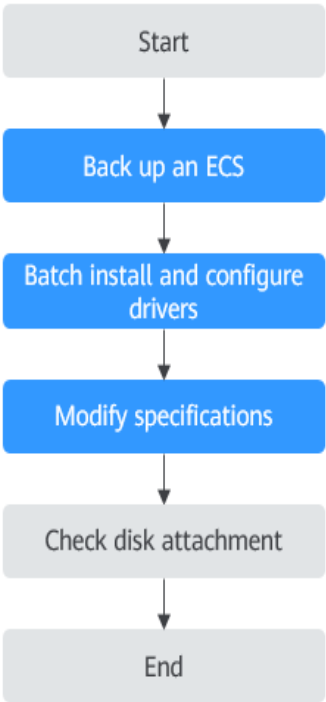
If the internal yum repository has been configured, the executor ECS does not require an EIP.

- You have obtained the IP addresses of the ECSs where drivers are to be batch installed, and the password of user **root** or the private key file for login.
- Only ECSs that use the same key pair support batch driver installation and configuration.

**Procedure**

**Figure 2-142** shows the flowchart for automatically changing Xen ECSs to KVM ECSs in a batch.

**Figure 2-142** Flowchart for automatically changing Xen ECSs to KVM ECSs in a batch



**Table 2-53** Process of automatically changing Xen ECSs to KVM ECSs in a batch

Step	Description
1	Step 1: Back Up an ECS
2	Step 2: Batch Install and Configure Drivers
3	Step 3: Modify Specifications
4	(Optional) Step 4: Check Disk Attachment

## Step 1: Back Up an ECS

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

1. Check the ECS.

Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.

For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

### NOTE

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Batch Install and Configure Drivers

Perform the operations described in this section if your ECSs support the batch configuration using a script.

If the ECSs cannot be automatically configured using a script, follow the instructions provided in [Manually Changing a Xen ECS to a KVM ECS \(Linux\)](#).

1. Log in to the executor ECS. The ECS must meet requirements in [Constraints](#).
2. Run the following command to install the dependency required to run a batch script:

```
yum install ansible -y
```

3. Run the following command to download the driver installation script to the **root** directory:

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

*URL* is the address for downloading the driver installation script.

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

- CN East-Shanghai 2: [https://cn-east-2-server-resize.obs.cn-east-2.myhuaweicloud.com/linux/server\\_resize/resize\\_ecs\\_modify\\_linux.sh](https://cn-east-2-server-resize.obs.cn-east-2.myhuaweicloud.com/linux/server_resize/resize_ecs_modify_linux.sh)
  - CN North-Beijing 1: [https://cn-north-1-server-resize.obs.cn-north-1.myhuaweicloud.com/linux/server\\_resize/resize\\_ecs\\_modify\\_linux.sh](https://cn-north-1-server-resize.obs.cn-north-1.myhuaweicloud.com/linux/server_resize/resize_ecs_modify_linux.sh)
  - CN South-Guangzhou: [https://cn-south-1-server-resize.obs.cn-south-1.myhuaweicloud.com/linux/server\\_resize/resize\\_ecs\\_modify\\_linux.sh](https://cn-south-1-server-resize.obs.cn-south-1.myhuaweicloud.com/linux/server_resize/resize_ecs_modify_linux.sh)
4. Run the following command to download the batch execution script to the **root** directory:
- ```
curl URL > ~/batch_resize_ecs_modify_linux.py
```
- URL* is the address for downloading the batch execution script.
- Select an address for downloading the script based on the region where the ECS is located:
- CN East-Shanghai 2: https://cn-east-2-server-resize.obs.cn-east-2.myhuaweicloud.com/linux/server_resize/batch_resize_ecs_modify_linux.py
 - CN North-Beijing 1: https://cn-north-1-server-resize.obs.cn-north-1.myhuaweicloud.com/linux/server_resize/batch_resize_ecs_modify_linux.py
 - CN South-Guangzhou: https://cn-south-1-server-resize.obs.cn-south-1.myhuaweicloud.com/linux/server_resize/batch_resize_ecs_modify_linux.py
5. Create **host_list.txt** and press **i** to enter editing mode.
- ```
vi host_list.txt
```

---

**CAUTION**

Place the driver installation script, batch execution script, and **host\_list.txt** file in the same directory.

6. Enter the information of the target ECSs in the **host\_list.txt** file.
- The information entered must match that of the target ECSs.
- If the target ECSs use a key pair for authentication, enter the following information:

---

**CAUTION**

- Upload the private key file saved during ECS creation to the folder in which **host\_list.txt** is stored.
- Ensure that the permission code of the private key file is 400.  
**chmod 400 Private key file**

---

Enter an ECS IP address in each line.

An example is provided as follows:

```
192.168.1.10
192.168.1.11
```

- If the target ECSs use a password for authentication, enter the following information:

Enter an ECS IP address and password of user **root** separated using a comma (,) in each line.

An example is provided as follows:

```
192.168.1.10,*****
192.168.1.11,*****
```

Press **Esc** to exit Insert mode and enter **:wq** to save the settings and exit.

7. Run the batch execution script **batch\_resize\_ecs\_modify\_linux.py** to automatically check the configured ECSs and install the native xen-pv and virtio drivers on them in a batch.
  - a. Install drivers on the configured ECSs.
    - ECSs using key pair authentication  
Since the private key file and the batch execution script are in the same directory, you only need to specify the private key file name.  
**python batch\_resize\_ecs\_modify\_linux.py {Name of the private key file or path in which the private key file is stored}**

**Figure 2-143** Executing the script

```
[root@allinone-centos ~]# python batch_resize_ecs_modify_linux.py id_rsa
2020-10-22 10:51:56 Start copying the scripts to all hosts.
2020-10-22 10:52:03 Start executing scripts on all hosts, it will take a while..
```

- ECSs using password authentication  
**python batch\_resize\_ecs\_modify\_linux.py**

**Figure 2-144** Executing the script

```
[root@allinone-centos ~]# python batch_resize_ecs_modify_linux.py
2020-10-22 10:53:44 Start copying the scripts to all hosts.
2020-10-22 10:53:52 Start executing scripts on all hosts, it will take a while..
```

8. Check execution results.

Wait until the script is executed. After checking that the drivers have been installed, the system automatically tags the ECSs. The specifications of only the tagged ECSs can be modified.

**Figure 2-145** shows the execution results. For details about the results, see **Table 2-54**.

**Figure 2-145** Successful script execution

```
2020-10-22 10:56:04 Please check the execution result.
status ip msg
[SUCCESS] 172.28.0.6 Success to set kvm meta
[SUCCESS] 172.28.0.8 Success to set kvm meta
[SUCCESS] 172.28.0.9 Success to set kvm meta
Total: 3 Success: 3 Failed: 0
You can check the logs/exec_origin.log for details.
```

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



**Table 2-54** Execution results

Command output	Description
Total: { <i>Quantity</i> } Success: { <i>Quantity</i> } Failed: 0	All ECSs have been checked, and the drivers have been installed on them.
Total: { <i>Quantity</i> } Success: { <i>Quantity</i> } Failed: { <i>Quantity</i> }	The value of <b>Failed</b> is not <b>0</b> , indicating that the driver check or installation failed on certain ECSs. In this case, view <b>logs/exec_origin.log</b> and identify the failure cause.
Please check the format of host_list.txt	The <b>host_list.txt</b> file does not meet the requirements. Follow the instructions provided in <b>6</b> to modify the file.
[Error] resize_ecs_modify_linux.sh not found.	The driver installation script described in <b>6</b> has not been downloaded. Download the script.
Host(s) is unreachable, please check the network or password of user root.	The target ECSs are unreachable from the executor. Check whether the network is connected or whether the password of user <b>root</b> configured in the <b>host_list.txt</b> file is correct.
[Error] host_list.txt not found.	The <b>host_list.txt</b> file is unavailable. Follow the instructions provided in <b>6</b> to create the file.
[Error] key-file { <i>Private key file</i> } not found.	The specified private key file is unavailable. Ensure that the private key file is available and perform <b>6</b> again.  Alternatively, use the password to log in to the target ECS and perform <b>6</b> again.
Please confirm that the ansible has been installed.	The dependency required in <b>2</b> has not been installed. Install the dependency.

 NOTE

- After the driver is installed, safely store the login password configured in the `host_list.txt` file and the private key file.
- Make sure that the ECSs have been configured successfully. Otherwise, the ECSs will be unavailable after the modification is performed. If the operation failed, follow the instructions provided in [Manually Changing a Xen ECS to a KVM ECS \(Linux\)](#).
- FAQs related to a script installation failure:
  - [What Should I Do If Executing a Driver Installation Script Failed on an ECS Running CentOS 5?](#)
  - [What Should I Do If Executing a Driver Installation Script Failed When I Attempted to Modify the Specifications of a Linux ECS?](#)

### Step 3: Modify Specifications

- **Using the management console**
  - a. Log in to the management console.
  - b. Click  in the upper left corner and select your region and project.
  - c. Click . Under **Compute**, click **Elastic Cloud Server**.
  - d. On the **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.
  - e. Choose **More > Modify Specifications** in the **Operation** column.  
The **Modify ECS Specifications** page is displayed.
  - f. Select the new ECS type, vCPUs, and memory.
  - g. (Optional) Set **DeH**.  
If the ECS is created on a DeH, the system allows you to change the DeH.  
To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that remaining DeH resources are insufficient and cannot be used to create the ECS with new specifications.
  - h. Select the check box to confirm that [Step 2: Batch Install and Configure Drivers](#) has been performed.
  - i. Click **Next**.
  - j. Confirm the new specifications, read and agree to the disclaimer, and then click **Submit Application**.

 NOTE

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

## (Optional) Step 4: Check Disk Attachment

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

- Linux ECS

For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

## Follow-up Procedure

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

### NOTE

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

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

## 2.8.6 Changing a KVM ECS to a QingTian ECS (Windows)

### Scenarios

This section describes how to change a KVM ECS that runs Windows to a QingTian ECS.

### NOTE

- KVM ECS flavors: See the **Virtualization** column in [A Summary List of x86 ECS Specifications](#).
- QingTian ECS flavors: See the **Virtualization** column in [A Summary List of x86 ECS Specifications](#).

### Constraints

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.
- The network type cannot be changed during the specifications modification.
- QingTian ECSs can only have SCSI disks attached, and the disks will use WWN identifiers.
- A KVM ECS can be changed to a QingTian ECS, but a QingTian ECS cannot be changed to a KVM ECS.
- A Xen ECS cannot be changed to a QingTian ECS.

### Procedure

[Figure 2-146](#) shows the flowchart for changing a KVM ECS to a QingTian ECS.

**Figure 2-146** Flowchart for changing a KVM ECS to a QingTian ECS**Table 2-55** Procedure for changing a KVM ECS to a QingTian ECS

Step	Description
1	<a href="#">Step 1: Back Up an ECS</a>
2	<a href="#">Step 2: Check the SCSI Driver</a>
3	<a href="#">Step 3: Modify Specifications</a>
4	<a href="#">(Optional) Step 4: Check Disk Attachment</a>

## Step 1: Back Up an ECS

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

### 1. Check the ECS.

Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

### 2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.  
For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

#### NOTE

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

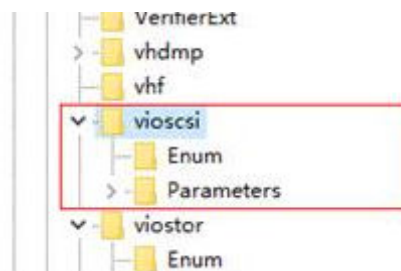
- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Check the SCSI Driver


1. Log in to the ECS.
2. In the **Run** dialog box, enter **regedit** to access the registry editor.
3. In the registry, check whether **StartOverride** exists in the following directory:  
Computer\HKEY\_LOCAL\_MACHINE\SYSTEM\ControlSet001\Services\vioscsi
  - If **StartOverride** does not exist and there are only **Enum** and **Parameters** in the directory, no changes need to be made.

Figure 2-147 Registry



- If **StartOverride** exists in the directory, go to step 3.
4. In **StartOverride**, double-click **REG\_DWORD** and change the value **3** to **0**.
  5. Stop the ECS.

## Step 3: Modify Specifications

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. 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.
4. Choose **More** > **Modify Specifications** in the **Operation** column.  
The **Modify ECS Specifications** page is displayed.
5. Select the new ECS type, vCPUs, and memory.
6. (Optional) Set **DeH**.  
If the ECS is created on a DeH, the system allows you to change the DeH.

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

7. Click **Next**.
8. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.

#### NOTE

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

## (Optional) Step 4: Check Disk Attachment

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

- Windows ECS

For details, see [Why Do the Disks of a Windows ECS Go Offline After I Modify the ECS Specifications?](#)

## Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, contact or reinstall the ECS OS to resolve this issue. For details, see [Reinstalling the OS](#).

#### NOTE

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

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

## 2.8.7 Changing a KVM ECS to a QingTian ECS (Linux)

### Scenarios

This section describes how to change a KVM ECS that runs Linux to a QingTian ECS.

#### NOTE

- KVM ECS flavors: See the **Virtualization** column in [ECS Specifications](#).

### Constraints

- The ECS needs to be stopped during the specification modification, so you are advised to perform this operation during off-peak hours.

- The network type cannot be changed during the specifications modification.
- If a Linux ECS uses LVM logical disks consisting of physical volumes or uses RAID arrays as the system or data disk, the ECS specifications cannot be modified to prevent data loss.
- QingTian ECSs can only have SCSI disks attached, and the disks will use WWN identifiers.
- A KVM ECS can be changed to a QingTian ECS, but a QingTian ECS cannot be changed to a KVM ECS.
- A Xen ECS cannot be changed to a QingTian ECS.

Procedure

Figure 2-148 shows the flowchart for changing a KVM ECS to a QingTian ECS.

Figure 2-148 Flowchart for changing a Linux ECS from KVM to QingTian




Table 2-56 Procedure for changing a KVM ECS to a QingTian ECS

Step	Description
1	Step 1: Back Up an ECS
2	Step 2: Execute the Specifications Modification Optimization Script
3	Step 3: Modify Specifications
4	(Optional) Step 4: Check Disk Attachment

## Step 1: Back Up an ECS

If you modify the specifications of an ECS without installing the driver, the ECS may become unavailable and the data on the system disk may be lost. You are advised to back up the ECS first to prevent data loss.

1. Check the ECS.

Before backing up the ECS, stop and then start the ECS to ensure that services can run properly after the ECS is started.

2. Back up the ECS.

Backing up ECSs generate storage costs. There are two types of EVS snapshots: standard snapshots and legacy snapshots. Legacy snapshots are in OBT and free of charge. Standard snapshots are billed. Select the one type as required.

- Method 1: Create a backup for the ECS.

For details, see [ECS Backup Procedure](#).

- Method 2: Create a system disk snapshot and a data disk snapshot.

For details about how to create a snapshot, see [Creating an EVS Snapshot](#) in *Elastic Volume Service User Guide*.

### NOTE

Backups and snapshots created for the ECS are used to restore data. If the specifications fail to be modified, you can use the backups or snapshots to restore data.

- To use backups to restore data, see [Restoring from a Cloud Server Backup](#).
- To use snapshots to roll back data, see [Rolling Back Disk Data from a Snapshot](#).

After the specifications are modified, if services are running properly, delete the ECS backups or snapshots from the corresponding service console.

## Step 2: Execute the Specifications Modification Optimization Script

1. Remotely log in to the ECS.
2. Download the script to the **root** directory.

```
curl URL > ~/offload_check_blockdevice.sh
```

The *URL* is the address for downloading the specifications modification optimization script:

Download address:

- CN East-Shanghai1: [https://sdi-resize-check-cn-east-3.obs.cn-east-3.myhuaweicloud.com:443/offload\\_check\\_blockdevice.sh](https://sdi-resize-check-cn-east-3.obs.cn-east-3.myhuaweicloud.com:443/offload_check_blockdevice.sh)
- CN North-Beijing4: [https://sdi-resize-check-cn-north-4.obs.cn-north-4.myhuaweicloud.com:443/offload\\_check\\_blockdevice.sh](https://sdi-resize-check-cn-north-4.obs.cn-north-4.myhuaweicloud.com:443/offload_check_blockdevice.sh)
- CN South-Guangzhou: [https://sdi-resize-check-cn-south-1.obs.cn-south-1.myhuaweicloud.com:443/offload\\_check\\_blockdevice.sh](https://sdi-resize-check-cn-south-1.obs.cn-south-1.myhuaweicloud.com:443/offload_check_blockdevice.sh)
- CN North-Ulanqab1: [https://sdi-resize-check-cn-north-9.obs.cn-north-9.myhuaweicloud.com:443/offload\\_check\\_blockdevice.sh](https://sdi-resize-check-cn-north-9.obs.cn-north-9.myhuaweicloud.com:443/offload_check_blockdevice.sh)
- CN Southwest-Guiyang1: [https://sdi-resize-check-cn-southwest-2.obs.cn-southwest-2.myhuaweicloud.com:443/offload\\_check\\_blockdevice.sh](https://sdi-resize-check-cn-southwest-2.obs.cn-southwest-2.myhuaweicloud.com:443/offload_check_blockdevice.sh)

 **NOTE**

If an abnormal output is returned, check whether the server has an EIP bound. In other regions, the script can be obtained only after the EIP is bound.

3. Execute the specifications modification optimization script. The script automatically checks whether the ECS can be changed to a QingTian ECS.

**bash offload\_check\_blockdevice.sh**

**Figure 2-149** Executing the script


```
root@ecs-10e0 ~]# bash offload_check_blockdevice.sh
2023-06-27 09:36:52 Info: -----
2023-06-27 09:36:52 Info: Begin to check and correct grub and fstab file
2023-06-27 09:36:52 Info: OK fstab file looks fine and does not contain any device names.
2023-06-27 09:36:52 Info: Search grub file...
2023-06-27 09:36:52 Info: Find grub file: /boot/grub2/grub.cfg
2023-06-27 09:36:52 Info: Check grub file: /boot/grub2/grub.cfg
2023-06-27 09:36:52 Info: OK /boot/grub2/grub.cfg file looks fine and does not contain any device names.
2023-06-27 09:36:52 Info: -----
root@ecs-10e0 ~]#
```

Wait until the script is executed. If the message "fstab file looks fine" is displayed, the script is executed successfully, and the KVM ECS can be changed to a QingTian ECS.

 **NOTE**

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

## Step 3: Modify Specifications

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. 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.
4. Choose **More > Modify Specifications** in the **Operation** column.  
The **Modify ECS Specifications** page is displayed.
5. Select the new ECS type, vCPUs, and memory.
6. (Optional) Set **DeH**.  
If the ECS is created on a DeH, the system allows you to change the DeH.  
To do so, select the target DeH from the drop-down list. If no DeH is available in the drop-down list, it indicates that remaining DeH resources are insufficient and cannot be used to create the ECS with new specifications.
7. Select the checkbox to confirm the ECS configuration.
8. Click **Next**.
9. Confirm the new specifications, read and agree to the agreement, and then click **Submit Application**.

 NOTE

- The cloud platform automatically creates a system disk snapshot for you. After the specifications are modified, manually delete the snapshot on the snapshot page if you have verified that services are running properly.
- If ECS specifications failed to be modified and the ECS becomes unavailable, reinstall the OS. This operation will clear the data on the system disk while the data on data disks is retained.

## (Optional) Step 4: Check Disk Attachment

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

- Linux ECS

For details, see [Why Does the Disk Attachment of a Linux ECS Fail After I Modify the ECS Specifications?](#)

## Follow-up Procedure

If the ECS specifications have been modified but the OS cannot be started after remote login, contact or reinstall the ECS OS to resolve this issue. For details, see [Reinstalling the OS](#).

 NOTE

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

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

## 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 in 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 quota must be greater than 0.
- If the target ECS is created using a private image, ensure that the private image is available.
- If the target ECS is billed on a pay-per-use basis, ensure that your account has sufficient balance.
- If the target ECS is billed on a yearly/monthly basis, ensure that the subscribed resources are within the subscription term.
- 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 a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image > Reinstall OS** in the **Operation** column.

Before reinstalling the OS, stop the ECS first 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.

**Figure 2-150** Reinstalling an OS

Reinstall OS

Note the following points before you reinstall the OS:  
1. An OS reinstallation has no effect on data disks, but all data on and all snapshots created for the system disk will be lost. [Back up](#) the data before you continue.  
2. The ECS will be automatically restarted after the OS reinstallation, and custom settings (such as the DNS and hostname) will be reset.  
[Hide](#)

Current Configuration

ECS Name	IP address	Specifications	Image	Syste...
ecs- 	192.168.0.30 (Private IP)	1 vCPU   2 GiB  ...	CentOS 7.5 64bit(64-bit)	40 GiB

☐ Stop the ECS (The ECS must be stopped before its OS can be reinstalled.)

System disk ☐ Encrypted ⓘ

Login Mode 

Key pair Password Set password later

Password 

Enter a password. ⓘ

You can use the original password or enter a new one.

Confirm Password 

Enter the password again. ⓘ

Cancel

OK

6. Click **OK**.
7. In the **Reinstall OS** dialog box, confirm the settings, read and agree to 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 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.
- For a yearly/monthly ECS, the system disk capacity may be insufficient if you change the image type. You need to detach the system disk, expand the disk capacity, attach the expanded disk, and then change the OS.
- The OS of a yearly/monthly ECS can be changed:
  - Only changes between free OSs are supported.
  - If an ECS is created from a private image in KooGallery and is billed on a yearly/monthly basis, the OS cannot be changed.
  - OS change between Windows and Linux is only supported in the Chinese mainland regions.
- The EVS disk quota must be greater than 0.

- The system disk type cannot be changed.
- H2 ECSs do not support OS change.
- For details about the change between different OSs, see [Notes on Change Between Windows and Linux](#).
- An ISO image created from an ISO file cannot be used to change the OS of an ECS. You need to install an OS and drivers on the ECS and use the ECS to create a system disk image first.
- The boot mode (BIOS or UEFI) cannot be changed.
- The OS cannot be changed between an x86 ECS and a Kunpeng ECS.

## Notes

- After the OS is changed, the original OS is not retained, and the original system disk is deleted, including the data in all partitions of the system disk.
- 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 status is **Changing OS**.

## Notes on Change Between Windows and Linux

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

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

### NOTE

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



## Billing Rules

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

## Prerequisites

- The 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 a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image > 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.

### NOTE

For yearly/monthly ECSs, if the system disk capacity of the current OS is less than that of the new OS, you need to detach the system disk, expand its capacity, and attach the system disk to the ECS again. Then, you can change the OS.

For instructions about how to expand the system disk capacity, see [Disk Capacity Expansion](#).

Figure 2-151 Changing an OS

**Change OS**

Note the following points before you change the OS:

1. All the data on the system disk, and any snapshots, will be lost. [Back up the data before you continue.](#)
2. Not all OSs support SCSI disks. If the new OS does not support SCSI disks, any SCSI disks attached to the ECS will become unavailable. [Learn more](#)
3. The ECS will be automatically restarted after the OS change. Any custom settings (such as the DNS or hostname) will be reset to their default settings.

[Hide](#)

**Current Configuration**

ECS Name	IP address	Specifications	Image	System Disk
ecs-444	192.168.0.30 (Private IP)	1 vCPU   2 GiB ...	CentOS 7.5 64bit (64-bit)	40 GiB

☐ Stop the ECS (The ECS must be stopped before its OS can be changed.)

**Image**

--Select OS--  --Select OS version--

**Login Mode**

**Password**

Enter a password.

You can use the original password or enter a new one.

**Confirm Password**

Enter the password again.

6. Select 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, read and agree to the agreement or disclaimer, 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 \(Less Than or Equal to 2 TiB\)](#).
  - 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 fails, perform steps [3](#) to [8](#) again to retry the OS change.
- If the attempt still fails, 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. View the ECS status in the **Status** column after purchasing an ECS.

#### NOTE

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

### 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 2-57](#) lists the types of failures that can be recorded in the **Failures** area.

**Table 2-57** 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, <b>Ecs.0013 Insufficient EIP quota</b> .
Operation failures	<ul style="list-style-type: none"><li>Modifying ECS specifications</li></ul> If an ECS specifications modification failed, this operation is recorded in <b>Failures</b> .

## Procedure

- Log in to the management console.
- Click  in the upper left corner and select a region and project.
- Click . Under **Compute**, click **Elastic Cloud Server**.
- View **Failures** on the right side of common operations.

**Figure 2-152** Failures

- 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

- Log in to the management console.

2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **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 contains the following information of an ECS:

- Basic information: includes the ECS name, region, AZ, status, flavor, image, and billing mode.
- Network information: includes the private IP address, EIP, VPC, subnet, and security group.
- Disk information: includes the number of disks, disk properties, and disk capacity.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **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 2-58](#).

 **NOTE**

- For certain properties, if you enter complete values, the system can automatically identify 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 2-58** Search syntax

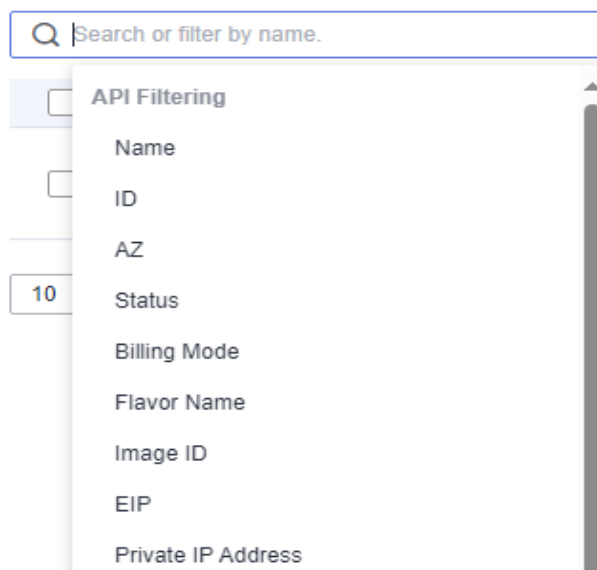
Search Type	Supported Property	Format	Example	Description
Property value Automatic property matching	ID Flavor name EIP Private IP address	Complete property value	ID: 4a79dfec-f0d8-4181-9bef-495b8b7220e1 Flavor Name: s2.xlarge.4 Private IP Address: 192.168.9.231	When you search by keyword, enter a complete 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.

Search Type	Supported Property	Format	Example	Description
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: 192.168.9 9.231	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: 192.168.9 9.231 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.
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.

Search Type	Supported Property	Format	Example	Description
	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 <ul style="list-style-type: none"><li>Multiple properties are in AND relationship.</li><li>Multiple values of the same property are in OR relationship.</li></ul> 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 2-153** 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 2-154** Entering a complete ECS ID**Figure 2-155** Automatic property matching

- Searching by multiple values  
Enter multiple complete flavor names and press **Enter** to search.

**Figure 2-156** Entering multiple complete flavor names


**Figure 2-157** Automatic property matching for search

## 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
  - Select **Private IP Address** in the search box.
  - 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 2-158** Searching by private IP address

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

**Figure 2-159** Searching for an ECS by ID

## 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.

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

**Figure 2-160** Searching by name and private IP address

Name/ID	Monitoring	Security	Status	AZ	Specifications/Image	IP Address	Billing Mode	Enterprise Project
ecs-98c8 1984237e-cdb8-4e3c-8...			Running	AZ2	1 vCPU   1 GiB   s6 small 1 CentOS 7.5 64bit	10.0.0.146 (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	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 2-161** 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 2-162** Searching by private IP address



<input type="checkbox"/>	Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
<input type="checkbox"/>	ecs-976d f964337e-cc8b-4e3c-8...			Running	AZ2	1 vCPU   1 GiB   s6.small.1 CentOS 7.5 64bit	(EIP)... 10.0.0.146 (Private ...
<input type="checkbox"/>	ecs-98c8 2ef2311b-c5b3-4273-b...			Running	AZ2	1 vCPU   2 GiB   s6.medium.2 CentOS 7.5 64bit	(... 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 2-163** Searching by status and flavor name



- Exact search

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

**Figure 2-164** Searching by status and ID

Q Status: Stopped Status: Running ID: f984337e-cc8b-4e3c-8752-3bf13e7d111d,2ef2311b-c5b3-4273-b0b4-a3f527193393 Add filter						
<input type="checkbox"/> Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
<div><input type="checkbox"/> ecs-976d</div> <div>f984337e-cc8b-4e3c-8752-3bf13e7d111d</div>			<div>Stopped</div>	AZ2	1 vCPU   1 GiB   s6_small.1 CentOS 7.5 64bit	10.0.0.146 (Pr...
<div><input type="checkbox"/> ecs-98c8</div> <div>2ef2311b-c5b3-4273-b0b4-a3f527193393</div>			<div>Running</div>	AZ2	1 vCPU   2 GiB   s6_medium.2 CentOS 7.5 64bit	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 2-165** Searching by tag

- 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 2-166** Searching by tag

Q

\_type=Empty value

z

a

b

Add filter

Name/ID

ecs-976d

f984337e-cc8b-4e3c-8...

Monitor

Security

Stopped

AZ2

Specifications/Image

1 vCPU | 1 GiB | s6\_small.1

CentOS 7.5 64bit

IP Address

10.0.0.146 (Private ...)

Billing Mode

Pay-per-use

Created on Oct 18, 202...

Enterprise Project

default

Tag

\_type=Empty value &#39;b

# 3 Recycle Bin (OBT)

---

## 3.1 Recycle Bin Overview

### What Is Recycle Bin?

Recycle bin is a resource recovery feature that enables you to restore ECSs that have been deleted. When using recycle bin, if you delete pay-per-use ECSs or unsubscribe from unexpired yearly/monthly ECSs, they are retained in the recycle bin for a time period that you specify before being permanently deleted. This can help protect your ECSs from accidental deletions.

### Constraints

- If an IAM user wants to use recycle bin, the user must have the following permissions:
  - `ecs:recycleBin:get`: Query the recycle bin.
  - `ecs:recycleBin:update`: Update (enable or disable) the recycle bin.
  - `ecs:recycleBin:updatePolicy`: Update the recycle bin policy.

You can assign these permissions to the IAM user by referring to [ECS Custom Policies](#).

- A deleted or unsubscribed ECS cannot be moved to the recycle bin if:
  - Your account is in arrears, restricted, or frozen.
  - The ECS is faulty.
  - The number of days that has passed since the ECS was created is less than that specified in the recycle bin policy.
  - The ECS is in the retention period.
  - The ECS was deleted after the ECS retention period ends.
  - A scheduled deletion time has been set for the ECS.
  - The ECS is a spot ECS.
- Quota considerations:
  - ECSs in the recycle bin still occupy the ECS resource quotas.

- If ECSs deployed on a DeH are moved to the recycle bin, they still occupy the DeH resource quotas.

 **NOTE**

If ECS quotas are insufficient, delete the ECSs from the recycle bin in a timely manner.

- When ECSs are used as compute resources for gPaaS & AI DaaS services, the ECSs that meet the recycle bin policy will be moved to the recycle bin. This may lead to the failure of clearing resources of gPaaS & AI DaaS services. In this case, you need to delete the ECSs from the recycle bin first.

**Pricing**

- After ECSs are moved to the recycle bin, they will be stopped and billed on a pay-per-use basis based on the billing rules in [Table 3-1](#). (For more details, see [Billing Mode Overview](#).)

**Table 3-1** Billing for a stopped ECS


Instance Type	Description	Billing Item	Billed or Not	Action on Resources
Common instances	Common instances include: <ul style="list-style-type: none"><li>• Non-bare-metal instance</li><li>• Instances without local disks</li><li>• Instances without FPGA cards</li></ul>	ECS (compute resources including vCPUs and memory)	No	<ul style="list-style-type: none"><li>• If ECSs are created from public resource pools, basic resources (vCPUs and memory) are no longer retained. The ECS may fail to be started due to insufficient resources. Wait patiently and try again later.</li><li>• If ECSs are created on a DeH or in an edge AZ, basic resources (vCPUs and memory) will be retained.</li></ul>
		Image	No	Resources are retained but are not billed.
		GPU	No	After GPU-accelerated ECSs without local disks attached are stopped, GPU resources are no longer retained.

Instance Type	Description	Billing Item	Billed or Not	Action on Resources
		EVS disks (system and data disks)	Yes	The resources are still billed based on the billing rules.
		EIP bandwidth	Yes	The resources are still billed based on the billing rules. <ul style="list-style-type: none"><li>EIP bandwidth price: pay-per-use EIPs (by bandwidth)</li><li>Shared bandwidth price</li></ul> For details, see <a href="#">EIP Billing</a> .
Special instances	Special instances include: <ul style="list-style-type: none"><li>Bare metal instances</li><li>Instances with local disks such as disk-intensive, ultra-high I/O ECSs.</li><li>ECSs with FPGA cards</li></ul>	ECS (compute resources including vCPUs and memory)	Yes	The ECS resources are still billed based on the billing rules. To stop the billing, you need to delete the instance and its associated resources.
		Image	Yes	
		GPU	Yes	
		EVS disks (system and data disks)	Yes	
		EIP bandwidth	Yes	

- If your account is in arrears, the ECSs kept in the recycle bin will enter a grace period and then a retention period. These ECSs may not be kept as long as the duration you specify in the recycle bin policy and may be deleted earlier.
- ECSs recovered from the recycle bin are billed on a pay-per-use basis. If you want to change the billing mode to yearly/monthly, see [From Pay-per-Use to Yearly/Monthly](#).

## Recycling Process

[Figure 3-1](#) shows the process of recycling ECSs.

**Figure 3-1** ECS recycling process

## Recycle Bin Policy Configuration Suggestions

When using ECS recycle bin, you need to configure a recycle bin policy. For details, see [Configuring a Recycle Bin Policy](#).

A recycle bin policy includes:

- The minimum number of days that must have passed after an ECS was created (the minimum ECS age) before it can be moved to the recycle bin upon deletion or unsubscription
- A duration that deleted or unsubscribed ECSs can be kept in the recycle bin

### NOTE

This rule only applies to the ECSs that are moved to the recycle bin after the rule is set.

Both ECS and EVS support recycle bins. Their recycle bin policies are configured separately. If you enable recycle bin for both ECS and EVS, you are advised to configure the same minimum number of days for moving ECSs and EVS disks to the recycle bin upon deletion to avoid issues caused by different lifecycles.

In special cases, you may need to configure different recycle bin policies for ECS and EVS. For details, see [Table 3-2](#).

Assume that an ECS (including system and data disks) was created 8 days ago. If you delete or unsubscribe from the ECS, the resources would be handled as follows.

**Table 3-2** Scenarios and rules for recycling resources

Scenario	Recycling Rule	ECS	EVS Disk
The minimum ECS age for moving to the recycle bin is less than that of EVS disks, for example: <ul style="list-style-type: none"><li>• Minimum ECS age for moving to the recycle bin: 7 days</li><li>• Minimum EVS disk age for moving to the recycle bin: 15 days</li></ul>	Both ECSs and EVS disks comply with the ECS recycle bin policy.	Moved to the recycle bin	Moved to the recycle bin
The minimum ECS age for moving to the recycle bin is greater than that of EVS disks, for example: <ul style="list-style-type: none"><li>• Minimum ECS age for moving to the recycle bin: 15 days</li><li>• Minimum EVS disk age for moving to the recycle bin: 7 days</li></ul>	ECSs comply with the ECS recycle bin policy. EVS disks comply with the EVS recycle bin policy.	Deleted or unsubscribed	Moved to the recycle bin

## Associated Resource Recycling

After an ECS is moved to the recycle bin, its associated resources are handled as follows:

- [Table 3-3](#) describes how EIPs bound to ECSs are handled.

**Table 3-3** Handling of bound EIPs

ECS Billing Mode	Operation	Moved to the Recycle Bin	After ECS Is Recovered	After ECS Is Permanently Deleted
Pay-per-use	Select <b>Release the EIPs bound to the ECSs</b> when deleting an ECS.	No	The EIP has been deleted and cannot be recovered.	The EIP has already been deleted.
	Deselect <b>Release the EIPs bound to the ECSs</b> when deleting an ECS.	No	The EIP has been unbound from the ECS and will not be recovered.	The EIP has been unbound from the ECS and will not be deleted.
Yearly/ Monthly	Unsubscribe from the EIP when unsubscribing from an ECS.	No	The EIP has been deleted and cannot be recovered.	The EIP has already been deleted.
	Do not unsubscribe from the EIP when unsubscribing from an ECS.	No	The EIP has been unbound from the ECS and will not be recovered.	The EIP has been unbound from the ECS and will not be deleted.

- **Table 3-4** describes how EVS disks attached to ECSs are handled.

**Table 3-4** Handling of attached EVS disks

ECS Billing Mode	Disk Type	Operation	Moved to the Recycle Bin	After ECS Is Recovered	After ECS Is Permanently Deleted
Pay-per-use	System disk	Select <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	Yes	The disk will also be recovered.	The disk will also be permanently deleted.

EC S Bill ing Mode	Disk Type	Operation	Moved to the Recycle Bin	After ECS Is Recovered	After ECS Is Permanentl y Deleted
		Deselect <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	Yes	The disk will also be recovered .	The disk will also be permanently deleted.
	Exclusi ve data disk	Select <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	Yes	The disk will also be recovered .	The disk will also be permanently deleted.
		Deselect <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	No	The disk will still be attached to the ECS, unless it is manually detached.	The disk has been detached from the ECS and will not be deleted.
	Shared data disk	Select <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	No (when the shared disk is attached to multiple ECSs)	The disk will still be attached to the ECS, unless it is manually detached.	The disk has been detached from the ECS and will not be deleted.
			Yes (when the shared disk is attached to only one ECS)	The disk will also be recovered .	The disk will also be permanently deleted.

EC S Billing Mode	Disk Type	Operation	Moved to the Recycle Bin	After ECS Is Recovered	After ECS Is Permanently Deleted
		Deselect <b>Delete all data disks attached to the ECSs</b> when deleting an ECS.	No	The disk will still be attached to the ECS, unless it is manually detached.	The disk has been detached from the ECS and will not be deleted.
Yearly/ Monthly	System disk	Unsubscribe from an ECS.	Yes	The disk will also be recovered.	The disk will also be permanently deleted.
	Exclusive data disk	Unsubscribe from an ECS with an exclusive data disk attached in the same order.	Yes	The disk will also be recovered.	The disk will also be permanently deleted.
		Unsubscribe from an ECS with an exclusive data disk attached in another order.	No	The disk will still be attached to the ECS, unless it is manually detached.	The disk has been detached from the ECS and will not be deleted.
	Shared data disk	Unsubscribe from an ECS with a shared data disk attached.	No	The disk will still be attached to the ECS, unless it is manually detached.	The disk has been detached from the ECS and will not be deleted.

- [Table 3-5](#) describes how network interfaces attached to ECSs are handled.

**Table 3-5** Handling of attached network interfaces

Network Interface Type	Operation	Moved to the Recycle Bin	After ECS Is Recovered	After ECS Is Permanently Deleted
Primary network interface	Delete or unsubscribe from an ECS.	No	The network interface will still be attached to the ECS.	The network interface will also be permanently deleted.
Extension network interface	Delete or unsubscribe from an ECS.	No	The network interface will still be attached to the ECS, unless it is manually detached.	<ul style="list-style-type: none"><li>• The network interface will not be deleted if it is manually detached from the ECS.</li><li>• The network interface will also be permanently deleted if it is not manually detached from the ECS.</li></ul>

## 3.2 Enabling Recycle Bin

### Scenarios

When using recycle bin, if you delete pay-per-use ECSs or unexpired yearly/monthly ECSs, they are retained in the recycle bin for a time period that you specify before being permanently deleted. This can help protect your ECSs from accidental deletions.

Recycle bin is disabled by default. This section describes how to enable recycle bin.

 NOTE



Recycle bin is in the OBT. If you want to use it, [submit a service ticket](#) to apply for the OBT.

The ECS recycle bin is enabled by IAM project. If multi-project management is used, you need to enable recycle bin for each project.

## Prerequisites

To enable ECS recycle bin, you need to enable EVS recycle bin first. For details, see [Enabling the Recycle Bin](#).

## 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 **Recycle Bin** tab.
5. Click **Enable Recycle Bin**.

Deleted or unsubscribed ECSs are displayed on the **Recycle Bin** tab.

## 3.3 Disabling Recycle Bin

### Scenarios

You can disable recycle bin as needed. After recycle bin is disabled, the deleted or unsubscribed ECSs will be immediately deleted and cannot be recovered.

 NOTE



If you disable the ECS recycle bin, that does not affect the EVS recycle bin.

## Prerequisites

Before disabling ECS recycle bin, you need to recover or permanently delete any ECSs currently in the recycle bin.

- [Recovering an ECS from the Recycle Bin](#)
- [Permanently Deleting ECSs from the Recycle Bin](#)

## 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 **Recycle Bin** tab.
5. In the upper right corner of the **Recycle Bin** tab, click **Disable Recycle Bin**.

A dialog box is displayed, requiring you to empty the recycle bin.

6. Click **OK**.

Recycle bin has been disabled.

## 3.4 Configuring a Recycle Bin Policy

### Scenarios

After Recycle Bin is enabled, you can configure a recycle bin policy that includes:

- The minimum number of days that must have passed after an ECS was created (the minimum ECS age) before it can be moved to the recycle bin upon deletion or unsubscription
- A duration that deleted or unsubscribed ECSs can be kept in the recycle bin



#### NOTE

This rule only applies to the ECSs that are moved to the recycle bin after the rule is set.

### Constraints

- The number of days you can set for the recycle bin policy ranges from 1 to 1,000 days and defaults to 7 days.
- ECSs in the recycle bin are kept for at least 1 hour and no more than 720 hours.

### 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 **Recycle Bin** tab.
5. In the upper right corner of the **Recycle Bin** tab, click **Modify Policy**.  
The **Configure Recycle Bin Policy** panel slides out.
6. Configuring a recycle bin policy.
  - The minimum number of days that must have passed after an ECS was created (the minimum ECS age) before it can be moved to the recycle bin upon deletion or unsubscription. The default value is 7 days.  
For example, if you set the minimum ECS age to 7 days, ECSs created 7 days ago will be moved to the recycle bin upon deletion or unsubscription and be billed on a pay-per-use basis. ECSs created within 7 days will not be moved to the recycle bin. They will be permanently deleted.

 NOTE

If you use both ECS and EVS recycle bins, you are advised to configure the same minimum number of days for moving ECSs and EVS disks to the recycle bins to avoid issues brought by different lifecycles. For details, see [Recycle Bin Policy Configuration Suggestions](#).

- A duration that deleted or unsubscribed ECSs can be kept in the recycle bin. The default value is 168 hours.

If you set the retention duration to 168 hours, ECSs can be kept in the recycle bin for a maximum of 168 hours, after which they will be permanently deleted and cannot be recovered.

7. Click **OK**.

## 3.5 Recovering an ECS from the Recycle Bin



### Scenarios

You can recover ECSs from the recycle bin as needed within the kept duration.

### Constraints

If your account is frozen or restricted, the recycle bin is unavailable. ECSs in the recycle bin cannot be recovered.

### 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 **Recycle Bin** tab.
5. In the recycle bin list, locate the ECS to be recovered and click **Recover** in the **Operation** column.  
The **Recover ECS** page is displayed.
6. Click **Submit**.
  - If the recovery is successful, the ECS is displayed in the ECS list and the ECS status is **Running**.  
To view the details of the ECS associated resources that have been recovered, see [Viewing ECS Details](#).
  - If the recovery fails, the ECS remains in the recycle bin, and the ECS status is **Stopped**.

 NOTE

After an ECS is recovered, its associated resources are processed as described in [Associated Resource Recycling](#).

ECSs recovered from the recycle bin are billed on a pay-per-use basis. If you want to change the billing mode to yearly/monthly, see [From Pay-per-Use to Yearly/Monthly](#).

## 3.6 Permanently Deleting ECSs from the Recycle Bin

### Scenarios

You can permanently delete ECSs from the recycle bin within the kept duration.

#### NOTE



Permanently deleted ECSs cannot be recovered. Exercise caution when performing this operation.

After an ECS is permanently deleted, its associated resources are processed as described in [Associated Resource Recycling](#).

### Constraints

If your account is in arrears, ECSs in the recycle bin will enter a grace period and then a retention period. Based on the grace period and retention period policies, these ECSs may not be kept as long as the duration you specify in the recycle bin policy. They may be deleted earlier.

### 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 **Recycle Bin** tab.
5. Locate the ECS to be deleted and click **Delete** in the **Operation** column.
6. In the displayed dialog box, click **OK**.

If the ECS is no longer displayed in the list, the deletion is successful.



## 3.7 Exporting ECS Information from the Recycle Bin

### Scenarios

The information of all your account's ECSs in the recycle bin can be exported to an XLSX file locally.

The file contains the ECS name, ID, status, AZ, specifications, image, IP address, remaining duration, and DeH name.

### 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. Click the **Recycle Bin** tab.
5. Above the ECS list, click **Export** to export ECS information.
  - Export all data to an XLSX file: The system automatically exports information of all your account's ECSs in the current region from the recycle bin to a local directory.
  - Export selected data to an XLSX file: The system automatically exports information of the selected ECSs from the recycle bin to a local directory.
6. In the default download path, view the exported ECS information.

# 4 Images

---

## 4.1 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 4-1** 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](#)
- [Image Source Management](#)

## 4.2 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 a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, choose **More** > **Manage Image** > **Create Image** in the **Operation** column.

5. Configure the following information:

[Table 4-2](#) and [Table 4-3](#) list the parameters in the **Image Type and Source** and **Image Information** areas, respectively.

**Table 4-2** Image type and source

Parameter	Description
Region	Select a region close to your business. If you select an incorrect region here, you can replicate the created image to your desired region. For details, see <a href="#">Replicating Images Across Regions</a> .
Type	Select <b>Create Image</b> .
Image Type	Select <b>System disk image</b> .
Source	Click the <b>ECS</b> tab and select an ECS with required configurations.

**Table 4-3** 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. <ul style="list-style-type: none"><li>Only an unencrypted private image can be created from an unencrypted ECS.</li><li>Only an encrypted private image can be created from an encrypted ECS.</li></ul>
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.

# 5 Disks


## 5.1 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.

Just like the physical disks in local PC need to be installed before they can be used, EVS disks need to be attached to servers before they can be used. They cannot be used alone. You also need to partition and create file systems on them before they can be used for persistent data storage.

**Figure 5-1** EVS architecture



- A system disk runs the server OS. It is like drive C in a PC.  
When a server is purchased, a system disk is automatically purchased and attached. You cannot purchase a system disk separately. The maximum size of a system disk is 1,024 GiB.

- Data disks store the server data. They are like drive D, drive E, and drive F in a PC.

Data disks can be purchased during or after the server purchase. If you purchase data disks during the server purchase, the system will automatically attach the data disks to the server. If you purchase data disks after the server purchase, you need to manually attach the data disks. The maximum size of a data disk is 32,768 GiB.

## 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 [Disk Types and Performance](#).

## 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?](#)

## Device Types Supported by ECS

Device types supported by ECS are determined by the ECS type and scenarios.

**Table 5-1** Device types supported by ECS

Scenario	Device Type (System Disk)	Device Type (Data Disk)
Purchasing an ECS	<ul style="list-style-type: none"><li>• General computing-plus C7: SCSI</li><li>• Memory-optimized M7: SCSI</li><li>• Bare metal ECS (with the additional identifier "physical"): SCSI</li><li>• Other: VBD</li></ul>	<ul style="list-style-type: none"><li>• General computing-plus C7: SCSI</li><li>• Memory-optimized M7: SCSI</li><li>• Bare metal ECS (with the additional identifier "physical"): SCSI</li><li>• Other: VBD or SCSI</li></ul>

Scenario	Device Type (System Disk)	Device Type (Data Disk)
Attaching a disk to an existing ECS	N/A	<ul style="list-style-type: none"><li>• General computing-plus C7: VBD or SCSI</li><li>• Memory-optimized M7: VBD or SCSI</li><li>• Bare metal ECS (with the additional identifier "physical"): SCSI</li><li>• Other: VBD or SCSI</li></ul>

## Helpful Links

- [Attaching an EVS Disk to an ECS](#)
- [Initializing EVS Data Disks](#)
- [Why Can't I Find My Newly Purchased Data Disk After I Log In to My Windows ECS?](#)
- [How Can I Adjust System Disk Partitions?](#)
- [Can I Attach Multiple Disks to an ECS?](#)
- [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)

## 5.2 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.

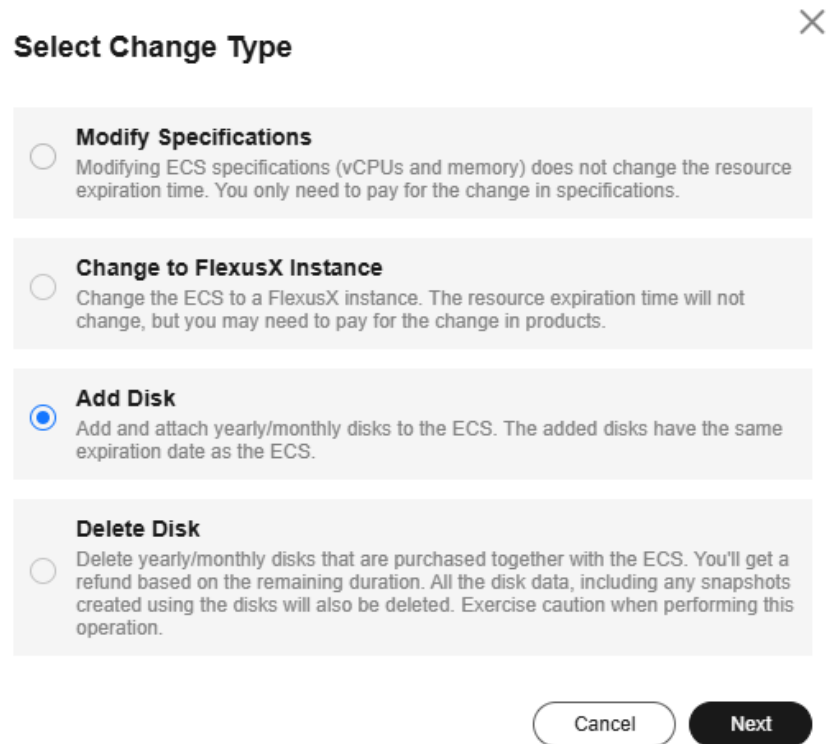
This section describes how to add a data disk after an ECS is purchased.

### 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 do as follows:
  - For a pay-per-use ECS:  
Choose **More** > **Manage Disk/Backup** > **Add Disk** in the **Operation** column.

- For a yearly/monthly ECS:
  - i. Choose **More > Change** in the **Operation** column.
  - ii. In the displayed **Select Change Type** dialog box, select **Add Disk** and click **Next**.

**Figure 5-2** Adding a disk to an ECS



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.
  - The expiration time of a new disk billed on a yearly/monthly basis is the same as that of the ECS.
6. Click **Next** to confirm the order and click **Submit** to complete the payment.  
The system automatically switches back to the **Disks** tab on the ECS management console. Then, you can view the information of the new disk.

## Follow-up Procedure

The system automatically attaches the new disk to the ECS, but the disk can be used only after it is initialized. To do so, log in to the ECS and initialize the disk.

For details about how to initialize a data disk, see [Initializing EVS Data Disks](#).

## 5.3 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.

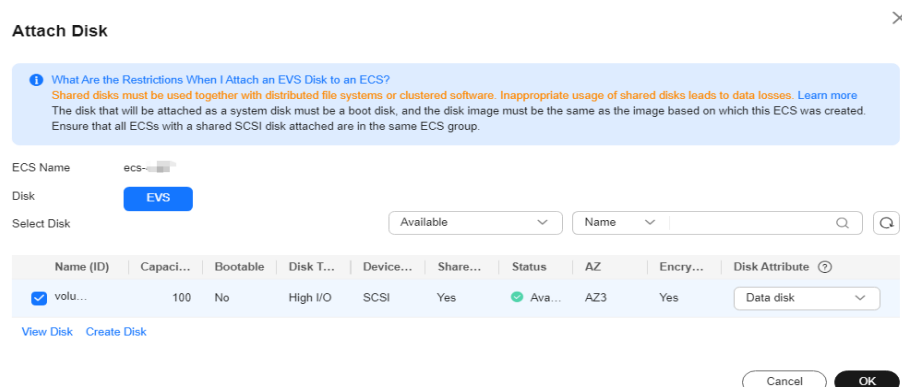
When attaching EVS disks to an existing ECS, their billing modes can be different from the ECS. You can select appropriate billing modes for these EVS disks based on your requirements.

### 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 a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
6. Click the **Disks** tab. Then, click **Attach Disk**.  
The **Attach Disk** dialog box is displayed.

**Figure 5-3 Attach Disk (KVM)**

7. Select the target disk and specify the disk as the system disk or data disk
  - For KVM ECSs, you can specify a disk as a system disk or data disk but cannot specify a device name for the disk.
  - For Xen ECSs, you can specify the device name of a disk, such as **/dev/vdb**.

**NOTE**

- If no EVS disks are available, click **Create Disk** in the lower part of the list.
- For 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 EVS Data Disks](#).

## 5.4 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.

### Scenarios

- **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 5-2](#).

- **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 5-2](#).

## Impact on the System

- An initialization operation includes partitioning, which deletes all the data on the disk.
- If you change the partition style of a disk, data on the disk will be erased. Select an appropriate partition style when initializing disks.
- 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

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.

**Table 5-2** 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: <ul style="list-style-type: none"><li>• Four primary partitions</li><li>• Three primary partitions and one extended partition</li></ul> 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	<a href="#">Initializing a Linux Data Disk (Less Than or Equal to 2 TiB)</a>
			Windows	<a href="#">Initializing a Windows Data Disk</a>
Capacity > 2 TiB	GPT	GPT partitions are not classified, and there is no limit on the number of GPT partitions.	Linux	<a href="#">Initializing a Linux Data Disk (Greater Than 2 TiB)</a>
			Windows	<a href="#">Initializing a Windows Data Disk</a>

## 5.5 Adding a Yearly/Monthly EVS Disk

### Scenarios

You are allowed to add yearly/monthly EVS disks to a yearly/monthly ECS. The expiration time of the newly added EVS disks is the same as that of the ECS.

### Procedure

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

4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS.  
The page providing details about the ECS is displayed.
6. Click the **Disks** tab. Then, click **Add Disk**.  
The system switches to the EVS disk purchase page.
7. Configure parameters for the new EVS disk as prompted.
8. Click **Next**.
9. Verify that the disk is correctly configured, select the agreement, and click **Submit**.

The new EVS disk is automatically attached to the target ECS.

#### NOTE

After the new disk is detached, it can only be attached to the original ECS.

## 5.6 Detaching an EVS Disk from a Running ECS

### Scenarios

You can detach EVS disks from an ECS.

- System disks (mounted to **/dev/sda** or **/dev/vda**) can only be detached offline. They must be stopped before being detached.
- Data disks (mounted to points other than **dev/sda**) can be detached online if the attached ECS is running certain OSs. You can detach these data disks without stopping the ECS.

This section describes how to detach a disk from a running ECS.

### Constraints

- The EVS disk to be detached must be mounted to a point other than **/dev/sda** or **/dev/vda**.  
EVS disks mounted to **/dev/sda** or **/dev/vda** are system disks and cannot be detached from running ECSs.
- Before detaching an EVS disk from a running Windows ECS, make sure that UVP VMTools have been installed on the ECS and that the tools are running properly.
- Before detaching an EVS disk from a running Windows ECS, ensure that no programs are reading data from or writing data to the disk. Otherwise, data will be lost.
- SCSI EVS disks cannot be detached from running Windows ECSs.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the **umount** command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

## Notes

- On a Windows ECS, if the disk is in non-offline state, the system forcibly detaches the EVS disk. If this occurs, the system may generate a xenvbd alarm. You can ignore this alarm.

### NOTE

To view the status of an EVS disk, perform the following operations:

- Click **Start** in the task bar. In the displayed **Start** menu, right-click **Computer** and choose **Manage** from the shortcut menu.  
The **Server Manager** page is displayed.
  - In the navigation pane on the left, choose **Storage > Disk Management**.  
The EVS disk list is displayed in the right pane.
  - View the status of each EVS disk.
- Do not detach an EVS disk from an ECS that is being started, stopped, or restarted.
  - Do not detach an EVS disk from a running ECS whose OS does not support this feature. OSs supporting EVS disk detachment from a running ECS are listed in [OSs Supporting EVS Disk Detachment from a Running ECS](#).
  - For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and then attached to it again. This is a normal case due to the drive letter allocation mechanism of the Linux system.
  - For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and the ECS is restarted. This is a normal case due to the drive letter allocation mechanism of the Linux system.

## OSs Supporting EVS Disk Detachment from a Running ECS

OSs supporting EVS disk detachment from a running ECS include two parts:

- For the first part, see [External Image File Formats and Supported OSs](#).
- [Table 5-3](#) lists the second part of supported OSs.

**Table 5-3** OSs supporting EVS disk detachment from a running ECS

OS	Version
CentOS	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Debian	8.6.0 64bit
	8.5.0 64bit
Fedora	25 64bit
	24 64bit
SUSE	SUSE Linux Enterprise Server 12 SP2 64bit

OS	Version
	SUSE Linux Enterprise Server 12 SP1 64bit
	SUSE Linux Enterprise Server 11 SP4 64bit
	SUSE Linux Enterprise Server 12 64bit
OpenSUSE	42.2 64bit
	42.1 64bit
Oracle Linux Server release	7.3 64bit
	7.2 64bit
	6.8 64bit
	6.7 64bit
Ubuntu Server	16.04 64bit
	14.04 64bit
	14.04.4 64bit
Windows	Windows Server 2008 R2 Enterprise 64bit
	Windows Server 2012 R2 Standard 64bit
	Windows Server 2016 R2 Standard 64bit
Red Hat Linux Enterprise	7.3 64bit
	6.8 64bit

 **NOTE**

Online detachment is not supported by the ECSs running OSs not listed in the preceding table. For such ECSs, stop the ECSs before detaching disks from them to prevent any possible problems from occurring.

## Procedure

1. On the **Elastic Cloud Server** page, click the name of the ECS from which the EVS disk is to be detached. The page providing details about the ECS is displayed.
2. Click the **Disks** tab. Locate the row containing the EVS disk to be detached and click **Detach**.

## 5.7 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:

- Apply for an EVS disk and attach it to an ECS.
- Expand the capacity of an existing EVS disk. The capacities of both system disks and data disks can be expanded.

For details 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.

## Related Operations

For a Windows ECS, if you want to expand the disk capacity by clearing disk files, you can reduce the size of the WinSxS folder using tools built into Windows. For details, see [Clean Up the WinSxS Folder](#).

# 5.8 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.



## 5.9 Enabling Advanced Disk

### Scenarios

- Disk functions have been upgraded on the platform. Newly created ECSs can have up to 60 attached disks. However, an existing ECS can still have a maximum of 24 attached disks (40 for certain ECSs). To allow such ECSs to have up to 60 attached disks, enable advanced disk.
- After advanced disk is enabled, you can view the mapping between device names and disks. For details, see "What Is the Mapping Between Device Names and Disks?"

This section describes how to enable advanced disk on an ECS.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS. The page providing details about the ECS is displayed.
5. Click the **Disks** tab.
6. View the current number of disks that can be attached to the ECS and enable advanced disk as prompted.  
The **Enable Advanced Disk** dialog box is displayed.
7. Click **OK**.
8. Stop and then start the target ECS.  
This operation allows advanced disk to take effect.

9. Switch to the page providing details about the ECS again, click the **Disks** tab, and check whether the number of disks that can be attached to the ECS has been changed.
  - If yes, advanced disk has been enabled.
  - If no, enabling advanced disk failed. In such a case, try again later or contact customer service.

# 6 Elastic Network Interfaces

---

## 6.1 Network Interface 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 cloud servers (such as 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.

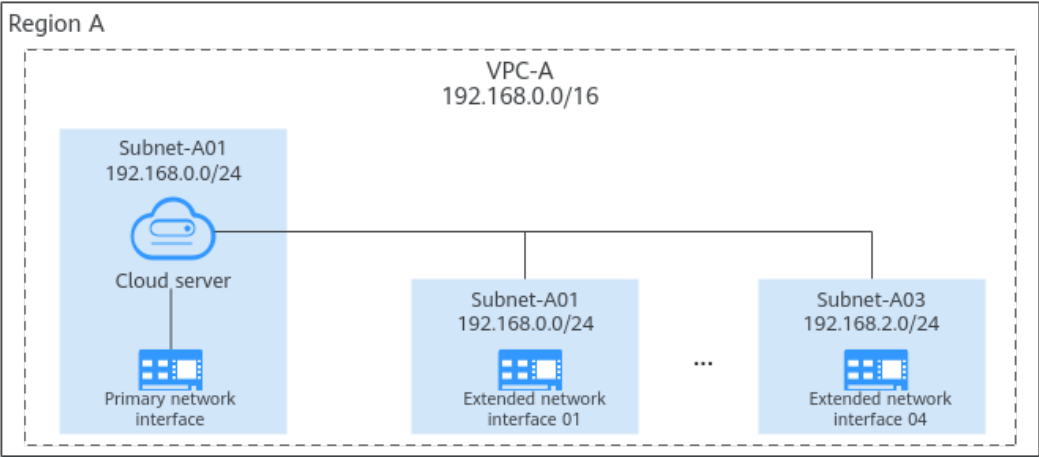
### Application Scenarios

Elastic network interfaces help you flexibly migrate and separate services.

- **Flexible migration:** You can detach an **extended** network interface from a cloud server and attach it to another one. The private IP address, EIP, and security group rules of the original cloud server can be migrated together, so you do not need to reconfigure them. This allows the service traffic on the faulty cloud server to be quickly switched to the standby one, achieving quick service recovery.
- **Service separation:** You can configure multiple network interfaces for a cloud server. These network interfaces can be in different subnets of the same VPC and process the internal, external, and management traffic of the cloud server respectively. You can configure access control policies and routes for each subnet, and define security group rules for each network interface to isolate networks and service traffic.

In [Figure 6-1](#), the cloud server has one primary network interface and four extended network interfaces. These network interfaces can be in different subnets. In this example, extended network interface 01 and the primary network interface are in Subnet-A01, and extended network interface 04 is in Subnet-A03.

Figure 6-1 Cloud server network interfaces



Each cloud server can have a limited number of elastic network interfaces attached. If the cloud server specifications support supplementary network interfaces, you can attach supplementary network interfaces to the elastic network interfaces.

For details, see [Supplementary Network Interface Application Scenarios](#).

Constraints on Using Network Interfaces

- The number of network interfaces that can be attached to an ECS is determined by the ECS specifications. For details, see [ECS Specifications](#).

Table 6-1 Constraints on using different types of network interfaces

Net wor k Inte rfac e Type	Creation	Attachment	Communication with External Networks	Commun ication with Public Service Zone
Prim ary netw ork inter face	Created along with the instance by default and cannot be created separately.	Cannot be detached from the instance.	Supported	Supporte d

Net wor k Inte rfac e Type	Creation	Attachment	Communication with External Networks	Commun ication with Public Service Zone
Extensio n netw ork inter face	Can be created separately on the Network Interface console.	Can be attached to or detached from the instance.	Not supported. Policy-based routes need to be configured for external network access. <a href="#">How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?</a>	Not supported



- Extended network interfaces cannot be used to directly access Huawei Cloud services, such as DNS. You can use VPC Endpoint (VPCEP) to access these services. For details, see [Buying a VPC Endpoint for Accessing Interface VPC Endpoint Services](#).

## 6.2 Attaching a Network Interface

### Scenarios

If your ECS requires multiple network interfaces, you can attach them to your ECS.

### Procedure

- Log in to the management console.
- Click  in the upper left corner and select a region and project.
- Click . Under **Compute**, click **Elastic Cloud Server**.
- Click the name of the ECS to which you want to attach a network interface. The page providing details about the ECS is displayed.
- On the **Network Interfaces** tab, click **Attach Network Interface**.
- Select either of the following methods to attach the network interface.
  - Use an existing network interface.
    - (Optional) Search for the network interface by name, ID, or private IP address.
    - 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 6-2** Configuring the subnet and security group

**Attach Network Interface**

ECS Name: ecs-XXXX

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

VPC: vpc-XXXX

Subnet: subnet-03(192.168.3.0/24) [View Subnet](#)

New Private IP Address: User-defined IP address [View In-Use IP Address](#)

Security Group:  [View Security Group](#)

<input type="checkbox"/> Security Group Name	Description
<input type="checkbox"/> default	Default security group
<input type="checkbox"/> sg-a2b4	
<input type="checkbox"/> sg-test	

- **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 ethX** 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.

## 6.3 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 network 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.

## 6.4 Changing a VPC

### Scenarios

This section describes how to change a VPC.

You can change the VPC of an individual ECS or use job management of Cloud Operations Center (COC) to batch change the VPCs of multiple ECSs.

### Constraints

- Only running or stopped ECSs support VPC change.
- VPC change is only supported when the ECS has a single NIC attached.
- 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.
- You can change the VPCs of a maximum of 50 ECSs in a batch.

## 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, such as ELB, VPN, NAT, traffic mirroring, and DNS.

## Prerequisites

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


If you want to change VPCs for multiple ECSs in a batch, do as follows:

- Enable and authorize COC.  
For IAM users, permissions for COC operations need to be granted. For details, see [Configuring Custom Policies for ECS Self-Service O&M](#).
- Create an agency to grant ECS access to COC.  
For details, see [Creating an Agency to Grant ECS Access to COC](#).

## Procedure

You can change the VPC of an individual ECS or batch change VPCs of multiple ECSs as needed.

### Changing the VPC of an Individual ECS

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 > Change VPC** in the **Operation** column.  
The **Change VPC** dialog box is displayed.

**Figure 6-3** Change VPC

×

Change VPC

Changing the VPC will interrupt ECS network connections and change the subnet, IP address, and MAC address of the ECS.  
During the change process, do not perform operations on the ECS, including its EIP.  
After the VPC is changed, reconfigure source/destination check and virtual IP address.  
After the VPC is changed, reconfigure network-related application software and services, such as ELB, VPN, NAT, and DNS.

ECS Name

ecs-764c

VPC

vpc1( )

View In-Use VPCs

Subnet

subnet-A( )

View Subnet

Private IP Address

User-defined IP address

View In-Use IP Address

Security Group

Sys-default

View Security Group

OK

Cancel


- Specify the VPC, subnet, private IP address, and security group.  
You can select multiple security groups. 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.

- Click **OK**.

## Batch Changing VPCs of Multiple ECSs

- Log in to the management console.
- Click . Under **Compute**, click **Elastic Cloud Server**.
- In the ECS list, select the ECSs for which you want to change VPCs.
- Above the ECS list, choose **More > Change VPC**.

**Figure 6-4** Changing VPCs

**Change VPC**

• Changing the VPC will interrupt ECS network connections and change the subnet, IP address, and MAC address of the ECS.  
• During the VPC change, do not perform operations on the ECS or its EIP.  
• After the VPC is changed, you will have to reconfigure the source/destination check, virtual IP address, and network-related application software and services (such as ELB, VPN, NAT Gateway, Traffic Mirroring, Global EIP, and DNS).

**Resources to Be Operated**

The VPC will be changed for 2 ECSs.

Name	ID	VPC
vpc-		vpc- 172.16.0.0/16)
ecs-		VPC 172.16.0.0/12)

**Agency**

IAM Agency

--Select-- [Create Agency](#)

Batch VPC change requires that you create an agency for COC and grant required ECS permissions to the agency. [Learn more](#)

**Network**

VPC [?](#)

vpc- 172.16.0.0/16) [Create VPC](#)

Subnet

subnet- 172.16.1.0/24) [Create Subnet](#)

Available IP Addresses: 247

IPv4 Address

[Automatically assign](#) [View In-Use IP Address](#)

**Security Group**

[Cancel](#) [OK](#)

5. In the displayed box, from the **IAM Agency** drop-down list, select an agency that has granted ECS access and related permissions to COC.

If no agency is available, create one by clicking **Create Agency**. For details, see [Creating an Agency to Grant ECS Access to COC](#).

6. Specify the VPC, subnet, private IP address, and security group.  
You can select multiple security groups. The access rules of all the selected security groups will apply to the ECSs.

#### NOTE

Using multiple security groups may deteriorate ECS network performance. You are advised to select no more than five security groups.

7. Click **OK**.  
The batch VPC change is requested. You can view the results on the **Job Tickets** tab of COC.  
For more information about job management, see [Executing Public Jobs](#).

## 6.5 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.

## 6.6 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. The configurations for ECSs will not be lost after restart.

- 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:

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:

```
[root@ecs-X-ubuntu:~]# 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, virtual IP address 172.16.0.125 is bound to network interface eth0.

```
[root@ecs-X-ubuntu:~]# 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:d5:c5: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:a5b3:62c: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-ubuntu:~# 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**:
- 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:
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-ubuntu:/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:

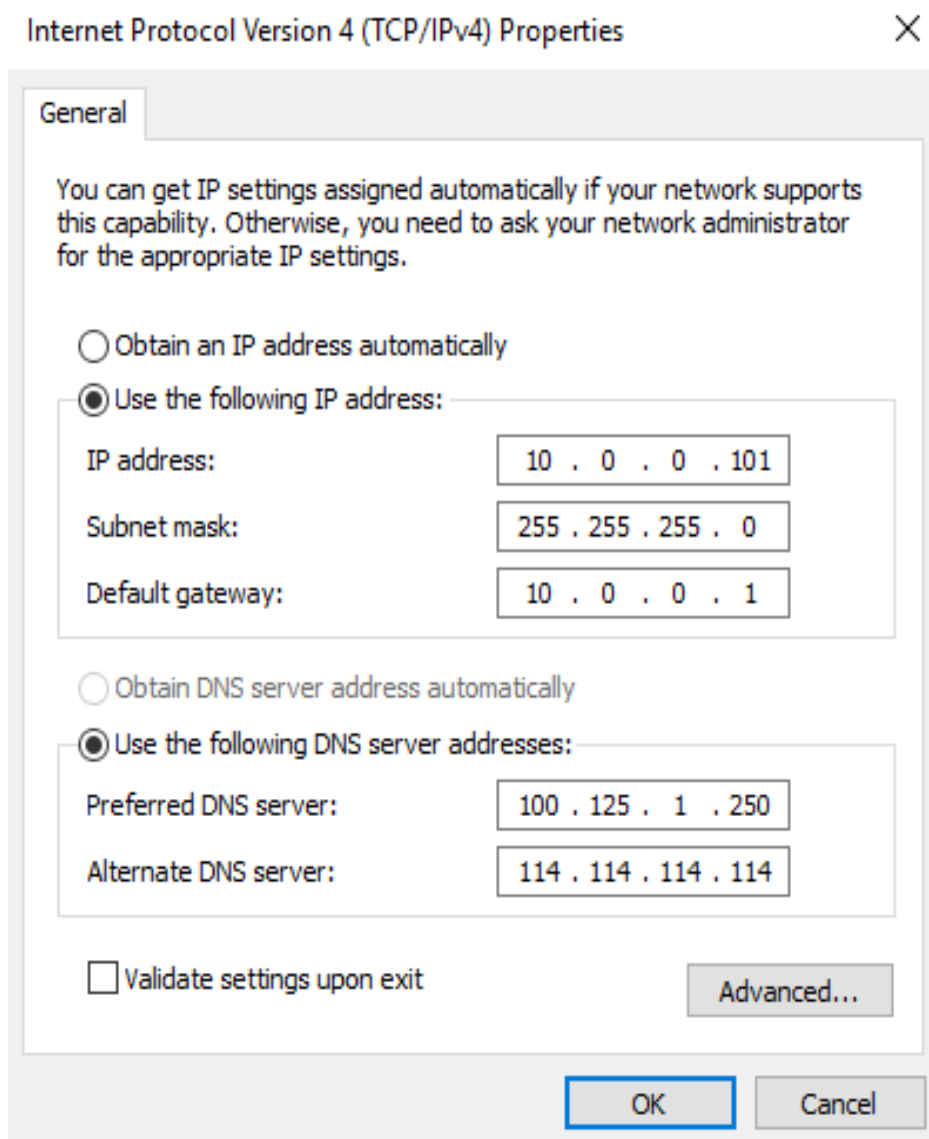
1. Open the configuration file **01-netcfg.yaml** and delete the virtual IP address of the corresponding network interface by referring to **3**.
2. Make the deletion take effect by referring to **4**.

## Windows: Windows Server

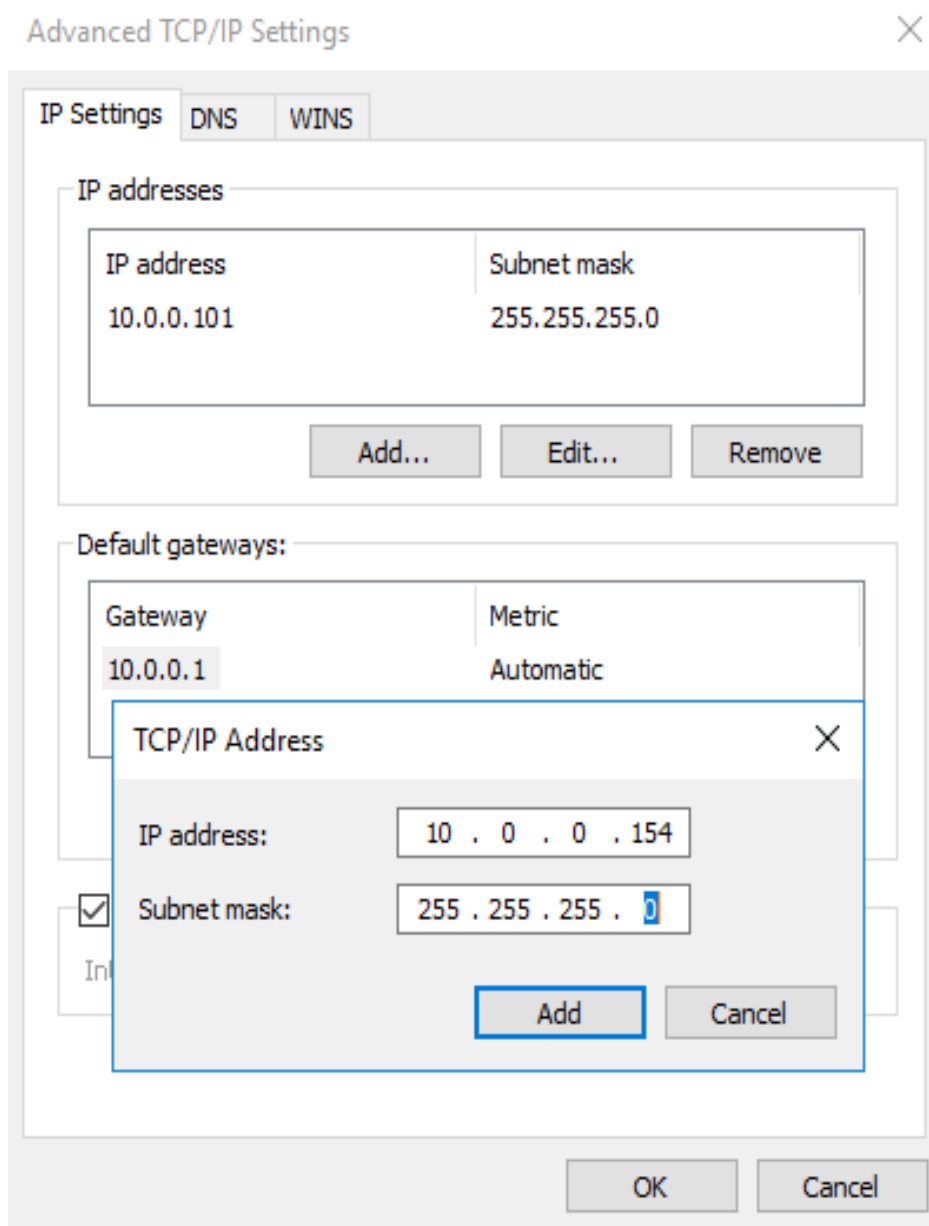
The following operations use Windows Server as an example.

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 6-5** Configuring private IP address



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 6-6** Configuring virtual IP address

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.

## 6.7 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 6-3](#) support NIC multi-queue.

#### NOTE

- The PV driver of a Windows ECS dynamically adjusts the number of NIC queues based on the number of vCPUs of the ECS, and you do not need to set the number of Windows NIC multi-queues.
- Public images that contain Windows Server 2008 are no longer available. However, you can still use private images that contain Windows Server 2008.
- It is a good practice to upgrade the kernel version of the Linux ECS to 2.6.35 or later. Otherwise, NIC multi-queue is not supported.

Run the **uname -r** command to obtain the kernel version. If the kernel version is earlier than 2.6.35, contact customer service to upgrade the kernel.

**Table 6-2** Support of NIC multi-queue for Windows ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Windows Server 2008 R2 Standard/Enterprise/DataCenter 64bit	Yes	Yes
Windows Server 2008 Enterprise SP2 64bit	Yes	Yes
Windows Server 2008 Web R2 64-bit	Yes	Yes
Windows Server 2008 R2 Enterprise 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard 64bit_WithGPUdriver	Yes	Yes
Windows Server 2012 R2 Standard/DataCenter 64 bit	Yes	Yes
Windows Server 2016 Standard/DataCenter 64 bit	Yes	Yes
Windows Server 2019 DataCenter 64 bit	Yes	Yes

**Table 6-3** Support of NIC multi-queue for Linux ECSs

Image	Support of NIC Multi-Queue	NIC Multi-Queue Enabled by Default
Ubuntu 14.04/16.04/18.04/20.04 server 64bit	Yes	Yes
OpenSUSE 42.2/15.* 64bit	Yes	Yes
SUSE Enterprise 12 SP1/SP2 64bit	Yes	Yes
CentOS 6.8/6.9/7.*/8.* 64bit	Yes	Yes
Debian 8.0.0/8.8.0/8.9.0/9.0.0/10.0.0/10.2.0 64bit	Yes	Yes
Fedora 24/25/30 64bit	Yes	Yes
EulerOS 2.2/2.3/2.5 64bit	Yes	Yes

## Importing the External Image File to the IMS Console

For details, see "Registering an Image File as a Private Image" in *Image Management Service User Guide*. After the image file is imported, view the value of **NIC Multi-Queue** on the page providing details about the image.


- If the value is **Supported**, go to [Creating an ECS Using a Private Image](#).
- If the value is **Not supported**, go to [Setting NIC Multi-Queue for the Image](#).

## Setting NIC Multi-Queue for the Image


Windows OSs have not commercially supported NIC multi-queue. If you enable NIC multi-queue in a Windows image, starting an ECS created using such an image may be slow.

Use one of the following methods to set the NIC multi-queue attribute:

### Method 1:

1. Log in to the management console.
2. Click . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab, locate the row containing the target image, click **Modify** in the **Operation** column.
4. Set the NIC multi-queue attribute of the image.

### Method 2:

1. Log in to the management console.
2. Click . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab. In the image list, click the name of the target image to switch to the page providing details about the image.
4. Click **Modify** in the upper right corner. In the displayed **Modify Image** dialog box, set the NIC multi-queue attribute.

### Method 3: Add `hw_vif_multiqueue_enabled` to an image through the API.

1. For instructions about how to obtain the token, see [Authentication](#).
2. For instructions about how to call an API to update image information, see [Updating Image Information \(Native OpenStack API\)](#).
3. Add **X-Auth-Token** to the request header.  
The value of **X-Auth-Token** is the token obtained in step 1.
4. Add **Content-Type** to the request header.

The value of **Content-Type** is **application/openstack-images-v2.1-json-patch**.

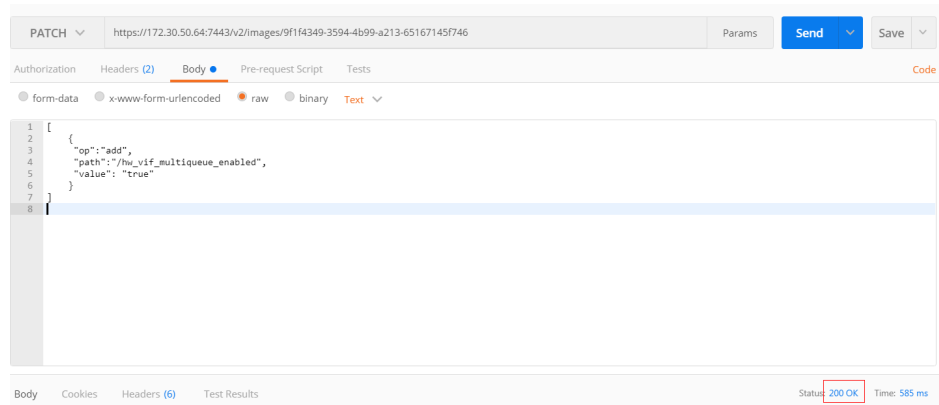
The request URI is in the following format:

PATCH /v2/images/{image\_id}

The request body is as follows:

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

**Figure 6-7** shows an example request body for modifying the NIC multi-queue attribute.

**Figure 6-7** Example request body

## 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

The PV driver of a Windows ECS dynamically adjusts the number of NIC queues based on the number of vCPUs of the ECS, and you do not need to set the number of Windows NIC multi-queues.

A script for automatically enabling NIC multi-queue on a Linux ECS is available. After the script is configured, the ECS supports NIC multi-queue.

### 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.

### NOTE

For Windows ECSs, you can call the API for [querying NICs of an ECS](#).

The response parameter **multiqueue\_num** indicates 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.

## 6.8 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. For details, see [Network](#).

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.

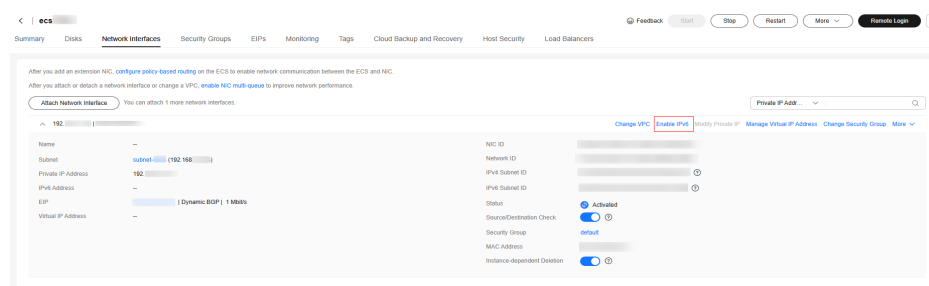
This section describes how to enable IPv6 for a created ECS.

### 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 C7 and M7

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.

Figure 6-8 Enabling IPv6



- Only one IPv6 address can be bound to a network interface.

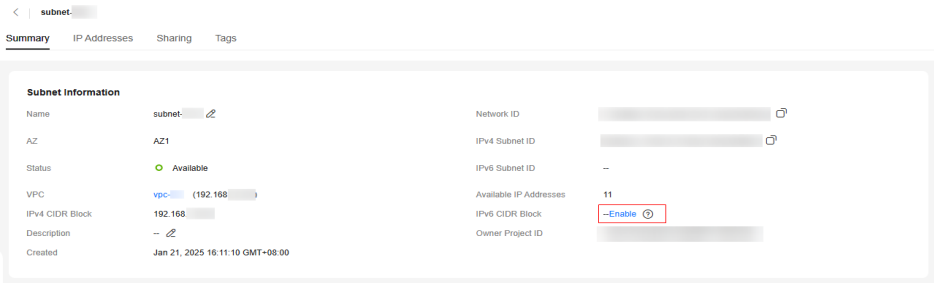
### (Optional) Step 1: Enabling IPv6 for a Subnet

#### 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.
- Log in to the management console and go to the [Elastic Cloud Server](#) page.
  - 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**.

Figure 6-9 Enabling IPv6 for a subnet

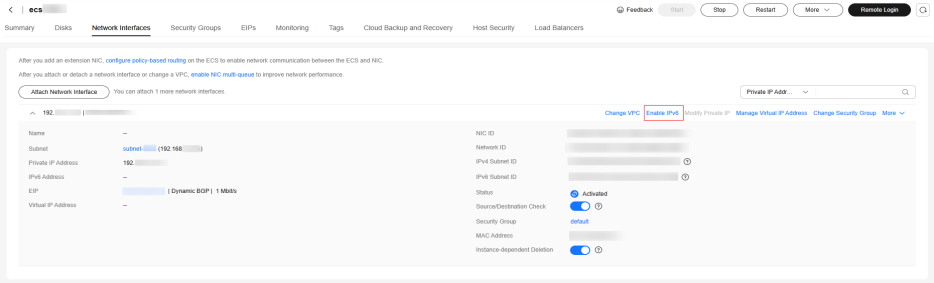


5. Click **OK** to enable IPv6 for the subnet.

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.

Figure 6-10 Enabling IPv6



 **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](#).
  - After IPv6 is enabled, if you need to enable public IPv6 communication, you need to create an EIP shared bandwidth and add the IPv6 address of the ECS to the shared bandwidth. For details, see [Setting Up an IPv4/IPv6 Dual-Stack Network in a VPC](#).
  - 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.

## 6.9 Dynamically Assigning IPv6 Addresses

### Scenarios

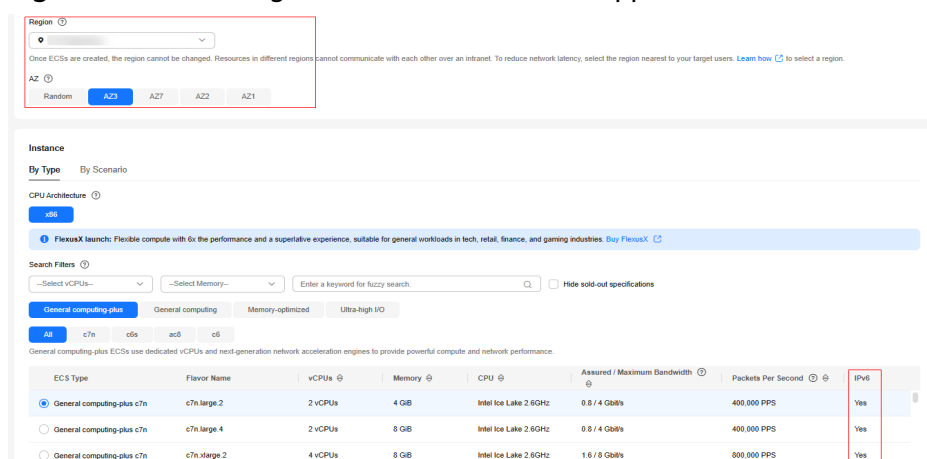
IPv6 addresses are used to deal with IPv4 address exhaustion. If an ECS uses an IPv4 address, the ECS can run in dual-stack mode after IPv6 is enabled for it. Then, the ECS will have two IP addresses to access the intranet and Internet: an IPv4 address and an IPv6 address.

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.

**Figure 6-11** Checking whether an ECS flavor supports IPv6


If the value of **IPv6** is **Yes** for an ECS flavor, the flavor supports IPv6.

#### NOTE

**AZ** and **Flavor** determine whether IPv6 is supported.

After you select an AZ, if **IPv6** is not displayed or the value of **IPv6** is **No**, IPv6 is not supported by any or certain flavors in the AZ.

- Ensure that **Automatically-assigned IPv6 address** is selected during ECS creation.

**Figure 6-12** Automatically-assigned IPv6 address

- 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 2012/2008 is used as an example to describe how to enable dynamic assignment of IPv6 addresses in Windows, as shown in [Table 6-4](#).
- Linux: Dynamic assignment of IPv6 addresses can be enabled automatically (recommended) or manually, as shown in [Table 6-4](#).

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 6-4** Enabling dynamic assignment of IPv6 addresses for different OSs

OS	Auto/Manual	Reference
Windows Server 2012	Auto	<a href="#">Windows Server 2012</a>
Windows Server 2008	Auto	<a href="#">Windows Server 2008</a>
Linux	Auto (recommended)	<a href="#">Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)</a>
Linux	Manual	<a href="#">Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses)</a>

## Windows Server 2012

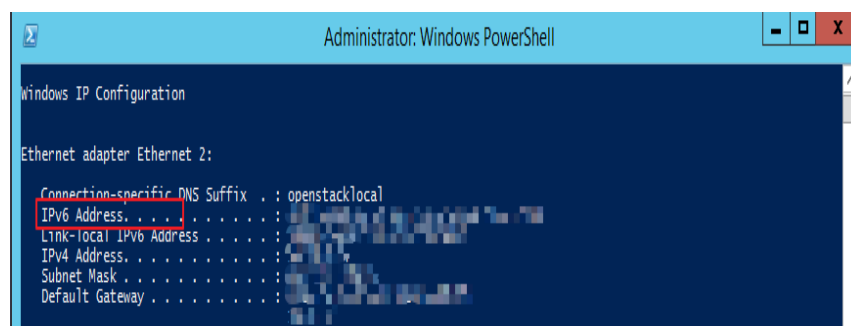
**Step 1** Check whether IPv6 is enabled for the ECS.

Run the following command in the CMD window:

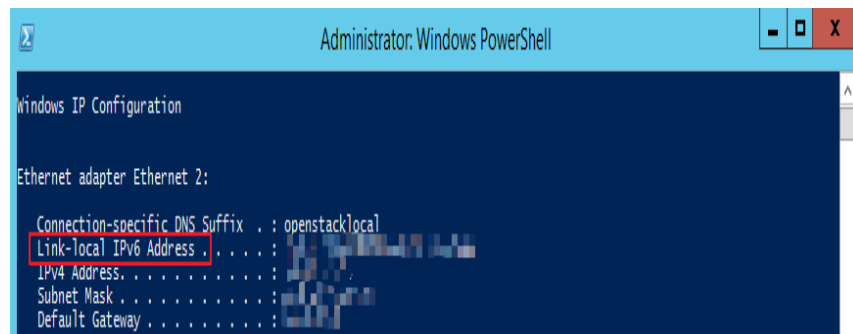
### ipconfig

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

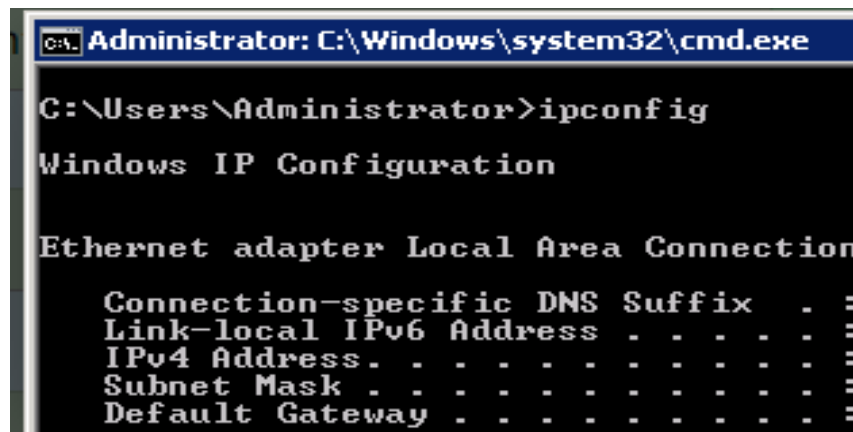
**Figure 6-13** Querying the IPv6 address



- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to [Step 2](#).

**Figure 6-14** Link-local IPv6 address


- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to [Step 3](#).

**Figure 6-15** IPv6 disabled**NOTE**

By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 6-13](#). No additional configuration is required.

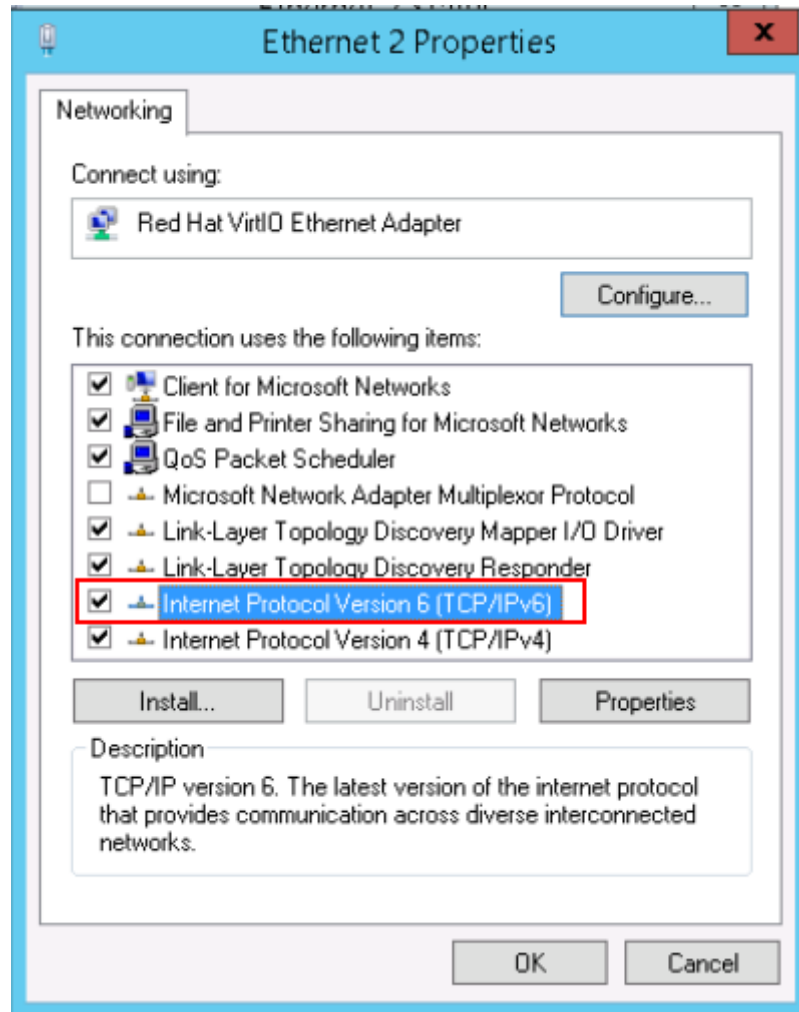
**Step 2** Enable dynamic IPv6 address assignment.

1. Choose **Start > Control Panel**.
2. Click **Network and Sharing Center**.
3. Click the Ethernet connection.

**Figure 6-16** Ethernet connection

4. In the **Ethernet Status** dialog box, click **Properties** in the lower left corner.
5. Select **Internet Protocol Version 6 (TCP/IPv6)** and click **OK**.

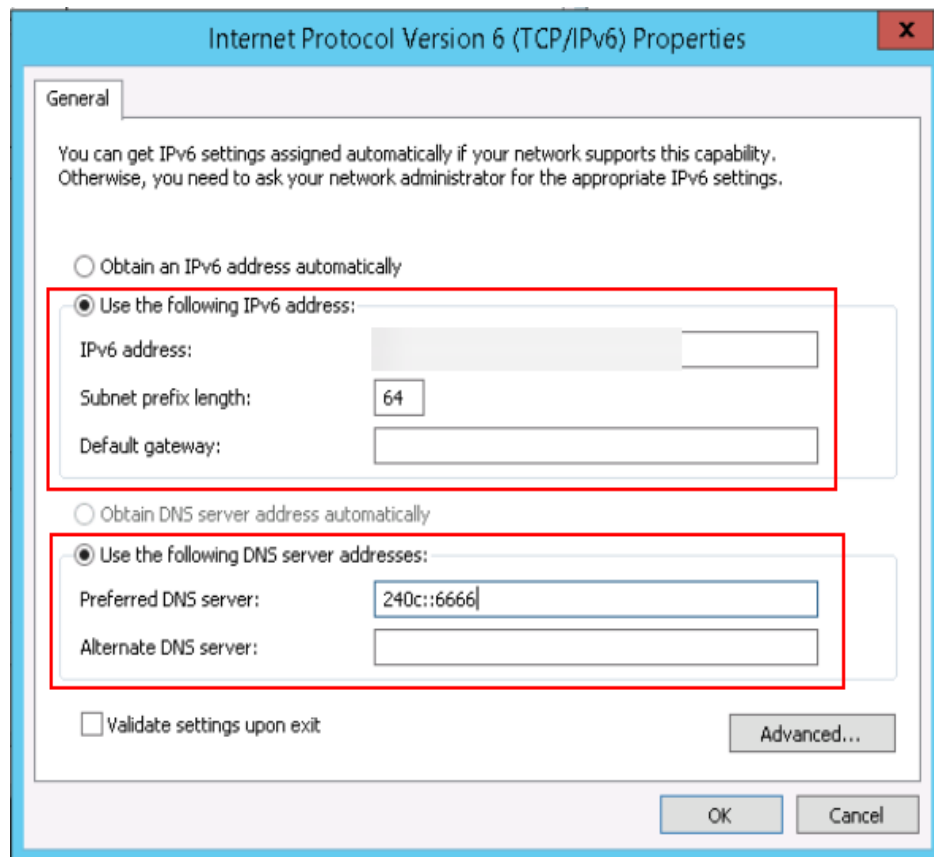
**Figure 6-17** Configuring dynamic IPv6 address assignment



6. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

**Step 3** Enable and configure IPv6.

1. 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 6-18** Configuring an IPv6 address and a DNS server address

- (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**
- Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

----End

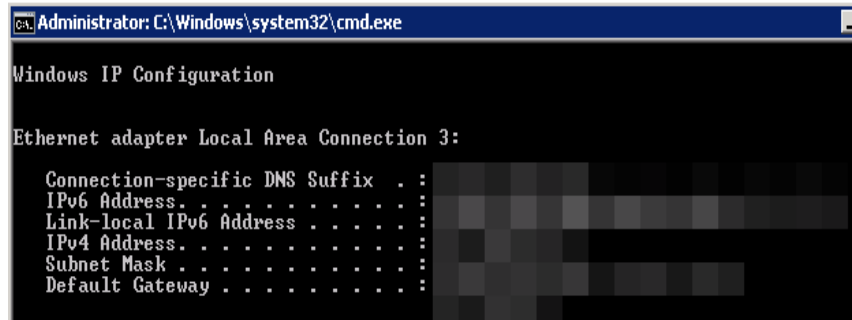
## Windows Server 2008

**Step 1** Check whether IPv6 is enabled for the ECS.

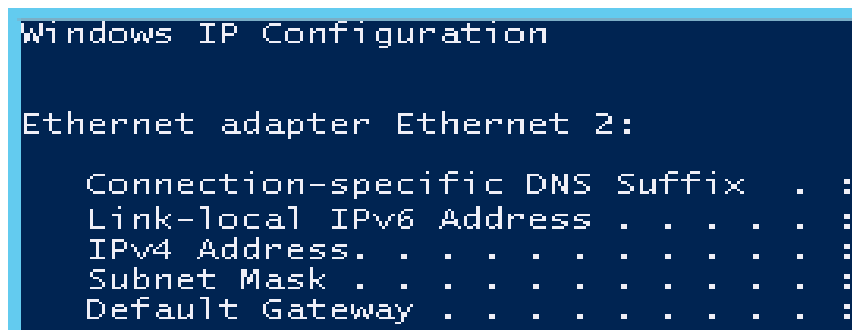
Run the following command in the CMD window:

### ipconfig

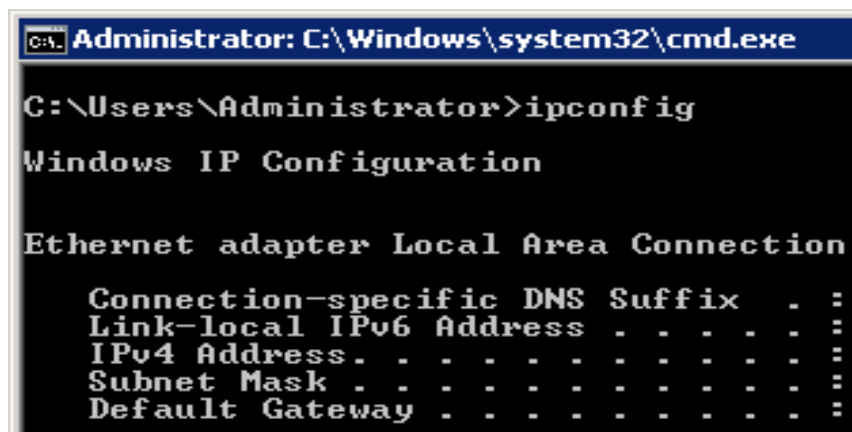
- If an IPv6 address and a link-local IPv6 address are displayed, IPv6 is enabled and dynamic IPv6 assignment is also enabled.

**Figure 6-19** Querying the IPv6 address

- If only a link-local IPv6 address is displayed, IPv6 is enabled but dynamic IPv6 assignment is not enabled. Go to [Step 2](#).

**Figure 6-20** Link-local IPv6 address

- If neither an IPv6 address nor link-local IPv6 address is displayed, IPv6 is disabled. Go to [Step 3](#).

**Figure 6-21** IPv6 disabled**NOTE**


By default, dynamic IPv6 address assignment is enabled for Windows public images, as shown in [Figure 6-19](#). 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 6-22** Configuring dynamic IPv6 address assignment

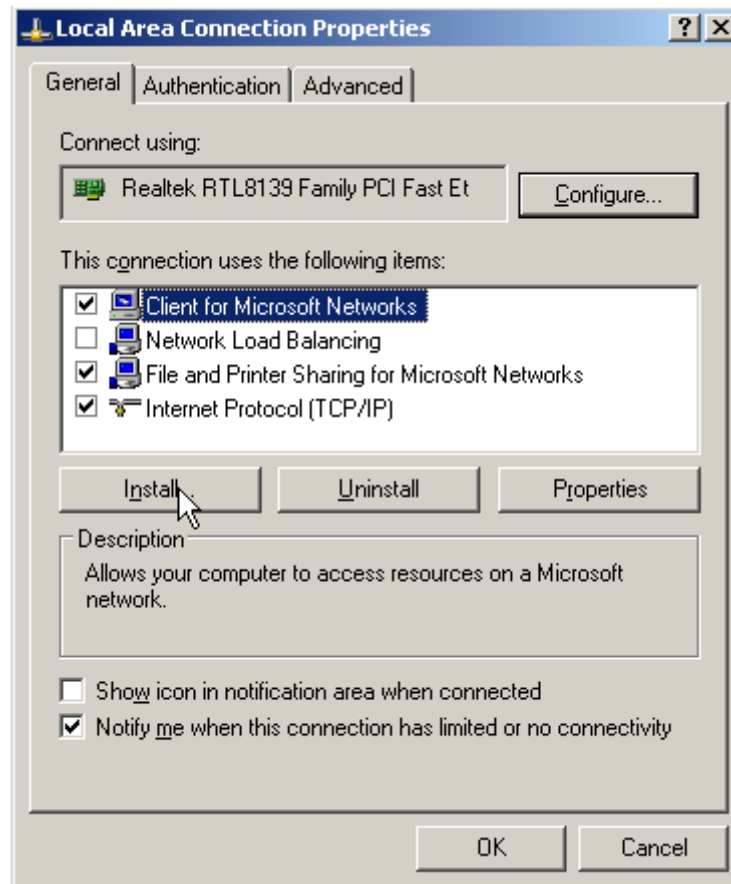


6. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

**Step 3** Enable and configure IPv6.


1. Choose **Start > Control Panel > Network Connection > Local Connection**.
2. Select **Properties**, select the following options, and click **Install**.

**Figure 6-23** Enabling and configuring IPv6




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

**Figure 6-24** Adding the protocol



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

**Figure 6-25** Network protocols

5. (Optional) Run the following commands depending on your ECS OS.  
For Windows Server 2008, run the following command in PowerShell or CMD:  
**netsh interface ipv6 set global randomizeidentifiers=disable**  
Disable the local connection and then enable it again.  
To disable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Right-click the local connection and choose **Disable** from the shortcut menu.  
To enable the local connection, choose **Start > Control Panel > Network and Internet > Network and Sharing Center > Change Adapter Options**. Right-click the local connection and choose **Enable** from the shortcut menu.
6. Perform [Step 1](#) to check whether dynamic IPv6 address assignment is enabled.

----End

## Linux (Automatically Enabling Dynamic Assignment of IPv6 Addresses)

The **ipv6-setup-xxx** tool can be used to enable Linux OSs to automatically acquire IPv6 addresses. xxx indicates a tool, which can be rhel or debian.

You can also enable dynamic IPv6 address assignment by following the instructions in [Linux \(Manually Enabling Dynamic Assignment of IPv6 Addresses\)](#).

**CAUTION**

- When you run **ipv6-setup-xxx**, the network service will be automatically restarted. As a result, the network is temporarily disconnected.
- If a private image created from a CentOS 6.x or Debian ECS with automatic IPv6 address assignment enabled is used to create an ECS in an environment that does not support IPv6, the ECS may start slow because of IPv6 address assignment timeout. Set the timeout duration for assigning IPv6 addresses to 30s by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#) and try to create a new private image again.

**Step 1** Run the following command to check whether IPv6 is enabled for the ECS:

**ip addr**

- If only an IPv4 address is displayed, IPv6 is disabled. Enable it by referring to [Setting the Timeout Duration for IPv6 Address Assignment](#).

**Figure 6-26** IPv6 disabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
link/ether fa:16:3e: b8:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
inet 10.0.0.1 brd 10.0.0.255 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 6-27** IPv6 enabled

```
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 10.0.0.1 brd 10.0.0.255 scope global noprefixroute dynamic eth0
 valid_lft 76391sec preferred_lft 76391sec
inet6 fe80::f816:3eff:fe00:0000/64 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 6-28** 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 10.0.0.1 brd 10.0.0.255 scope global noprefixroute dynamic eth0
 valid_lft 86395sec preferred_lft 86395sec
inet6 2407:c080:802::1/128 scope global dynamic
 valid_lft 7496sec preferred_lft 7196sec
inet6 fe80::f816:3eff:fe00:0000/64 scope link noprefixroute
 valid_lft forever preferred_lft forever
```

**NOTE**

IPv6 is enabled for Linux public images by default, as shown in [Figure 6-27](#).

**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 6-5](#) lists the paths for obtaining **ipv6-setup-rhel** and **ipv6-setup-debian**.

**Table 6-5** Download paths of **ipv6-setup-rhel** and **ipv6-setup-debian**

Series	Release Version	How to Obtain
RHEL	<ul style="list-style-type: none"><li>– CentOS 6/7</li><li>– EulerOS 2.2/2.3</li><li>– Fedora 25</li></ul>	<a href="https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-rhel">https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-rhel</a>
Debian	<ul style="list-style-type: none"><li>– Ubuntu 16/18</li><li>– Debian 8/9/10</li></ul>	<a href="https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-debian">https://ecs-instance-driver.obs.cn-north-1.myhuaweicloud.com/ipv6/ipv6-setup-debian</a>

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 6-29** IPv6 disabled

```
eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
link/ether fa:16:3e: b8:ff:ff:ff:ff:ff
inet b8:ff:ff:ff:ff:ff 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 6-30** IPv6 enabled

```
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 b8:ff:ff:ff:ff:ff scope global noprefixroute dynamic eth0
 valid_lft 76391sec preferred_lft 76391sec
inet6 fe80::f816: b8:ff:ff:ff:ff:ff /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 6-31** 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 b8:ff:ff:ff:ff:ff scope global noprefixroute dynamic eth0
 valid_lft 86395sec preferred_lft 86395sec
inet6 2407:c080:802: b8:ff:ff:ff:ff:ff /128 scope global dynamic
 valid_lft 7496sec preferred_lft 7196sec
inet6 fe80::f816:3eff: b8:ff:ff:ff:ff:ff /64 scope link noprefixroute
 valid_lft forever preferred_lft forever
```

### NOTE

IPv6 is enabled for Linux public images by default, as shown in [Figure 6-30](#).

**Step 2** Enable IPv6 for the ECS.


- 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 6-32** Configuration file name




```
root@ecs-: /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 6-33** Edited configuration file



```
Let NetworkManager manage all devices on this 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 6-34** 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 6-35 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 6-36 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 6-37 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/ip6tablef 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 6-38 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 6-39 Network service status

```
Welcome to Huawei Cloud Service

[root@ecs-~]# 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 18:32:01 CST; 24h ago
     Docs: man:NetworkManager(8)
    Main PID: 581 (NetworkManager)
    CGroup: /system.slice/NetworkManager.service
            └─ 581 /usr/sbin/NetworkManager --no-daemon
               └─ 555 /usr/sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -pf /var/run/dhclient-eth0.pid -lf /var/lib/NetworkManager/dhclient-5fb8b0d-8bb8-7ffb-45f1-d5edd55f3e83-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 6-40 Method for ECS NIC to obtain an IP address

```
[root@ecs-~]# 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 6-41 Checking the DHCP process

```
[root@ecs-~]# 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 18:32:01 CST; 1 day 4h ago
     Docs: man:NetworkManager(8)
    Main PID: 581 (NetworkManager)
    CGroup: /system.slice/NetworkManager.service
            └─ 581 /usr/sbin/NetworkManager --no-daemon
               └─ 555 /usr/sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -pf /var/run/dhclient-eth0.pid -lf /var/lib/NetworkManager/dhclient-5fb8b0d-8bb8-7ffb-45f1-d5edd55f3e83-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 6-42 Checking the static configuration

```
[root@localhost ~]# 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 fc00:f816:3eff:fcfc:14/64 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 6-43 Modification 1

```
115 case "$1" in
116 start)
117     if init_is_upstart; then
118         exit 1
119     fi
120     process_options
121     check_ifstate
122
123     if [ "$CONFIGURE_INTERFACES" = no ]
124     then
125         log_action_msg "Not configuring network interfaces, see /etc/default/networking"
126         exit 0
127     fi
128     set -f
129     exclusions=$(process_exclusions)
130     log_action_begin_msg "Configuring network interfaces"
131     if /usr/bin/timeout 30 ifup -a $exclusions $verbose && ifup_hotplug $exclusions $verbose
132     then
133         log_action_end_msg $?
134     else
135         log_action_end_msg $?
136     fi
137     ;;
138 stop)
139 if init_is_upstart; then
140     exit 0
141 fi
142 check_network_file_systems
143 check_network_swap
144
145 log_action_begin_msg "Deconfiguring network interfaces"
146 if /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose; then
147     log_action_end_msg $?
```

Figure 6-44 Modification 2

```
154 reload)
155     process_options
156
157     log_action_begin_msg "Reloading network interfaces configuration"
158     state=$(cat /run/network/ifstate)
159     /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
160     if /usr/bin/timeout 30 ifup --exclude=lo $state $verbose ; then
161         log_action_end_msg $?
162     else
163         log_action_end_msg $?
164     fi
165     ;;
166
167 force-reload|restart)
168     if init_is_upstart; then
169         exit 1
170     fi
171     process_options
172
173     log_warning_msg "Running $0 $1 is deprecated because it may not re-enable some interfaces"
174     log_action_begin_msg "Reconfiguring network interfaces"
175     /usr/bin/timeout 30 ifdown -a --exclude=lo $verbose || true
176     set -f
177     exclusions=$(process_exclusions)
178     if /usr/bin/timeout 30 ifup -a --exclude=lo $exclusions $verbose && ifup_hotplug $exclusions $verbose
179     then
180         log_action_end_msg $?
181     else
182         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 **/** 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/systemd/system/network-online.target.wants/networking.service
 - b. Press **i** to enter editing mode and change **TimeoutStartSec=5min** to **TimeoutStartSec=30**.

6.10 Creating an Agency to Grant ECS Access to COC

Scenarios

COC's job management is required for batch VPC change. To use job management, you need to create an agency to grant ECS access and related permissions to COC.

If you need to batch change VPCs of multiple ECSs, you can refer to this section to create an agency first.

Procedure

1. Log in to the [IAM console](#).
2. On the IAM console, choose **Agencies** from the left navigation pane and click **Create Agency** on the displayed page.
3. Configure agency parameters.
 - **Agency Name:** Enter an agency name.
 - **Agency Type:** Select **Cloud service**.
 - **Cloud Service:** Select the cloud service to be authorized, for example, **CloudOperationsCenter**.

- **Validity Period:** Set this parameter based on service requirements.
- **Description (Optional):** Enter a description.

Figure 6-45 Creating an agency

Agencies / Create Agency

* Agency Name

* Agency Type ☐ Account
Delegate another Huawei Cloud account to perform operations on your resources.
☒ Cloud service
Delegate a cloud service to access your resources in other cloud services.

* Cloud Service

* Validity Period

Description

0/255

OK Cancel

4. Click **OK**.
5. In the displayed dialog box, click **Authorize**.
6. Select the permissions to be granted to the agency.

If no policy is available in the list, click **Create Policy** in the upper right corner of the list to create one. For details, see [Creating a Custom Policy](#).

A custom policy must contain the following content:

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

7. Click **Next** and specify the authorization scope.
8. Click **OK**. The agency is created.

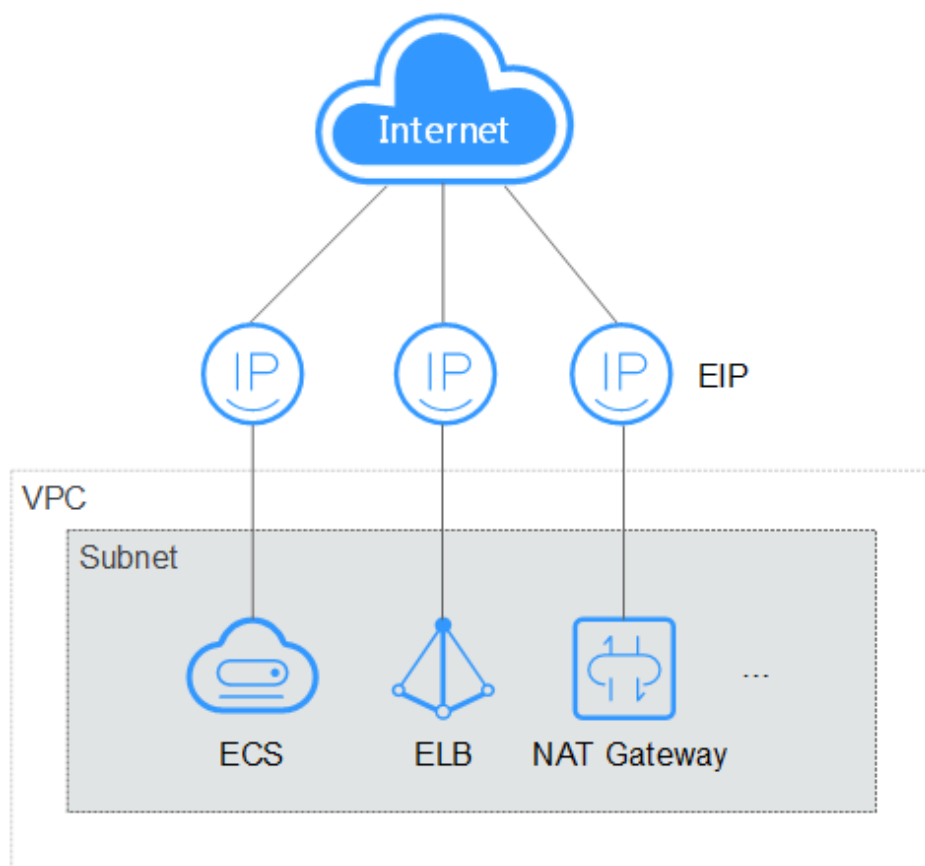
7 EIPs

7.1 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 7-1 Accessing the Internet using an EIP

Helpful Links



- [Binding an EIP](#)
- [Unbinding an EIP](#)
- [Modifying an EIP Bandwidth](#)
- [EIP FAQ](#)
- [Internet Inaccessible](#)

7.2 Binding an EIP

Scenarios

You can assign an EIP and bind it to an ECS so that the ECS can access the public network.

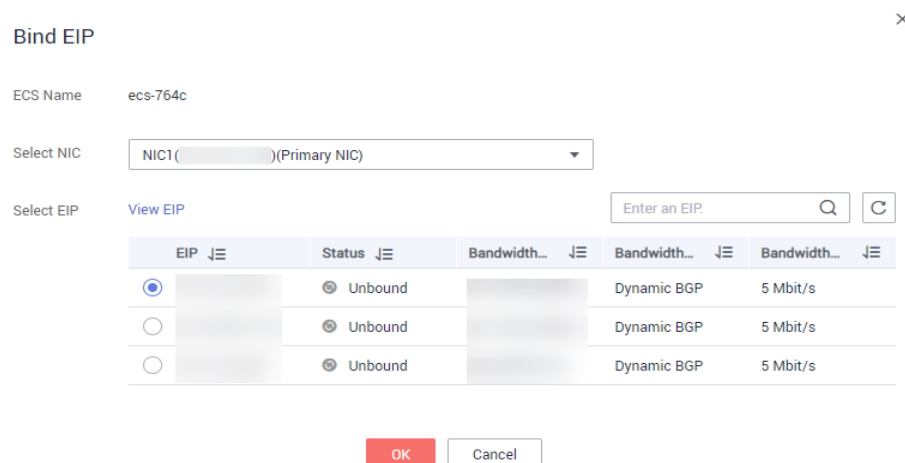
Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select a 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.

Figure 7-2 Binding an EIP

6. Click **OK**.
After the EIP is bound, view it in the ECS list on the **Elastic Cloud Server** page.

7.3 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 a 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**.

7.4 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 > 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** > **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.

Figure 7-3 Binding a new EIP

Bind EIP

ECS Name: ecs-764c

Select NIC: NIC1 (Primary NIC)

Select EIP: [View EIP](#)

| EIP | Status | Bandwidth... | Bandwidth... | Bandwidth... |
|----------------------------------|---------|--------------|--------------|--------------|
| <input checked="" type="radio"/> | Unbound | | Dynamic BGP | 5 Mbit/s |
| <input type="radio"/> | Unbound | | Dynamic BGP | 5 Mbit/s |
| <input type="radio"/> | Unbound | | Dynamic BGP | 5 Mbit/s |

7.5 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.

7.6 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 choose from the following solutions:

- **(Recommended) Using NAT Gateway**

A NAT gateway provides SNAT and DNAT to enable multiple ECSs in the same VPC to share an EIP to access the Internet or provide services for the Internet.

- **Using a Linux Proxy ECS**

If you already have an ECS that has an EIP bound, you can use the ECS as the proxy to provide an Internet access channel for other ECSs that are in the same network segment, in the same security group, and have no EIPs bound.

(Recommended) Using NAT Gateway

NOTE

Prerequisites for this solution:

- You have purchased an EIP and a public NAT gateway.
- The public NAT gateway and the ECSs that need to access the Internet are in the same VPC.

Table 7-1 NAT gateway operations

| Scenario | Gateway Rule | Reference |
|---|--------------|---|
| Enabling ECSs without EIPs bound to access the Internet | SNAT | Using a Public NAT Gateway to Enable Servers to Share One or More EIPs to Access the Internet |
| Enabling ECSs without EIPs bound to provide services for the Internet | DNAT | Using a Public NAT Gateway to Enable Servers to Be Accessed by 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.

For other OSs and versions, see the help documentation of the corresponding official website or [\(Recommended\) Using NAT Gateway](#).




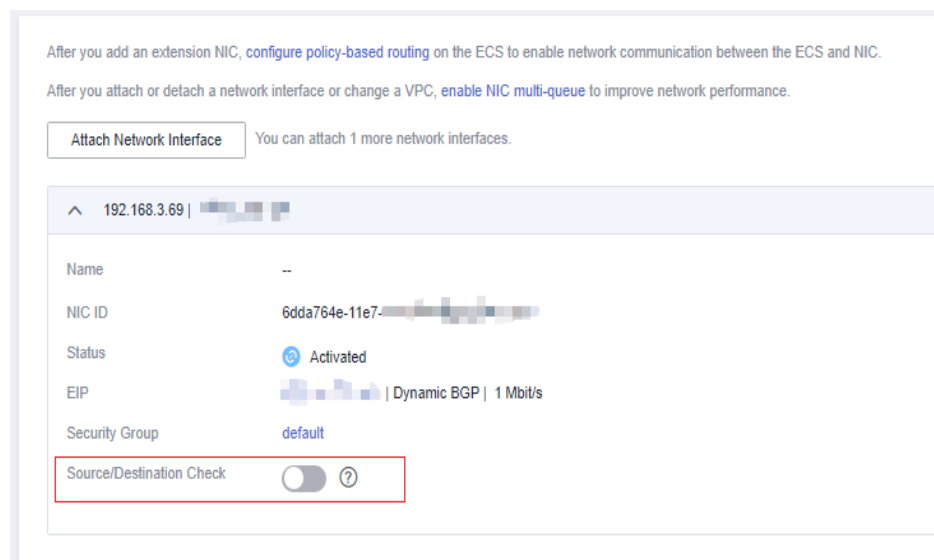
1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
5. Click the name of the proxy ECS. The page providing details about the ECS is displayed.
6. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

Figure 7-4 Disabling source/destination check



By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. Therefore, disable the source/destination check.

7. Log in to the proxy ECS.


For more details, see [Login Overview \(Linux\)](#).

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 7-5 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 7-6](#) is displayed.

Figure 7-6 Successful SNAT configuration



20. Add a route.

- a. Log in to the management console.
- b. Click  in the upper left corner and select a region and project.
- c. Under **Networking**, click **Virtual Private Cloud**.
- d. Choose **Route Tables** in the left navigation pane. In the route table list, click a target route table. On the displayed page, click **Add Route**.
- e. Set route information on the displayed page.
 - **Destination:** indicates the destination network segment. The default value is **0.0.0.0/0**.
 - **Next Hop:** indicates the private IP address of the proxy ECS.
You can obtain the private IP address of the ECS on the **Elastic Cloud Server** page.

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
```

8 Security

8.1 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 8-1 Methods for improving ECS security

| Type | Description | Protection Method |
|-------------------|---|---|
| External security | DDoS attacks and Trojan horses or other viruses are common external security issues. To address these issues, you can choose services such as Host Security Service (HSS) based on your service requirements: | <ul style="list-style-type: none">• Enabling HSS• Configuring Cloud Bastion Host• 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. | <ul style="list-style-type: none">• 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.

- **Scenario 1: An ECS is to be created.**

When you use certain public images to create ECSs, you are advised to use HSS to protect your ECSs.

Select one of the following options:

- **Basic edition (one-month free trial):** After this function is enabled, the HSS basic edition can be used free of charge for 30 days. The HSS basic edition supports detection of OS vulnerabilities, weak passwords, and brute force cracking to improve the overall security for your ECSs.

 **NOTE**

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

If you want to retain or upgrade HSS security capabilities, you are advised to enable the advanced HSS edition. For details, see [What Should I Do When the Free Trial of HSS Basic Edition Expires?](#)

This option is selected by default.


- **Advanced HSS edition (paid):** You can choose from HSS basic, professional, enterprise, premium, and Web Temper Protection (WTP) editions and you need to pay for it.

After ECSs are purchased, you can switch between different editions on the HSS console after **Advanced HSS edition (paid)** is enabled. For details about the differences among different editions, see [Specifications of Different Editions](#).

- **None:** HSS is disabled and servers are not protected.

After you select an HSS edition, the system automatically installs the HSS agent, enables account cracking prevention, and offers host security functions. If the basic or enterprise edition does not meet service requirements, you can [purchase an HSS quota](#) and switch the edition on the HSS console to obtain advanced protection without reinstalling the agent.

Figure 8-1 Enabling HSS



- **Scenario 2: An ECS is already created and HSS is not configured for it.**

For an existing ECS without HSS configured, you can manually install an Agent on it.

For details, see [Installing the Agent on Huawei Cloud Servers](#) and [Enabling Protection](#).

Configuring Cloud Bastion Host

CBH can monitor the usage of the CBH system, monitor O&M activities of each managed resource, and identify suspicious O&M actions in real time. This protects resources and data from being accessed or damaged by external or internal users. CBH reports alarms to customers, who can then more easily handle or audit O&M issues in a timely, centralized manner.

After logging in to the CBH instance, you need to create users, resources, and access policies successively, and then you can manage assets and perform O&M operations.

- You can log in to your CBH system through a web browser, MSTSC client, or SSH client.

[Using a Web Browser to Log In to a CBH System](#) supports system management and resource O&M. This method is recommended for system user **admin** or administrators to manage the CBH system and assign O&M permissions.

[Using an SSH Client to Log In to a CBH System](#) enables you to count on your experiences in using SSH clients to manage O&M on authorized resources. You can use an SSH client to directly log in to the CBH system for resource O&M.

[Using the MSTSC Client to Log In to a CBH System](#) enables you to count on your experience in using MSTSC to manage O&M on authorized resources. You can use an MSTSC client to directly log in to the CBH system for resource O&M.

- Before using CBH for system management and O&M, administrators need to [create a system user](#) and assign system roles to them.

System users then can access the modules within the permissions.

 NOTE

Only the **admin** user has the permissions to manage system roles.

- A CBH system allows you to centrally manage cloud resources as well as their accounts and permissions. To manage resources centrally, you need to [add resources to your CBH system](#) first.

A host or application resource may have multiple accounts for logins.

CBH allows you to log in to managed resources through managed accounts without having to repeatedly enter the usernames and passwords.

The default account for each managed resource is **Empty**. If you use the **Empty** account, enter the account username and password for accessing the host resource.

- To use CBH to maintain resources, [configure access control policies](#), associate system users with resources, and assign resource permissions to system users.
- After obtaining the resource access control permission, you can log in to resources and [perform O&M](#) through the system. The entire O&M process is monitored and recorded.

You can select different login methods based on resource types.

- You can log in to many types of managed resources, including databases, within the granted permissions for further O&M in the CBH system.

Administrators can audit logins and operations of other system users, and [audit their O&M sessions](#) on the system web page.

Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

Server monitoring includes basic monitoring, OS monitoring, and process monitoring for servers.

- Basic monitoring

Basic monitoring does not require the agent to be installed and automatically reports ECS metrics to Cloud Eye. Basic monitoring for KVM ECSs is performed every 5 minutes.

- OS monitoring

By installing the Agent on an ECS, OS monitoring provides system-wide, active, and fine-grained monitoring. OS monitoring for KVM ECSs is performed every minute.

To enable OS monitoring when purchasing an ECS:

Select **Enable Detailed Monitoring** when purchasing an ECS. After this option is selected, the cloud platform automatically installs the agent required for OS monitoring.

NOTE

Currently, you can enable OS monitoring only when you purchase ECSs running specific OSs in specific regions.

Figure 8-2 Enabling OS monitoring when purchasing an ECS**To enable OS monitoring for a created ECS:**

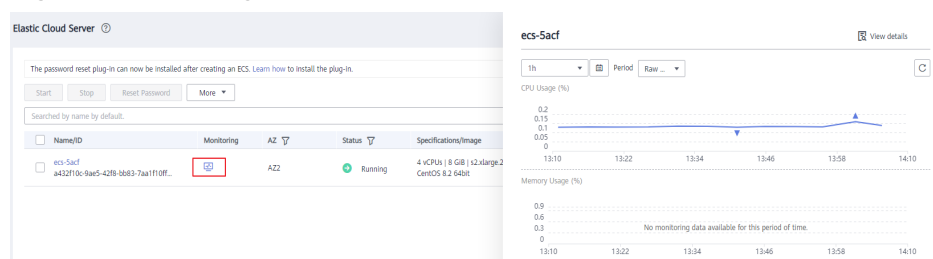
You need to manually install the agent if **Enable Detailed Monitoring** is not selected during the creation.

For instructions about how to install and configure the Agent, see [Agent Installation and Configuration](#).

- Process monitoring
Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. Processes are monitored at an interval of 1 minute (for KVM ECSs).

After server monitoring is enabled, you can set ECS alarm rules to customize the monitored objects and notification policies and learn about the ECS running status at any time.

On the ECS console, click  to view monitoring metrics.

Figure 8-3 Viewing ECS metrics

Enabling Anti-DDoS

To defend against DDoS attacks, Huawei Cloud provides multiple security solutions. You can select an appropriate one based on your service requirements. Anti-DDoS Service on Huawei Cloud provides three sub-services: Cloud Native Anti-DDoS (CNAD) Basic (also known as Anti-DDoS), CNAD Pro, and Advanced Anti-DDoS (AAD).

Anti-DDoS is free while CNAD Pro and AAD are paid services.

For details about CNAD Pro and AAD, see [What Is Anti-DDoS?](#)

If you choose to purchase an EIP when purchasing an ECS, the console will display a message indicating that you have enabled free-of-charge Anti-DDoS protection.

Figure 8-4 Enabling anti-DDoS protection

EIP ☒ Auto assign ☐ Use existing ☐ Not required ?

EIP Type **Dynamic BGP** Premium BGP

Greater than or equal to 99.95% service availability rate

Billed By **Bandwidth** Traffic Shared bandwidth

Bandwidth Size 5 10 20 50 100 Custom 1 + The bandwidth can be from 1 to 300 Mbit/s.

Free 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.


To enable CBR when purchasing an ECS:

Set CBR when purchasing an ECS. The system will associate the ECS with a cloud backup vault and the selected backup policy to periodically back up the ECS.

- Create new
 - a. Enter the name of the cloud backup vault. The name consists of 1 to 64 characters. Only letters, digits, underscores (_), and hyphens (-) are allowed. For example, **vault-f61e**. The default naming rule is **vault_XXXX**.
 - b. Enter the vault capacity, which is required for backing up the ECS. The vault capacity cannot be smaller than that of the ECS to be backed up. Its value ranges from the total capacity of the ECS to 10,485,760 in the unit of GB.
 - c. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Use existing
 - a. Select an existing cloud backup vault from the drop-down list.
 - b. Select a backup policy from the drop-down list, or log in to the CBR console and configure a desired one.
- Not required

Skip this configuration if CBR is not required. If you need to enable CBR after purchasing an ECS, log in to the CBR console, locate the target vault, and associate it with the ECS.

Figure 8-5 Setting CBR



To back up data for a created ECS:

You can use the cloud server backup and cloud disk backup to [back up your ECS data](#).

- Cloud server backup (recommended): Use this backup method if you want to back up the data of all EVS disks (system and data disks) attached to an ECS. This prevents data inconsistency caused by the time difference in creating a backup.
- Cloud disk backup: Use this backup method if you want to back up the data of one or more EVS disks (system or data disk) attached to an ECS. This minimizes backup costs on the basis of data security.

Enhancing the Login Password Strength

Key pair authentication is recommended because it is more secure than password-based authentication.

If you select **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 8-2 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 8-2 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](#).

8.2 Security Groups

8.2.1 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 8-3](#) 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 8-3 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 |

- If you specify an IP address group or inconsecutive ports for a security group rule, the rule is only applied for certain ECSs. For details, see [Table 8-4](#).

Table 8-4 Scenarios that security group rules do not take effect

| Rule Configuration | ECS Type |
|--|--|
| Source or Destination is set to IP address group . | The following x86 ECS types are not supported: <ul style="list-style-type: none">• General computing (S1, C1, and C2 ECSs)• Memory-optimized (M1 ECSs)• High-performance computing (H1 ECSs)• Disk-intensive (D1 ECSs)• GPU-accelerated (G1 and G2 ECSs)• Large-memory (E1, E2, and ET2 ECSs) |
| Port is set to non-consecutive ports. | The following x86 ECS types are not supported: <ul style="list-style-type: none">• General computing (S1, C1, and C2 ECSs)• Memory-optimized (M1 ECSs)• High-performance computing (H1 ECSs)• Disk-intensive (D1 ECSs)• GPU-accelerated (G1 and G2 ECSs)• Large-memory (E1, E2, and ET2 ECSs) |
| | All Kunpeng ECS flavors do not support inconsecutive ports. <p>If you use inconsecutive port numbers in a security group rule of a Kunpeng ECS, this rule and rules configured after this one do not take effect.</p> <p>If you configure security group rule A with inconsecutive ports 22,24 and then configure security group rule B with port 9096, both rule A and rule B do not take effect.</p> |

 **NOTE**

- For details about x86 ECSs, see [ECS Specifications \(x86\)](#).
- For details about Kunpeng ECSs, see [ECS Specifications \(Kunpeng\)](#).
- **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?](#)

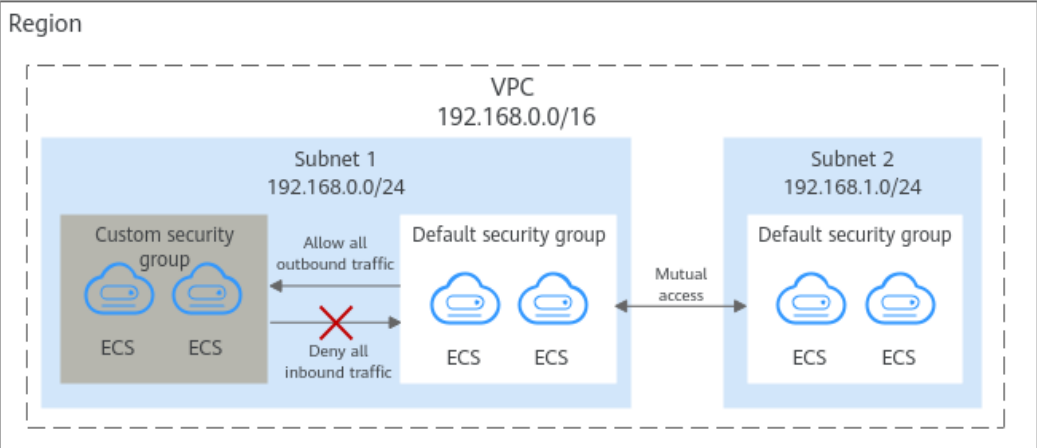
8.2.2 Default Security Groups and Rules

Note the following when using default security group rules:

- **Inbound rules** control incoming traffic to instances in the default security group. The instances can 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 8-6](#) shows the default security group.

Figure 8-6 Default security group



NOTE

- Both default and custom security groups are free of charge. The name of a default security group is **default**.
- You cannot delete the default security group, but you can modify existing rules or add rules to the group.
- The default security group is automatically created to simplify the process of creating an instance for the first time. The default security group denies all external requests. To log in to an instance, add a security group rule by referring to [Remotely Logging In to an ECS from a Local Server](#).

Table 8-5 describes the rules in the default security group.

Table 8-5 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 8-6 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 8-7 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. |

| Direction | Action | Type | Protocol & Port | Source/Destination | Description |
|-----------|--------|------|-----------------|--|--|
| 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. |
| 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 8-8 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. |

| Direction | Action | Type | Protocol & Port | Source/Destination | Description |
|-----------|--------|------|-----------------|---|--|
| 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. |
| 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. |

8.2.3 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.

- **If the source is set to 0.0.0.0/0 or ::/0, then the access from all external IP addresses are either allowed or denied, depending on if the action is Allow or Deny.** If the access is allowed, exposing **high-risk ports**, such as port 22, 3389, or 8848, to the public network will leave your instances vulnerable to network intrusions, causing service interruptions, data leakage, or ransomware attacks. You should only configure known IP addresses for the security group rule.
- **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 8-9](#).

Table 8-9 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 8-10](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 8-11](#).

Table 8-10 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 8-11 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 |

⚠ CAUTION

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 8-12](#).

Table 8-12 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 8-13 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 |

CAUTION

- 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 8-14](#).
- You must first install the FTP server program on the ECSs and then check whether ports 20 and 21 are working properly.

Table 8-14 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 8-15 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 8-16 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 8-17 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

In this example, the source is for reference only. Set the source based on actual requirements.

Table 8-18 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. |

| Direction | Priority | Action | Type | Protocol & Port | Source | Description |
|-----------|----------|--------|------|-----------------|-----------------------------|---|
| 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. |
| 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. |

Allowing ECSs to Access Only Specific External Websites

By default, a security group allows all outbound traffic. [Table 8-20](#) 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 8-19 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 8-20 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. |

8.2.4 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 [ECS console](#).
2. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
3. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
4. Click the security group ID.
The system automatically switches to the security group details page.
5. Configure inbound rule parameters as prompted.
You can click  to add more inbound rules.

Figure 8-7 Add Inbound Rule

Add Inbound Rule [Learn more about security group configuration.](#)

Some security group rules may not be applied for cloud servers with certain specifications. [Learn more](#)

If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,), vertical bars (|), or spaces. Each IP address represents a different security group rule.

If the source is set to 0.0.0.0/0 or ::0, then all external IP addresses are either allowed or denied to access your instances, depending on if the action is Allow or Deny. If the access is allowed, exposing high-risk ports, such as port 22, 3389, or 8848, to the public network will leave your instances vulnerable to network intrusions, service interruptions, data leakage, or ransomware attacks. You should only configure known IP addresses for the security group rule.

Security Groupsg-default-smb

You can import multiple rules in a batch.

| Priority | Action | Type | Protocol & Port | Source | Description | Operation |
|----------|--------|------|--------------------------|----------------|-------------|--|
| 1-100 | Allow | IPv4 | Protocols / TCP (Cust... | IP address | | Replicate Delete |
| | | | 22 | 192.168.0.0/24 | | |

+ Add Rule

Cancel

OK


Table 8-21 Inbound rule parameter description

| Param eter | 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 . <ul style="list-style-type: none">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: <ul style="list-style-type: none">IPv4IPv6 | IPv4 |
| Protoc ol & 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 |
|-----------|--|---------------|
| | <p>Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.</p> <p>Inbound rules control incoming traffic over specific ports to instances in the security group.</p> <p>Enter ports in any of the following formats:</p> <ul style="list-style-type: none">• Individual port: Enter a port, such as 22.• Consecutive ports: Enter a port range, such as 22-30.• Non-consecutive ports: Enter ports and port ranges, such as 22,23-30. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.• All ports: Leave it empty or enter 1-65535. | 22, 22-30 |

| Parameter | Description | Example Value |
|-------------|--|--|
| Source | <p>Used to match the source of an external request. The source can be:</p> <ul style="list-style-type: none">• IP address: You can enter multiple IP addresses, separated by commas (,). Each IP address defines a different security group rule.<ul style="list-style-type: none">– Single IP address: IP address/mask
Example IPv4 address: 192.168.10.10/32
Example IPv6 address: 2002:50::44/128– IP address range in CIDR notation: IP address/mask
Example IPv4 address range: 192.168.52.0/24
Example IPv6 address range: 2407:c080:802:469::/64– Any IP addresses
0.0.0.0/0 represents any IPv4 addresses.
::/0 represents any IPv6 address.• Security group: The source is from another security group. You can select a security group in the same region from the drop-down list. If there is instance A in security group A and instance B in security group B, and the inbound rule of security group A allows traffic from security group B, traffic is allowed from instance B to instance A.• IP address group: An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in an easier way. If no IP address groups are available, create one by referring to Creating an IP Address Group. | IP address:
192.168.52.0/24,10.0.0.0/24 |
| Description | <p>(Optional) Supplementary information about the security group rule.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).</p> | N/A |

6. Configure outbound rule parameters as prompted.

You can click  to add more outbound rules.

Add Outbound Rule [Learn more about security group configuration.](#)

1

Some security group rules may not be applied for cloud servers with certain specifications. [Learn more](#)

If you select IP address for Destination, you can enter multiple IP addresses, separated with commas (,), vertical bars (|), or spaces. Each IP address represents a different security group rule.

Security Group **sg-default-smb**

You can [import multiple rules in a batch.](#)

| Priority | Action | Type | Protocol & Port | Destination | Description | Operation |
|----------|--------|------|---|-------------------------|-------------|--|
| 1-100 | Allow | IPv4 | Protocols / TCP (Cust...
Example: 22 or 22,24 or 22-30 | IP address
0.0.0.0/0 | | Replicate Delete |

[+ Add Rule](#)

Cancel

OK

| Parameter | Description | Example Value |
|-----------------|--|---------------|
| Priority | <p>The security group rule priority.</p> <p>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.</p> | 1 |
| Action | <p>The value can be Allow or Deny.</p> <ul style="list-style-type: none"> 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. <p>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.</p> | Allow |
| Type | <p>Destination IP address version. You can select:</p> <ul style="list-style-type: none"> IPv4 IPv6 | IPv4 |
| Protocol & Port | <p>The network protocol used to match traffic in a security group rule. The protocol can be All, TCP, UDP, GRE, or ICMP.</p> | TCP |

| Parameter | Description | Example Value |
|-----------|---|---------------|
| | <p>Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.</p> <p>Outbound rules control outgoing traffic over specific ports from instances in the security group.</p> <p>Enter ports in any of the following formats:</p> <ul style="list-style-type: none">• Individual port: Enter a port, such as 22.• Consecutive ports: Enter a port range, such as 22-30.• Non-consecutive ports: Enter ports and port ranges, such as 22,23-30. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.• All ports: Leave it empty or enter 1-65535. | 22, 22-30 |

| Parameter | Description | Example Value |
|-------------|---|--|
| Destination | <p>Used to match the destination of an internal request. The destination can be:</p> <ul style="list-style-type: none">• IP address: You can enter multiple IP addresses, separated by commas (,). Each IP address defines a different security group rule.<ul style="list-style-type: none">– Single IP address: IP address/mask
Example IPv4 address: 192.168.10.10/32
Example IPv6 address: 2002:50::44/128– IP address range in CIDR notation: IP address/mask
Example IPv4 address range: 192.168.52.0/24
Example IPv6 address range: 2407:c080:802:469::/64– Any IP addresses
0.0.0.0/0 represents any IPv4 addresses.
::/0 represents any IPv6 address.• Security group: The destination is another security group. You can select a security group in the same region under the current account from the drop-down list. If there is instance A in security group A and instance B in security group B, and the outbound rule of security group A allows traffic to security group B, traffic is allowed from instance A to instance B.• IP address group: An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in an easier way. If no IP address groups are available, create one by referring to Creating an IP Address Group. | IP address:
192.168.52.0/24,10.0.0.0/24 |
| Description | <p>(Optional) Supplementary information about the security group rule.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).</p> | N/A |

7. 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 8-23](#) shows the rule.

Table 8-23 Security group rule

| Direction | Priority | Action | Type | Protocol & Port | Source |
|-----------|----------|--------|------|-----------------|-----------------------|
| Inbound | 1 | Allow | IPv4 | 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 8-9 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**
 - i. Choose **Start > Run**. Type **cmd** to open the Command Prompt.
 - ii. 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 8-10 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.

8.2.5 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.



Figure 8-11 Change Security Group



Change Security Group

ECS Name ecs-e498

NIC (primary)

Security Group Enter a security group name.   [Create Security Group](#)

| Security Group Name | Description |
|---|------------------------|
| <input checked="" type="checkbox"/> default | Default security group |
| <input type="checkbox"/> sg-0228 | |

Selected security groups: default

OK Cancel

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**.

8.3 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. The installation method varies depending on whether your ECS is to be created or already exists.

- **Scenario 1: An ECS is to be created.**

When you use certain public images to create ECSs, you are advised to use HSS to protect your ECSs.

Select one of the following options:

- **Basic edition (one-month free trial):** After this function is enabled, the HSS basic edition can be used free of charge for 30 days. The HSS basic edition supports detection of OS vulnerabilities, weak passwords, and brute force cracking to improve the overall security for your ECSs.

 **NOTE**

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

If you want to retain or upgrade HSS security capabilities, you are advised to enable the advanced HSS edition. For details, see [What Should I Do When the Free Trial of HSS Basic Edition Expires?](#)

This option is selected by default.

- **Advanced HSS edition (paid):** You can choose from HSS basic, professional, enterprise, premium, and Web Temper Protection (WTP) editions and you need to pay for it.


After ECSs are purchased, you can switch between different editions on the HSS console after **Advanced HSS edition (paid)** is enabled. For details about the differences among different editions, see [Specifications of Different Editions](#).

- **None:** HSS is disabled and servers are not protected.

After you select an HSS edition, the system automatically installs the HSS agent, enables account cracking prevention, and offers host security functions.

If the basic or enterprise edition does not meet service requirements, you can [purchase an HSS quota](#) and switch the edition on the HSS console to obtain advanced protection without reinstalling the agent.

Figure 8-12 Enabling HSS



The screenshot shows a configuration interface for HSS. At the top, there are four tabs: 'Public image', 'Private image', 'Shared image', and 'Marketplace image'. Below these tabs, there are two dropdown menus. The first dropdown menu is labeled 'CentOS' and the second is labeled 'CentOS 7.5 64bit(40GB)'. Below these dropdowns, there is a section labeled 'Protection' with three radio button options: 'HSS basic edition (free)', 'Advanced HSS edition (paid)', and 'None'. The 'HSS basic edition (free)' option is selected. A red box highlights the 'Protection' section.

- **Scenario 2: An ECS is already created and HSS is not configured for it.**

For an existing ECS without HSS configured, you can manually install an Agent on it.

For details, see [Installing 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 & Compliance > Host Security Service**.
3. Choose **Asset Management > Servers & Quota** and go to the **Servers** tab to view the protection status of the target servers.

Figure 8-13 ECS security statuses




Table 8-24 Statuses

| Parameter | Description |
|-------------------|---|
| Agent Status | <ul style="list-style-type: none">● 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 | <ul style="list-style-type: none">● 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 | <ul style="list-style-type: none">● 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?](#)

8.4 CBH

What Is CBH?

Cloud Bastion Host (CBH) is a unified security management and control platform. It provides account, authorization, authentication, and audit management services that enable you to centrally manage cloud computing resources.

CBH provides various functional modules, such as department, user, resource, policy, operation, and audit modules. It integrates functions such as single sign-on (SSO), unified asset management, multi-terminal access protocols, file transfer, and session collaboration. With the unified O&M login portal, protocol-based forward proxy, and remote access isolation technologies, CBH enables centralized, simplified, secure management and maintenance auditing for cloud resources such as servers, cloud hosts, databases, and application systems.

How to Configure a CBH Instance Quickly

CBH can monitor the usage of the CBH system, monitor O&M activities of each managed resource, and identify suspicious O&M actions in real time. This protects resources and data from being accessed or damaged by external or internal users. CBH reports alarms to customers, who can then more easily handle or audit O&M issues in a timely, centralized manner. To do all these, you only need to configure your CBH instance first.


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

Step 2 On the management console, choose **Security and Compliance > Cloud Bastion Host**. In the upper right corner of the page, click **Buy CBH Instance**. When your CBH instance is ready, click **Remote Login** in the **Operation** column to go to the CBH system login page.

NOTE

- When the first time you log in to a CBH system as user **admin**, enter the login password you configure when you purchase the corresponding CBH instance. System administrator **admin** is the default user. It is the first account that can be used to log in to a CBH system and has the highest operation permissions. Its permissions cannot be deleted or changed.
- After logging in to a CBH system for the first time, all users need to change the password and bind the mobile number as prompted.


Figure 8-14 Logging in to a CBH instance



Step 3 After logging in to a CBH system, choose **User > User**. In the upper right corner of the displayed page, click **New**. In the displayed dialog box, **create a user**.

 NOTE

- By default, there are four roles: system administrator, policy administrator, audit administrator, and O&M personnel. The **admin** user can [create a custom role](#) to assign system operation permissions.
- You need to set **LoginName** to a unique name in the CBH system.
- After a user is created, you can enable [multi-factor authentication](#) for the user to log in to the CBH system.

Figure 8-15 New User**New User**

The dialog box titled "New User" contains the following fields and instructions:

- * LoginName**: A text input field. Below it, the text reads: "The value contains 1 to 64 characters and must start with a letter or digit. The following characters are not supported: / \ [] ; : | = , + ' ' ? < > @ * and Spaces".
- * Verification Type**: A dropdown menu with "Local" selected.
- * Password**: A text input field with a password icon (eye with a slash) on the right.
- * Confirm Password**: A text input field with a password icon (eye with a slash) on the right. Below it, the text reads: "The password is 8-32 characters long and must contain at least four of the following character types: uppercase letters, lowercase letters, digits, and special characters (!@#\$%^&*_+=[]{};.,/?~#*). It cannot contain the username or the username spelled backwards".
- * UserName**: A text input field. Below it, the text reads: "1-255 length of characters, allowed characters including letter, digit, '@', '.', '_', or '-'".

At the bottom of the dialog box are two buttons: "OK" (in a red box) and "Cancel" (in a white box with a gray border).

Step 4 After creating a user, add a host resource. To do so, choose **Resource > Host**. On the displayed page, click **New** in the upper right corner. In the dialog box displayed, [complete basic settings and network settings](#).

 NOTE

- **Host Address** indicates the IP address used for communication between the host and CBH instance. You can select the EIP or private IP address assigned to the host. You are advised to select an available private IP address.
- You can use enhanced editions to manage databases in the **Host** module. Currently, four types of databases are supported: MySQL, SQL Server, Oracle, and DB2.
- Application resources are managed through the Windows remote access function. You need to [configure an application server](#) first.
- After a resource is added to CBH, you still need to add a resource account to log in to the resource O&M system. You can use any of the following login modes:
Automatic login: You use CBH to manage resource account usernames and passwords. In this mode, you do not need to enter the username and password for logging in to a specific resource.
Manual login: The **Empty** account is automatically generated when adding resources to CBH. When logging in to a resource, you need to enter the account username and password.
Sudo login: When a user logs in to a specific resource as a sudoer, the user is automatically switched to a privileged account.

Figure 8-16 New Host

New Host

* Host Name
1-128 length of characters

* Protocol

* Host Address
IP address or domain name

* Port
Digits of 1-65535

OS Type

Options

| | |
|--|--|
| <input checked="" type="checkbox"/> File Manage | <input checked="" type="checkbox"/> X11 forward |
| <input checked="" type="checkbox"/> Uplink Clipboard | <input checked="" type="checkbox"/> Downlink Clipboard |
| <input checked="" type="checkbox"/> Keyboard Audit | |

* Department

Step 5 Click **Next**, configure host account information, and click **OK**.

Step 6 Choose **Policy > ACL Rules** and click **New** in the upper right corner. In the displayed dialog box, [configure an access control rule](#).

 NOTE

- Access control rules are used to associate users with resources by granting specific permissions for certain resources to a specific user. CBH system users can operate and maintain resources only after being authorized.
- **IP Limit** is used to set the local IP address of a user to restrict or allow the user from the IP address to access resources.

Figure 8-17 New ACL Rule**New ACL rule**

* Rule Name
1-64 length of characters, including letters, digit or "-"

Period of validity

File Transmission ☒ Upload ☒ Download

Options ☒ File Manage ☒ Uplink Clipboard ☒ Downlink Clipboard
☐ Watermark ☒ Keyboard Audit

Logon Time Limit ☒ Permit ☐ Forbid

| | | | | | | | | | | | | | | | | | | | | | | | | |
|-----|---|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| Mon | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 |
| Tue | | | | | | | | | | | | | | | | | | | | | | | | |
| Wed | | | | | | | | | | | | | | | | | | | | | | | | |
| Thu | | | | | | | | | | | | | | | | | | | | | | | | |
| Fri | | | | | | | | | | | | | | | | | | | | | | | | |
| Sat | | | | | | | | | | | | | | | | | | | | | | | | |
| Sun | | | | | | | | | | | | | | | | | | | | | | | | |

ID Limit

Step 7 Click **Next**, associate the user with the host resource, and click **OK**.

Step 8 Log in to the CBH system using the created user and choose **Operation > Host Operation**.

Step 9 Select the target host resource and click **Login** and perform O&M as needed. For details, see [Logging In to Host Resources](#).

NOTE

- For host resources with SSH, Telnet, or Rlogin protocol configured, you can use an SSH client for O&M.
- For host resources with FTP, SFTP, or SCP protocol configured, you need to use the FTP, SFTP, or SCP client for O&M, respectively.
- For MySQL, SQL Server, Oracle, and DB2 host resources, you need to configure an SSO tool and database management tool first. Then you can use the SSO tool to call the database client and implement resource O&M.
- For host resources with SSH, RDP, VNC, or Telnet protocol configured, you can use web browsers for O&M. For application resources, you can use only web browsers for remote access and O&M.

----End

8.5 Configuring Secure Boot for an ECS

Scenarios

Secure boot is used to check the integrity of each component (driver loader, kernel, and kernel driver) during device startup. This prevents security threats to the system and user data caused by loading and running unauthenticated components.

Secure boot is a feature of the Unified Extensible Firmware Interface (UEFI) that requires all low-level firmware and software components to be verified before

being loaded. During the boot process, UEFI secure boot checks the signature of each boot software, including the UEFI firmware driver (also called ROM option), Extensible Firmware Interface (EFI) application, OS driver, and binary files. If the signature is valid or recognized by the original equipment manufacturer (OEM), the device will boot, and the firmware will hand over control to the OS.

Secure boot uses key pairs to sign and verify boot devices, ensuring that ECSs only boot software signed by encryption keys. This protects software from threats during the boot process. The secure boot key is stored in the key database of the UEFI non-volatile variable storage.

- After secure boot is enabled, the system verifies component integrity during system startup, which does not affect services.
- After secure boot is disabled, the system does not verify component integrity during system startup, which does not affect services.

Background

UEFI is a standard that describes a new type of interface. It enables an OS to automatically load from a pre-boot operating environment to an OS.

The OSs that use the UEFI boot mode support secure boot. Secure boot provides verification of the boot chain status to ensure that only encrypted and verified UEFI binary files are executed after the firmware is initialized. These binary files include UEFI drivers, primary boot loaders, and chain loading components.

By default, secure boot is disabled and the system is in SetupMode. When the system is in SetupMode, all key variables can be updated without a cryptographic signature. After secure boot is enabled, the system exits SetupMode.

NOTICE

After secure boot is enabled, the system enforces signature validation on any UEFI binary files.

Secure boot uses key databases in a chain of trust, as shown in [Table 8-25](#). These databases are stored in the UEFI variable storage.

Table 8-25 Key databases

| Key Data base | Abbreviation | Mandatory | Description | Format Constraints |
|----------------------------|--------------|-----------|--|--|
| Platform key data base | PK | Yes | <p>The PK database is the root of trust for the secure boot instance.</p> <p>The PK database contains a public PK key, which is used in the chain of trust to update the key exchange key (KEK) database.</p> <p>To change the PK database, you must have the private PK key to sign an update request. This includes deleting the PK database by writing an empty PK key.</p> | <ul style="list-style-type: none">• DER, ESL, and AUTH certificate formats are supported.• Only one certificate is supported.• Only X.509 signatures are supported.• Certificate digest is not supported. |
| Key exchange key data base | KEK | Yes | <p>The KEK database is a list of public KEK keys used in the chain of trust to update the signature (DB) and denylist (DBX) databases.</p> <p>The private KEK is used to add keys to the DB, which is a list of authorized signatures to be booted on the system.</p> <p>To change the public KEK database, you must have the private PK key to sign an update request.</p> | <ul style="list-style-type: none">• DER, ESL, and AUTH certificate formats are supported.• Only X.509 signatures are supported for ESL and AUTH files.• Certificate digest is not supported. |

| Key Data base | Abbreviation | Mandatory | Description | Format Constraints |
|--------------------------------|--------------|-----------|---|---|
| Signature data base | DB | Yes | <p>The DB database is a list of public keys and hash values used in the chain of trust to verify all UEFI boot binary files.</p> <p>The DB list contains authorized keys that are authorized to boot on the system. To modify the list, you must have a private KEK.</p> <p>To change the DB database, you must have either the private PK key or any of the private KEK keys to sign an update request.</p> | <ul style="list-style-type: none">• DER, ESL, and AUTH certificate formats are supported.• Only X.509 signatures are supported for ESL and AUTH files.• Certificate digest is not supported. |
| Forbidden signatures data base | DBX | No | <p>The DBX database is a list of untrusted public keys and binary hashes used to revoke files in the chain of trust.</p> <p>The DBX database always takes precedence over all other key databases.</p> <p>To change the DBX database, you must have either the private PK key or any of the private KEK keys to sign an update request.</p> <p>You can obtain a publicly available DBX from the UEFI forum. For details, see Unified Extensible Firmware Interface Forum.</p> | <ul style="list-style-type: none">• ESL and AUTH certificate formats are supported.• Only X.509 signatures are supported.• SHA256, SHA284, SHA512, and SM3 certificate digests are supported. |

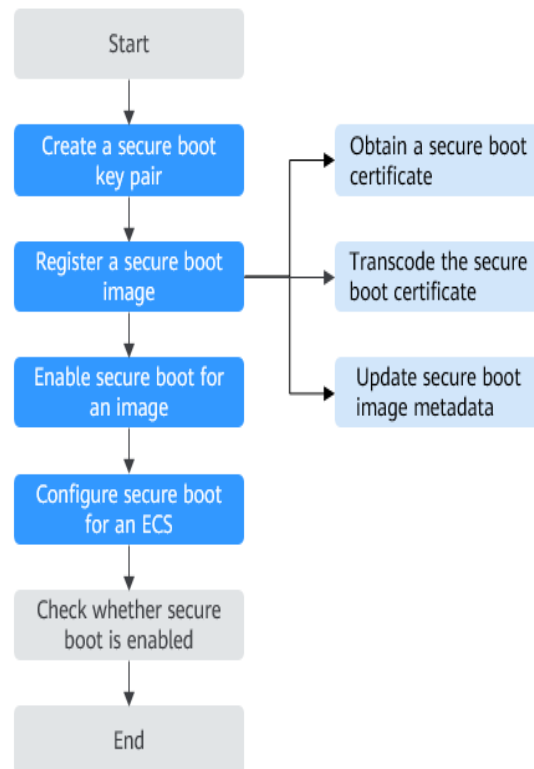
| Key Data base | Abbreviation | Mandatory | Description | Format Constraints |
|-------------------------------|--------------|-----------|---|--|
| Timestamp signatures database | DBT | No | The DBT database retains the signature of the code in the timestamp signature database. | <ul style="list-style-type: none">• DER, ESL, and AUTH certificate formats are supported.• Only X.509 signatures are supported for ESL and AUTH files.• Certificate digest is not supported. |

Constraints

- UEFI secure boot must be supported by instance flavors, including C7t, kC2, and aC8.
- Keys can be imported only when a private image is being created. Keys cannot be imported on the BIOS boot interface of an ECS.
- Keys or related configurations cannot be imported to an ECS using the mokutil tool.
- BMSs do not support UEFI secure boot.
- For details about the OSs that support UEFI boot, see [OSs Supporting UEFI Boot Mode](#).
- For Windows, you need to download the corresponding certificate from the OS's official website. For details, see [Keys Required for Secure Boot on all PCs](#).

Process Flow

Figure 8-18 Configuring secure boot for an ECS



Creating Secure Boot Keys (Linux)

You need to create keys based on the UEFI secure boot standard, create a certificate for each key, convert each certificate into a UEFI signature list (binary file), and sign each certificate using the related key.

Step 1 Create a globally unique identifier (GUID).

Before creating a key pair, create a GUID for key generation.

1. [Connect to an ECS](#).
2. Create a GUID for key generation.
uuidgen --random > GUID.txt

Step 2 (Mandatory) Create a platform key (PK).

A PK is the root of trust for the UEFI secure boot instance.

1. Create a key and name the variable **PK**.
openssl req -newkey rsa:4096 -nodes -keyout PK.key -new -x509 -sha256 -days 3650 -subj "/CN=Platform key/" -out PK.crt
Parameters in this command are described as follows:

- **keyout PK.key**: the private key file
 - **days 3650**: the certificate validity period
 - **out PK.crt**: the certificate for creating a UEFI variable
 - **CN=Platform key**: common name (CN) of the key. You can enter the name of your enterprise instead of the platform key.
2. Create a certificate.
openssl x509 -outform DER -in PK.crt -out PK.cer
 3. Convert the certificate into a UEFI signature list.
cert-to-efi-sig-list -g "\$(< GUID.txt)" PK.crt PK.esl
 4. Use the private PK to sign the UEFI signature list.
sign-efi-sig-list -g "\$(< GUID.txt)" -k PK.key -c PK.crt PK PK.esl PK.auth

Step 3 (Mandatory) Create a key exchange key (KEK).

The private KEK is used to add keys to the signature database (DB), which is the list of authorized signatures to boot on the system.

1. Create a key.
openssl req -newkey rsa:4096 -nodes -keyout KEK.key -new -x509 -sha256 -days 3650 -subj "/CN=Key Exchange Key/" -out KEK.crt
2. Create a certificate.
openssl x509 -outform DER -in KEK.crt -out KEK.cer
3. Convert the certificate into a UEFI signature list.
cert-to-efi-sig-list -g "\$(< GUID.txt)" KEK.crt KEK.esl
4. Use the private PK to sign the signature list.
sign-efi-sig-list -g "\$(< GUID.txt)" -k PK.key -c PK.crt KEK KEK.esl KEK.auth

Step 4 (Mandatory) Create a signature database (DB).

The DB list contains authorized keys that are authorized to boot on the system. To modify the list, you must have a private KEK.

Boot images will be signed with the private key that is created in this step.

1. Create a key.
openssl req -newkey rsa:4096 -nodes -keyout db.key -new -x509 -sha256 -days 3650 -subj "/CN=Signature Database key/" -out db.crt
2. Create a certificate.
openssl x509 -outform DER -in db.crt -out db.cer
3. Convert the certificate into a UEFI signature list.
cert-to-efi-sig-list -g "\$(< GUID.txt)" db.crt db.esl
4. Use the private KEK to sign the signature list.
sign-efi-sig-list -g "\$(< GUID.txt)" -k KEK.key -c KEK.crt db db.esl db.auth

Step 5 (Optional) Create a forbidden signatures database (DBX).

The DBX disables specific DB certificates. It can be a certificate list or a hash list.

Step 6 (Optional) Create a timestamp signatures database (DBT).

The DBT database signs signature timestamps. If the DB certificate is in the DBX and the DBX specifies the revocation time, the DB certificate can be used for signature verification only if the signature time is earlier than the revocation time.

----End

Signing an ECS Image Using a Key

For Huawei Cloud EulerOS (HCE), sign the following images:

/boot/efi/EFI/hce/shimx64.efi

/boot/efi/EFI/hce/mmx64.efi

/boot/efi/EFI/hce/grubx64.efi

/boot/vmlinuz-5.10.0-60.18.0.50.r951_39_66.hce2.x86_64

NOTE

Run **uname -a** to check the kernel version of the image. For example:

```
linux-iXZPNN:~ # uname -a
```

```
Linux linux-iXZPNN 5.10.0-60.18.0.50.r951_39_66.hce2.x86_64 #1 SMP Sun Jul 30 15:05:32  
UTC 2023 x86_64 x86_64 x86_64 GNU/Linux
```

To sign an image

Use the following syntax to sign an image.

```
sbsign --key db.key --cert db.crt --output /boot/  
vmlinuz-5.10.0-60.18.0.50.r951_39_66.hce2.x86_64 /boot/  
vmlinuz-5.10.0-60.18.0.50.r951_39_66.hce2.x86_64
```

Refer to the documentation for your distribution to find out about your boot chain and required images.

¹ Thanks to the ArchWiki community for all of the work they have done. The commands for creating the PK, creating the KEK, creating the DB, and signing the image are from [Creating keys](#), authored by the ArchWiki Maintenance Team and/or the ArchWiki contributors.

Obtaining a Secure Boot Certificate

Contact the image provider to obtain a secure boot certificate.

Transcoding the Secure Boot Certificate

- Platform Key (PK) transcoding
__platform_key=\$(cat PK.der | base64 -w 0)
- Key Exchange Key (KEK) transcoding
__key_exchange_key=\$(cat KEK.der | base64 -w 0)
- Signature Database (DB) transcoding
__signature_database=\$(cat DB.der | base64 -w 0)

- Forbidden Signatures Database (DBX) transcoding
`__forbidden_signatures_database=$(cat DBX.der | base64 -w 0)`
- Timestamp Signatures Database (DBT) transcoding
`__timestamp_signatures_database=$(cat DBT.der | base64 -w 0)`

CAUTION

- **PK.der** is the platform key (PK) certificate.
 - **KEK.der** is the Key Exchange Key (KEK) certificate.
 - **DB.der** is the Signature Database (DB) certificate.
 - **DBX.der** is the Forbidden Signatures Database (DBX) certificate.
 - **DBT.der** is the Timestamp Signatures Database (DBT) certificate.
-

Creating a Private Image

You can create a private image using IMS. A private image can be created using either of the following methods:

- [Creating a Private Image from a Cloud Server or a Backup](#)
- [Creating a Private Image from an Image File](#)

Updating Secure Boot Image Metadata

You can call the following API to update the image metadata:

Updating an Image



```
hw_firmware_type=uefi
__support_secure_boot=true
__platform_key=${__platform_key}
__key_exchange_key=${__key_exchange_key}
__signature_database=${__signature_database}
__forbidden_signatures_database=${__forbidden_signatures_database}
__timestamp_signatures_database=${__timestamp_signatures_database}
```

Enabling Secure Boot for an Image

1. Log in to the management console.
2. Choose **Compute > Image Management Service** to access the IMS console.
3. Click the **Private Images** tab to show the image list.
4. Locate the row containing the target image and click **Modify** in the **Operation** column.
5. In the displayed dialog box, select the UEFI boot mode.
6. After secure boot is enabled, specify the key databases.
7. Click **OK**.

Configuring Secure Boot for an ECS

1. Log in to the management console.

2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Click **Buy ECS**.
5. Set **Trusted Settings** based on service requirements.
 - To enable secure boot, select **Secure boot**.
 - To disable secure boot, deselect **Secure boot**.

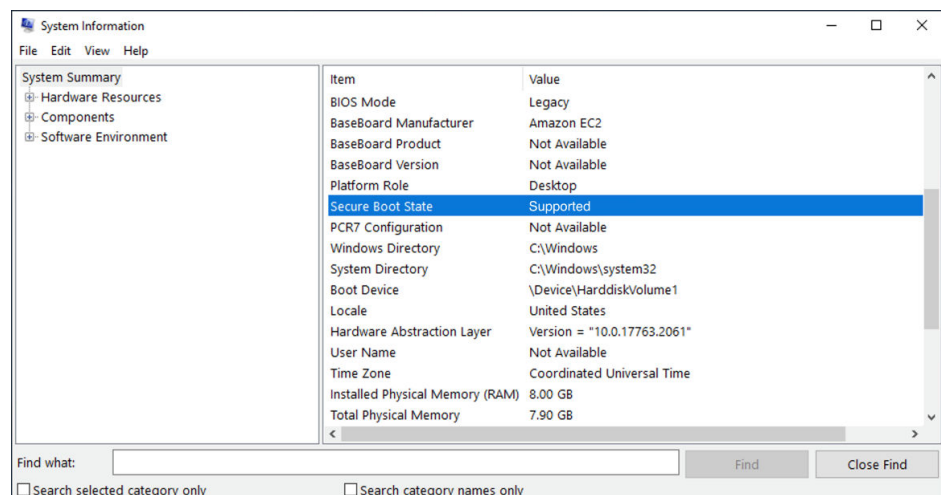
For details about how to configure basic, network, and advanced settings when purchasing an ECS, see [Purchasing an ECS in Custom Config Mode](#).

6. Click **OK**.

Follow-up Operations

- (Windows) Check whether UEFI secure boot is enabled.
 - a. Open the msinfo32 tool.
 - b. Check the **Secure Boot State** field.

If the value of **Secure Boot State** is **Supported**, UEFI secure boot is enabled.



- (Linux) Check whether UEFI secure boot is enabled.
 - a. Run the following command on the ECS:
mokutil --sb-state
 - b. View the command output.
 - If UEFI secure boot is enabled, the command output contains **SecureBoot enabled**.
 - If UEFI secure boot is not enabled, the command output contains **SecureBoot disabled** or **Failed to read SecureBoot**.

8.6 Configuring ECS QingTian TPM

Scenarios

QingTian TPM is a virtual device that complies with the TPM 2.0 specifications. QingTian TPM can be used as the root of trust of an ECS to build a trust chain that covers system boot and user-specified applications and implement remote attestation. In addition, QingTian TPM can securely store tenant identity authentication data, such as passwords, certificates, and encryption keys. QingTian TPM can generate keys and use them for cryptographic functions, such as hashing, signing, encryption, and decryption.

QingTian TPM provides measured boot. During the process, the bootloader and OS create a cryptographic hash for each boot binary file and combine them with the previous values in the Platform Configuration Registers (PCRs) of QingTian TPM. With measured boot, you can obtain signed PCR values from QingTian TPM and use them to prove the integrity of the boot software of an instance to a remote entity. This is called remote attestation.

With QingTian TPM, keys and secrets can be tagged with specific PCR values so that they can never be accessed if the PCR values and instance integrity change. This special form of conditional access is called sealing and unsealing. OS technologies, such as BitLocker, can use QingTian TPM to seal drive decryption keys so that drives can only be decrypted when the OS is correctly booted and in a known good state.



Constraints

- If you want to change the specifications of an ECS booted using an image with QingTian TPM enabled, the target specifications must also support QingTian TPM.
- The boot mode of the image with QingTian TPM enabled must be UEFI.
- BitLocker volumes encrypted using QingTian TPM keys can only be used on the original instance.
- The QingTian TPM status of an ECS is not displayed in the ECS list.
- The QingTian TPM status is not included in image snapshots.

Billing Rules

There is no additional cost for using QingTian TPM. You only pay for the ECS resources you use.

Configuring QingTian TPM on an ECS

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

5. Set **Trusted Settings** based on service requirements.
 - To enable QingTian TPM, select **TPM**.
 - To disable QingTian TPM, deselect **TPM**.

For details about how to configure basic, network, and advanced settings when purchasing an ECS, see [Purchasing an ECS in Custom Config Mode](#).

6. Click **OK**.

8.7 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](#).

8.8 Protection for Mission-Critical Operations

Scenarios

ECS protects against mission-critical operations. If you want to perform a mission-critical operation on the management console, you must enter a credential for identity verification. You can perform the operation only after your identity is verified. For account security, it is a good practice to enable operation protection. The setting will take effect for both the account and IAM users under the account.


Operation protection applies to the following operations: stop, restart, or delete an ECS, reset the password for logging in to an ECS, detach a disk from an ECS, and unbind an EIP from an ECS.

Enabling Operation Protection

Operation protection is disabled by default. Perform the following operations to enable it:

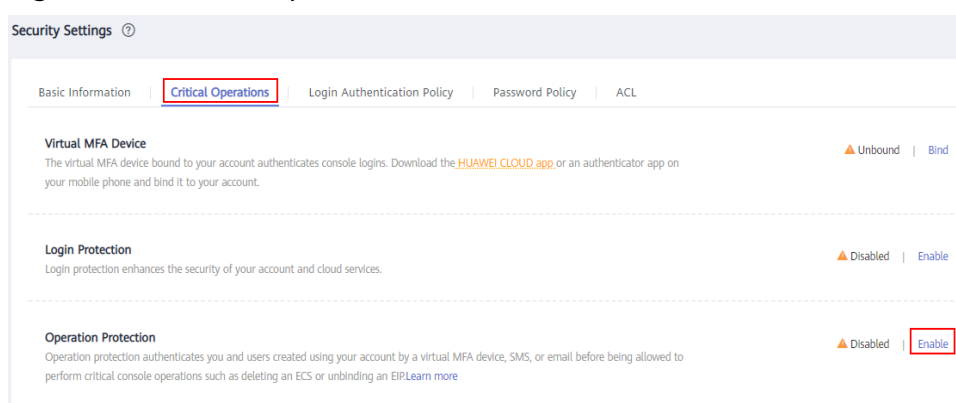
1. Log in to the management console.
2. Move the cursor to the username in the upper right corner of the page and select **Security Settings** from the drop-down list.

Figure 8-19 Security Settings



3. On the **Security Settings** page, choose **Critical Operations** > **Operation Protection** > **Enable**.

Figure 8-20 Critical Operations



4. On the **Operation Protection** page, select **Enable** to enable operation protection.

When you or the IAM users under your account perform critical operations, for example, deleting ECS resources, you are required to enter a verification code based on the selected verification method.

NOTE

- When performing a critical operation, you need to choose a verification method from email, SMS, and virtual MFA device.
 - If you have bound only a mobile number, only SMS verification is available.
 - If you have bound only an email address, only email verification is available.
 - If you have not bound an email address, mobile number, or virtual MFA device, you are required to bind one to continue with the critical operation.
- If you want to change the mobile number, email address, or virtual MFA device, see [Basic Information](#).

Verifying an Identity

After operation protection is enabled, when you perform a mission-critical operation, the system will verify your identity.

- If you have bound an email address, enter the email verification code.

- If you have bound a mobile number, enter the SMS verification code.
- If you have bound a virtual MFA device, enter a 6-digit dynamic verification code of the MFA device.

When you attempt to stop an ECS, select a verification method.

Figure 8-21 Identity Verification

Identity Verification ×

i You have enabled operation protection. If you do not require operation protection for critical operations, go to Security Settings > Critical Operations > Operation Protection to disable it. [Disable Identity Verification](#)

Verification Method ☒ SMS ☐ Email ☐ Virtual MFA device ?

Mobile Number [Change](#)

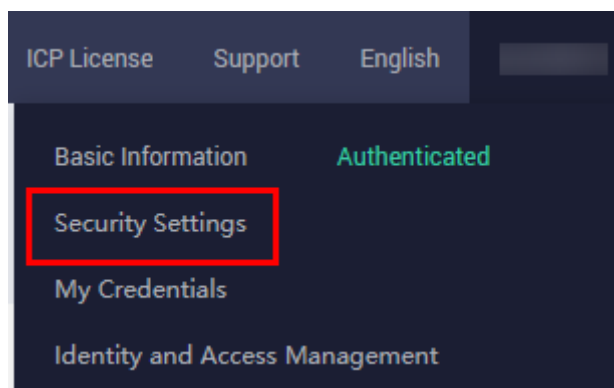
Verification Code

Disabling Operation Protection

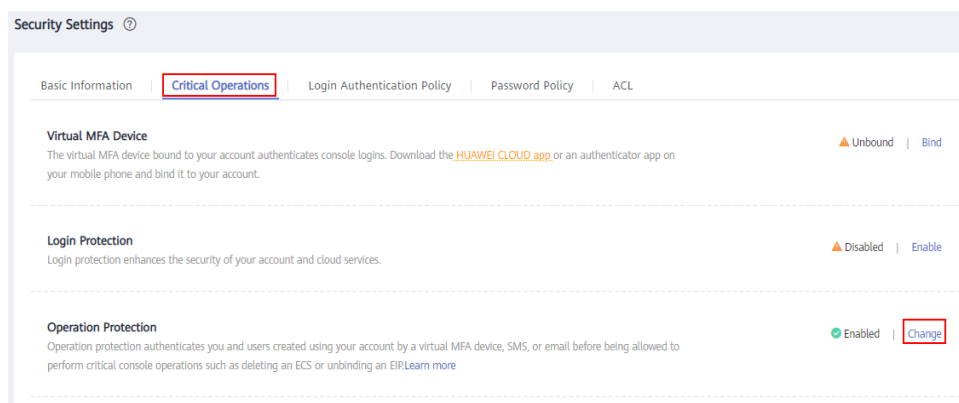
Perform the following operations to disable operation protection.

1. Log in to the management console.
2. Move the cursor to the username in the upper right corner of the page and select **Security Settings** from the drop-down list.

Figure 8-22 Security Settings



3. On the **Security Settings** page, choose **Critical Operations > Operation Protection > Change**.

Figure 8-23 Modifying operation protection settings

4. On the **Operation Protection** page, select **Disable** and click **OK**.

Helpful Links

- [How Do I Bind a Virtual MFA Device?](#)
- [How Do I Obtain a Virtual MFA Verification Code?](#)

9 Backup Using CBR

9.1 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.

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|---------------------|---|---|---|--|--|
| Cloud server backup | All disks (system and data disks) on an ECS | <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none">• Restoring Data Using a Cloud Server Backup• How Do I Restore Data on the Original Server to a New Server? |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|-------------------|--|---|--|--|---|
| Cloud disk backup | One or more specified disks (system or data disks) | <ul style="list-style-type: none">• 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 Disk Backup | <ul style="list-style-type: none">• Restoring from a Cloud Disk Backup• Creating a Disk from a Cloud Disk Backup |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|--------------|--|--|---|-------------------------------------|---|
| Snapshot | One or more specified disks (system or data disks) | <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none">• The snapshot data is stored with the disk data to facilitate rapid data back up 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 |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|--------------|---------------|---|----------------------------|----------------|--------------------|
| | | <p>disks can be used as data resources for various services, for example data mining, report query, and development and testing.</p> <p>This method protects the initial data and creates disks rapidly, meeting the diversified service data requirements.</p> <p>NOTE</p> <ul style="list-style-type: none">• 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. | | | |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|-------------------|--------------------|---|--|--|--|
| System disk image | System disk | <ul style="list-style-type: none">• 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 | <ul style="list-style-type: none">• 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 | Purchasing an EVS Disk |

| Back up Type | Backup Object | Application Scenario | Differences and Advantages | Back up Method | Restoration Method |
|----------------|---|--|---|---|---|
| Full-ECS image | All disks (system and data disks) on an ECS | <ul style="list-style-type: none">• 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 9-1 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. |

9.2 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 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 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](#).

Follow-up Procedure

After the backup is successful, you can view the backup details.



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 to go to the details page.
5. View the backup details.
 - a. On the displayed page, view the backup name and time in the **Cloud Backup and Recovery** area in the lower right corner.
- b. Click a backup name to go to the backup details page of the CBR console.

Figure 9-1 Cloud backup and recovery area


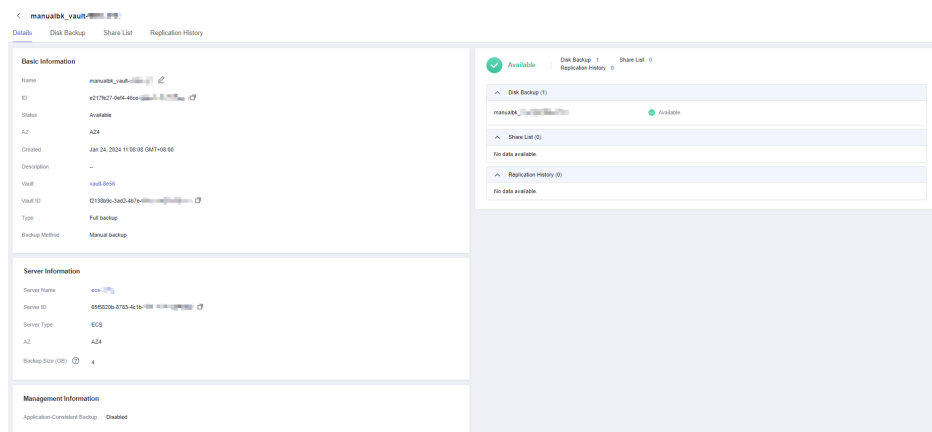


Figure 9-2 Backup details

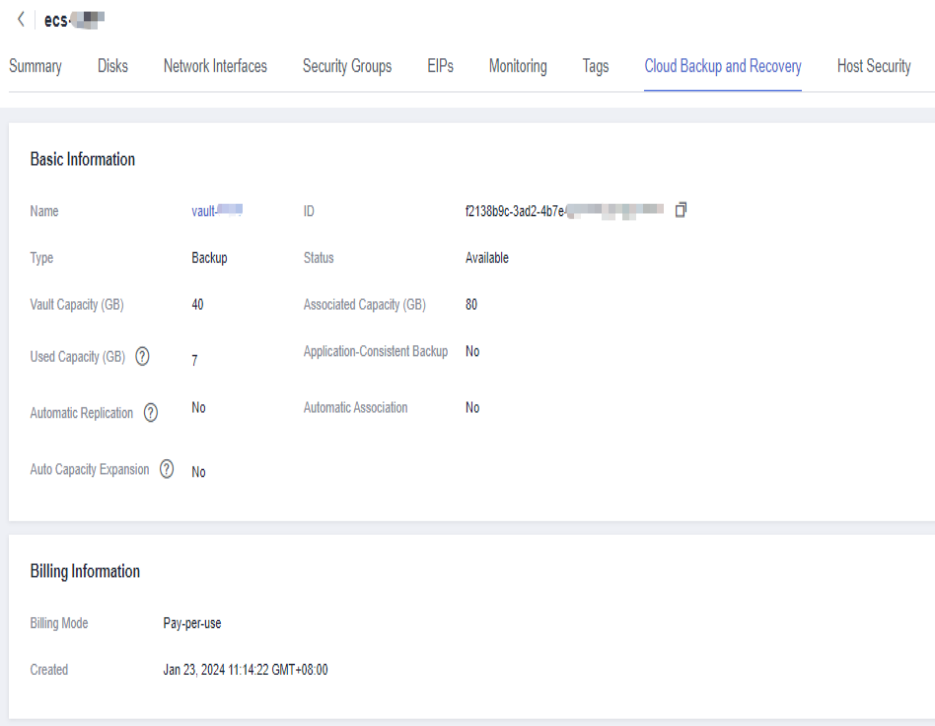


On the **Details** page, you can view the backup name, ID, status, AZ, creation time, vault, as well as server details.

For more information, see [Viewing a Backup](#).

6. On the ECS details page, click the **Cloud Backup and Recovery** tab to view the backup details and billing information.

Figure 9-3 Cloud backup and recovery details



10 Passwords and Key Pairs

10.1 Password Reset

10.1.1 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.
- You have purchased an ECS in quick config mode.

Table 10-1 provides guidance on how to reset your password in different scenarios.

Table 10-1 Resetting a password

| Scenario | Prerequisites |
|---|---|
| Resetting the Password for Logging In to an ECS on the Management Console | The password reset plug-in has been installed.
NOTE <ul style="list-style-type: none">• If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default.• The reference is for Windows or Linux ECSs. |
| Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed | The password reset plug-in has not been installed. |

| Scenario | Prerequisites |
|---|--|
| Resetting the Password for Logging In to a Linux ECS Without the Password Reset Plug-in Installed | The password reset plug-in has not been installed. |

Background

Table 10-2 shows the password complexity requirements in two scenarios: ECS creation and password reset.

Table 10-2 Password complexity requirements

| Password complexity requirements | ECS Creation | Password Reset |
|----------------------------------|--|--|
| Length | Must contain 8 to 26 characters. | Must contain 8 to 26 characters. |
| Characters | Must contain at least three of the supported four character types. | Must contain at least three of the supported four character types. |
| Security | <ul style="list-style-type: none">• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.) | <ul style="list-style-type: none">• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the same sequence as they appear in the username. (This requirement applies only to Windows ECSs.) |
| Other | None | Cannot start with a slash (/) for Windows ECSs. |
| Supported Characters | Uppercase letters | Uppercase letters |
| | Lowercase letters | Lowercase letters |
| | Digits | Digits |

| Password complexity requirements | ECS Creation | Password Reset |
|----------------------------------|---|---|
| | <ul style="list-style-type: none">Special characters for Linux ECSs: !@\$%^_-=+[{ }];,./?~#*Special characters for Windows ECSs: !@\$%^_-=+[{ }];,./?~#* | <ul style="list-style-type: none">Online password reset:
Special characters for Linux ECSs: @%-_=[]:/^,{ }?
Special characters for Windows ECSs: \$@%-_=[]:/,?
Offline password reset:
Special characters for Linux ECSs: !@%-_=[]:/?
Special characters for Windows ECSs: !@%-_=[]:/? |

10.1.2 Resetting the Password for Logging In to an ECS on the Management Console

Scenarios

If you did not set a password when creating an ECS, or the password expires or is forgotten, reset the password by following the instructions provided in this section.

You can reset the password online or offline.

Notes and Constraints

Table 10-3 Notes and constraints

| Constraint | Online Password Reset | Offline Password Reset |
|--------------------|--|------------------------|
| Service dependency | Cloud Operations Center (COC) needs to be enabled and authorized.
For IAM users, permissions for COC operations need to be granted. For details, see Configuring Custom Policies for ECS Self-Service O&M . | N/A |

| Constraint | Online Password Reset | Offline Password Reset |
|--------------------|---|--|
| Plug-in dependency | UniAgent needs to be installed. UniAgent is a unified data collection agent and supports script delivery and execution.

To install UniAgent on an ECS, see Installing UniAgent on an ECS . | N/A |
| Supported status | Running | <ul style="list-style-type: none">RunningStopped |
| Effective time | The new password will be applied immediately after the reset. | Depending on the ECS status. <ul style="list-style-type: none">If the ECS is running, the new password is applied after the ECS is restarted. NOTICE <ul style="list-style-type: none">Auto restart may cause data loss. You are advised to stop the ECS first before resetting the password.An ECS restart may cause a service interruption. You are advised to perform the operation during off-peak hours.If the ECS is stopped, the new password is applied after the ECS is started. |

Prerequisites

- You have installed the password reset plug-in before your ECS password expires or is forgotten.
 - If your ECS was created using a public image, the password reset plug-in was installed on the ECS by default.
 - If your ECS was created using a private image and has no password reset plug-in installed, see [Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed](#) or [Resetting the Password for Logging In to a Linux ECS Without the Password Reset Plug-in Installed](#).

- Do not delete the CloudResetPwdAgent or CloudResetPwdUpdateAgent process. Otherwise, one-click password reset will not be available.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP needs to be enabled in the VPC which the ECS belongs to.
- The ECS network connectivity is normal.
- Ensure that the one-click password reset plug-in is not blocked by security software. Otherwise, the one-click password reset function is unavailable.
- After the password is reset, you must restart the ECS for the new password to be applied.

Procedure

You can reset the password online or offline.

- Online password reset: This function depends on Cloud Operations Center (COC) and UniAgents provided by Application Operations Management (AOM). The new password is applied immediately without restarting the ECSs.
- Offline password reset: This function does not depend on COC and UniAgents provided by AOM. The new password is applied only after the ECS is restarted or started.

Online Password Reset

1. Log in to the management console and go to the [Elastic Cloud Server](#) console.
2. In the ECS list, select the ECSs for which you want to reset the passwords.

NOTE

You can select one or more ECSs to batch reset their passwords. The passwords of the selected ECSs will be the same after being reset.

3. Click **Reset Password** above the list.
4. Click the **Online Password Reset** tab.
5. (Optional) On the **Enable COC and Grant Permissions** page, read and agree to the service statement, and click **Enable and Authorize**.
This page is displayed if COC is not enabled and authorized.
6. Configure the required parameters.

NOTICE

A UniAgent is required for online password reset. If a message is displayed indicating that the UniAgent is not installed or failed to be installed, install it first by referring to [Installing UniAgent on an ECS](#).

Figure 10-1 Online Password Reset

Reset Password

Online Password Reset

Offline Password Reset

1

Ensure that you have been granted the required permissions. [View permissions](#)

1

Resetting the password online takes effect immediately without restarting ECSs. You can use the new password to log in to the ECSs.

Resources to Be Operated

The password will be reset online for 1 ECS.

| Name | ID | OS |
|------|----|-------|
| ecs- | | Linux |

Total Records: 1

<1>

Login Name

root

Enter an existing login name in the OS.

New Password

Enter a value.

Confirm Password

Enter a value.

Table 10-4 Parameters for resetting the password online

| Parameter | Description |
|------------|---|
| Login Name | The login name of the ECS. This parameter is set by default and does not need to be specified. <ul style="list-style-type: none">Linux ECS login name: rootWindows ECS login name: Administrator |

| Parameter | Description |
|------------------|--|
| New Password | <p>The new password.</p> <p>The new password must comply with the following rules:</p> <ul style="list-style-type: none">• Must contain 8 to 26 characters.• Must contain at least three of the following character types:<ul style="list-style-type: none">– Uppercase letters– Lowercase letters– Digits– Special characters <p>Special characters for Linux ECSs: @%-_+=[]:./^,{}?</p> <p>Special characters for Windows ECSs: \$@%-_+=[]:./,?</p> <ul style="list-style-type: none">• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the username (applying only to Windows ECSs).• Cannot start with a slash (/) (applying only to Windows ECSs). |
| Confirm Password | Must be the same as the new password. |

7. Click **OK**.

Offline Password Reset

1. Log in to the management console and go to the [Elastic Cloud Server](#) console.
2. In the ECS list, select the ECSs for which you want to reset the passwords.

NOTE

You can select one or more ECSs to batch reset their passwords. The passwords of the selected ECSs will be the same after being reset.

3. Click **Reset Password** above the list.
4. Click the **Offline Password Reset** tab and configure the required parameters.

Figure 10-2 Offline Password Reset

Reset Password

Online Password Reset

Offline Password Reset

The new password will take effect after the ECS is restarted.

Resources to Be Operated

The password will be reset offline for 1 ECS.

| Name | ID | Login Name |
|------|----|------------|
| ecs- | | root |

Total Records: 1

Login Name

root

New Password

Enter a value.

Confirm Password

Enter a value.

☐ The new password will take effect after the preceding ECS is automatically restarted.

Ensure that you save data and then proceed with this operation. Otherwise, ECS data will be lost and cannot be recovered.

Table 10-5 Parameters for resetting the password offline

| Parameter | Description |
|------------|--|
| Login Name | <div>The login name of the ECS. This parameter is set by default and does not need to be specified.</div> <ul style="list-style-type: none">Linux ECS login name: rootWindows ECS login name: Administrator |

| Parameter | Description |
|------------------|--|
| New Password | <p>The new password.</p> <p>The new password must comply with the following rules:</p> <ul style="list-style-type: none">• Must contain 8 to 26 characters.• Must contain at least three of the following character types:<ul style="list-style-type: none">– Uppercase letters– Lowercase letters– Digits– Special characters <p>Special characters for Linux
ECSs: !@%-_+=[:./?</p> <p>Special characters for
Windows ECSs: !@%-_+=[:./?</p> <ul style="list-style-type: none">• Cannot contain the username or the username spelled backwards.• Cannot contain more than two consecutive characters in the username (applies only to Windows ECSs).• Cannot start with a slash (/)
(applying only to Windows ECSs). |
| Confirm Password | Must be the same as the new password. |

5. Select **The new password will take effect after the preceding ECS is automatically restarted.**

This checkbox is displayed only when an ECS is in the **Running** state. You must select the checkbox to proceed with the password reset.

6. Click **OK**.

Helpful Links

- [Why Does Login to My ECS Using the Reset Password Fail?](#)
- [Why Am I Seeing the Message Indicating That the Port Is Used by a One-Click Password Reset Plug-in?](#)
- [Why Cannot the Installation Script Be Downloaded When I Try to Install UniAgent on an ECS?](#)

10.1.3 Obtaining and Deleting the Password of a Windows ECS

10.1.3.1 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



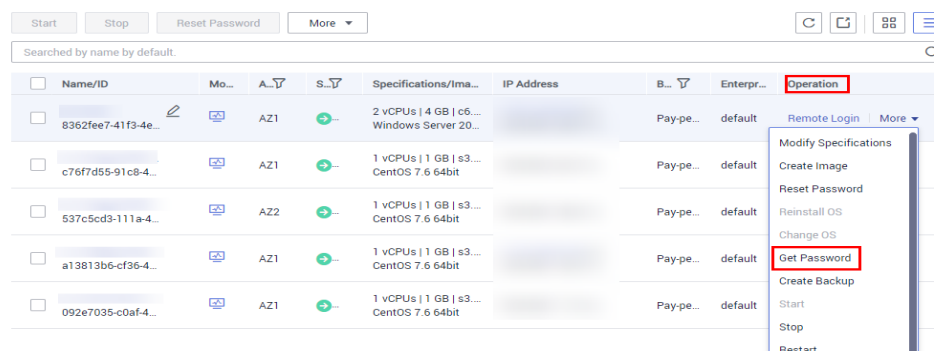
1. Obtain the private key file (.pem file) used when you created the ECS.
2. Log in to the management console.
3. Click  in the upper left corner and select a region and project.
4. Click . Under **Compute**, click **Elastic Cloud Server**.
5. On the **Elastic Cloud Server** page, select the target ECS.
6. In the **Operation** column, click **More** and select **Get Password**.

Figure 10-3 Obtaining a password



 **NOTE**

If **Get Password** is not displayed, the one-click password reset plug-in may not be installed.

In this case, you can reset the password by referring to [Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed](#).

7. Use either of the following methods to obtain the password through the 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.



10.1.3.2 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 a 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.

10.1.4 One-Click ECS Password Reset Plug-in

10.1.4.1 Obtaining the ECS One-Click Password Reset Plug-in

Scenarios

If the password failed to be reset, this may be because the one-click password reset plug-in has not been installed. You can install the plug-in and verify its integrity following the instructions provided in this section.

Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Linux)

1. Log in to the ECS as user **root**.
2. Download the one-click password reset plug-in and SHA256 checksum.
Obtain the download address from [Table 10-6](#) based on the region and OS (32- or 64-bit).

NOTE

To help you download the one-click password reset plug-in from the intranet faster, we provide the download addresses for different regions. The plug-in provided in different regions for the same OS is the same.

If the region where your ECS is located is not listed in [Table 10-6](#), bind an EIP to the ECS and select the nearest region to download the plug-in.

For details about how to bind an EIP, see [Binding an EIP](#).

Example command to download the plug-in for a 32-bit x86 ECS in the **CN North-Beijing1** region:

```
wget https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip
```

```
wget https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
```

3. Obtain the hash value of your local one-click password reset plug-in.
`sha256sum {software-package-directory}/CloudResetPwdAgent.zip`

Replace *software-package-directory* with the actual download directory, such as */root*.

4. Check whether the SHA256 hash value obtained in step 2 is consistent with that obtained in step 3.
 - If they are consistent, the verification is successful.
 - If they are inconsistent, download the one-click password reset plug-in of the corresponding version and repeat steps 2 to 4 to verify it.

Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity (Windows)

1. Log in to the ECS.
2. Download the one-click password reset plug-in and SHA256 checksum.

Obtain the download address from [Table 10-6](#) based on the region where the ECS resides.

NOTE

If the region where your ECS is located is not listed in [Table 10-6](#), bind an EIP to the ECS and select the nearest region to download the plug-in.

For details about how to bind an EIP, see [Binding an EIP](#).

3. Open Command Prompt as an administrator and run the following command to obtain the hash value of the local one-click password reset plug-in:

```
certutil -hashfile {software-package-directory}\CloudResetPwdAgent.zip  
SHA256
```

Replace *software-package-directory* with the actual download directory.

4. Check whether the SHA256 hash value obtained in step 2 is consistent with that obtained in step 3.
 - If they are consistent, the verification is successful.
 - If they are inconsistent, download the one-click password reset plug-in of the corresponding version and repeat steps 2 to 4 to verify it.

How to Obtain the One-Click Password Reset Plug-in and SHA256 Checksum

Table 10-6 Addresses for downloading the one-click password reset plug-in

| Region | OS | Name | How to Obtain |
|-------------------|---------------|------------------------|---|
| CN North-Beijing1 | Linux(x86_32) | CloudResetPwdAgent.zip | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |

| Region | OS | Name | How to Obtain |
|-------------------|---------------|------------------------|---|
| | Linux(x86_64) | CloudResetPwdAgent.zip | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(armv8) | CloudResetPwdAgent.zip | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Windows | CloudResetPwdAgent.zip | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| CN North-Beijing4 | Linux(x86_64) | CloudResetPwdAgent.zip | https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Windows | CloudResetPwdAgent.zip | https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| CN East-Shanghai2 | Linux(x86_32) | CloudResetPwdAgent.zip | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip |

| Region | OS | Name | How to Obtain |
|-----------------|--------------------|------------------------|---|
| | | SHA256 checksum | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(x86_64) | CloudResetPwdAgent.zip | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(armv8_64) | CloudResetPwdAgent.zip | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Windows | CloudResetPwdAgent.zip | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | CN South-Guangzhou | Linux(x86_32) | CloudResetPwdAgent.zip |
| SHA256 checksum | | | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| Linux(x86_64) | | CloudResetPwdAgent.zip | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |

| Region | OS | Name | How to Obtain |
|--------------|---------------|------------------------|---|
| | Linux(armv64) | CloudResetPwdAgent.zip | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Windows | CloudResetPwdAgent.zip | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| CN-Hong Kong | Linux(x86_32) | CloudResetPwdAgent.zip | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(x86_64) | CloudResetPwdAgent.zip | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(armv64) | CloudResetPwdAgent.zip | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |

| Region | OS | Name | How to Obtain |
|------------|---------------|------------------------|---|
| | Windows | CloudResetPwdAgent.zip | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| AP-Bangkok | Linux(x86_32) | CloudResetPwdAgent.zip | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/32/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(x86_64) | CloudResetPwdAgent.zip | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Linux(armv64) | CloudResetPwdAgent.zip | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | SHA256 checksum | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/arm/linux/64/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |
| | Windows | CloudResetPwdAgent.zip | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |
| | | CloudResetPwdAgent.zip | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip |

| Region | OS | Name | How to Obtain |
|--------|----|-----------------|---|
| | | SHA256 checksum | https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256 |

Related Operations

After obtaining the password reset plug-in, you can install the plug-in or update it as follows:

- [Installing the One-Click Password Reset Plug-in on an ECS](#)
- [Updating the One-Click Password Reset Plug-in for an ECS](#)

10.1.4.2 Installing the One-Click Password Reset Plug-in on an ECS

You can reset the password for logging in to an ECS with just a few clicks if you forgot the password or the password expires.

After you have created an ECS, it is a good practice to log in to it and install the password reset plug-in.

NOTE

The password reset plug-in has been installed on the ECSs created using a public image by default. To check whether the plug-in has been installed, see [Step 1](#).

Notes

1. The password reset plug-in is not installed by default. You can determine whether to install it.
2. After the installation is complete, do not uninstall the plug-in by yourself. Otherwise, the ECS password cannot be reset.
3. After you reinstall or change the OS of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
4. After you replace the system disk of an ECS, the one-click password reset function will become invalid. If you want to continue using this function, reinstall the password reset plug-in.
5. The password reset plug-in cannot be installed on ECSs running CoreOS or KylinOS.
6. To reset the password, the one-click password reset plug-in must be installed before the ECS password is lost or expires.
7. The one-click password reset plug-in can be installed only after an EIP is bound to the ECS.

Prerequisites

- The available space in drive C of a Windows ECS is greater than 300 MB, and data can be written to it.
- The available space in the root directory of a Linux ECS is greater than 300 MB, and data can be written to it.
- For Linux ECSs, [disable SELinux](#) if it has been enabled.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP is enabled in the VPC which the ECS belongs to.
- The ECS network connectivity is normal.
- The NIC has been set to DHCP so that the ECS can dynamically obtain an IP address.

NOTE

For details about how to set the NIC to DHCP for a Linux ECS, see [Setting the NIC to DHCP](#).

For details about how to set the NIC to DHCP for a Windows ECS, see [Setting the NIC to DHCP](#).

- The ECS security group rule in the outbound direction meets the following requirements:

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

If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:

- Protocol & Port: All
- Port range: All ports
- Destination: 0.0.0.0/0

Installing the Password Reset Plug-in on a Linux ECS

Step 1 Check whether the one-click password reset plug-in has been installed on the ECS.

1. Log in to the ECS as user **root**.
2. Run the following command to check whether CloudResetPwdAgent has been installed:

ls -lh /Cloud*

Figure 10-4 Checking whether the plug-in has been installed

```
[root@ecs-test ~]# ls -lh /Cloud*
total 20K
drwx----- 2 root root 4.0K Jun 13 14:13 bin
drwxr-xr-x 2 root root 4.0K Jun 13 11:53 conf
drwx----- 3 root root 4.0K Jun 13 11:53 depend
drwx----- 2 root root 4.0K Jun 13 11:53 lib
drwx----- 2 root root 4.0K Jun 13 14:13 logs
[root@ecs-test ~]#
[root@ecs-test ~]#
```

Check whether the obtained information is similar to that shown in [Figure 10-4](#).

- If yes, the plug-in has been installed.
- If no, the plug-in has not been installed. Then, install it.

Step 2 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Linux\)](#).

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 3 Run the following command to decompress **CloudResetPwdAgent.zip**:

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

unzip -o -d *Decompressed directory* CloudResetPwdAgent.zip

For example:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

unzip -o -d /home/linux/test CloudResetPwdAgent.zip

Step 4 Install the one-click password reset plug-in.

1. Run the following command to open the **CloudResetPwdAgent.Linux** file:

**cd {*Plug-in decompressed directory*}/CloudResetPwdAgent/
CloudResetPwdAgent.Linux**

For example:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

cd /home/linux/test/CloudResetPwdAgent/CloudResetPwdAgent.Linux

2. Run the following command to add the execute permission for the **setup.sh** file:

chmod +x setup.sh

3. Run the following command to install the plug-in:

sudo sh setup.sh

If "cloudResetPwdAgent install successfully." is displayed and "Failed to start service cloudResetPwdAgent" is not displayed, the installation is successful.

NOTE

- You can also check whether the password reset plug-in has been installed using the methods provided in [Step 1](#).
- If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

Step 5 Modify the file permissions of the password reset plug-in.

chmod 700 /CloudResetPwdAgent/bin/cloudResetPwdAgent.script

chmod 700 /CloudResetPwdAgent/bin/wrapper

```
chmod 600 /CloudrResetPwdAgent/lib/*
```

----End

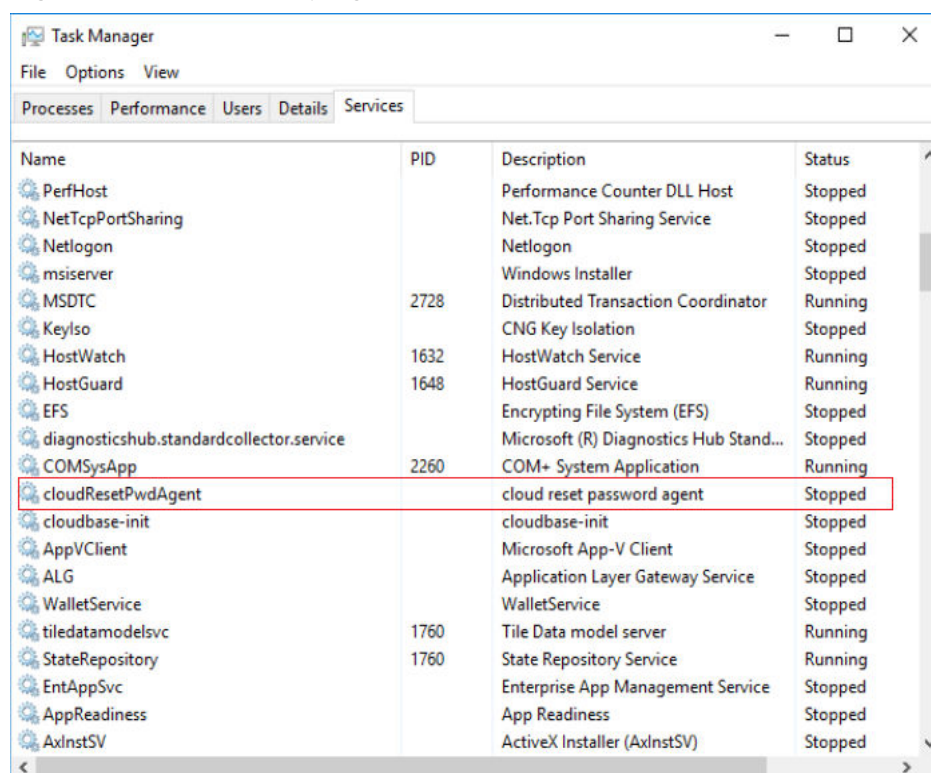
Installing the Password Reset Plug-in on a Windows ECS

Step 1 Log in to the ECS.

Step 2 Check whether the password reset plug-in CloudResetPwdAgent has been installed on the ECS. To check this, perform the following operations:

Start the **Task Manager** and check whether **cloudResetPwdAgent** is displayed on the **Services** tab. As shown in the [Figure 10-5](#), the password reset plug-in has been installed on the ECS.

Figure 10-5 Successful plug-in installation



- If yes, no further action is required.
- If no, go to [Step 3](#).

Step 3 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Windows\)](#).

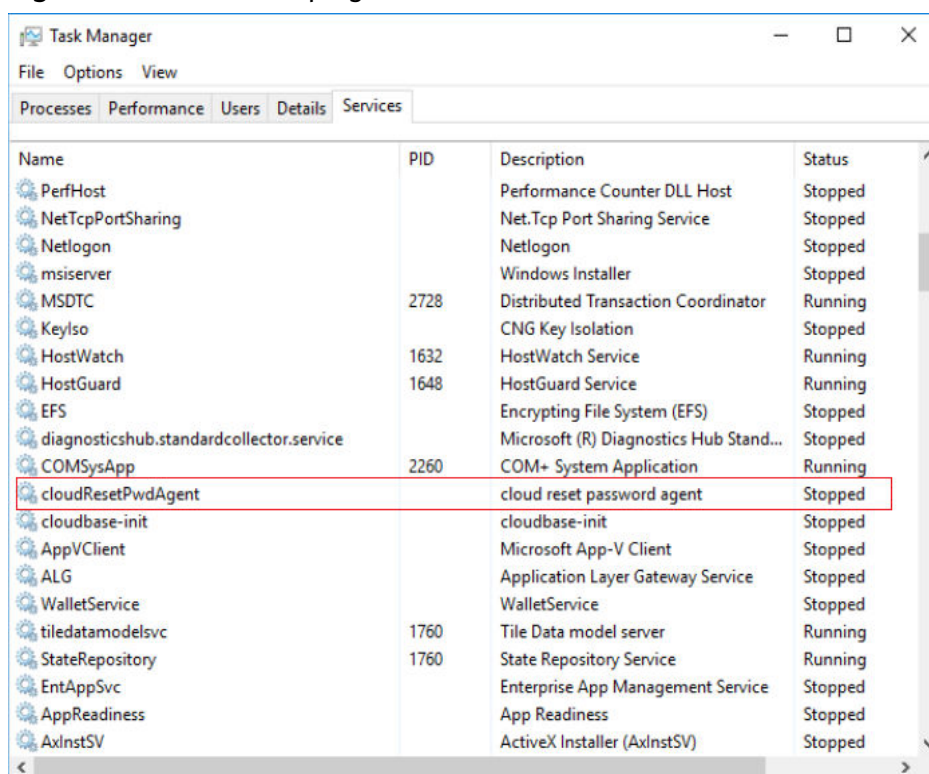
There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 4 Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

Step 5 Install the plug-in.

1. Double-click **setup.bat** in **CloudResetPwdAgent.Windows**.
The password reset plug-in starts to be installed.
2. View the **Task Manager** and check whether the installation was successful.
If **cloudResetPwdAgent** is displayed in the **Task Manager**, as shown in [Figure 10-6](#), the installation was successful. Otherwise, the installation failed.

Figure 10-6 Successful plug-in installation**NOTE**

If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

----End

Follow-up Procedure

- After the one-click password reset plug-in is installed, you can add it to the startup items if it cannot automatically start upon ECS startup. For details, see [What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?](#)
- After the one-click password reset plug-in is installed, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.
- If you have updated the one-click password reset plug-in, newly created ECSs work in PIPE mode by default to prevent the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in selects an idle port with the smallest port number from 31000 to 32999. The system will automatically select an idle port with the smallest port number.

Related Operations

- **Uninstall plug-in:** If you do not need to use the one-click password reset function, perform the following operations to uninstall the plug-in:
 - Linux ECS
 - i. Log in to the ECS.
 - ii. Run the following commands to switch to the **bin** directory and delete **cloudResetPwdAgent**:
cd /CloudrResetPwdAgent/bin
sudo ./cloudResetPwdAgent.script remove
 - iii. Delete the plug-in.
sudo rm -rf /CloudrResetPwdAgent
 - iv. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, go to the **bin** directory and delete the **cloudResetPwdUpdateAgent** service.
cd /CloudResetPwdUpdateAgent/bin
sudo ./cloudResetPwdUpdateAgent.script stop
sudo ./cloudResetPwdUpdateAgent.script remove
 - v. Delete the plug-in.
sudo rm -rf /CloudResetPwdUpdateAgent
 - Windows ECS
 - i. Uninstall and delete CloudResetPwdAgent.
 - 1) Switch to the **C:\CloudResetPwdAgent\bin** folder.
 - 2) Double-click **UninstallApp-NT.bat**.
 - 3) Delete the file in **C:\CloudResetPwdAgent**.
 - ii. (Optional) Uninstall and delete CloudResetPwdUpdateAgent.
The plug-in varies depending on the Windows version. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, perform the following operations to uninstall and delete it. If it does not exist, skip this step.
 - 1) Go to the **C:\CloudResetPwdUpdateAgent** folder.
 - 2) Double-click **UninstallApp-NT.bat**.
 - 3) Delete the file in **C:\CloudResetPwdUpdateAgent**.
If the deletion fails, delete **CloudResetPwdUpdateAgent** from Task Manager first and then delete the file in **C:\CloudResetPwdUpdateAgent**.
- **Stop plug-in:** If you do not need to use the one-click password reset function, perform the following operations to stop the plug-in:
 - LinuxECS
 - i. Log in to the ECS.
 - ii. Run the following commands to switch to the **bin** directory and delete **cloudResetPwdAgent**:
cd /CloudrResetPwdAgent/bin
sudo ./cloudResetPwdAgent.script stop

- iii. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, go to the **bin** directory and stop the **cloudResetPwdUpdateAgent** service.
cd /CloudResetPwdUpdateAgent/bin
sudo ./cloudResetPwdUpdateAgent.script stop
- WindowsECS
 - i. Stop **CloudResetPwdAgent**.
 - 1) Log in to the ECS.
 - 2) Open Task Manager.
 - 3) On the **Services** tab, right-click the **CloudResetPwdAgent** and select **Stop** from the shortcut menu.
 - ii. (Optional) Stop **CloudResetPwdUpdateAgent**.

The plug-in varies depending on the Windows version. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, perform the following operations to stop it. If it does not exist, skip this step.

 - 1) Log in to the ECS.
 - 2) Open Task Manager.
 - 3) On the **Services** tab, right-click **CloudResetPwdUpdateAgent** and select **Stop** from the shortcut menu.
- **Restart plug-in:** If the one-click password reset function is abnormal, perform the following operations to restart the plug-in:
 - LinuxECS
 - i. Log in to the ECS.
 - ii. Run the following commands to switch to the **bin** directory and restart **cloudResetPwdAgent**:
cd /CloudResetPwdAgent/bin
sudo ./cloudResetPwdAgent.script reboot
 - iii. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, go to the **bin** directory and restart the **cloudResetPwdUpdateAgent** service.
cd /CloudResetPwdUpdateAgent/bin
sudo ./cloudResetPwdUpdateAgent.script reboot
 - WindowsECS
 - i. Restart **CloudResetPwdAgent**.
 - 1) Log in to the ECS.
 - 2) Open Task Manager.
 - 3) On the **Services** tab, right-click **CloudResetPwdAgent** and select **Restart** from the shortcut menu.
 - ii. (Optional) Restart **CloudResetPwdUpdateAgent**.

The plug-in varies depending on the Windows version. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, perform the following operations to restart it. If it does not exist, skip this step.

 - 1) Log in to the ECS.

- 2) Open Task Manager.
- 3) On the **Services** tab, right-click **CloudResetPwdUpdateAgent** and select **Restart** from the shortcut menu.

10.1.4.3 Updating the One-Click Password Reset Plug-in for an ECS

You can reset the password for logging in to an ECS with just a few clicks if you forgot the password or the password expires.

This section describes how to update the one-click password reset plug-in for an ECS.

Notes

1. The one-click password reset plug-in can be updated only after an EIP is bound to the ECS.
2. By default, the one-click password reset plug-in has been installed on ECSs created using public images by default. Before updating the plug-in, uninstall it first.

Prerequisites

- The available space in drive C of a Windows ECS is greater than 300 MB, and data can be written to it.
- The available space in the root directory of a Linux ECS is greater than 300 MB, and data can be written to it.
- For Linux ECSs, [disable SELinux](#) if it has been enabled.
- One-click password reset can be used on the ECSs created using SUSE 11 SP4 only if their memory capacity is greater than or equal to 4 GiB.
- DHCP is enabled in the VPC which the ECS belongs to.
- The ECS network connectivity is normal.
- The NIC has been set to DHCP so that the ECS can dynamically obtain an IP address.

NOTE

For details about how to set the NIC to DHCP for a Linux ECS, see [Setting the NIC to DHCP](#).

For details about how to set the NIC to DHCP for a Windows ECS, see [Setting the NIC to DHCP](#).

- The ECS security group rule in the outbound direction meets the following requirements:
 - Protocol & Port: TCP
 - Port: 80
 - Destination: 169.254.0.0/16

If you use the default security group rules for the outbound direction, the preceding requirements are met, and the ECS can be initialized. The default security group rules for the outbound direction are as follows:

- Protocol & Port: All
- Port range: All ports

- Destination: 0.0.0.0/0

Updating the One-Click Password Reset Plug-in on a Linux ECS

Step 1 Uninstall the plug-in.

1. Log in to the ECS.
2. Run the following commands to switch to the **bin** directory and delete **cloudResetPwdAgent**:
cd /CloudResetPwdAgent/bin
sudo ./cloudResetPwdAgent.script remove
3. Delete the plug-in.
sudo rm -rf /CloudResetPwdAgent
4. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, go to the **bin** directory and delete the **cloudResetPwdUpdateAgent** service.
cd /CloudResetPwdUpdateAgent/bin
sudo ./cloudResetPwdUpdateAgent.script stop
sudo ./cloudResetPwdUpdateAgent.script remove
5. Delete the plug-in.
sudo rm -rf /CloudResetPwdUpdateAgent

Step 2 Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Linux\)](#).

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

Step 3 Run the following command to decompress **CloudResetPwdAgent.zip**:

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

unzip -o -d *Decompressed directory* CloudResetPwdAgent.zip

For example:

If the plug-in is decompressed to **/home/linux/test**, run the following command:

unzip -o -d /home/linux/test CloudResetPwdAgent.zip

Step 4 Install the one-click password reset plug-in.

1. Run the following command to open the **CloudResetPwdAgent.Linux** file:
cd {Plug-in decompressed directory}/CloudResetPwdAgent/CloudResetPwdAgent.Linux
For example:
If the plug-in is decompressed to **/home/linux/test**, run the following command:
cd /home/linux/test/CloudResetPwdAgent/CloudResetPwdAgent.Linux
2. Run the following command to add the execute permission for the **setup.sh** file:

```
chmod +x setup.sh
```

3. Run the following command to install the plug-in:

```
sudo sh setup.sh
```

If "cloudResetPwdAgent install successfully." is displayed and "Failed to start service cloudResetPwdAgent" is not displayed, the installation is successful.

 **NOTE**

- You can also check whether the password reset plug-in has been installed using the methods provided in [Step 1](#).
- If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

- Step 5** Modify the file permissions of the password reset plug-in.

```
chmod 700 /CloudrResetPwdAgent/bin/cloudResetPwdAgent.script
```

```
chmod 700 /CloudrResetPwdAgent/bin/wrapper
```

```
chmod 600 /CloudrResetPwdAgent/lib/*
```

----End

Updating the One-Click Password Reset Plug-in on a Windows ECS

- Step 1** Uninstall the plug-in.

1. Uninstall and delete CloudResetPwdAgent.
 - a. Switch to the **C:\CloudResetPwdAgent\bin** folder.
 - b. Double-click **UninstallApp-NT.bat**.
 - c. Delete the file in **C:\CloudResetPwdAgent**.
2. (Optional) Uninstall and delete CloudResetPwdUpdateAgent.

The plug-in varies depending on the Windows version. Check whether **CloudResetPwdUpdateAgent** exists. If it exists, perform the following operations to uninstall and delete it. If it does not exist, skip this step.

- a. Go to the **C:\CloudResetPwdUpdateAgent** folder.
- b. Double-click **UninstallApp-NT.bat**.
- c. Delete the file in **C:\CloudResetPwdUpdateAgent**.

If the deletion fails, delete **CloudResetPwdUpdateAgent** from Task Manager first and then delete the file in **C:\CloudResetPwdUpdateAgent**.

- Step 2** Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Windows\)](#).

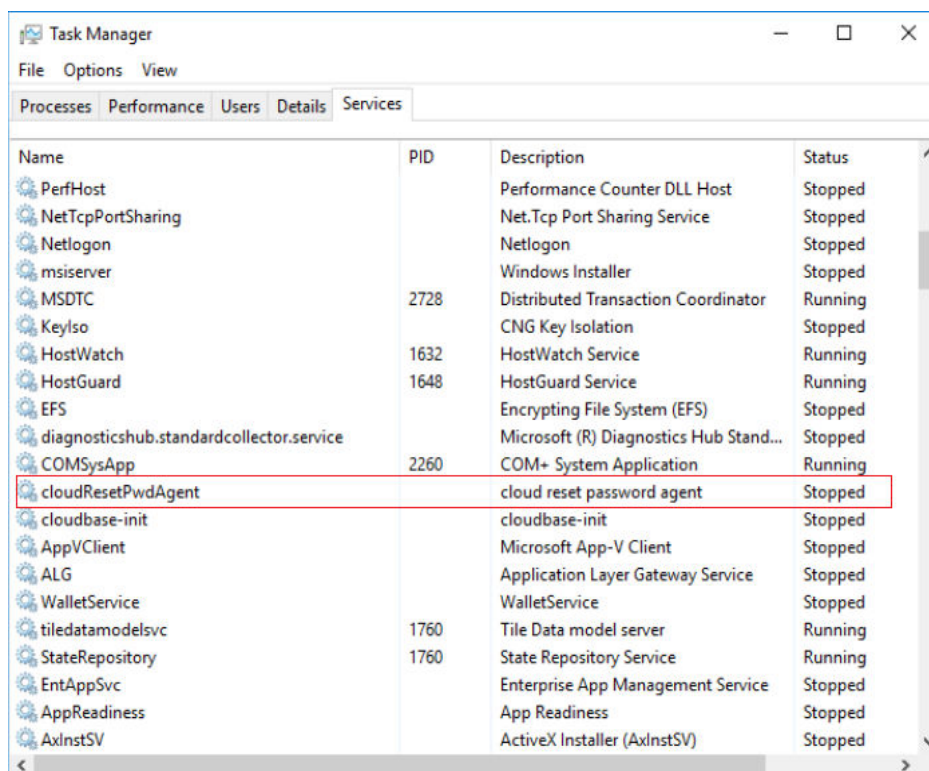
There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

- Step 3** Decompress **CloudResetPwdAgent.zip**.

There is no special requirement for the directory that stores the decompressed **CloudResetPwdAgent.zip**. Use any directory.

Step 4 Install the plug-in.

1. Double-click **setup.bat** in **CloudResetPwdAgent.Windows**.
The password reset plug-in starts to be installed.
2. View the **Task Manager** and check whether the installation was successful.
If **cloudResetPwdAgent** is displayed in the **Task Manager**, as shown in [Figure 10-7](#), the installation was successful. Otherwise, the installation failed.

Figure 10-7 Successful plug-in installation**NOTE**

If the installation failed, check whether the installation environment meets requirements and install the plug-in again.

----End

Follow-up Procedure

- After the one-click password reset plug-in is updated, you can add it to the startup items if it cannot automatically start upon ECS startup. For details, see [What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?](#)
- After the one-click password reset plug-in is updated, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.
- If you have updated the one-click password reset plug-in, newly created ECSs work in PIPE mode by default to prevent the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in selects an idle port with the smallest port number from 31000 to 32999.

10.1.4.4 Using Scripts to Batch Update the One-Click Password Reset Plug-ins for Linux ECSs

Scenarios

You can use scripts to batch update the one-click password reset plug-ins for multiple Linux servers.

Prerequisites

- An executor ECS meeting the requirements in [Constraints](#) is available.
- You have obtained the IP addresses of the ECSs where the plug-ins are to be batch installed, and the password of user **root** or the private key file for login.
- The executor ECS and the ECSs whose password reset plug-ins are to be updated must be in the same VPC.
- The EIP can be unbound only after you perform step 6.

Constraints

- The executor ECS must use the public image CentOS 7, has an EIP bound, and can communicate with the ECSs where the plug-ins are to be batch installed.

NOTE

If the internal yum repository has been configured, the executor ECS does not require an EIP.

- Only ECSs that use the same key pair support the batch installation of plug-ins.

Procedure

1. Log in to the executor ECS as user **root**.
2. Install the dependency required for batch script execution.
yum install ansible -y
3. Download the plug-in package **CloudResetPwdAgent.zip** and verify its integrity by referring to [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Linux\)](#).

There is no special requirement for the directory that stores **CloudResetPwdAgent.zip**.

4. Download the batch execution script to the **root** directory.

curl URL > ~/batch_update_log4j_version.py

URL is the address for downloading the batch execution script.

Select an address for downloading the script based on the region where the ECSs are located.

- CN North-Beijing1: https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py
- CN North-Beijing4: https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py

- CN East-Shanghai2: https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py
 - CN South-Guangzhou: https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py
 - CN-Hong Kong: https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py
 - AP-Bangkok: https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/batch_update_resetpwd/batch_update_log4j_version.py
5. Download the plug-in update script to the **root** directory.
- curl URL > ~/update_log4j_version_for_resetpwdagent.sh**
- URL** is the address for downloading the plug-in update script.
- Select an address for downloading the script based on the region where the ECSs are located.
- CN North-Beijing1: https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
 - CN North-Beijing4: https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
 - CN East-Shanghai2: https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
 - CN South-Guangzhou: https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
 - CN-Hong Kong: https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
 - AP-Bangkok: https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/linux/batch_update_resetpwd/update_log4j_version_for_resetpwdagent.sh
6. Check whether the following scripts are in the **root** directory:
- batch_update_log4j_version.py
 - update_log4j_version_for_resetpwdagent.sh
 - CloudResetPwdAgent.zip
7. Create **host_list.txt** and press **i** to enter editing mode.
- vi host_list.txt**
- Enter the information of the target ECSs in the **host_list.txt** file.
- The file format must match the login mode to be switched.
- If the target ECSs use a key pair for authentication, enter the following information:

CAUTION

- Upload the private key file saved during ECS creation to the folder in which **host_list.txt** is stored.
- Ensure that the permission code of the private key file is 400.

chmod 400 *Private key file*

Enter an ECS IP address in each line.

An example is provided as follows:

```
192.168.1.10  
192.168.1.11
```

- If the target ECSs use a password for authentication, enter the following information:

Enter an ECS IP address and password of user **root** separated using a comma (,) in each line.

An example is provided as follows:

```
192.168.1.10, !*****!  
192.168.1.11, !*****!
```

8. Run **batch_update_log4j_version.py**.

- For ECSs authenticated using key pairs

If the private key file and the batch execution script are in the same directory, you can simply use the private key file name.

python batch_update_log4j_version.py {*Private key file path/Private key file name*}

Figure 10-8 Successful script execution

```
2024-04-11 16:59:04 Start copying the file /root/CloudResetPwdAgent.zip to all hosts.  
2024-04-11 16:59:09 Start copying the file /root/update_log4j_version_for_resetpwdagent.sh to all hosts.  
2024-04-11 16:59:10 Start executing scripts on all hosts, it will take a while..  
2024-04-11 17:00:03 Please check the execution result.  
status ip msg  
[SUCCESS] [redacted] [redacted] install resetpwd successful  
Total: 1 Success: 1 Failed: 0  
You can check the logs/exec_origin.log for details.
```

If information shown in **Figure 10-8** is displayed, the execution is successful.

- For ECSs authenticated using passwords

python batch_update_log4j_version.py

Figure 10-9 Successful script execution

```
2024-04-11 16:22:27 Start copying the file /root/CloudResetPwdAgent.zip to all hosts.  
2024-04-11 16:22:32 Start copying the file /root/update_log4j_version_for_resetpwdagent.sh to all hosts.  
2024-04-11 16:22:33 Start executing scripts on all hosts, it will take a while..  
2024-04-11 16:23:47 Please check the execution result.  
status ip msg  
[SUCCESS] [redacted] [redacted] install resetpwd successful  
Total: 1 Success: 1 Failed: 0  
You can check the logs/exec_origin.log for details.
```

If information shown in **Figure 10-9** is displayed, the execution is successful.

9. View the execution result log in the last line of **/root/logs/exec_origin.log**.

vim /root/logs/exec_origin.log

If information shown in [Figure 10-10](#) is displayed, the one-click password reset plug-ins are batch updated.

Figure 10-10 Execution result log

```
begin install CloudResetPwdAgent
=====cp LinuxCloudResetPwdAgent=====
=====v1 /etc/profile=====
=====begin install CloudResetPwdAgent=====
Detected RHEL or Fedora:
Installing the cloudResetPwdAgent daemon using systemd...
creating default service file...
Reading file /etc/systemd/system/cloudResetPwdAgent.service
/home/CloudResetPwdAgent/CloudResetPwdAgent.Linux
cloudResetPwdAgent install successfully.
2024-04-11 17:08:03 Info:sh setup.sh successful
2024-04-11 17:08:03 Info:install resetpwd successful
2024-04-11 17:08:03 Info:update file permission successfulRemoved symlink /etc/systemd/system/multi-user.target.wants/cloudResetPwdAgent.service.
ls: cannot access /CloudResetPwdUpdateAgents: No such file or directory
Created symlink from /etc/systemd/system/multi-user.target.wants/cloudResetPwdAgent.service to /etc/systemd/system/cloudResetPwdAgent.service.
Redirecting to /bin/systemctl status cloudResetPwdAgent.service
```

Follow-up Procedure

- After the one-click password reset plug-in is updated, you can add it to the startup items if it cannot automatically start upon ECS startup. For details, see [What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?](#)
- After the one-click password reset plug-in is updated, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.
- If you have updated the one-click password reset plug-in, newly created ECSs work in PIPE mode by default to prevent the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in selects an idle port with the smallest port number from 31000 to 32999.

10.1.4.5 Using Scripts to Batch Update the One-Click Password Reset Plug-ins for Windows ECSs

Scenarios

You can batch update one-click password reset plug-ins for multiple Windows servers.

Prerequisites

- An executor ECS meeting the requirements in [Constraints](#) is available.
- You have obtained the IP addresses of the ECSs where plug-ins are to be batch installed, and the password of user **Administrator**.
- The executor ECS and the ECSs whose password reset plug-ins are to be updated must be in the same VPC.
- The EIP can be unbound only after you perform step [7](#).

Constraints

The executor ECS must use the public image CentOS 8.2, has an EIP bound, and can communicate with the ECSs where the plug-ins are to be batch installed.

NOTE

If the internal yum repository has been configured, the executor ECS does not require an EIP.

Procedure

1. Log in to the executor ECS as user **root**.
2. Install the required dependencies.

- a. Install **epel**.

yum install epel-release -y

- b. Install **ansible**.

yum install ansible -y --skip-broken

Run **ansible --version** to check whether Ansible is successfully installed.

NOTE

If Ansible cannot be installed due to yum repository configuration issues, run the following commands to install Ansible:

yum install python3 python3-pip

pip3 install --upgrade pip

pip3 install ansible

- c. Install **pip**.

**python3.6 -m pip install bcrypt==3.2.0 paramiko==3.3.1
cryptography==3.3.0 pywinrm PyYAML Jinja2 httpplib2 six**

If an error shown in [Figure 10-11](#) is displayed, perform the following operations:


- i. Install the dependency.

dnf install python3-devel

- ii. Run the following command again:

**python3.6 -m pip install bcrypt==3.2.0 paramiko==3.3.1
cryptography==3.3.0 pywinrm PyYAML Jinja2 httpplib2 six**

Figure 10-11 Error message



```

running build_ext
generating cffi module 'build/temp.linux-x86_64-3.6/_sodium.c'
building '_sodium' extension
creating build/temp.linux-x86_64-3.6/build
creating build/temp.linux-x86_64-3.6/build/temp.linux-x86_64-3.6
gcc -pthread -DDYNAMIC_ANNOTATIONS_ENABLED=1 -DNDEBUG -O2 -g -pipe -Wall -Werror=format-security -Wp,-D_FORTIFY_SOURCE=2 -Wp,-D_GLIBCXX_ASSERTIONS -fexceptions -fstack-protector-strong -grecord-gcc-switches -m64 -mtune=generic -fasynchronous-unwind-tables -fstack-clash-protection -fcf-protection -D_GNU_SOURCE -fPIC -fwrapv -O2 -g -pipe -Wall -Werror=format-security -Wp,-D_FORTIFY_SOURCE=2 -Wp,-D_GLIBCXX_ASSERTIONS -fexceptions -fstack-protector-strong -grecord-gcc-switches -m64 -mtune=generic -fasynchronous-unwind-tables -fstack-clash-protection -fcf-protection -D_GNU_SOURCE -fPIC -fwrapv -O2 -g -pipe -Wall -Werror=format-security -Wp,-D_FORTIFY_SOURCE=2 -Wp,-D_GLIBCXX_ASSERTIONS -fexceptions -fstack-protector-strong -grecord-gcc-switches -m64 -mtune=generic -fasynchronous-unwind-tables -fstack-clash-protection -fcf-protection -D_GNU_SOURCE -fPIC -fwrapv -O2 -g -pipe -Wall -Werror=format-security -Wp,-D_FORTIFY_SOURCE=2 -Wp,-D_GLIBCXX_ASSERTIONS -fexceptions -fstack-protector-strong -grecord-gcc-switches -m64 -mtune=generic -fasynchronous-unwind-tables -fstack-clash-protection -fcf-protection -D_GNU_SOURCE -fPIC -fwrapv -fPIC -I/usr/include/python3.6m -Ibuild/temp.linux-x86_64-3.6/ -c build/temp.linux-x86_64-3.6/_sodium.c -o build/temp.linux-x86_64-3.6/build/temp.linux-x86_64-3.6/_sodium.o
# include <pyconfig.h>
# include <pyconfig.h>
compilation terminated.
error: command 'gcc' failed with exit status 1

Command "/usr/bin/python3.6 -m -c "import setuptools, tokenize; file _ = '/tmp/pip-build-qdglubhd/pynacl/setup.py';f=getattr(tokenize, 'open', open);_file_ =codef.read().replace('\n\n', '\n');f.close();exec(compilecode, _file_, 'exec'))" install --record /tmp/pip-0u7xt6-record/install-record.txt --single-version-externally-managed --compile" failed with error code 1 in /tmp/pip-build-qdglubhd/pynacl/

```

3. Download the one-click password reset plug-in **CloudResetPwdAgent.zip** for Windows ECSs in the corresponding region by referring to [Table 10-6](#) and verify the plug-in integrity.

For example, to download the plug-in for Windows ECSs in the **CN North-Beijing4** region, run the following command:

```
wget https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip
```

```
wget https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/reset_pwd_agent/CloudResetPwdAgent.zip.sha256
```

For details about how to verify the plug-in integrity, see 4 in [Obtaining the One-Click Password Reset Plug-in and Verifying Its Integrity \(Linux\)](#).

The password reset plug-in can be stored in any directory.

4. Download the Windows installation package of the corresponding version and OS architecture to the **root** directory.

Note: Run the `pwd` command to check whether the current directory is **root**. If not, switch to the **root** directory.

- 32-bit OS, x86 architecture

```
wget https://www.7-zip.org/a/7z2107.exe '--no-check-certificate'
```

- 64-bit OS, x86 architecture

```
wget https://www.7-zip.org/a/7z2107-x64.exe '--no-check-certificate'
```

- 64-bit OS, Kunpeng architecture

```
wget https://www.7-zip.org/a/7z2107-x64.exe '--no-check-certificate'
```

5. Download the batch execution script to the **root** directory.

```
curl URL > ~/batch_update_log4j_version_for_windows.py
```

URL is the address for downloading the batch execution script.

Select an address for downloading the script based on the region where the ECSs are located.

- CN North-Beijing1: https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

- CN North-Beijing4: https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

- CN East-Shanghai2: https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

- CN South-Guangzhou: https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

- CN-Hong Kong: https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

- AP-Bangkok: https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/batch_update_resetpwd/batch_update_log4j_version_for_windows.py

6. Download the plug-in update script to the **root** directory.

```
curl URL > ~/update_log4j_version_for_resetpwdagent_windows.bat
```

URL is the address for downloading the plug-in update script.

Select an address for downloading the script based on the region where the ECSs are located.

- CN North-Beijing1: https://cn-north-1-cloud-reset-pwd.obs.cn-north-1.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat
- CN North-Beijing4: https://cn-north-4-cloud-reset-pwd.obs.cn-north-4.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat
- CN East-Shanghai2: https://cn-east-2-cloud-reset-pwd.obs.cn-east-2.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat
- CN South-Guangzhou: https://cn-south-1-cloud-reset-pwd.obs.cn-south-1.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat
- CN-Hong Kong: https://ap-southeast-1-cloud-reset-pwd.obs.ap-southeast-1.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat
- AP-Bangkok: https://ap-southeast-2-cloud-reset-pwd.obs.ap-southeast-2.myhuaweicloud.com/windows/batch_update_resetpwd/update_log4j_version_for_resetpwdagent_windows.bat

7. Check whether the following files are in the **root** directory:

- batch_update_log4j_version_for_windows.py
- update_log4j_version_for_resetpwdagent_windows.bat
- CloudResetPwdAgent.zip
- 7z*.exe

8. Create **host_list.txt** and press **i** to enter editing mode.

vi host_list.txt

Enter the information of the target ECSs in the **host_list.txt** file.

In each line, enter the IP address and the password of the **Administrator** user and separate them with a comma (,).

For example:

```
192.168.1.10,'*****'  
192.168.1.11,'*****'
```

9. Add the Ansible configuration file.

mkdir -p /etc/ansible

touch /etc/ansible/ansible.cfg

10. Execute the batch execution script

batch_update_log4j_version_for_windows.py.

python3.6 batch_update_log4j_version_for_windows.py

Figure 10-12 Executing the script


```
2022-01-08 17:00:36 Start copying the file /root/CloudResetPwdAgent.zip to all hosts.  
2022-01-08 17:00:43 Start copying the file /root/update_log4j_version_for_resetpwdagent_windows.bat to all hosts.  
2022-01-08 17:00:45 Start copying the file /root/7z.exe to all hosts.  
2022-01-08 17:00:46 Start executing scripts on all hosts, it will take a while..  
2022-01-08 17:00:57 Please check the execution result.  
status ip msg  
[SUCCESS] 192.168.96.118 started successfully  
Total: 1 Success: 1 Failed: 0  
You can check the logs/exec_origin.log for details.
```

11. View the execution result log in the last line of `/root/logs/exec_origin.log`.

`vim /root/logs/exec_origin.log`

If the following information is displayed, the execution is successful.

Figure 10-13 Successful script execution



```
C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\json-20160810.jar",
"C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\log4j-api-2.17.0.jar",
"C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\log4j-core-2.17.0.jar",
"C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\resetpwdupdateagent.jar",
"C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\wrapper.dll",
"C:\temp\CloudResetPwdAgent\CloudResetPwdUpdateAgent.Windows\CloudResetPwdUpdateAgent\lib\wrapper.jar",
"205 File(s) copied",
"The cloud reset password update agent service is starting.",
"The cloud reset password update agent service was started successfully.",
""
2022/01/08 17:00:56.63 \Info:"install ResetPwdUpdateAgent success\\"",
2022/01/08 17:00:56.83 \Info:"*****Update Success*****\\"",
"Press any key to continue . . ."
```

Follow-up Procedure

- After the one-click password reset plug-in is updated, you can add it to the startup items if it cannot automatically start upon ECS startup. For details, see [What Do I Do If the One-Click Password Resetting Plug-In Failed to Start?](#)
- After the one-click password reset plug-in is updated, do not delete the CloudResetPwdAgent process. Otherwise, one-click password reset will not be available.
- If you have updated the one-click password reset plug-in, newly created ECSs work in PIPE mode by default to prevent the plug-in from using service ports. Existing ECSs still work in AUTO mode, in which the plug-in selects an idle port with the smallest port number from 31000 to 32999.

10.2 Key Pairs

10.2.1 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.

This section describes how to create and import a key pair. For details about other operations, see [Managing Key Pairs](#).

Scenarios

When purchasing an ECS, you are advised to select the key pair login mode. For Windows ECSs, key pairs are required to decrypt the passwords so that you can use the decrypted password to log in.

- Logging in to a Linux ECS
You can directly use a key pair to log in 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](#).
 - After the ECS is created, [bind a key pair](#).
- Logging in to a Windows ECS
You can use the key pair to obtain a password for login. The password is randomly generated and is more secure.
For details, see [Obtaining the Password for Logging In to a Windows ECS](#).

Creating a Key Pair

You can create a key pair or use an existing one for remote login authentication.

- Creating a key pair
You can create a key pair using either of the following 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.

10.2.2 (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**.

NOTE

Key pairs include private key pairs and account key pairs. Private key pairs are only available to the user itself. Account key pairs are available to all users under the account.

You can create key pairs based on your needs.

6. Configure the following parameters:
 - a. Enter a key pair name.
 - b. Select a type.
 - c. Confirm KMS encryption.

This option is displayed only if you select **I agree to host the private key of the key pair**.

NOTE

- If you do not select **I agree to host the private key of the key pair**, you can download the private key only once for security purposes. Please keep the private key secure.

If the key pair is lost, you can [reset 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](#).
- d. Select **I have read and agree to the Key Pair Service Disclaimer**.

7. Click **OK**.

Related Operations

- If your private key file is lost, you can [reset the key pair](#).
- If your private key file is leaked, you can [use a new key pair to replace the public key of the ECS](#).

10.2.3 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.

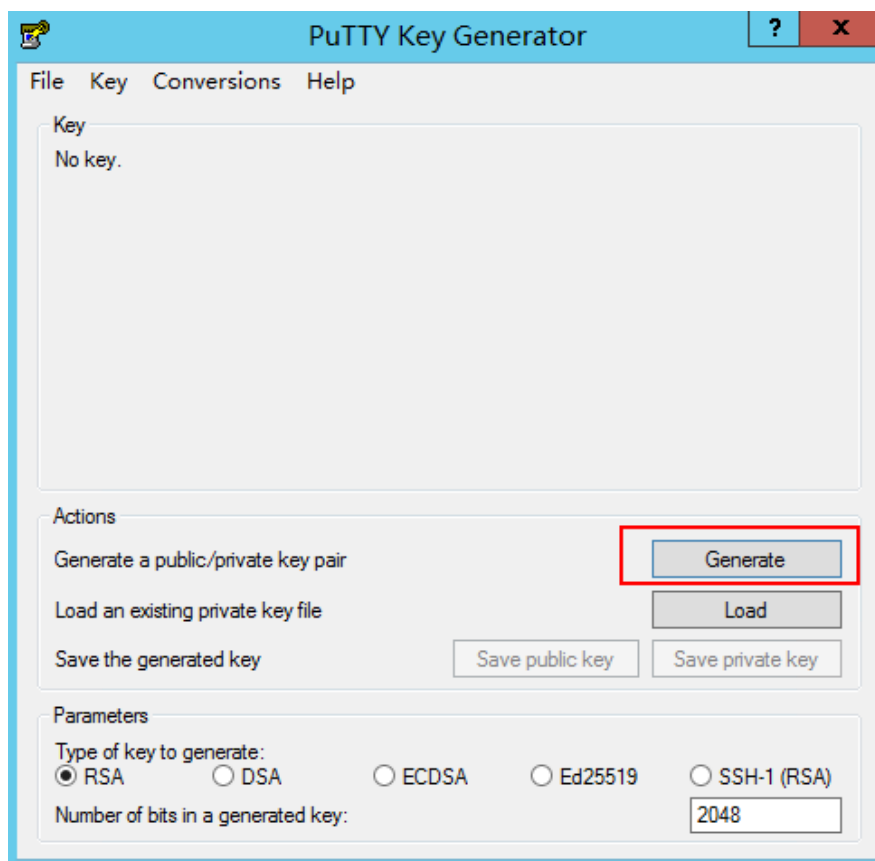
Procedure

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 10-14 PuTTY Key Generator

- b. Click **Generate**.

The key generator automatically generates a key pair that consists of a public key and a private key. The content shown in the red box in [Figure 10-15](#) is the public key.

Figure 10-15 Generating the public and private keys

3. Copy the public key to a .txt file and save it to a local directory.

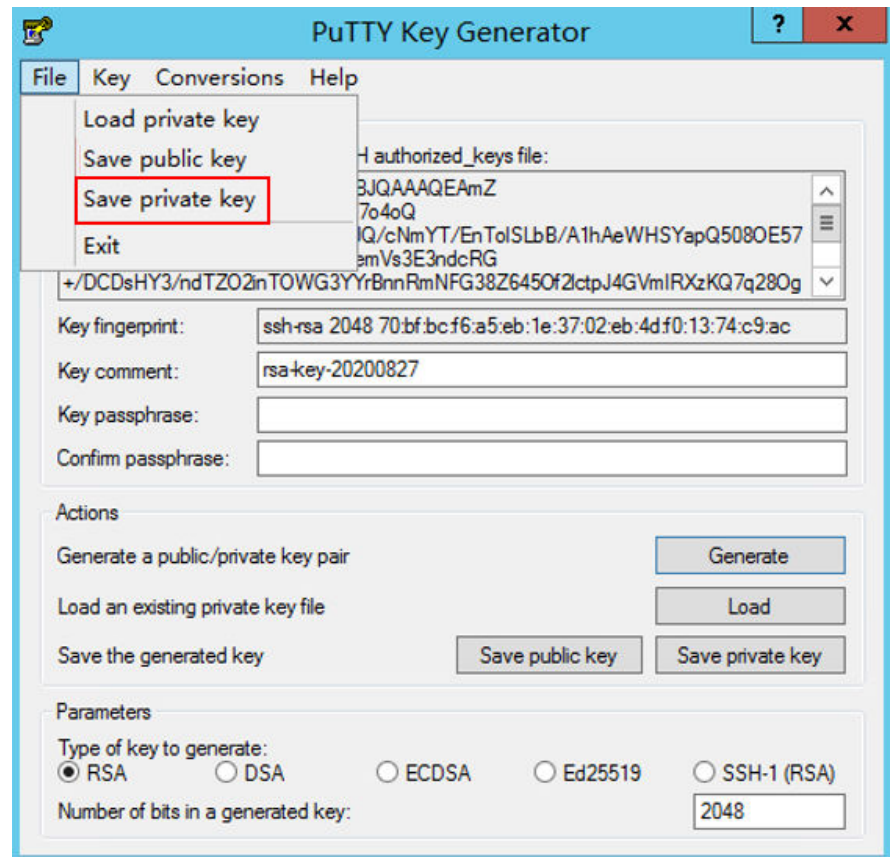
NOTE

Do not save the public key by clicking **Save public key** because this operation will change the format of the public key content and cause the public key to fail to be imported to the management console.

4. Save the private key and keep it secure. The private key can be downloaded only once.

The format in which to save your private key file varies depending on application scenarios.

- When using PuTTY 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 10-16 Saving a private key


- ii. Save the converted private key file, such as **kp-123.ppk**, locally.
- When using Xshell to log in to a Linux ECS or obtaining the password for logging in to a Windows ECS:

Save the private key file in the **.pem** format.

- i. Choose **Conversions > Export OpenSSH key**.

NOTE

If you use this private file to obtain the password for logging in to a Windows ECS, do not specify **Key passphrase** for **Export OpenSSH key** so that you can obtain the password successfully.

Figure 10-17 Saving a private key

- ii. Save the private key, for example, **kp-123.pem**, locally.
5. After you have saved the key pair, import your public key to the ECS by referring to [Importing a Key Pair](#).

Related Operations

- If your private key file is lost, you can [reset the key pair](#).
- If your private key file is leaked, you can [use a new key pair to replace the public key of the ECS](#).

10.2.4 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


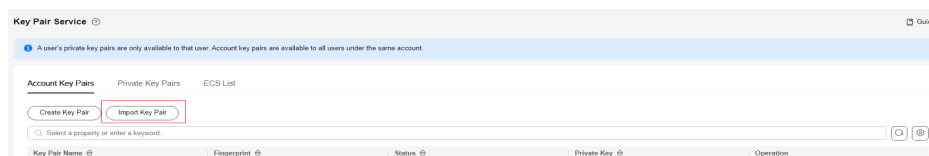
1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Key Pair**.
4. On the **Key Pair Service** page, click **Import Key Pair**.

Figure 10-18 Importing a key pair



5. In the displayed **Import Key Pair** dialog box, import the public and private key content.

Click **Select File** to import the local public key file (for example, the .txt file saved in step 3 of [Creating a Key Pair Using PuTTY Key Generator](#)).

For details about how to configure parameters, see [Importing a Key Pair](#).


Figure 10-19 Configuring parameters

×

Import Key Pair

1. Click Select File to import a public key file. You can change the key name if necessary.
2. Copy the content of a public key file to the Public Key Content field and enter a name in the Name field.
Notes: Only RSA keys are supported. The key file size must be 1024, 2048, or 4096 bits.

Key Pair Name

KeyPair-

Public Key

Select File

No file is selected.


Public Key Content

Private Key Content

KMS Encryption

kps/default

▼


 [View Key List](#) [Enter](#)

Key ID

4d358bbe-6d71-4095-8eca-51772de9c78e

☒

I agree to host the private key of the key pair. [Learn more](#)



What you use beyond the free API request quota given by KMS will be billed.
[Pricing details](#)

☒

I have read and agree to the [Key Pair Service Disclaimer](#)

Cancel

OK

6. 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?](#)

11 Launch Templates

11.1 Overview

What Is a Launch Template?

A launch template contains the configuration information to launch an ECS, for example, the ECS specifications, network settings, and a key pair (excluding the password). You can launch an ECS quickly without specifying the configuration parameters every time.

A launch template cannot be modified once it is created. You can create multiple launch templates and configure different parameters for each template.

Creating a Launch Template

Create a launch template on the console.

For details, see [Creating a Launch Template](#).

11.2 Creating a Launch Template

Scenarios

This section describes how to create a launch template on the management console.



Constraints

- Each account can have a maximum of 30 launch templates in each region.
- The parameters you can configure when you create a launch template are optional.

However, if your launch template does not include parameters, such as the flavor and image, you need to set them when you use the template to create an ECS.




- The configuration cannot be saved as a launch template if the billing mode is yearly/monthly or spot block, host security is enabled, or the login mode is password.
- After a launch template is created, it cannot be modified.
- Currently, launch templates are supported in AP-Singapore and CN-Hong Kong.

Creating a Launch Template on the Launch Templates Console

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Launch Templates**.
5. On the **Launch Templates** page, click **Create Launch Template**.
6. On the **Create Launch Template** page, configure the required parameters.
For details about the parameters, see [Purchasing an ECS in Custom Config Mode](#).
7. In the **Configuration Summary** panel on the right side, review the ECS configuration details.
Unspecified mandatory fields are displayed in red. You need to set them in the parameter configuration area.
8. Read and agree to the agreement, and click **Create Now**.
Go back to the launch template list and you can view the created launch template.

Creating a Launch Template When Buying an ECS

You can save the ECS configurations as a launch template when creating an ECS.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Click **Buy ECS**.
5. On the **Custom Config** tab, set the required parameters.
For details about the parameters, see [Purchasing an ECS in Custom Config Mode](#).
6. In the **Configuration Summary** panel on the right side, review the ECS configuration details.
Unspecified mandatory fields are displayed in red. You need to set them in the parameter configuration area.
7. In the **Configuration Summary** panel, click  and select **Save as Launch Template**.
8. Read the agreement and click **Agree and Submit**.

9. In the displayed dialog box, enter a launch template name and description, and click **OK**.

You can view the created template on the launch template list page.

11.3 Managing Launch Templates

Scenarios

You can:

- [Viewing Details About a Launch Template](#)
- [Deleting a Launch Template](#)

Viewing Details About a Launch Template




1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Launch Templates**.
5. On the **Launch Templates** page, click the name of the launch template to view its details.

Table 11-1 Launch template details

| Parameter | Description |
|---------------------|---|
| Name | The name of the launch template. |
| ID | The ID of the launch template. |
| Created | The time when the launch template is created. |
| Description | The description of the launch template. |
| Version Information | The version information contains the configuration information about the launch template of the current version, such as the region, specifications, and image. |

Deleting a Launch Template

1. Log in to the management console.
2. Click . Under **Compute**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Launch Templates**.
4. Locate row that contains the launch template to be deleted and click **Delete** in the **Operation** column.
5. Click **OK**.

12 Auto Launch Groups

12.1 Overview

What Is an Auto Launch Group?

An auto launch group lets you rapidly create ECSs distributed across multiple AZs, using a combination of different types of spot and pay-per-use ECSs to meet capacity targets at the lowest price possible.

Application Scenarios

Auto launch groups are applicable to scenarios such as image rendering, stateless web services, DNA sequencing, offline analysis, function computing, batch computing, sample analysis, CI/CD, and test.

Notes

- An auto launch group can create ECSs across AZs but cannot create ECSs across regions.
- The target capacity of each auto launch group is limited.
 - If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.
 - If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.
- You can specify one launch template for each auto launch group.

Advantages

- Spot ECSs and pay-per-use ECSs
Spot ECSs are much less expensive than regular pay-per-use ECSs, but they can be reclaimed suddenly. Spot ECSs are a great way to save money when running stateless, fault-tolerant instances that are not sensitive to interruptions. Pay-per-use ECSs can be created and deleted at any time, and the inventory is sufficient and stable, but the price is higher than that of spot ECSs.

An auto launch group lets you rapidly create both spot and pay-per-use ECSs to meet capacity targets at the lowest price possible.

- ECSs from different AZs

An auto launch group can create ECSs across AZs to improve the disaster recovery capability.

- ECSs of different types

An auto launch group can create ECSs of different types to meet your requirements of different scenarios.

- Flexible allocation strategies

You can specify your desired target capacity and how much of that must be pay-per-use ECSs.

You can also let your auto launch group continue to create ECSs until the total target capacity is reached or delete ECSs when the target capacity is exceeded.

- Cost effectiveness

If your set **Optimize for** to **Lowest price**, the auto launch group will create the least expensive ECSs possible.

Pricing Details

Auto launch groups are free, but you will be billed for the ECSs created by the group.

For details, see [Elastic Cloud Server Pricing Details](#).

12.2 Creating an Auto Launch Group



Scenarios

This section describes how to create an auto launch group on the management console.

Constraints

Currently, auto launch groups are supported in AP-Singapore and CN-Hong Kong regions.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Auto Launch Groups**.
5. On the **Auto Launch Groups** page, click **Create Group**.
6. Set the name of the auto launch group.

The name can contain 2 to 64 characters, including letters, digits, underscores (_), and hyphens (-).

7. Set the total target capacity.

You can specify the number of ECSs or vCPUs.

If you choose to include pay-per-use ECSs, set the quantity of pay-per-use ECSs or vCPUs.

The target capacity of each auto launch group is limited.

 - If the number of ECSs is used as the target capacity, a maximum of 500 ECSs can be created.
 - If the number of vCPUs is used as the target capacity, a maximum of 40,000 vCPUs can be created.
8. Select a launch template.

You can select a launch template and its corresponding version as the configuration source. You can also select other required ECS configurations.
9. Set the allocation strategy.
 - **Lowest price:** The auto launch group will create the least expensive ECSs possible.
 - **Compute balancing:** The auto launch group will prioritize balancing compute loads by creating ECSs distributed across multiple AZs as evenly as possible.
 - **High specifications:** The auto launch group creates ECSs with the highest specifications possible.

If you have configured a target number of ECSs, ECSs with more vCPUs will be prioritized and if the target is vCPUs, then that target will be met with as few ECSs as possible.
10. Select a delivery type.
 - **Single use:** The auto launch group only attempts to create ECSs to meet the target capacity when it is started, but will not create ECSs again even if the target capacity is not reached.
 - **Continuous:** The auto launch group continues to create ECSs until the total target capacity is reached.
11. Set the start time.

Set the time when the auto launch group starts to launch ECSs. You can set both the start time and the end time to determine the validity period of the group.

 - **Immediately:** The auto launch group starts to launch ECSs immediately after the group is created.
 - **Custom:** You can specify when the auto launch group starts to launch ECSs.
12. Set the end time.

Set the time when the auto launch group expires. You can set both the start time and the end time to determine the validity period of the group.

 - **Never expire:** The auto launch group does not expire.
 - **Custom:** You can specify when the auto launch group expires.
13. Set the global maximum price.

Set the allowed maximum price of a single spot ECS in the auto launch group. If the market price of a spot ECS in the group exceeds the global maximum price, the spot ECS will be deleted.

If both the specific maximum price of a spot ECS and the global maximum price are set, the specific maximum price of the spot ECS will be used.

The price cannot be less than 0. If the price is set to be greater than the pay-per-use ECS price, there is no upper limit on the spot ECS price.

14. Configure ECS deletion settings.

- **Delete ECSs When Auto Launch Group Expires:** ECSs in the auto launch group will be deleted when the group expires.
- **Delete ECSs When Target Capacity Is Exceeded:** When the number of ECSs or vCPUs in the auto launch group exceeds the target capacity, the ECSs or vCPUs that exceed the target capacity will be deleted.

 **NOTE**

If you do not select **Delete ECSs When Target Capacity Is Exceeded**, the ECSs that exceed the target capacity will be removed from the group but not deleted. The removed ECSs will be displayed in the ECS list. To avoid charges on such ECSs, manually delete them.

15. Click **Create Now**.

Execution Result

After an auto launch group is created, the group starts to create ECSs at the specified time. If you select the **Continuous** delivery mode, the auto launch group monitors the target and current capacity in real time, and automatically creates a new ECS if a spot ECS is reclaimed.


12.3 Managing Auto Launch Groups

Scenarios

You can:

- [Viewing Details About an Auto Launch Group](#)
- [Modifying an Auto Launch Group](#)
- [Deleting an Auto Launch Group](#)

Viewing Details About an Auto Launch Group


1. Log in to the management console.
2. Click . Under **Compute**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. On the **Auto Launch Groups** page, click the name of the target auto launch group to view its details.

You can view the basic information and capacity overview of the auto launch group.

The basic information includes the name, launch template, delivery type, and allocation strategy of the group.

In the **Capacity Overview** area, you can view the total current and target capacity, and the current and target capacity of the spot ECS and pay-per-use ECS.


Modifying an Auto Launch Group

1. Log in to the management console.
2. Click  . Under **Compute**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. Locate the row that contains the target auto launch group and click **Modify** in the **Operation** column.

You can modify the name, target capacity, quantity of pay-per-use ECSs, allowed maximum price of a spot ECS, and whether to delete ECSs when the auto launch group expires or the target capacity is exceeded.

5. Click **OK**.

Deleting an Auto Launch Group

1. Log in to the management console.
2. Click  . Under **Compute**, choose **Elastic Cloud Server**.
3. In the navigation pane on the left, choose **Auto Launch Groups**.
4. Locate the row that contains the target auto launch group and click **Delete** in the **Operation** column.

5. Determine whether to delete the ECSs in the auto launch group after the group is deleted.

If you do not want to delete the ECSs, you can view the ECSs on the ECS list page. To avoid charges on such ECSs, manually delete them.

6. Click **OK**.

13 Events

13.1 Overview

Huawei Cloud can predict and proactively prevent hardware or software faults of hosts accommodating ECSs.

If host failures cannot be avoided, the system will generate and report events for affected ECSs to minimize impacts of instance unavailability or performance deterioration. These events include instance redeployment and local disk replacement. For details, see [Event Type](#). The system does not frequently report events.

You can view events details on the ECS console, including the event type, instance name/ID, and event status. You can also check ECS events details on the **Event Monitoring** page on the Cloud Eye console. For details, see [Viewing Event Monitoring Data](#).

Event Type

[Table 13-1](#) describes events that can be reported by the system.

Table 13-1 Events

| Event Type | Generated When | Impact | Handling Suggestion |
|-----------------------|--|--|--|
| Instance redeployment | The system detects that the host accommodating ECSs is faulty and it plans to deploy the ECSs on a new host. | <p>During the instance redeployment, ECSs will be temporarily unavailable for a short period of time.</p> <p>The system will send the event notification 24 to 72 hours earlier than the scheduled execution time.</p> <p>NOTICE
For ECSs using local disks, all data stored on the local disks will be lost.</p> | <p>Refer to the following to rectify the fault. After the fault is rectified, check the impacts on services. If any problems occur, contact technical support.</p> <p>Handling an Instance Redeployment Event</p> <p>You are advised to select off-peak time as the scheduled start time during authorization. If you do not specify the start time, the current time is used as the start time by default.</p> |

| Event Type | Generated When | Impact | Handling Suggestion |
|------------------------|--|---|--|
| Local disk replacement | The system detects that a disk of the host accommodating disk-intensive ECSs or bare metal ECSs is faulty. | Local disk replacement will cause data loss on local disks. | <p>Refer to the following to rectify the fault. After the fault is rectified, check the impacts on services. If any problems occur, contact technical support.</p> <p>NOTICE</p> <p>Local disk replacement will cause data loss on local disks. If you do not need to retain data on local disks, use one of the following methods:</p> <ul style="list-style-type: none">• Redeployment: All local disk data will be lost.• Authorizing Disk Replacement: Only data on the faulty local disk will be lost. <p>You are advised to select off-peak time as the scheduled start time during authorization. If you do not specify the start time, the current time is used as the start time by default.</p> <p>The local disk replacement will be completed within five working days generally after it is started. Please wait patiently.</p> |

| Event Type | Generated When | Impact | Handling Suggestion |
|--------------------|---|---|---|
| Instance migration | The system detects that the host accommodating ECSs is faulty and needs to be restarted, stopped, or brought offline, and it plans to migrate ECSs. | The system attempts to perform a live migration of ECSs first. The HA mechanism will be triggered if an exception occurs (ECSs will be unavailable temporarily during this period). | After the fault is rectified, check the impacts on services. If any problems occur, contact technical support. |
| System maintenance | The system detects that there are hardware or software faults in the host accommodating ECSs (including bare metal ECSs) and plans to perform maintenance operations on the affected instances. | During system maintenance, the host may be powered off, and ECSs running on it become unavailable. | <p>Refer to the following to rectify the fault. After the fault is rectified, check the impacts on services. If any problems occur, contact technical support.</p> <p>Handling a System Maintenance Event</p> <p>Ensure that services running on the instances have been stopped and select an off-peak time as the scheduled start time during authorization. If you do not specify the start time, the current time is used as the start time by default.</p> <p>The duration required for system maintenance varies depending on the faults. The system maintenance will be completed within five working days generally after the authorization is started. Please wait patiently.</p> |

Event Status

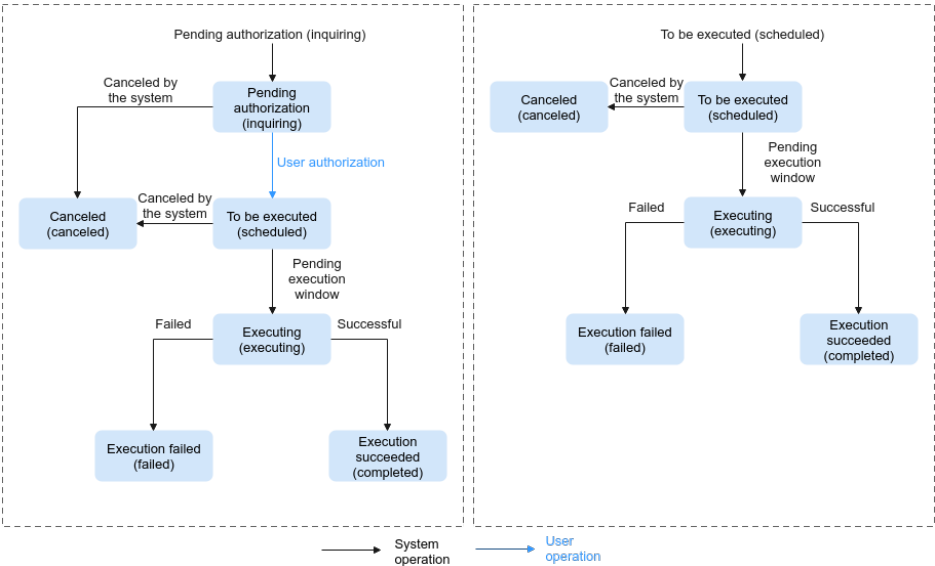
Table 13-2 lists statuses of the events reported by the system. You can check progresses of the events and filter events by status.

Table 13-2 Event statuses

| Type | Description |
|-----------------------|---|
| Pending authorization | An event is waiting to be authorized with the start time specified. The system will complete operations within a specified time. For details, see Handling an Event . |
| To be executed | The event is waiting for the system to schedule resources. |
| Executing | The system has scheduled resources and is rectifying the fault. |
| Execution succeeded | The system has completed event execution. Check the impacts on services. If any problems occur, contact technical support. |
| Execution failed | The system fails to automatically rectify the fault. |
| Canceled | The event has been canceled. |

The event status changes with the operations performed by users and the system.

Figure 13-1 Event statuses



13.2 Querying an Event

Scenarios

This section describes how to query events on the ECS console.

You can also check ECS events details on the **Event Monitoring** page on the Cloud Eye console. For details, see [Viewing Event Monitoring Data](#).

Prerequisites



If you need to perform operations as an IAM user, ensure that the IAM user has been granted the required permissions.

Event query and handling require the permissions defined in the following policies:

- Query events: **ecs:instanceScheduledEvents:list**
- Accept and authorize events: **ecs:instanceScheduledEvents:accept**
- Change the reservation time: **ecs:instanceScheduledEvents:update**

For details about how to grant permissions to IAM users, see [ECS Custom Policies](#).

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. Choose **Events** from the left navigation pane.

In the displayed list on the **Events** page, you can view the event ID, instance name/ID, event type, and more event details.

You can perform the operations displayed in the **Operation** column on events. For details, see [Handling an Event](#).

13.3 Handling an Event

13.3.1 Handling an Instance Redeployment Event

Scenarios

When the system reports an instance redeployment event, you need to authorize the redeployment for the instance.

This section describes how to handle an instance redeployment event in **Pending authorization** status and guides you through the authorization required for the instance redeployment.

You can authorize the redeployment now or later.

- **Now:** After receiving a request, the system immediately redeploys the instance.
- **Later:** After receiving a request, the system redeploys the instance at the scheduled time.

You can change the scheduled time as needed. For details, see [Changing the Reservation Time](#).

NOTICE

- Instance redeployment will restart the instance. Switch service traffic in advance and select an appropriate time for authorizing the redeployment to minimize the impact on your services.
- Instance redeployment will not affect the instance's system disk and EVS data disks.
- For instances using local disks, all data stored on the local disks will be deleted after the instance redeployment. To ensure data security, back up local disk data before authorizing the redeployment.
- For instances using local disks, before authorizing the redeployment, you need to [preprocess local disks](#) to prevent the instance from starting in emergency mode.

Choose a method to authorize redeployment based on the types of disks attached to instances.

- [Authorizing Redeployment for Instances that Not Using Local Disks](#)
- [Authorize Redeployment for Instances Using Local Disks](#)

Prerequisites

If you need to perform operations as an IAM user, ensure that the IAM user has been granted the required permissions.

Event query and handling require the permissions defined in the following policies:

- Query events: **ecs:instanceScheduledEvents:list**
- Accept and authorize events: **ecs:instanceScheduledEvents:accept**
- Change the reservation time: **ecs:instanceScheduledEvents:update**

For details about how to grant permissions to IAM users, see [ECS Custom Policies](#).

Authorizing Redeployment for Instances that Not Using Local Disks



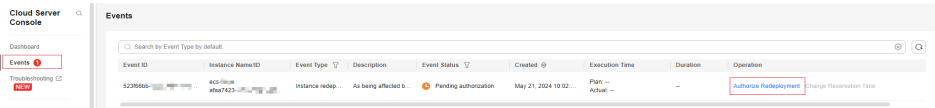

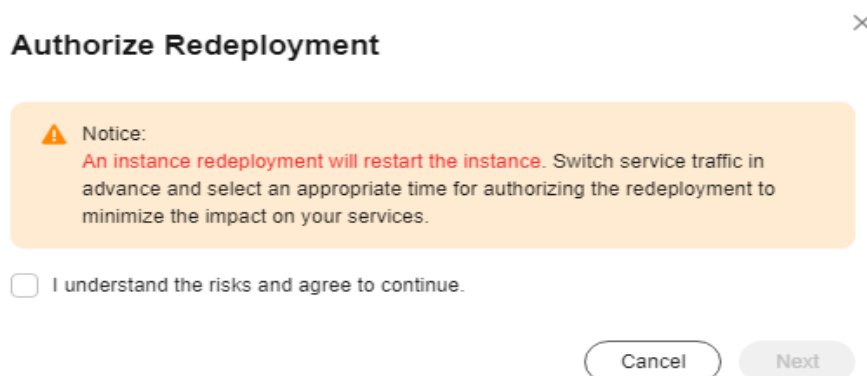
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Choose **Events** from the left navigation pane.

Figure 13-2 Events



5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the instance redeployment events in **Pending authorization** status.

- a. In the **Event Type** column, select **Instance redeployment**.
- b. In the **Event Status** column, select **Pending authorization**.
6. Click **Authorize Redeployment** in the **Operation** column.

Figure 13-3 Authorizing redeployment**NOTE**

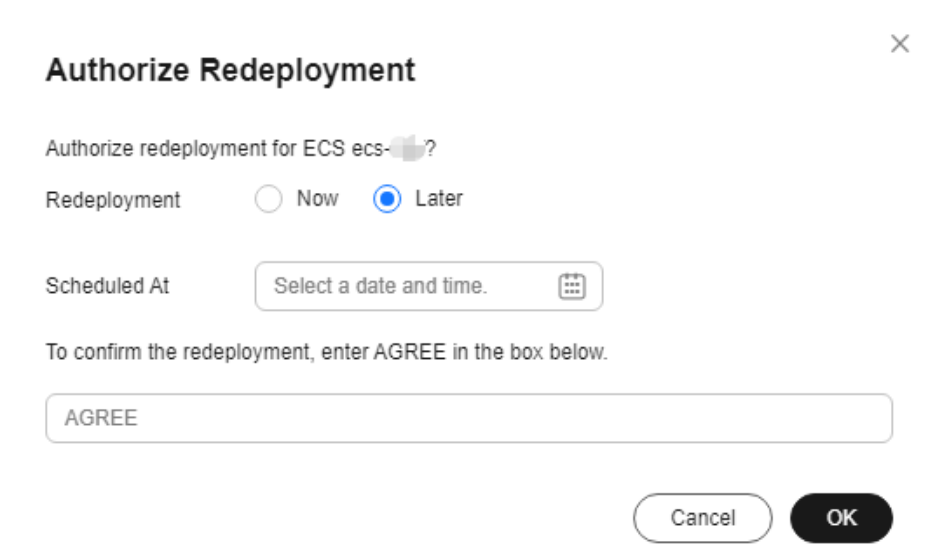
Before redeploying an instance, you need to learn about the notice.

7. Select the checkbox and click **Next**.
8. Specify a redeployment time.
 - Select **Now**.
The system redeploys the instance upon receiving the request.

Figure 13-4 Selecting Now

- Select **Later** and specify **Scheduled At**.
The system will redeploy the instance at the scheduled redeployment time.

Figure 13-5 Selecting Later



You can change the scheduled time as needed. For details, see [Changing the Reservation Time](#).

9.


Enter **AGREE** in the text box below to confirm the redeployment.
10.


Click **OK**.
- After you submit the authorization request, the event status becomes **To be executed**. After the system receives the request, it decides when to change the event status to **Executing** based on the redeployment time you specify and starts to redeploy the instance.
- When the status of the instance redeployment event becomes **Execution succeeded**, the instance redeployment is complete. Check the status of services running on the instance.

Authorize Redeployment for Instances Using Local Disks

1.

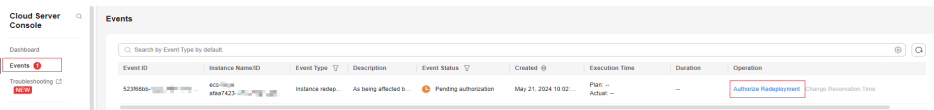
Log in to the management console.
2.

Click  in the upper left corner and select your region and project.
3.


Click . Under **Compute**, choose **Elastic Cloud Server**.
4.

Choose **Events** from the left navigation pane.

Figure 13-6 Events



5.

On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the instance redeployment events in **Pending authorization** status.
- a.

In the **Event Type** column, select **Instance redeployment**.

- b. In the **Event Status** column, select **Pending authorization**.
6. Click **Authorize Redeployment** in the **Operation** column.

Figure 13-7 Authorizing redeployment

Authorize Redeployment

Notice:

- An instance redeployment will restart the instance. Switch service traffic in advance and select an appropriate time for authorizing the redeployment to minimize the impact on your services.
- The instance redeployment will delete data from all local disks (EBS, Instance Store). Back up the data before authorizing the redeployment.
- Before authorizing the redeployment, you need to perform local disk preprocessing to prevent the instance from starting in emergency mode.

☐ I understand the redeployment will cause an instance restart.

☐ I understand data on the local disk (EBS, Instance Store) will be deleted.

☐ I have completed local disk preprocessing and agree to continue.

Cancel Next

 **NOTE**

Before redeploying an instance, you need to understand risks, complete local disk preprocessing, and add parameter **nofail** to all data disks.

For details, see [Preprocessing for Instance Redeployment](#).


7. Select the checkbox and click **Next**.
8. Specify a redeployment time.

- **Select Now.**

The system redeploys the instance upon receiving the request.

Figure 13-8 Selecting Now

Authorize Redeployment

Authorize redeployment for ECS ecs-?


Redeployment ☒ Now ☐ Later

Make sure you understand the data loss risks. To authorize the redeployment, enter AGREE in the box below.

Cancel

OK

- Select **Later** and specify **Scheduled At**.
The system will redeploy the instance at the scheduled redeployment time.

Figure 13-9 Selecting Later

You can change the scheduled time as needed. For details, see [Changing the Reservation Time](#).




9. Enter **AGREE** in the text box below to authorize the redeployment.
10. Click **OK**.

After you submit the authorization request, the event status becomes **To be executed**. After the system receives the request, it decides when to change the event status to **Executing** based on the redeployment time you specify and starts to redeploy the instance.

When the status of the instance redeployment event becomes **Execution succeeded**, the instance redeployment is complete. Check the status of services running on the instance.

Changing the Reservation Time

After you submit a scheduled authorization request, the event status changes to **To be executed**. You can change the reservation time as needed.

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. Choose **Events** from the left navigation pane.
5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the instance redeployment events in **To be executed** state.

- a. In the **Event Type** column, select **Instance redeployment**.
- b. In the **Event Status** column, select **To be executed**.
6. In the **Operation** column of the event to be executed, click **Change Reservation Time**.
7. Specify **New Reservation Time** and click **OK**.

 **NOTE**

The reservation time can be changed at least 30 minutes before the reservation deadline.

13.3.2 Handling a Local Disk Replacement Event

Scenarios

When the system detects that there are hardware or software faults in local disks of the host accommodating ECSs (including bare metal ECSs), it will automatically generate a local disk replacement event.

You can handle a local disk replacement event in **Pending authorization** status in either of the following ways:

- **Authorizing disk replacement:** You can authorize disk replacement to replace the faulty local disk.

You can authorize disk replacement now or later.

- **Now:** The disk replacement will be completed within five working days generally after the system receives the request.
- **Later:** The disk replacement will be completed within five working days after the scheduled replacement time.

You can change the reservation time as needed. For details, see [Changing the Reservation Time](#).

NOTICE

- Replacing local disks will not affect the instance's system disk, EVS data disks, and other normal local disks.
- During the local disk replacement, the faulty local disk will be detached, and all data stored on the faulty local disk will be deleted. To ensure data security, back up local disk data in advance.
- For disk-intensive (D series) ECSs, during the local disk replacement, the instances will not be stopped. Only the faulty local disk is unavailable.
- Replacing local disks for bare metal ECSs may restart or power off hosts. Make sure that services have been stopped or they will not be affected by the shutdown of instances.

If services on the instance cannot be stopped, contact technical support.

- Before authorizing the disk replacement, perform operations described in [Preprocessing for Local Disk Replacement](#) to prevent exceptions.

-
- **Redeployment:** You can choose to redeploy the instance.

NOTICE

- Instance redeployment will restart the instance. Switch service traffic in advance and select an appropriate time for the redeployment to minimize the impact on your services.
- Instance redeployment will not affect the instance's system disk and EVS data disks.
- Instance redeployment will delete data from all local disks. To ensure data security, back up data before the redeployment.
- Before the instance redeployment, you need to **preprocess local disks** to prevent the instance from starting in emergency mode.

Prerequisites

If you need to perform operations as an IAM user, ensure that the IAM user has been granted the required permissions.

Event query and handling require the permissions defined in the following policies:

- Query events: **ecs:instanceScheduledEvents:list**
- Accept and authorize events: **ecs:instanceScheduledEvents:accept**
- Change the reservation time: **ecs:instanceScheduledEvents:update**

For details about how to grant permissions to IAM users, see [ECS Custom Policies](#).

Authorizing Disk Replacement



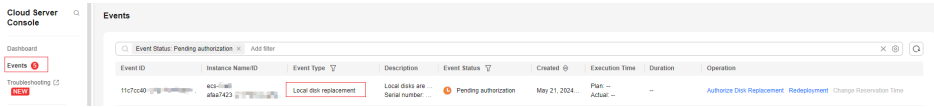
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Choose **Events** from the left navigation pane.

Figure 13-10 Events




5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the local disk replacement events in **Pending authorization** status.
 - a. In the **Event Type** column, select **Local disk replacement**.
 - b. In the **Event Status** column, select **Pending authorization**.
6. Click **Authorize Disk Replacement** in the **Operation** column.

Figure 13-11 Authorizing disk replacement (for disk-intensive ECSs)

Authorize Disk Replacement ×

Notice:

- During the local disk replacement, the faulty local disk (Serial number: [REDACTED];WWN: [REDACTED]) will be detached and unavailable. The instance will not be restarted and other disks will not be affected. [Learn how to query disk device names](#)
- During the local disk replacement, data on the faulty local disk (Serial number: [REDACTED];WWN: [REDACTED]) will be deleted. Back up the data in advance.
- Before authorizing the disk replacement, perform [disk replacement preprocessing](#) to prevent exceptions.

☐ I understand the faulty local disk is unavailable during the disk replacement.

☐ I understand data on the faulty local disk (Serial number: [REDACTED];WWN: [REDACTED]) will be deleted.

☐ I have completed local disk preprocessing and agree to continue.

Cancel

Next

NOTE

You need to understand the risks and complete the disk replacement preprocessing before authorizing disk replacement.

- If you want to retain the local disk data, back up data stored on the faulty local disk first.
- In the `/etc/fstab` configuration file, comment out the mount point of the faulty local disk to prevent the instance from entering the maintenance mode after the faulty disk is replaced.
- Run the `umount` command in the operating system to unmount the device or file system of the faulty local disk.

Complete the disk replacement preprocessing based on your instance type. For details, see [Preprocessing for Local Disk Replacement](#).


7. Select the checkbox and click **Next**.

8. Specify **Authorization**.

- Select **Now**.

The disk replacement will be completed within five working days generally after the system receives the request.

Figure 13-12 Selecting Now



Authorize Disk Replacement ✕

Are you sure you have completed the disk replacement preprocessing and you want to authorize the replacement of local disk (██████████) attached to instance "ecs-██████████"?

Authorization ☒ Now ☐ Later

The disk replacement will be completed within five working days generally after the system receives the request.


Make sure you understand the data loss risks. To authorize the redeployment, enter AGREE in the box below.

AGREE

Cancel OK

- Select **Later** and specify **Scheduled Time**.
The disk replacement will be completed within five working days after the scheduled replacement time.

Figure 13-13 Selecting Later




Authorize Disk Replacement ✕

Are you sure you have completed the disk replacement preprocessing and you want to authorize the replacement of local disk (██████████) attached to instance "ecs-██████████"?

Authorization ☐ Now ☒ Later

The disk replacement will be completed within five working days generally after the scheduled replacement time.

Scheduled Time 

Make sure you understand the data loss risks. To authorize the redeployment, enter AGREE in the box below.

AGREE

Cancel OK

You can change the reservation time as needed. For details, see [Changing the Reservation Time](#).

9. Enter **AGREE** in the text box below to authorize the local disk replacement.
10. Click **OK**.

After you submit the authorization request, the event status becomes **To be executed**. After the system receives the request, it decides when to change the event status to **Executing** based on the replacement time you specify.

When the status of the local disk replacement event becomes **Execution succeeded**, the local disk replacement is complete. Check the status of services running on the instance.

Redeployment



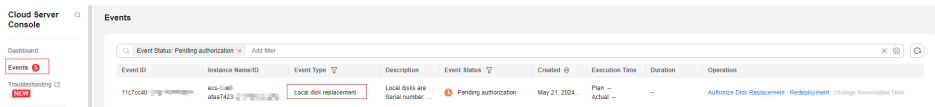
- 1. Log in to the management console.
- 2. Click  in the upper left corner and select your region and project.
- 3. Click . Under **Compute**, choose **Elastic Cloud Server**.
- 4. Choose **Events** from the left navigation pane.

Figure 13-14 Events




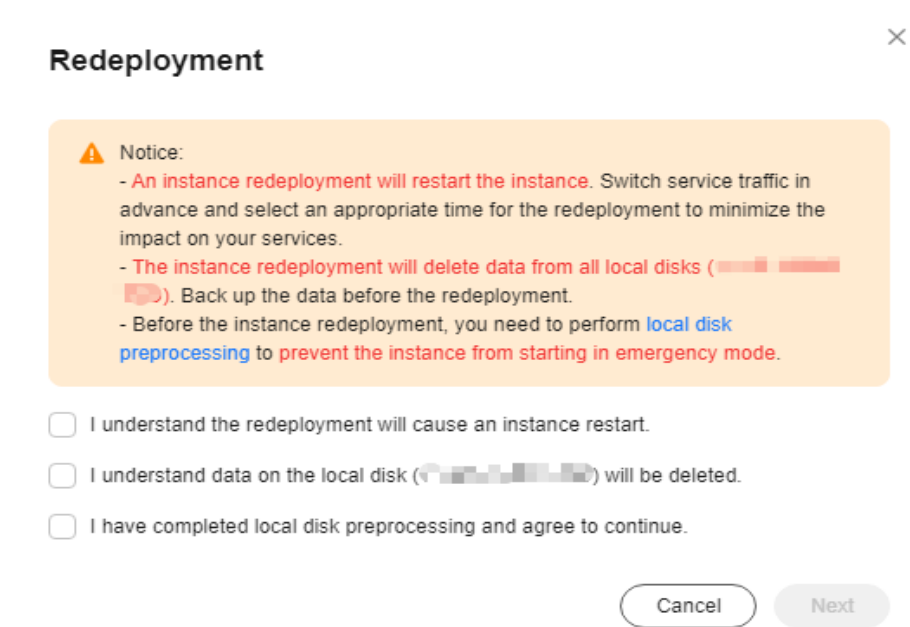
- 5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the local disk replacement events in **Pending authorization** status.
 - a. In the **Event Type** column, select **Local disk replacement**.
 - b. In the **Event Status** column, select **Pending authorization**.
- 6. Click **Redeployment** in the **Operation** column.

Figure 13-15 Redeployment



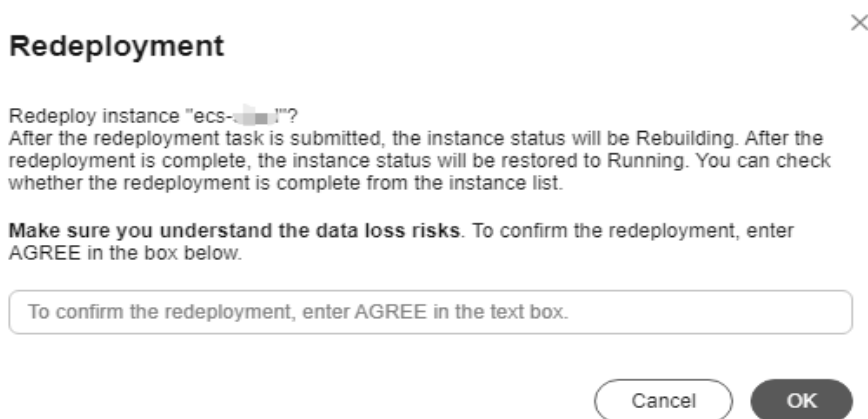
 **NOTE**

Before redeploying an instance, you need to understand risks of local disk data loss, preprocess the instance redeployment, and add parameter **nofail** to all data disks.

For details, see [Preprocessing for Instance Redeployment](#).

7. Select the checkbox and click **Next**.
8. Enter **AGREE** to confirm the redeployment.




Figure 13-16 Confirming the instance redeployment



9. Click **OK**.
After you submit the redeployment request, the event status becomes **To be executed**. After the system receives the request, it changes the event status to **Executing** and starts to redeploy the instance.
When the status of the local disk replacement event becomes **Execution succeeded**, the instance redeployment is complete. Check the status of services running on the instance.

Changing the Reservation Time

After you submit a scheduled authorization request, the event status changes to **To be executed**. You can change the reservation time as needed.

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. Choose **Events** from the left navigation pane.
5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the local disk replacement events in **To be executed** state.
 - a. In the **Event Type** column, select **Local disk replacement**.
 - b. In the **Event Status** column, select **To be executed**.
6. In the **Operation** column of the event to be executed, click **Change Reservation Time**.

7. Specify **New Reservation Time** and click **OK**.

 **NOTE**

The reservation time can be changed at least 30 minutes before the reservation deadline.

13.3.3 Handling a System Maintenance Event

Scenarios

When the system reports a system maintenance event, you need to authorize maintenance for ECSs (including bare metal ECSs).

This section describes how to handle a system maintenance event in **Pending authorization** status and guides you through the authorization required for the system maintenance.

You can authorize maintenance now or later.

- **Now:** The system maintenance will be completed within five working days generally after the system receives the request.
- **Later:** The system maintenance will be completed within five working days after the scheduled maintenance time.

You can change the reservation time as needed. For details, see [Changing the Reservation Time](#).

NOTICE

System maintenance requires the host to be powered off. The instance will be unavailable during this period. Make sure that services have been stopped or they will not be affected by the shutdown of the instance.

If services on the instance cannot be stopped, contact technical support.

Prerequisites


If you need to perform operations as an IAM user, ensure that the IAM user has been granted the required permissions.

Event query and handling require the permissions defined in the following policies:

- Query events: **ecs:instanceScheduledEvents:list**
- Accept and authorize events: **ecs:instanceScheduledEvents:accept**
- Change the reservation time: **ecs:instanceScheduledEvents:update**

For details about how to grant permissions to IAM users, see [ECS Custom Policies](#).

Authorizing Maintenance

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. Choose **Events** from the left navigation pane.

Figure 13-17 Events




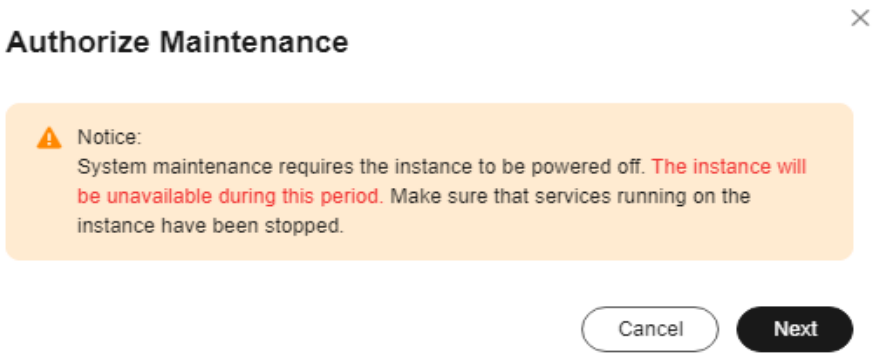

5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the system maintenance events in **Pending authorization** status.
- a. In the **Event Type** column, select **System maintenance**.
- b. In the **Event Status** column, select **Pending authorization**.
6. Click **Authorize Maintenance** in the **Operation** column.

Figure 13-18 Authorizing maintenance




7. Learn about the notice and click **Next**.
8. Specify **Maintenance**.
- Select **Now**.
- The system maintenance will be completed within five working days generally after the system receives the request.

Figure 13-19 Selecting Now

The screenshot shows a dialog box titled "Authorize Maintenance" with a close button (X) in the top right corner. The text inside reads: "Make sure that services have been stopped and authorize maintenance for cloud server ecs-? ?". Below this, there are two radio buttons: "Now" (which is selected) and "Later". A message states: "To confirm the maintenance authorization, enter AGREE in the box below." Below the message is a text input field containing the word "AGREE". At the bottom right, there are two buttons: "Cancel" and "OK".

- Select **Later** and specify **Scheduled At**.
The system maintenance will be completed within five working days after the scheduled maintenance time.

Figure 13-20 Selecting Later

The screenshot shows the same "Authorize Maintenance" dialog box, but with the "Later" radio button selected. Below the radio buttons, there is a "Scheduled At" label followed by a date and time picker button that says "Select a date and time." with a calendar icon. The rest of the dialog, including the "AGREE" text box and "Cancel/OK" buttons, remains the same.

You can change the reservation time as needed. For details, see [Changing the Reservation Time](#).

9. Enter **AGREE** in the text box below to confirm the maintenance authorization.
10. Click **OK**.

After you submit the authorization request, the event status becomes **To be executed**. After the system receives the request, it decides when to change the event status to **Executing** based on the maintenance time you specify.

When the status of the system maintenance event becomes **Execution succeeded**, the system maintenance is complete. Check the status of services running on the instance.

Changing the Reservation Time

After you submit a scheduled authorization request, the event status changes to **To be executed**. You can change the reservation time as needed.





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. Choose **Events** from the left navigation pane.
5. On the **Events** page, in the event list, click  in table headers and set the following criteria to filter the system maintenance events in **To be executed** state.
 - a. In the **Event Type** column, select **System maintenance**.
 - b. In the **Event Status** column, select **To be executed**.
6. In the **Operation** column of the event to be executed, click **Change Reservation Time**.

Figure 13-21 Changing the reservation time



The dialog box titled "Change Reservation Time" has a close button (X) in the top right corner. It contains three rows of information: "Current Reservation Time" with the value "May 22, 2024 10:18:26 GMT+08:00", "New Reservation Time" with a text input field containing "May 22, 2024 10:18:26" and a calendar icon, and "Reservation Deadline" with the value "May 29, 2024 10:18:26 GMT+08:00". At the bottom right, there are two buttons: "Cancel" and "OK".

7. Specify **New Reservation Time** and click **OK**.

NOTE

The reservation time can be changed at least 30 minutes before the reservation deadline.

13.4 Preprocessing for Local Disk Replacement

13.4.1 Preprocessing for Disk-intensive ECSs

Scenarios

Before authorizing the replacement of a faulty local disk, you need to preprocess the local disk.

This section describes how to preprocess a faulty local disk of a disk-intensive ECS (D series) before authorizing the replacement of this disk.



NOTICE

Replacing a local disk will lose data on it. If you do not need these data, perform operations in this section to preprocess the disk to be replaced.

If you need to retain the data, stop your operation here and contact technical support.

Procedure (Linux ECS)

Obtain the WWN of the faulty local disk.

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, choose **Events**.
5. In the event list, check **Description** of the **Local disk replacement** event to obtain the WWN of the faulty local disk.

For example, the WWN is **wwn-0x5000c500e01a4930**.

Obtain the mount point of the faulty local disk.

1. [Remotely log in to](#) the ECS whose disk needs to be preprocessed.
2. Obtain the drive letter matching the WWN of the faulty disk.

ll /dev/disk/by-id/ | grep WWN

For example, if the WWN is **wwn-0x5000c500e01a4930**, run the following command:

ll /dev/disk/by-id/ | grep wwn-0x5000c500e01a4930

The following information is displayed:

```
lrwxrwxrwx 1 root root 9 May 13 14:05 wwn-0x5000c500e01a4930 -> ../../sdb
```

3. Obtain the mount point of the faulty local disk.

df -Th | grep drive-letter

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

df -Th | grep /dev/sdb

The following information is displayed:

```
/dev/sdb ext4 3.6T 28K 3.4T 1% /data
```

NOTICE

Preprocessing a local disk to be replaced will lose data on it. If you need to retain the data, stop your operation here and contact technical support.

Unmount the faulty local disk.

1. Unmount the faulty local disk.

umount *mount-point*

For example, if the mount point is **/data**, run the following command:

umount /data

Check whether automatic mounting of the faulty disk is configured in `/etc/fstab` of the ECS.

 **NOTE**

If automatic mounting is configured for the faulty disk in **`/etc/fstab`** of the ECS, delete or comment out the configuration so that the ECS will not enter the emergency mode after its local disk is replaced.

1. Obtain the UUID of the disk partition.

blkid *disk-partition*

For example, if the disk partition is **`/dev/sdb`**, run the following command:

blkid /dev/sdb

The following information is displayed:

```
/dev/sdb: UUID="626c4774-e60e-4d86-bbe6-031bac126e4c" TYPE="ext4"
```

2. Check whether **`/etc/fstab`** contains information about automatic mounting of the disk partition.

cat /etc/fstab | grep *disk-partition-UUID*

For example, if the UUID is **`626c4774-e60e-4d86-bbe6-031bac126e4c`**, run the following command:

cat /etc/fstab | grep 626c4774-e60e-4d86-bbe6-031bac126e4c

If the following information is displayed, the **`/etc/fstab`** file contains information about automatic mounting of the disk partition:

```
UUID=626c4774-e60e-4d86-bbe6-031bac126e4c /mnt ext4 defaults 0 0
```

3. In the **`/etc/fstab`** file, delete or comment out the configuration of automatic mounting for the disk partition to prevent the ECS from entering the emergency mode upon startup after its local disk is replaced.

- a. Open the **`/etc/fstab`** file.

vi **`/etc/fstab`**

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

- c. Delete or comment out the configuration of automatic mounting found in [2](#).


To comment out the configuration, add **#** in front of it. For example:


```
# UUID=626c4774-e60e-4d86-bbe6-031bac126e4c /mnt ext4 defaults 0 0
```

- d. Press **Esc**, enter **:wq**, and press **Enter** to save the change and exit the editing mode.

Procedure (Windows ECS)

Obtain the SN of the faulty local disk.

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, choose **Events**.
5. In the event list, check **Description** of the **Local disk replacement** event to obtain the SN of the faulty local disk.

For example, the SN is **WS22LX16**.

Obtain the disk ID to confirm the disk to be replaced.

1. [Remotely log in to](#) the ECS whose disk needs to be preprocessed.
2. Open **Windows PowerShell** as an administrator and obtain the disk ID based on the SN.

Get-Disk | select Number, SerialNumber

Figure 13-22 Checking the mapping between the disk ID and SN

```
PS C:\Users\Administrator> Get-Disk | select Number, SerialNumber
Number SerialNumber
-----
1      WS22LX16
2      WS22LWC7
0 55a3514f-5288-444c-9
```

NOTICE

Replacing a local disk will lose data on it. If you need to retain the data, stop your operation here and contact technical support.

13.4.2 Preprocessing for Bare Metal ECSs

Scenarios

Before authorizing the replacement of a faulty local disk, you need to preprocess the local disk.

This section describes how to preprocess a faulty local disk of a bare metal ECS before authorizing the replacement of this disk.

NOTICE



Replacing a local disk will lose data on it. If you do not need these data, perform operations in this section to preprocess the disk to be replaced.

Replacing a local disk of a bare metal ECS may restart or power off the host. Before you preprocess a local disk, ensure that services on the bare metal ECS are stopped or that stopping the bare metal ECS during the replacement will not affect services running on it.

If you need to retain data on the local disk or if services on the bare metal ECS cannot be stopped, stop your operation here and contact technical support.

Procedure (Linux ECS)

Obtain the SN of the faulty local disk.

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, choose **Events**.
5. In the event list, check **Description** of the **Local disk replacement** event to obtain the SN of the faulty local disk.

For example, the SN is **99K0A046FL3G**.

Obtain the mount point of the faulty local disk.

1. [Remotely log in](#) to the bare metal ECS whose disk needs to be preprocessed.
2. Obtain the drive letter matching the SN of the faulty disk.

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

For example, if the SN is **99K0A046FL3G**, run the following command:

```
ll /dev/disk/by-id/ | grep 99K0A046FL3G
```

The following information is displayed:

```
/lrxxrrxx 1 root root 9 Sep 18 19: 20 ata-MG06ACA10TE_99K0A046FL3G ->././sda  
lrwxrwxrwx 1 root root 10 Sep 18 19: 20 ata-MG06ACA10TE_99K0A046FL3G-part1->././sda1
```

3. Obtain the mount point of the faulty local disk.

```
df -Th | grep drive-letter
```

For example, if the drive letter is **/dev/sda1**, run the following command:

```
df -Th| grep /dev/sda1
```

The following information is displayed:

```
/dev/sdal ext4 9.1T 7.5T 1.6T 83% /srv/data
```

NOTICE

Preprocessing a local disk to be replaced will lose data on it. If you need to retain the data, stop your operation here and contact technical support.

Unmount the faulty local disk.

1. Unmount the faulty local disk.

```
umount mount-point
```

For example, if the mount point is **/data**, run the following command:

```
umount /data
```

Check whether automatic mounting of the faulty disk is configured in `/etc/fstab` of the bare metal ECS.

 NOTE

If automatic mounting is configured for the faulty disk in **/etc/fstab** of the bare metal ECS, delete or comment out the configuration so that the bare metal ECS will not enter the emergency mode after its local disk is replaced.

1. Obtain the UUID of the disk partition.

blkid *disk-partition*

For example, if the disk partition is **/dev/sda1**, run the following command:

blkid /dev/sda1

The following information is displayed:

```
/dev/sda1: UUID="e7100f3e-af65-49da-a013-c4ace8e5aba7" TYPE
="ext4"RTLABEL="logical"PARTUUID="cd358d32-e02b-4b23-bbd9-8a8bdae0e070"
```

2. Check whether **/etc/fstab** contains information about automatic mounting of the disk partition.

cat /etc/fstab | grep disk-partition-UUID

For example, if the UUID is **e7100f3e-af65-49da-a013-c4ace8e5aba7**, run the following command:

cat /etc/fstab | grep e7100f3e-af65-49da-a013-c4ace8e5aba7

If the following information is displayed, the **/etc/fstab** file contains information about automatic mounting of the disk partition:

```
UUID=e7100f3e-af65-49da-a01 3-c4ace8e5aba7 /srv/data ext4 defaults,noatime, nodiratime 1 0
```

3. In the **/etc/fstab** file, delete or comment out the configuration of automatic mounting for the disk partition to prevent the ECS from entering the emergency mode upon startup after its local disk is replaced.

- a. Open the **/etc/fstab** file.

vi /etc/fstab

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

- c. Delete or comment out the configuration of automatic mounting found in 2.

To comment out the configuration, add **#** in front of it. For example:

```
# UUID=e7100f3e-af65-49da-a01 3-c4ace8e5aba7/srv/data ext4 defaults,noatime, nodiratime 1
0
```

- d. Press **Esc**, enter **:wq**, and press **Enter** to save the change and exit the editing mode.

Stop the bare metal ECS.



1. In the event list, check the name or ID of the bare metal ECS (**Instance Name/ID** in the **Local disk replacement** event).
2. In the navigation pane, choose **Elastic Cloud Server**. Locate the bare metal ECS based on its name or ID. In the **Operation** column, choose **More > Stop**.

 NOTE

After you complete preprocessing the local disk, restart the bare metal ECS to synchronize the disk information to the virtualization layer.

Procedure (Windows ECS)

Obtain the SN of the faulty local disk.

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, choose **Events**.
5. In the event list, check **Description** of the **Local disk replacement** event to obtain the SN of the faulty local disk.
For example, the SN is **WS22LX16**.

Obtain the disk ID to confirm the disk to be replaced.

1. **Remotely log in** to the bare metal ECS whose disk needs to be preprocessed.
2. Open **Windows PowerShell** as an administrator and obtain the disk ID based on the SN.

Get-Disk | select Number, SerialNumber

Figure 13-23 Checking the mapping between the disk ID and SN

```
PS C:\Users\Administrator> Get-Disk | select Number, SerialNumber
Number SerialNumber
-----
1      WS22LX16
2      WS22LWC7
0 55a3514f-5288-444c-9
```

NOTICE

Replacing a local disk will lose data on it. If you need to retain the data, stop your operation here and contact technical support.

Stop the bare metal ECS.

1. In the event list, check the name or ID of the bare metal ECS (**Instance Name/ID** in the **Local disk replacement** event).
2. In the navigation pane, choose **Elastic Cloud Server**. Locate the bare metal ECS based on its name or ID. In the **Operation** column, choose **More > Stop**.

NOTE

After you complete preprocessing the local disk, restart the bare metal ECS to synchronize the disk information to the virtualization layer.

13.5 Preprocessing for Instance Redeployment

Scenarios

When you authorize the redeployment for an ECS with local disks attached in local disk replacement and instance redeployment events, you need to preprocess local disks before the instance redeployment to prevent risks caused by ECS exceptions.

This section describes how to preprocess local disks for disk-intensive (D series) and ultra-high I/O (I series) ECSs running Linux.

NOTICE

Instance redeployment will cause all data loss on the ECS local disks. If you do not need to retain data on the local disks, you can perform the operations below.

If you need to retain data on the local disks, do not authorize the redeployment. Instead, contact technical support.

Procedure

Perform the following operations to add **nofail** to all data disks in the **/etc/fstab** configuration file:

1. **Remotely log in to** the instance to be preprocessed.
2. Run the following command to open and edit the **/etc/fstab** configuration file:

vim /etc/fstab

3. Press **i** to enter editing mode and add parameter **nofail** to all data disks in the **/etc/fstab** configuration file.

Example:

```
UUID-8232fee7-f20a-416c-a2e0-cbc8c85a01a2 /mnt/nvme0n1 ext4 defaults,nofail 0 2
```

- **UUID-8232fee7-f20a-416c-a2e0-cbc8c85a01a2**: UUID of a data disk
- **/mnt/nvme0n1**: data disk mount point, which can be queried using **mount | grep drive-letter**.
- **ext4**: file system type of the data disk, which can be queried using **blkid drive-letter**.
- **nofail**: allows the system to ignore mounting of a missing drive during boot. When local disks contained in the file system are missing, the system continues booting.

4. Press **:wq** to exit editing mode.
5. Run the following command for the configuration to take effect:

systemctl daemon-reload

14 OS Dump

14.1 Configuring OS Dump

Scenarios

If your ECS OS becomes faulty, a system crash or blue screen of death (BSOD) may occur. In this case, you can configure OS dump to enable the memory data to be stored in a file. You can then use this information to locate the fault.

This section describes how to configure an OS dump.

- [Configuring an OS Dump \(Windows\)](#)
- [Configuring an OS Dump \(Linux\)](#)

Background

An OS dump is also known as a kernel dump or crash dump. When a system crash (usually a stop error) occurs, the memory data is completely backed up in real time. A complete memory dump is produced from this event.

The OS dump function relies on the OS kernel's capability of processing hardware exceptions. The OS kernel can detect hardware exceptions and determine the handling method based on the exception type. If the system encounters an exception or error that cannot be handled, some or all information will be dumped to disks. The information includes the CPU register, physical memory, process status, file system status, and hardware device status.

An OS dump can be triggered in the following ways:

- System crash: When an exception or error that cannot be handled occurs, such as invalid memory access or kernel panic, the system automatically triggers a dump.
- Manual triggering: The system administrator can use commands or operations to manually trigger a dump. For example, in Linux, the system administrator can add the parameter **c** to **/proc/sysrq-trigger** to trigger a dump.
- Scheduled triggering: The system administrator can schedule a dump at a specified time to simulate a system crash and obtain the system information at that time.

To obtain a dump file, you need to configure the OS of your ECS. For example, in the Linux OS, you need to configure kdump. The following describes how to manually trigger a dump.

Constraints

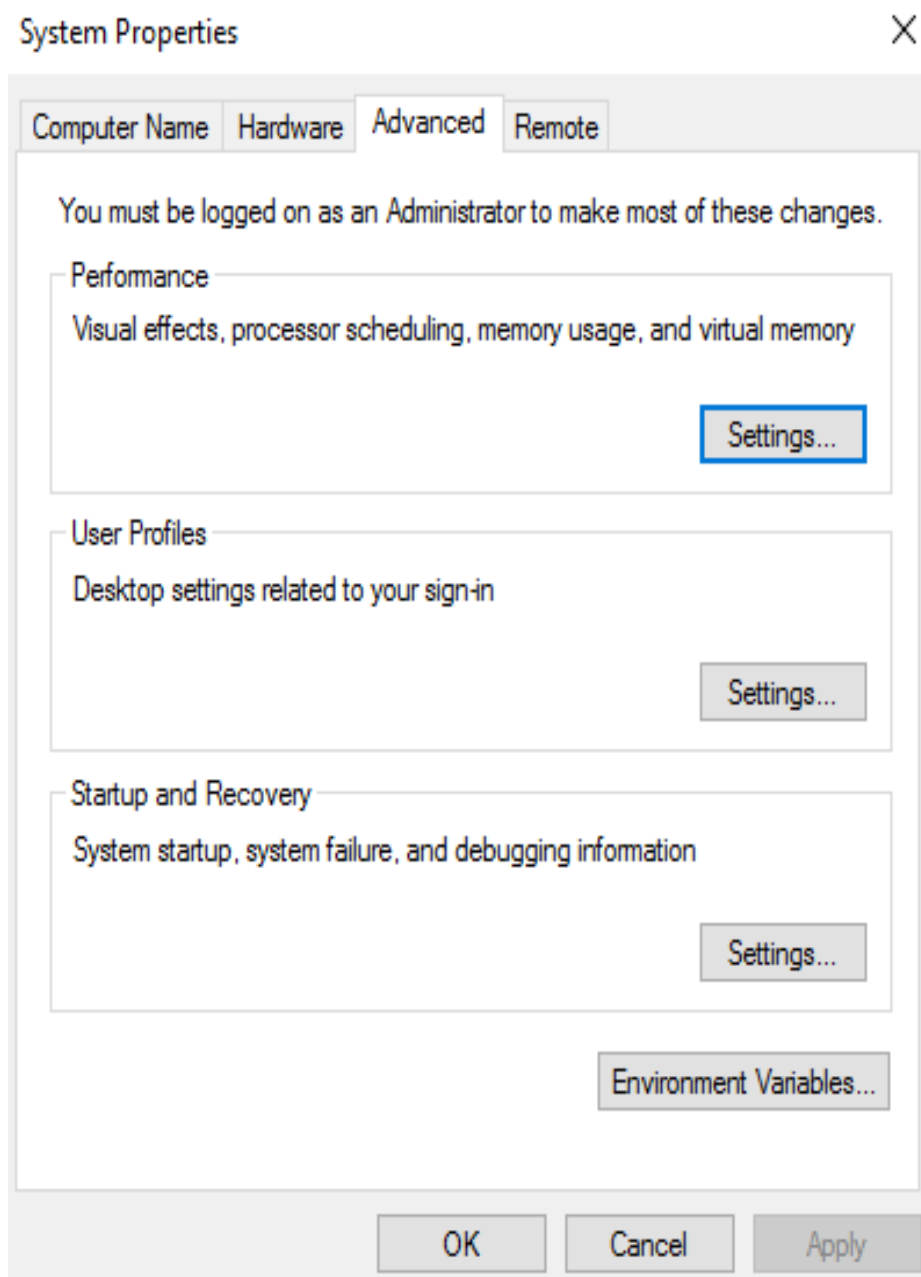
- An OS dump will cause the ECS to restart. Back up data in advance to prevent data loss.
- Kunpeng ECSs do not support OS dump.
- OS dump is only available for ECSs in **Running**, **Stopped**, or **Faulty** status.
- OS dump is only available for ECSs whose specifications are 4 vCPUs and 8 GiB of memory, or above.
- The reserved memory must be greater than the memory of the ECS.

Configuring an OS Dump (Windows)

The Windows Server 2016 is used as an example. For operations on more versions, see the help documents provided on the corresponding official website.

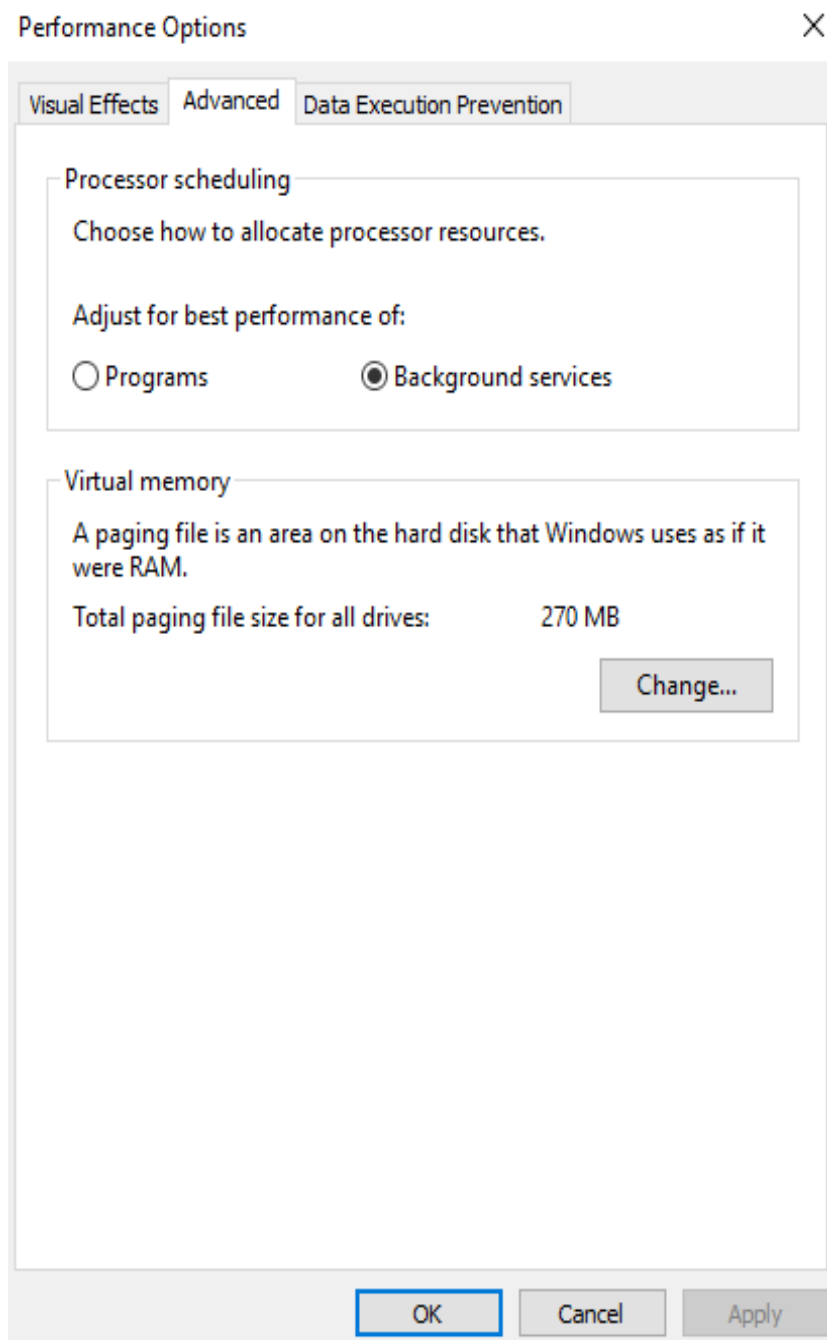
1. [Log in to a Windows ECS](#).
2. Open **Control Panel** and choose **System** > **Advanced system settings**.

Figure 14-1 Advanced system settings

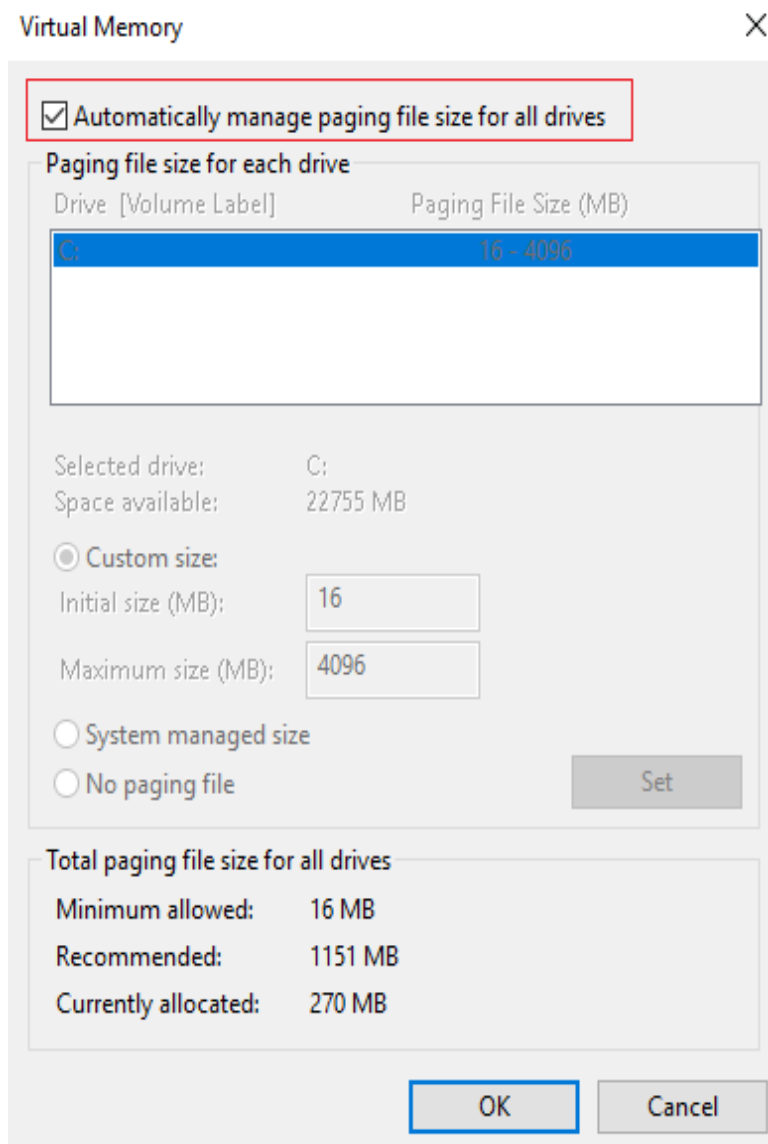


3. In the **Performance** area of the **System Properties** dialog box, click **Settings**.
4. In the **Virtual memory** area on the **Advanced** tab, click **Change**.

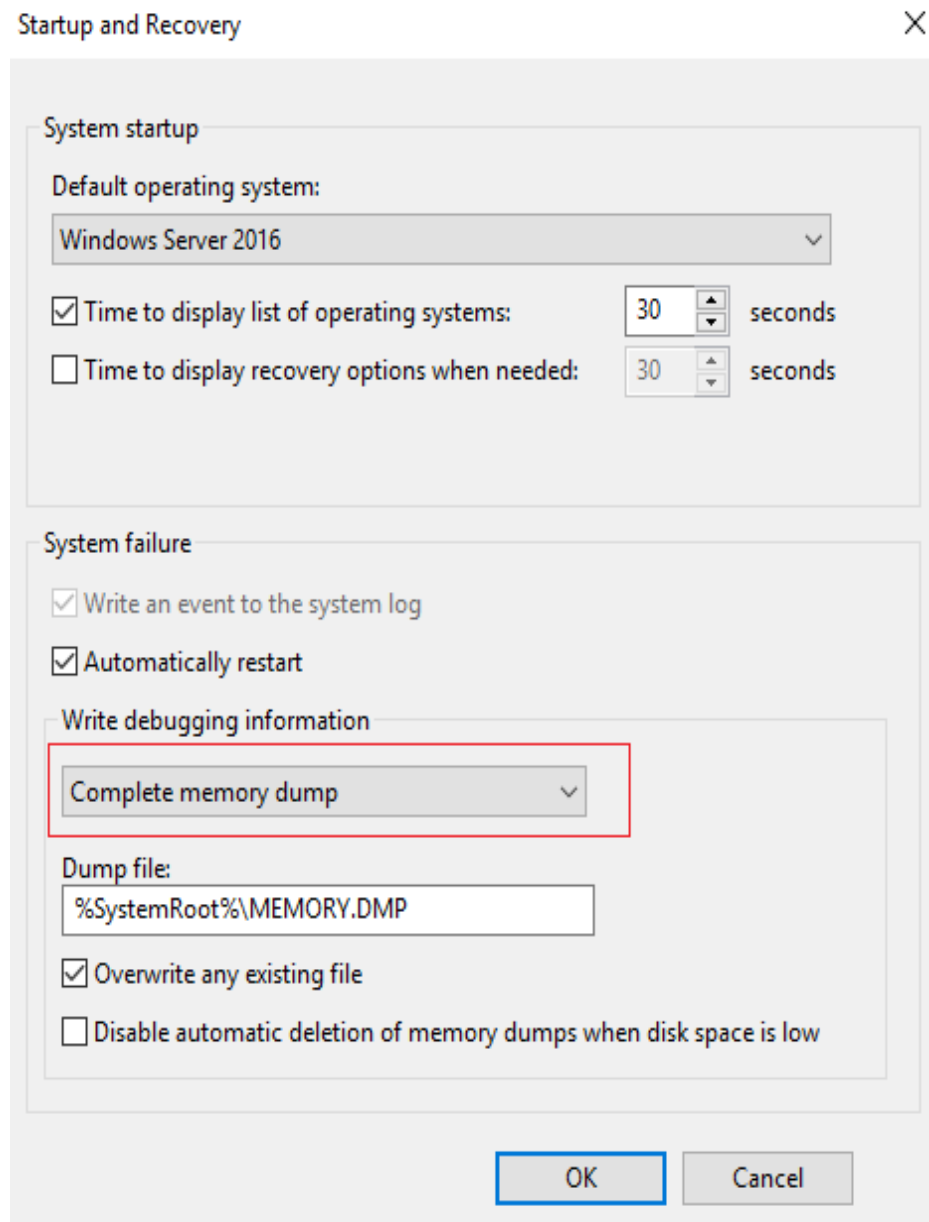
Figure 14-2 Performance options



5. In the **Virtual Memory** dialog box, select **Automatically manage paging file size for all drives**.

Figure 14-3 Virtual memory

6. Click **OK** to complete the virtual memory settings.
7. In the **Startup and Recovery** area of the **System Properties** dialog box, click **Settings**.
8. Configure required parameters.
 - **Write debugging information:** Select **Complete memory dump**.
 - **Dump file:** Change `%SystemRoot%\MEMORY.DMP` to a local drive with sufficient disk space, for example, `E:\Memory.dmp`.

Figure 14-4 Startup and recovery settings

9. Click **OK**.
10. Restart the ECS to complete the settings.

Configuring an OS Dump (Linux)

The CentOS 7.5 is used as an example. For operations on more versions, see the help documents provided on the corresponding official website.

1. [Log in to a Linux ECS](#).
2. Install kexec-tools.
yum install kexec-tools
3. Check the size of the memory reserved for **crashkernel**.
cat /etc/default/grub

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="net.ifnames=0 consoleblank=600 console=tty0 console=ttyS0,115200n8
spectre_v2=off nopti crashkernel=auto "
GRUB_DISABLE_RECOVERY="true"
```

4. Adjust the reserved memory size by changing the value of **crashkernel**.

vim /etc/default/grub

The reserved memory size for **crashkernel** is **auto**. You can adjust the reserved memory size by changing the value of **crashkernel** in the **GRUB_CMDLINE_LINUX** field.

```
GRUB_TIMEOUT=5
GRUB_DISTRIBUTOR="$(sed 's, release .*$,g' /etc/system-release)"
GRUB_DEFAULT=saved
GRUB_DISABLE_SUBMENU=true
GRUB_TERMINAL_OUTPUT="console"
GRUB_CMDLINE_LINUX="net.ifnames=0 consoleblank=600 console=tty0 console=ttyS0,115200n8
spectre_v2=off nopti crashkernel=256M"
GRUB_DISABLE_RECOVERY="true"
```

NOTE

You can determine the reserved memory size for **crashkernel** based on the system architecture. If the size is too small, the coredump file may fail to be generated.

In this example, the system memory is 1 GiB, and the reserved memory is 256 MB.

5. Generate a new grub configuration file.

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

6. Edit the kernel parameter file **/etc/sysctl.conf**.

vim /etc/sysctl.conf

Add the following parameters to the **/etc/sysctl.conf** file to switch the OS to the emergency mode and generate a memory dump when an NMI is received:

```
kernel.panic_on_unrecovered_nmi=1
kernel.unknown_nmi_panic=1
```

7. Reboot the system.

reboot

8. Open **/etc/kdump.conf** and modify the default configuration.

vim /etc/kdump.conf

```
# Specify the path where the coredump file is stored.
path /var/crash
# Add the parameter -c to compress the coredump file.
core_collector makedumpfile -c -l --message-level 1 -d 31
# Reboot the system after the coredump file is generated.
default reboot
```

9. Start kdump and enable auto-start upon system boot.

systemctl start kdump.service

systemctl enable kdump.service

10. Check whether kdump is started.

service kdump status



```
Redirecting to /bin/systemctl status kdump.service
● kdump.service - Crash recovery kernel arming
   Loaded: loaded (/usr/lib/systemd/system/kdump.service; enabled; vendor preset: enabled)
   Active: active (exited) since Wed 2024-05-22 21:19:47 CST; 11min ago
```

```
Main PID: 591 (code=exited, status=0/SUCCESS)
CGroup: /system.slice/kdump.service
```

```
systemctl is-active kdump.service
```

```
active
```

Triggering an OS Dump

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the ECS list, locate the row containing the target ECS and choose **More > Troubleshoot > Trigger OS Dump** in the **Operation** column.
5. Learn about the OS dump risks and select the checkbox.
6. Click **OK** to trigger an OS dump.

14.2 Viewing a Dump File

Scenarios

After triggering an OS dump, you can view the generated memory dump file on your ECS.



This section describes how to trigger an OS dump and view the generated dump file.

- [Triggering an OS Dump](#)
- [Viewing a Dump File \(Windows\)](#)
- [Viewing a Dump File \(Linux\)](#)

Prerequisites

Operations described in [Configuring OS Dump](#) have been completed.

Triggering an OS Dump

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the ECS list, locate the row containing the target ECS and choose **More > Troubleshoot > Trigger OS Dump** in the **Operation** column.
5. Learn about the OS dump risks and select the checkbox.
6. Click **OK** to trigger an OS dump.

Viewing a Dump File (Windows)

The Windows Server 2016 is used as an example. For operations on more versions, see the help documents provided on the corresponding official website.

1. [Log in to a Windows ECS](#).
2. Go to the path specified in [Configuring an OS Dump \(Windows\)](#) to view the generated dump file.
For example, view the generated **memory.dmp** file in **C:\Windows**.
You can use tools to check and analyze the dump files. For details, see [Analyzing a Kernel-Mode Dump File](#).

Viewing a Dump File (Linux)

The CentOS 7.5 is used as an example. For operations on more versions, see the help documents provided on the corresponding official website.

1. [Log in to a Linux ECS](#).
2. Check whether a dump file is generated.

ls /var/crash/

```
127.0.0.1-2024-05-22-21:35:26
```

You can use the crash tool in Linux to analyze the dump files. For details, see [Crash](#).

15 Self-Service O&M

15.1 Configuring Custom Policies for ECS Self-Service O&M

Scenario

ECS self-service O&M depends on Cloud Operations Center (COC). After COC is enabled and authorized, two agencies "ServiceLinkedAgencyForCOC" and "ServiceAgencyForCOC" will be automatically created and granted permissions required to perform operations on behalf of COC. For details, see [Enabling COC](#).

If you want to perform self-service O&M on ECSs as an IAM user, contact the account which is used to create the IAM user to assign COC permissions to the user.

This section describes how to assign COC permissions to an IAM user.

Prerequisites

The IAM user has been granted ECS permissions.

NOTE

If the assigned ECS permissions do not meet your requirements, assign permissions to the IAM user by referring to [Creating a User and Granting ECS Permissions](#).

Procedure

1. Log in to the IAM console and access the [Policies/Roles](#) page using the account.
2. In the upper right corner of the page, click **Create Custom Policy** and assign COC permissions.

For details, see [Creating a Custom Policy](#).

The following example describes how to add a custom policy for enabling COC and performing operations on COC.

- a. (Optional) Create a custom policy named **Enable COC** and assign the following permissions:

 **NOTE**

If the account used to create the IAM user has enabled COC, you can skip this step.

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:agencies:list*",
        "iam:agencies:createAgency",
        "iam:agencies:createServiceLinkedAgencyV5",
        "coc:agency:get",
        "coc:agency:create",
        "iam:permissions:grantRoleToAgency",
        "iam:permissions:grantRoleToAgencyOnDomain",
        "iam:roles:listRoles"
      ]
    }
  ]
}
```


- b. Create a custom policy named **COC Operations** and assign the following permissions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "coc:instance:listResources",
        "coc:application:listResources",
        "coc:schedule:list",
        "coc:schedule:enable",
        "coc:schedule:update",
        "coc:schedule:disable",
        "coc:schedule:approve",
        "coc:schedule:create",
        "coc:schedule:delete",
        "coc:schedule:count",
        "coc:schedule:get",
        "coc:schedule:getHistories",
        "coc:application:GetDiagnosisTaskDetails",
        "coc:application:CreateDiagnosisTask",
        "coc:document:create",
        "coc:document:listRunbookAtomics",
        "coc:document:getRunbookAtomicDetails",
        "coc:document:list",
        "coc:document:delete",
        "coc:document:update",
        "coc:document:get",
        "coc:document:analyzeRisk",
        "coc:instance:autoBatchInstances",
        "coc:instance:executeDocument",
        "coc:instance:start",
        "coc:instance:reboot",
        "coc:instance:stop",
        "coc:job:action",
        "coc:instance:reinstallOS",
        "coc:instance:changeOS",
        "coc:hostAccount:describe",
        "coc:instance:syncResources"
      ]
    }
  ]
}
```

```
}  
}
```

3. Create a user group and assign COC permissions to it.
For details, see [Creating a User Group and Assigning Permissions](#).
Assign the following permissions to the user group:
 - COC ReadOnlyAccess
 - Custom policy created in 2 for enabling COC and performing operations on COC

Figure 15-1 Assigning permissions to the user group



4. On the **Users** page of the IAM console, locate the target user and click **Authorize** in the **Operation** to assign COC permissions to the user.
For details, see [Assigning Permissions to an IAM User](#).
On the displayed page, select the user group created in 3 so that the user inherits permissions from the group.

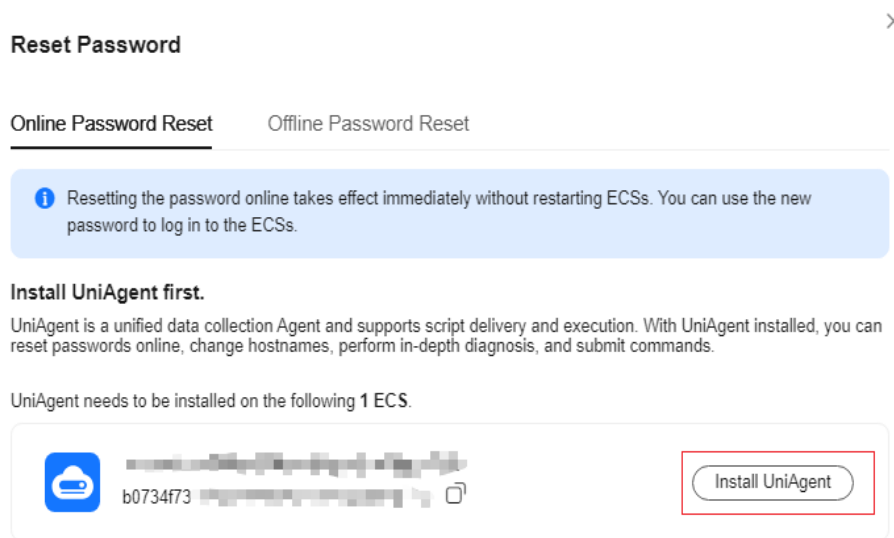
15.2 Installing UniAgent on an ECS

Scenario

Self-service O&M operations on ECSs such as hostname change and online password reset depend on UniAgent provided by Application Operations Management (AOM).

If UniAgent is not installed or the installation failed, you need to log in to the ECS and install UniAgent before performing self-service O&M.

Figure 15-2 Displayed message (online password reset for an individual ECS used as an example)



You can install UniAgent in either of the following ways:

- **Installing UniAgent on the ECS Console:** When you intend to perform self-service O&M operations on an individual ECS and a message is displayed indicating that UniAgent is not installed or the installation failed, install UniAgent on the ECS console.
- **Installing UniAgent on the Collection Management Console:** When you intend to perform self-service O&M operations on multiple ECSs and a message is displayed indicating that UniAgent is not installed or the installation failed, install UniAgent on the Collection Management console.

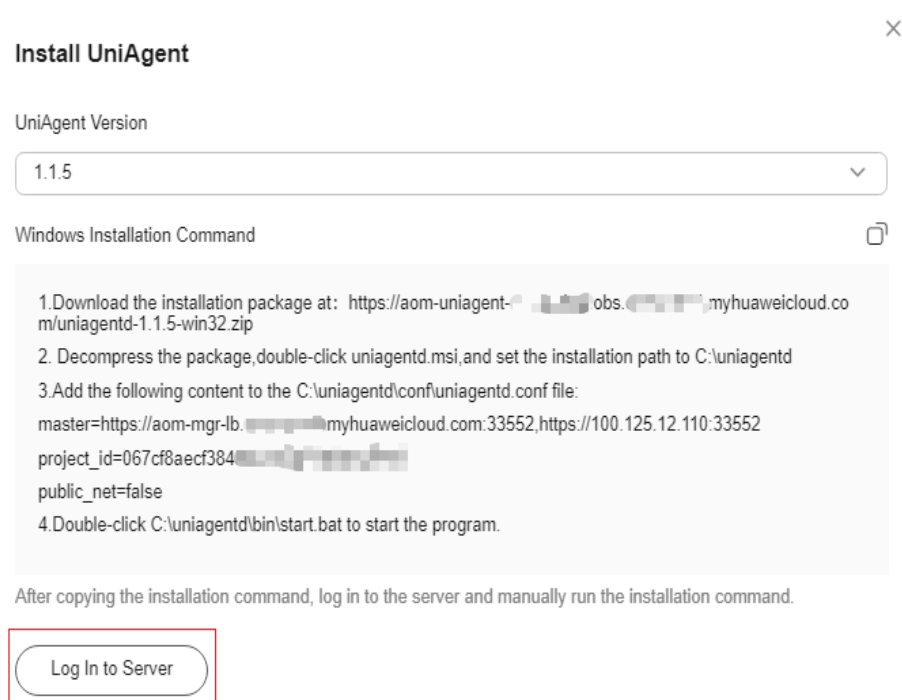
Prerequisites

AOM 2.0 has been subscribed.

Installing UniAgent on the ECS Console


When you intend to perform self-service O&M operations on an individual ECS and a message is displayed indicating that UniAgent is not installed or the installation failed, do as follows to install UniAgent:

1. In the slide-out panel of the ECS console, click **Install UniAgent**. The **Install UniAgent** panel slides out from the right.

Figure 15-3 Installing UniAgent (using a Windows ECS as an example)

2. Select a UniAgent version. You are advised to select the latest version.
3. Click the copy button in the upper right corner of the installation command area to copy the installation command.
4. Click **Log In to Server** to log in to the ECS.
5. Run the UniAgent installation command on the ECS.
 - Windows ECSs: download the installation package and install UniAgent as instructed.
 - Linux ECSs: paste and execute the installation command copied in [3](#).
6. Return to the ECS console and click **Check** in the panel to check the UniAgent installation status.

After UniAgent is installed, wait for several minutes until the status is synchronized and click **Check** to update and check the UniAgent installation status. If the installation is complete, you can perform self-service O&M operations as needed.

Figure 15-4 Checking UniAgent installation

Installing UniAgent on the Collection Management Console

When you intend to perform self-service O&M operations on multiple ECSs and a message is displayed indicating that UniAgent is not installed or the installation failed, do as follows to install UniAgent:

1. In the slide-out panel of the ECS console, click **Install UniAgent**. Then you will be redirected to the **Install UniAgent** page of the Collection Management console.

Figure 15-5 Installing UniAgent on the Collection Management console

| Host IP Address | OS | Login Account | Login Port | Authentication Mode | Password | Connectivity Test Result | Operation |
|-----------------|-------|---------------|------------|---------------------|----------|--------------------------|-------------------------------|
| | Linux | root | 22 | Password | | -- | Delete Copy Test Connectivity |

2. Set parameters required to install UniAgent.
For details, see [Installing a UniAgent](#).
3. Return to the ECS console and click **Check UniAgent Installation Status** in the panel.

After UniAgent is installed, wait for several minutes until the status is synchronized and click **Check UniAgent Installation Status** to update and check the UniAgent installation status. After the installation is complete, you can perform self-service O&M operations as needed.

15.3 Setting Scheduled Tasks for ECSs

Scenario

ECS supports scheduled O&M. You can schedule tasks such as starting and stopping multiple ECSs.

For example, if your services have regular peaks and lulls, you can start ECSs during peak hours and stop them during off-peak hours to reduce costs. Manually starting or stopping your ECSs is time-consuming and can be accident prone. You can set scheduled tasks for automatic O&M on multiple ECSs, reducing labor and time costs.

This section describes how to set scheduled tasks for ECSs.

Notes and Constraints

- ECS scheduled O&M is dependent on Cloud Operations Center (COC). COC must be enabled and authorized.
For IAM users, permissions for COC operations need to be granted. For details, see [Configuring Custom Policies for ECS Self-Service O&M](#).

- Each account can create a maximum of 100 scheduled tasks.
- Scheduled tasks only apply to ECSs in the same region.

Procedure

You can create scheduled start or stop tasks for one or more ECSs.

Creating a Scheduled Task for an Individual ECS

1. Log in to the management console and go to the [Elastic Cloud Server](#) console.
2. In the ECS list, locate the target ECS and choose **More > Set Scheduled Tasks > Set Scheduled Start** or **Set Scheduled Stop** in the **Operation** column.
3. (Optional) On the **Enable COC and Grant Permissions** page, read and agree to the service statement, and click **Enable and Authorize**.
This page is displayed if COC is not enabled and authorized.
4. On the **Create Task** page, set parameters required for the scheduled task.

Table 15-1 Scheduled task parameters

| Parameter | Description |
|--------------------|--|
| Task Name | Name of the scheduled task to be created.
The value can contain 3 to 100 characters, including letters, digits, hyphens (-), and underscores (_). |
| Enterprise Project | Enterprise project that the scheduled task belongs to. It is used to manage resources and members in the project.
The default enterprise project is default . |
| Time Zone | Time zone of the scheduled task. |
| Scheduled Type | Scheduled type of the task. It can be one-time or periodic execution. <ul style="list-style-type: none">• One-time execution: A scheduled task is executed once at a specified time.• Periodic execution: The scheduled task is executed periodically on a specific schedule. |

| Parameter | Description |
|-----------------------|---|
| Execute Time | <p>Time when the scheduled task is executed.</p> <ul style="list-style-type: none">• When One-time execution is selected for Scheduled Type, you specify the start time for the scheduled task.• When Periodic execution is selected for Scheduled Type, you specify the execution period and time for the scheduled task.<ul style="list-style-type: none">– Simple: Specify the date and time for the scheduled task to be automatically executed every week.– Cron: Specify the execution period and time for the scheduled task by setting the second, minute, hour, day, month, and week of a cron expression. |
| Rule Expired | <p>End time of a periodic scheduled task.</p> <p>This parameter is displayed only when Periodic execution is selected for Scheduled Type.</p> |
| Task Type | <p>Type of the scheduled task.</p> <ul style="list-style-type: none">• If you choose Set Scheduled Start in step 2, the type is ECS startup.• If you choose Set Scheduled Stop in step 2, the type is ECS shutdown. |
| IAM Permission Agency | <p>IAM agency assumed by COC to execute the scheduled task.</p> |
| Target Instance | <p>The instance where the scheduled task is to be executed. An instance is selected by default.</p> |

| Parameter | Description |
|--------------|---|
| Batch Policy | <p>The batch policy for executing a scheduled task, which can be any of the following:</p> <ul style="list-style-type: none">• Automatic: The system automatically sets the batches for executing the scheduled task on target instances.• Manual: You set the batches for executing the scheduled task on target instances.• No Batch: The scheduled task is executed on all target instances at the same time. <p>The maximum number of target instances in each batch is 100.</p> |

5. Click **Submit**.

Creating a Scheduled Task for Multiple ECSs

1. Log in to the management console and go to the [Elastic Cloud Server](#) console.
2. In the ECS list, select the ECSs for which you want to create a scheduled task.
3. Above the ECS list, choose **More > Set Scheduled Tasks > Set Scheduled Start** or **Set Scheduled Stop**.
4. (Optional) On the **Enable COC and Grant Permissions** page, read and agree to the service statement, and click **Enable and Authorize**.
This page is displayed if COC is not enabled and authorized.
5. On the **Create Task** page, set parameters required for the scheduled task.

Table 15-2 Parameters for configuring a scheduled task

| Parameter | Description |
|--------------------|---|
| Task Name | <p>Name of the scheduled task to be created.</p> <p>The value can contain 3 to 100 characters, including letters, digits, hyphens (-), and underscores (_).</p> |
| Enterprise Project | <p>Enterprise project that the scheduled task belongs to. It is used to manage resources and members in the project.</p> <p>The default enterprise project is default.</p> |

| Parameter | Description |
|-----------------------|---|
| Time Zone | Time zone of the scheduled task. |
| Scheduled Type | <p>Scheduled type of the task. It can be one-time or periodic execution.</p> <ul style="list-style-type: none">• One-time execution: A scheduled task is executed once at a specified time.• Periodic execution: The scheduled task is executed periodically at a specified time. |
| Execute Time | <p>Execution time of the scheduled task.</p> <ul style="list-style-type: none">• When One-time execution is selected for Scheduled Type, you specify the start time for the scheduled task.• When Periodic execution is selected for Scheduled Type, you specify the execution period and time for the scheduled task.<ul style="list-style-type: none">– Simple: Specify the date and time for the scheduled task to be automatically executed every week.– Cron: Specify the execution period and time for the scheduled task by setting the second, minute, hour, day, month, and week of a cron expression. |
| Rule Expired | <p>End time of a periodic scheduled task.</p> <p>This parameter is only displayed when Periodic execution is selected for Scheduled Type.</p> |
| Task Type | <p>Type of the scheduled task.</p> <ul style="list-style-type: none">• If you choose Set Scheduled Start in step 3, the type is ECS startup.• If you choose Set Scheduled Stop in step 3, the type is ECS shutdown. |
| IAM Permission Agency | IAM agency assumed by COC to execute the scheduled task. |

| Parameter | Description |
|-----------------|--|
| Target Instance | Target instances where the scheduled task is to be executed. Instances selected in step 2 are selected by default. |
| Batch Policy | <p>Batch policy for executing a scheduled task. It can be any of the following:</p> <ul style="list-style-type: none">• Automatic: The system automatically sets the batches for executing the scheduled task on target instances.• Manual: You set the batches for executing the scheduled task on target instances.• No Batch: The scheduled task is executed on all target instances at the same time. <p>The maximum number of target instances in each batch is 100.</p> |

6. Click **Submit**.

Follow-up Procedure

On the [Scheduled O&M](#) page, you can enable, disable, modify, and delete scheduled tasks, and configure review and notification settings for scheduled tasks.

For details, see [Managing Scheduled Tasks](#).

16 Resources and Tags

16.1 Tag Management

16.1.1 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.

If your organization has created a tag policy for ECS, you need to add tags for ECS based on the tag policy. If a tag does not comply with the tag rules, the creation may fail. Contact the organization administrator to learn details about the tag policy.

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 16-1 Example tags




Figure 16-1 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 16-1** lists the tag key and value requirements.

Table 16-1 Tag key and value requirements

| Parameter | Requirement | Example Value |
|-----------|--|---------------|
| Key | <ul style="list-style-type: none">• Cannot be left blank.• The key value must be unique for an ECS.• Can contain a maximum of 36 characters. | Organization |
| Value | <ul style="list-style-type: none">• Can contain a maximum of 43 characters. | Apache |

16.1.2 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](#).

Notes and Constraints

If your organization has configured tag policies for cloud resources, add tags to resources based on the tag policies. If a tag does not comply with the tag rules, the creation may fail. Contact the organization administrator to learn details about the tag policy.

Adding Tags During ECS Creation


1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Choose Compute > Elastic Cloud Server. The **Elastic Cloud Server** page is displayed.
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 16-1](#).

NOTE

For details about other parameters, see [Purchasing an ECS](#).

Figure 16-2 Adding tags



Advanced Options ☒ Configure now

Tag

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View Predefined Tags](#)

Tag key Tag value

You can add 10 more tags.

Adding Tags on the ECS Details Page



1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Choose Compute > Elastic Cloud Server. The **Elastic Cloud Server** page is displayed.
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 16-1](#).

Figure 16-3 Adding tags on the Tags tab

6. Click **OK**.

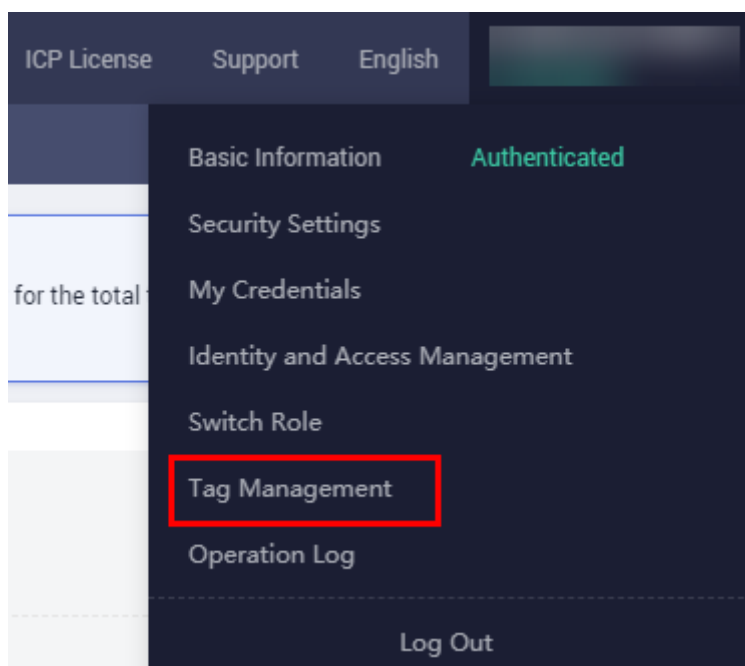
After tags are added, you can click **Edit** in the **Operation** column to edit them.


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.

Figure 16-4 Tag Management

3. On the displayed **Resource Tags** page, select the region where the resource is located, select **ECS-ECS** for **Resource Type**, and click **Search**.
All ECSs matching the search criteria are displayed.
4. In the **Search Result** area, click **Create Key**. In the displayed dialog box, enter a key (for example **project**) and click **OK**.
After the tag is created, the tag key is added to the resource list. If the key is not displayed in the resource list, click  and select the created key from the drop-down list.

By default, the value of the tag key is **Not tagged**. You need to set a value for the tag of each resource to associate the tag with the resource.

Figure 16-5 Resource list

Search Result Last Updated: Aug 08, 2019 19:44:04 GMT+08:00 (1 min ago) Invalid Tag

Create Key Manage Tag ? Enter a resource name. Q Edit View ⚙ ↺

| <input type="checkbox"/> | Resource Type | Resource Name | Region (Project) | Total Tags | Organization | |
|----------------------------|---------------|-------------------|-------------------------------|------------|--------------|---|
| ▼ <input type="checkbox"/> | ECS-ECS | ecs-bf2e | CN North-Beijing1(cn-north-1) | 0 | Not tagged | + |
| ▼ <input type="checkbox"/> | ECS-ECS | as-config-lb8v... | CN North-Beijing1(cn-north-1) | 0 | Not tagged | + |


5. Click **Edit** to make the resource list editable.
6. Locate the row containing the target ECS, click , and enter a value (for example **A**).
After a value is set for a tag key, the number of tags is incremented by 1.
Repeat the preceding steps to add tag values for other ECSs.

Figure 16-6 Setting a tag value

Search Result Last Updated: Aug 08, 2019 19:50:01 GMT+08:00 (1 min ago)

Create Key Manage Tag ? Enter a resource name. Q Edit View ⚙ ↺

| <input type="checkbox"/> | Resource Type | Resource Name | Region (Project) | Total Tags | Organization | |
|----------------------------|---------------|---------------------|-------------------------------|------------|--------------|-----|
| ▼ <input type="checkbox"/> | ECS-ECS | ecs-bf2e | CN North-Beijing1(cn-north-1) | 0 | A | ✓ ✗ |
| ▼ <input type="checkbox"/> | ECS-ECS | as-config-lb8v-0... | CN North-Beijing1(cn-north-1) | 0 | Not tagged | + |

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.

16.1.3 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 key-value pair.

1. Log in to the management console.
2. Click  in the upper left corner and select a region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box, select a tag key under **Resource Tag** and then a tag value, and press **Enter**.

You can search for ECSs by multiple tags and they are automatically joined with AND.

Figure 16-7 Searching by tag





Figure 16-8 Searching by tag



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.

16.1.4 Deleting Tags

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

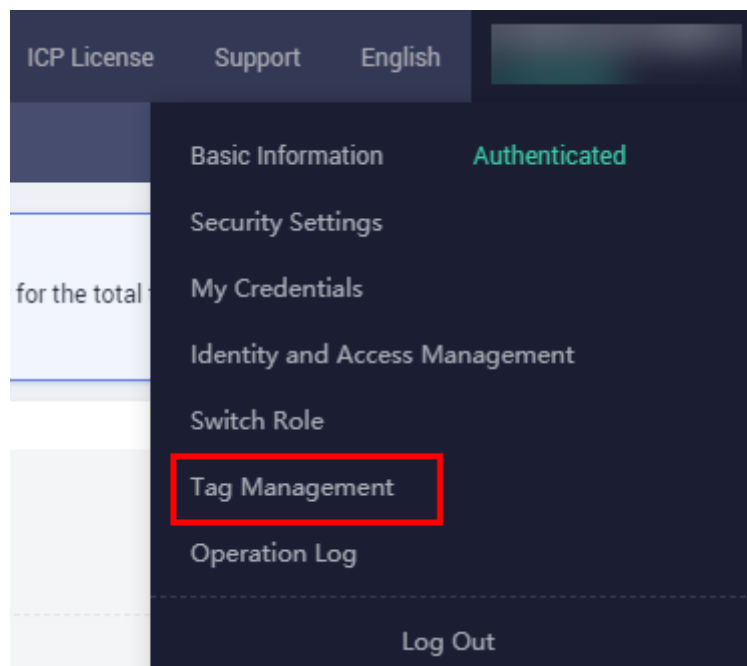
1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Choose **Compute > Elastic Cloud Server**. The **Elastic Cloud Server** page is displayed.

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.

Figure 16-9 Tag Management



3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. In the **Search Result** area, click **Edit** to make the resource tag list editable.


If the key of a tag you want to delete is not contained in the list, click  and select the tag key from the drop-down list. It is a good practice to select at most 10 keys to display.

5. Locate the row containing the target ECS and click .
6. (Optional) Click  in the upper right of the **Search Result** area.
The resource list is refreshed and the refresh time is updated.

Batch Deleting Tags on the TMS Console

NOTICE

Exercise caution when deleting tags in a batch. After you delete the tags, they will be removed from all the associated ECSs and cannot be recovered.

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. Select the target ECSs.
5. Click **Manage Tag** in the upper left corner of the list.
6. In the displayed **Manage Tag** dialog box, click **Delete** in the **Operation** column. Click **OK**.
7. (Optional) Click  in the upper right of the **Search Result** area. The resource list is refreshed and the refresh time is updated.

16.2 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?


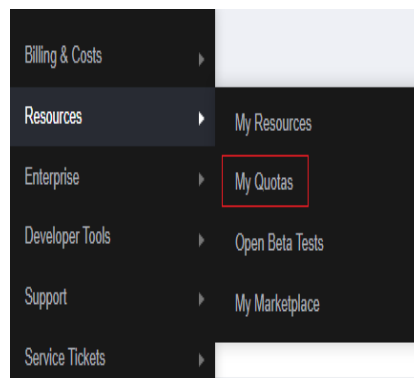
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources > My Quotas**. The **Quotas** page is displayed.

Figure 16-10 My Quotas

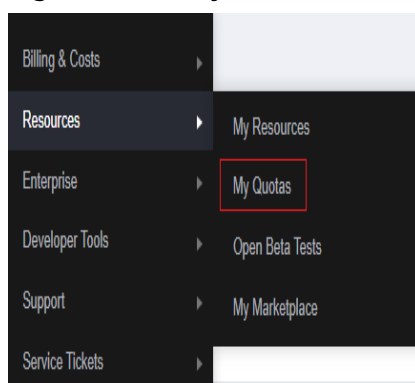


4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Quotas** page is displayed.

Figure 16-11 My Quotas



3. Click **Increase Quota** in the upper right corner of the page.

Figure 16-12 Increasing quota

| Service Quota ⓘ | | | | Increase Quota |
|-----------------------------------|--------------------------|------------|-------------|----------------|
| Service | Resource Type | Used Quota | Total Quota | |
| Auto Scaling | AS group | 0 | | |
| | AS configuration | 0 | | |
| Image Management Service | Image | 0 | | |
| Cloud Container Engine | Cluster | 0 | | |
| FunctionGraph | Function | 0 | | |
| Elastic Volume Service | Code storage(MB) | 0 | | |
| | Disk | 3 | | |
| | Disk capacity(GB) | 120 | | |
| Storage Disaster Recovery Service | Snapshots | 4 | | |
| | Protection group | 0 | | |
| Cloud Server Backup Service | Replication pair | 0 | | |
| | Backup Capacity(GB) | 0 | | |
| | Backup | 0 | | |
| Scalable File Service | File system | 0 | | |
| | File system capacity(GB) | 0 | | |
| CDN | Domain name | 0 | | |
| | File URL refreshing | 0 | | |
| | Directory URL refreshing | 0 | | |
| | URL preheating | 0 | | |

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**.

17 Cloud Eye Monitoring

17.1 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](#).

Helpful Links

- [Why Is My Windows ECS Running Slowly?](#)
- [Why Is My Linux ECS Running Slowly?](#)

17.2 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 17-1](#).

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 Windows ECS, see "Installing and Configuring the Agent (Windows)" in *Cloud Eye User Guide*.
 - For details about how to install the Agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in *Cloud Eye User Guide*.

Table 17-1 Basic ECS metrics

| Metric ID | Metric | Windows | | Linux | |
|-----------|-----------|-----------|-----------------|-----------|-----------------|
| - | None | Xen | KVM or QingTian | Xen | KVM or QingTian |
| cpu_util | CPU Usage | Supported | Supported | Supported | Supported |

| Metric ID | Metric | Windows | | Linux | |
|--------------------------|----------------------|---|---|---|---------------|
| mem_util | Memory Usage | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable. QingTian ECSs do not support this metric.) | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Not supported |
| disk_util_inband | Disk Usage | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Not supported |
| disk_read_bytes_rate | Disk Read Bandwidth | Supported | Supported | Supported | Supported |
| disk_write_bytes_rate | Disk Write Bandwidth | Supported | Supported | Supported | Supported |
| disk_read_requests_rate | Disk Read IOPS | Supported | Supported | Supported | Supported |
| disk_write_requests_rate | Disk Write IOPS | Supported | Supported | Supported | Supported |

| Metric ID | Metric | Windows | | Linux | |
|---------------------------------------|-----------------------|--|------------------------------|--|------------------------------|
| network_incoming_bytes_rate_inband | Inband Incoming Rate | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Supported | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Not supported |
| network_outgoing_bytes_rate_inband | Inband Outgoing Rate | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Supported | Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.) | Not supported |
| network_incoming_bytes_aggregate_rate | Outband Incoming Rate | Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.) | Supported | Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.) | Supported |
| network_outgoing_bytes_aggregate_rate | Outband Outgoing Rate | Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.) | Supported | Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.) | Supported |
| cpu_credit_usage | CPU Credit Usage | Not supported | Supported (only for T6 ECSs) | Not supported | Supported (only for T6 ECSs) |

| Metric ID | Metric | Windows | | Linux | |
|---------------------------|---------------------|---------------|------------------------------|---------------|------------------------------|
| cpu_credit_balance | CPU Credit Balance | Not supported | Supported (only for T6 ECSs) | Not supported | Supported (only for T6 ECSs) |
| network_vm_connections | Network Connections | Not supported | Supported | Not supported | Supported |
| network_vm_bandwidth_in | Inbound Bandwidth | Not supported | Supported | Not supported | Supported |
| network_vm_bandwidth_out | Outbound Bandwidth | Not supported | Supported | Not supported | Supported |
| network_vm_pps_in | Inbound PPS | Not supported | Supported | Not supported | Supported |
| network_vm_pps_out | Outbound PPS | Not supported | Supported | Not supported | Supported |
| network_vm_newconnections | New Connections | Not supported | Supported | Not supported | Supported |

Table 17-2 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 17-2 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 | <p>CPU usage of an ECS</p> <p>This metric is used to show the CPU usage of the physical server accommodating the monitored ECS, which is not as 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?</p> <p>Formula: CPU usage of an ECS/Number of vCPUs in the ECS</p> | 0-100 | % | N/A | ECS | 5 minutes |
| mem_util | Memory Usage | <p>Memory usage of an ECS</p> <p>This metric is unavailable if the image has no UVP VMTTools installed.</p> <p>Formula: Used memory of an ECS/ Total memory of the ECS</p> <p>NOTE</p> <p>The memory usage of QingTian ECSs cannot be monitored.</p> | 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

$\text{byte_out} = (\text{rd_bytes} - \text{last_rd_bytes}) / \text{Time difference}$ | ≥ 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
$\text{req_out} = (\text{rd_req} - \text{last_rd_req}) / \text{Time difference}$ | ≥ 0 | Requests/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
$\text{req_in} = (\text{wr_req} - \text{last_wr_req}) / \text{Time difference}$ | ≥ 0 | Requests/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 and 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 and 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 | Inband 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_vmpps_in | Inbound PPS | Number of public and private packets received by the ECS per second | ≥ 0 | Packets/s | N/A | ECS | 5 minutes |
| network_vmpps_out | Outbound PPS | Number of public and private packets sent by the ECS per second | ≥ 0 | Packets/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 | connections/s | N/A | ECS | 5 minutes |

Dimensions

| Dimension | Key | Value |
|-----------|-------------|-----------------------|
| ECS | instance_id | Specifies the ECS ID. |

17.3 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 17-3 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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">Linux: Other Process CPU Usage = 1- Idle CPU Usage - Kernel Space CPU Usage - User Space CPU UsageWindows: 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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 | <p>Percentage of time that the CPU is servicing software interrupts</p> <ul style="list-style-type: none">Linux: Check metric value changes in file /proc/stat in a collection period. Run the top command to check the %Cpu(s) value.Windows is not supported currently. | 0-100 | % | N/A | ECS | 1 minute |

OS Metric: CPU Load

Table 17-4 CPU load metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|--------------------------------|--|-------------|------|-----------------|------------------------------|------------------------------|
| load_average 1 | (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_average 5 | (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_average 15 | (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 |

 NOTE

The Windows OS does not support the CPU load metrics.

OS Metric: Memory

Table 17-5 Memory metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|--------------------------|--|-------------|------|-----------------|------------------------------|------------------------------|
| mem_available | (Agent) Available Memory | <p>Amount of memory that is available and can be given instantly to processes</p> <ul style="list-style-type: none">Linux: Obtain the metric value from /proc/meminfo.<ul style="list-style-type: none">If MemAvailable is displayed in /proc/meminfo, obtain the value.If MemAvailable is not displayed in /proc/meminfo, MemAvailable = MemFree + Buffers + CachedWindows: 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 | <p>Memory usage of the monitored object</p> <ul style="list-style-type: none">Linux: Obtain the metric value from the /proc/meminfo file:
(MemTotal - MemAvailable)/ MemTotal<ul style="list-style-type: none">If MemAvailable is displayed in /proc/meminfo,
MemUsedPercent = (MemTotal - MemAvailable)/MemTotalIf MemAvailable is not displayed in /proc/meminfo,
MemUsedPercent = (MemTotal - MemFree - Buffers - Cached)/ MemTotalWindows: 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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">Linux: Use the <code>/proc/{pid}/fd</code> 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 17-6 Disk metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------|------------------------------|---|-------------|------|-----------------|------------------------------|------------------------------|
| disk_free | (Agent) Available Disk Space | <p>Free space on the disks</p> <ul style="list-style-type: none">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 | <p>Total space on the disks, including used and free</p> <ul style="list-style-type: none">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 | <p>Used space on the disks</p> <ul style="list-style-type: none">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 <code>GetDiskFreeSpaceExW</code> 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 | <p>Percentage of total disk space that is used, which is calculated as follows:
Disk Usage = Used Disk Space/Disk Storage Capacity</p> <ul style="list-style-type: none">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 17-7 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 | <p>Number of bytes read from the monitored disk per second</p> <ul style="list-style-type: none">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:<ul style="list-style-type: none">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) | <ul style="list-style-type: none">EC S - DiskEC S - Mount point | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|---|-------------|------|-----------------|------------------------------|------------------------------|
| | | <p>instantaneous value returned by the object indicates the metric value in a collection period.</p> <ul style="list-style-type: none">- 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_avg_read_requests_rate | (Agent) Disks Read Requests | <p>Number of read requests sent to the monitored disk per second</p> <ul style="list-style-type: none">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:<ul style="list-style-type: none">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 | <ul style="list-style-type: none">ECS - DiskECS - Mount point | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|---|-------------|------|-----------------|------------------------------|------------------------------|
| | | <p>value returned by the object indicates the metric value in a collection period.</p> <ul style="list-style-type: none">- 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 | <p>Number of bytes written to the monitored disk per second</p> <ul style="list-style-type: none">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:<ul style="list-style-type: none">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) | <ul style="list-style-type: none">ECS - DiskECS - Mount Point | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|--|-------------|------|-----------------|------------------------------|------------------------------|
| | | <p>by the object indicates the metric value in a collection period.</p> <ul style="list-style-type: none">- 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_avg_write_requests_rate | (Agent) Disks Write Requests | <p>Number of write requests sent to the monitored disk per second</p> <ul style="list-style-type: none">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:<ul style="list-style-type: none">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 | <ul style="list-style-type: none">EC2 - DiskEC2 - Mount point | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------|-----------|---|-------------|------|-----------------|------------------------------|------------------------------|
| | | <p>value returned by the object indicates the metric value in a collection period.</p> <ul style="list-style-type: none">- 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 | <p>Average amount of time that read requests have waited on the disks</p> <ul style="list-style-type: none">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 | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>Average amount of time that write requests have waited on the disks</p> <ul style="list-style-type: none">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 | <ul style="list-style-type: none">ECS - DiskECS - Mount point | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|------------------------|---|-------------|------|-----------------|--|------------------------------|
| disk_io Utils | (Agent) Disk I/O Usage | <p>Percentage of the time that the disk has had I/O requests queued to the total disk operation time</p> <ul style="list-style-type: none">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 | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>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.</p> <ul style="list-style-type: none">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, hyphens (-), periods (.), and swung dashes (~).Windows is not supported currently. | ≥ 0 | count | N/A | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>Average number of bytes in an I/O write for the monitored disk in the monitoring period</p> <ul style="list-style-type: none">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 <code>/proc/diskstats</code> 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/operation | N/A | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>Average number of bytes in an I/O read for the monitored disk in the monitoring period</p> <ul style="list-style-type: none">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/operation | N/A | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>Average time in an I/O read or write for the monitored disk in the monitoring period</p> <ul style="list-style-type: none">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 | <ul style="list-style-type: none">ECS - DiskECS - 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 | <p>Percentage of the physical disk usage of the monitored object. Calculation formula: Used storage space of all mounted disk partitions/Total disk storage space</p> <ul style="list-style-type: none">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 17-8 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 . | <ul style="list-style-type: none">0: read able and writ able1: 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) Percent age 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 17-9 NIC metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------------|----------------------------|--|-------------|-------|-----------------|------------------------------|------------------------------|
| net_bit Recv | (Agent) Outbound Bandwidth | Number of bits sent by this NIC per second
<ul style="list-style-type: none">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_bit Sent | (Agent) Inbound Bandwidth | Number of bits received by this NIC per second
<ul style="list-style-type: none">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 <ul style="list-style-type: none">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 | Cou nts/s | N/A | ECS | 1 minute |
| net_packetSent | (Agent) NIC Packet Send Rate | Number of packets sent by this NIC per second <ul style="list-style-type: none">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 | Cou nts/s | N/A | ECS | 1 minute |
| net_errin | (Agent) Receive Error Rate | Percentage of receive errors detected by this NIC per second <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 17-10 NTP metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------|--------------------|---|-------------|------|-----------------|------------------------------|------------------------------|
| ntp_offset | (Agent) NTP Offset | <div>NTP offset of the monitored object</div> <ul style="list-style-type: none">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 17-11 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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 | <p>Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client</p> <ul style="list-style-type: none">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 | <p>Number of TCP connections in the FIN_WAIT2 state</p> <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 <ul style="list-style-type: none">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 | <p>Percentage of packets that are resent</p> <ul style="list-style-type: none">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 17-12 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 <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | <ul style="list-style-type: none">0: The GPU is healthy.1: The GPU is subhealthy.2: The GPU is faulty. | N/A | N/A | <ul style="list-style-type: none">EC SEC S - 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 <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | 0-100 | % | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_usage_decoder | Decoding Usage | Decoding capability usage of the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | 0-100 | % | N/A | <ul style="list-style-type: none">EC SEC S - 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 | <p>Number of correctable ECC errors since the GPU is reset. The value is reset to 0 each time the GPU is reset.</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | Count | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_volatile_uncorrectable | Volatile Uncorrectable ECC Errors | <p>Number of uncorrectable ECC errors since the GPU is reset. The value is reset to 0 each time the GPU is reset.</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | Count | N/A | <ul style="list-style-type: none">EC SEC S - 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 <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-mi.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 | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_aggregate_uncorrectable | Aggregate Uncorrectable ECC Errors | Aggregate uncorrectable ECC Errors on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-mi.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 | <ul style="list-style-type: none">EC SEC S - 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 | <p>Number of retired page single bit errors, which indicates the number of single-bit pages blocked by the graphics card</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | Count | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_retired_page_double_bit | Retired Page Double Bit Errors | <p>Number of retired page double bit errors, which indicates the number of double-bit pages blocked by the graphics card</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | Count | N/A | <ul style="list-style-type: none">EC SEC S - 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 <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-mL.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | P0-P15, P32 <ul style="list-style-type: none">P0: indicates the maximum performance status.P15: indicates the minimum performance status.P32: indicates | N/A | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------|--------------------------|--|---------------------------------------|------|-----------------|---|------------------------------|
| | | | tests the unknown performance status. | | | | |
| gpu_usage_mem | (Agent) GPU Memory Usage | GPU memory usage of the monitored object <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | 0-100 | % | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_usage_gpu | (Agent) GPU Usage | GPU usage of the monitored object <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | 0-100 | % | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|---------------------|---------------------|--|-------------|------|-----------------|---|------------------------------|
| gpu_free_mem | GPU Free Memory | Free Memory on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MB | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_graphics_clocks | GPU Graphics Clocks | Current Graphics Clocks on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_memory_clocks | GPU Memory Clocks | Current Memory Clocks on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------------|---------------------------|--|-------------|-----------|-----------------|---|------------------------------|
| gpu_power_draw | GPU Draw Power | Draw Power on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | NA | W | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_rx_throughput_pci | GPU PCI Inbound Bandwidth | Current PCI Rx Throughput on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MB byte/s | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_sm_clocks | GPU SM Clocks | Current SM Clocks on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------------|-----------------------|--|-------------|-----------|-----------------|---|------------------------------|
| gpu_temperature | GPU Temperature | Current Temperature on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | °C | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_tx_throughput_pci | GPU PCI Tx Throughput | Current PCI Tx Throughput on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MB byte/s | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |
| gpu_used_memory | GPU Used Memory | Memory Used on the GPU <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-ml.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MB | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|------------------|--|-------------|------|-----------------|---|------------------------------|
| gpu_video_clocks | GPU Video Clocks | <p>Current Video Clocks on the GPU</p> <ul style="list-style-type: none">Linux: Obtain the metric value using the libnvidia-mL.so.1 library file of the graphics card.Windows: Obtain the metric value using the nvml.dll library of the graphics card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">EC SEC S - GPU | 1 minute |

OS Metrics: NPU

Table 17-13 NPU metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------------|-------------------|--|---|------|-----------------|---|------------------------------|
| npu_device_health | NPU Device Health | Overall measurement of the NPU health
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | <ul style="list-style-type: none">• 0: healthy• 1: minor alarms• 2: major alarms• 3: critical alarms | N/A | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_util_rate_mem | NPU Memory Usage | The memory usage of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------------------|----------------------------|--|-------------|------|-----------------|---|------------------------------|
| npu_util_rate_ai_core | NPU AI Core Usage | The AI core usage of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_util_rate_ai_cpu | NPU AI CPU Usage | The AI CPU usage of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_util_rate_ctrl_cpu | NPU Control CPU Usage | The CPU control usage of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_util_rate_mem_bandwidth | NPU Memory Bandwidth Usage | The memory bandwidth usage of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------------|-----------------------|--|-------------|-------|-----------------|---|------------------------------|
| npu_freq_mem | NPU Memory Frequency | Clock frequency of the NPU memory

Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_freq_ai_core | NPU AI Core Frequency | Clock frequency of the NPU's AI core

Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | MHz | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_usage_mem | Used NPU Memory | Used size of the NPU memory

Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | MB | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_single_bit_errors | NPU Single-Bit Errors | Numbers of single-bit errors of the NPU

Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | count | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-----------------|-----------------------|--|-------------|-------|-----------------|---|------------------------------|
| npu_double | NPU Double-Bit Errors | Numbers of double-bit errors of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | count | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_power | NPU Power | The power of the NPU (Rated power displayed for 310 only, actual power displayed for other cards)
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | W | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |
| npu_temperature | NPU Temperature | Current temperature of the NPU
Linux: Obtain the metric value from the libdcmi.so library file of the NPU card. | ≥ 0 | °C | N/A | <ul style="list-style-type: none">• ECS• ECS - NPU | 1 minute |

 NOTE

The Windows OS does not support NPU metrics.

OS Metrics: DAVP

Table 17-14 DAVP metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------------|--------------------|---|--|------|-----------------|--|------------------------------|
| davp_device_health | DAVP Device Health | An overall measurement of the DAVP health

Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | <ul style="list-style-type: none">• 0: healthy• 1: abnormal | N/A | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_util_rate_mem | DAVP Memory Usage | The memory usage of the DAVP

Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_usage_mem | Used DAVP Memory | Used size of the DAVP memory

Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | ≥ 0 | MB | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_util_rate_ai_core | DAVP AI Core Usage | The AI core usage of the DAVP

Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------------------|--------------------------------|---|-------------|------|-----------------|--|------------------------------|
| davp_util_rate_vdsp_core | DAVP VDSP Core Usage | The VDSP core usage of the DAVP
Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_util_rate_encoding_core | DAVP Encoding Core Usage | The encoding core usage of the DAVP
Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_util_rate_decoding_core | DAVP Decoding Core Usage | The decoding core usage of the DAVP
Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | 0-100 | % | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |
| davp_sysc_temperature | DAVP System Module Temperature | The temperature of DAVP system module
Linux: Obtain the metric value from the libvaml.so library file in the VAtools tool of the DAVP card. | ≥ 0 | °C | N/A | <ul style="list-style-type: none">• ECS• ECS - DAVP | 1 minute |

 **NOTE**

The Windows OS does not support DAVP metrics.

Dimensions

| Dimension | Key | Value |
|-------------------|-------------|---|
| ECS | instance_id | Specifies the ECS ID. |
| ECS - Disk | disk | Specifies the disks attached to an ECS. |
| ECS - Mount point | mount_point | Specifies the mount point of a disk. |
| ECS - GPU | gpu | Specifies the graphics card of an ECS. |
| ECS - NPU | npu | Specifies the NPU graphics card of an NPU-based ECS. |
| ECS - DAVP | davp | Specifies the DaoCloud DAVP1 video acceleration card of a DAVP-based ECS. |

17.4 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 17-15 Process metrics

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------|---------------------------|--|------------------------|------|-----------------|------------------------------|------------------------------|
| proc_pHashId_cpu | (Agent) Process CPU Usage | <p>CPU consumed by a process. pHashId (process name and process ID) is the value of md5.</p> <ul style="list-style-type: none">Linux: Check metric value changes in file /proc/pid/stat.Windows: Call the Windows API GetProcessTimes to obtain the CPU usage of the process. | 0-1 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 | <p>Memory consumed by a process. pHashId (process name and process ID) is the value of md5.</p> <ul style="list-style-type: none">Linux: $RSS * PAGE_SIZE / MemTotal$
Obtain the RSS value by checking the second column of file /proc/pid/statm.
Obtain the PAGE_SIZE value by running the getconf PAGE_SIZE command.
Obtain the MemTotal value by checking file /proc/meminfo.Windows: Call the Windows API <code>procGlobalMemoryStatusEx</code> to obtain the total memory size. Call <code>GetProcessMemoryInfo</code> to obtain the used memory size. Use the used memory size to divide the total memory size to get the memory usage. | 0-100 | % | N/A | ECS | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|--------------------|----------------------------|---|-------------|-------|-----------------|------------------------------|------------------------------|
| 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 . <ul style="list-style-type: none">Linux: Run the ls -l /proc/pid/fd command to view the number of opened files.Windows is not supported currently. | ≥0 | Count | N/A | ECS | 1 minute |
| proc_running_count | (Agent) Running Processes | Number of processes that are running <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently. | ≥0 | Count | N/A | ECS | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|-------------------|--------------------------|--|-------------|-------|-----------------|------------------------------|------------------------------|
| proc_idle_count | (Agent) Idle Processes | Number of processes that are idle <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently. | ≥ 0 | Count | N/A | ECS | 1 minute |
| proc_zombie_count | (Agent) Zombie Processes | Number of zombie processes <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently. | ≥ 0 | 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 <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently. | ≥ 0 | Count | N/A | ECS | 1 minute |
| proc_sleeping_count | (Agent) Sleeping Processes | Number of processes that are sleeping <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows is not supported currently. | ≥ 0 | Count | N/A | ECS | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|----------------------|-----------------------------|---|-------------|-------|-----------------|---|------------------------------|
| proc_total_count | (Agent) Total Processes | Total number of processes on the monitored object <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows: Obtain the total number of processes from the modules supported by the psapi.dll system process statuses. | ≥ 0 | Count | N/A | ECS | 1 minute |
| proc_specified_count | (Agent) Specified Processes | Number of specified processes <ul style="list-style-type: none">Linux: You can obtain the state of each process by checking the Status value in the /proc/pid/status file, and then collect the total number of processes in each state.Windows: Obtain the total number of processes from the modules supported by the psapi.dll system process statuses. | ≥ 0 | Count | N/A | <ul style="list-style-type: none">ECSECS - Process | 1 minute |

| Metric | Parameter | Description | Value Range | Unit | Conversion Rule | Monitored Object & Dimension | Monitoring Period (Raw Data) |
|------------------------|-----------------------------------|---|-------------|-------|-----------------|---|------------------------------|
| specified_process_file | (Agent) Files Opened by a Process | Number of files opened by a specific process of the monitored object <ul style="list-style-type: none">Linux: Run the ls -l /proc/pid/fd command to view the number of opened files.Windows is not supported currently. | ≥0 | Count | N/A | <ul style="list-style-type: none">ECSECS - Process | 1 minute |

Dimensions

| Dimension | Key | Value |
|---------------|-------------|----------------------------|
| ECS | instance_id | Specifies the ECS ID. |
| ECS - Process | proc | Specifies the ECS process. |


17.5 Creating Alarm Rules

Scenarios

You can set alarm rules and notifications for ECSs. This helps you learn about ECS statuses and receive alarm notifications (if any) in a timely manner.

This section describes how to set ECS alarm rules.

Procedure

- Log in to the management console.
- Click  in the upper left corner and select a region and project.
- Under **Management & Governance**, choose **Cloud Eye**.
- In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.

5. On the displayed **Alarm Rules** page, click **Create Alarm Rule**.
You can also select an existing alarm rule and modify it as needed.
6. Configure basic information of the alarm rule.

Figure 17-1 Alarm rule information

★ Name

alarm-cprn

Description

0/256

Table 17-16 Basic parameters for alarm rules

| Parameter | Description | Example Value |
|-------------|--|---------------|
| Name | Name of the alarm rule. The system generates a random name, and you can change it as needed. | alarm-cprn |
| Description | Description of the alarm rule. This parameter is optional. | - |

7. Select an object to be monitored and configure alarm parameters.
- If you select **Metric** for **Alarm Type**, you can create an alarm rule for ECS metrics.

Figure 17-2 Metric alarm configuration

Alarm Type

MetricEvent

Cloud Product

Elastic Cloud Server - ECSs

Resource Level

Cloud productSpecific dimension

Monitoring Scope

Add resourcesResource groupsSpecific resources

An alarm will be triggered anytime a resource, including resources that will be purchased, in this dimension meets the alarm rule.

Select Resources to Exclude

Method

Associate templateConfigure manually

Alarm Policy

Trigger an alarm whenany policyis notUse Template

Metric Name

Enter a keyword

Alarm Policy

Rate data

3 times (consec...)

TriggerDaily

Add Alarm PolicyYou can add 40 more.

Key parameters are described below. For details, see [Creating an Alarm Rule and Notifications](#).

- **Alarm Type:** The type of the alarm that the alarm rule applies to. You can select **Metric** or **Event**.
- **Cloud Product:** The cloud product you want to monitor. This parameter is only available if you select **Metric** for **Alarm Type**. Example: Elastic Cloud Server - ECSs

- **Resource Level:** Select the resource level of the monitored object. You are advised to select **Specific dimension** and then select different dimensions such as disk, mount point, and process from the drop-down list based on service requirements.

 **NOTE**

To create an alarm rule for a GPU-accelerated ECS, select **ECSs - GPU** for **Specific dimension**.


- **Monitoring Scope:** The scope that the alarm rule applies to. You can select **All resources**, **Resource groups**, or **Specific resources**.
- **Method:** Select **Associate template** or **Configure manually**.

 **NOTE**

After an associated template is modified, the policies contained in this alarm rule to be created will be modified accordingly.

- **Alarm Policy:** The policy for triggering an alarm.
- If you select **Event** for **Alarm Type**, you can create an alarm rule for ECS events.

Figure 17-3 Event alarm configuration



Key parameters are described below. For details, see [Creating an Alarm Rule and Notifications](#).

- **Event Type:** You can select **System event** or **Custom event**.
 - **Event Source:** Select the name of the cloud service from which events are reported.
 - **Method:** Select **Associate template** or **Configure manually**.
 - **Alarm Policy:** The policy for triggering an alarm.
8. Set alarm notification parameters.
To send alarm notifications by email, SMS, HTTP, or HTTPS, enable **Alarm Notifications**.
For details about related parameters, see [Creating an Alarm Rule and Notifications](#).
9. Click **Create**.

 **NOTE**

For more information about ECS alarm rules, see [Cloud Eye User Guide](#).

17.6 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 is 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 [Creating 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 a 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. You can also select a custom time range to view the historical monitoring data of the last six months.

18 Audit Using CTS

18.1 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 18-1 ECS operations recorded by CTS

| Operation | Resource Type | Trace |
|-------------------|---------------|---|
| Creating an ECS | ecs | createServer
createServerV2
createServerV21 |
| Deleting an ECS | ecs | deleteServer
deleteServerV2
deleteServerV21
deleteVmByNative |
| Starting an ECS | ecs | startServer |
| Restarting an ECS | ecs | rebootServer |
| Stopping an ECS | ecs | stopServer |
| Adding an ECS NIC | ecs | addNic |

| Operation | Resource Type | Trace |
|--|---------------|------------------------------------|
| Deleting an ECS NIC | ecs | deleteNic
delNic |
| Attaching a disk | ecs | attachVolume
attachVolumeV2 |
| Attaching a disk (on the EVS console) | ecs | attachVolume |
| Detaching a disk | ecs | detachVolume |
| Reinstalling an OS | ecs | reinstallOs
reInitOs |
| Changing an OS | ecs | changeOs
reInitOs |
| Modifying specifications | ecs | resizeServer |
| Enabling automatic recovery on an ECS | ecs | addAutoRecovery |
| Disabling automatic recovery on an ECS | ecs | deleteAutoRecovery |
| Updating metadata or setting metadata of a specified key (excluding the agency name) | ecs | updateMetadata |
| Updating metadata or setting metadata of a specified key (including the agency name) | ecs | updateMetadataAnd-
ServerAgency |
| Updating an agency | ecs | updateServerAgency |
| Logging in to an ECS using VNC | ecs | remoteConsole |
| Modifying ECS information | ecs | updateServer |
| Managing the ECS status (OpenStack Nova API) | server | operateServer |
| Adding, deleting, or editing tags | ecsTags | dealUnifiedTags |

18.2 Viewing Traces

Scenarios

After you enable Cloud Trace Service (CTS) and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, CTS starts recording operations on data in Object Storage Service (OBS) buckets. 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


- Traces of a single account can be viewed on the CTS console. Multi-account traces can be viewed only on the **Trace List** page of each account, or in the OBS bucket or the **CTS/system** log stream configured for the management tracker with the organization function enabled.
- You can only query operation records of the last seven days on the CTS console. To store operation records for longer than seven days, 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 1 minute later and query data traces 5 minutes later.
- 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.

- **Enterprise Project ID:** Enter an enterprise project ID.
 - **Access Key:** Enter a temporary or permanent access key ID.
 - **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 **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 **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 | | Nov 16, 2023 10:54:04 GMT+08:00 | View Trace |

request

trace_id

code

trace_name

resource_type

trace_rating

api_version

message

source_ip

domain_id

trace_type

9. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace

```
{
  "request": "",
  "trace_id": " ",
  "code": "200",
  "trace_name": "createDockerConfig",
  "resource_type": "dockerlogincmd",
  "trace_rating": "normal",
  "api_version": "",
  "message": "createDockerConfig, Method: POST Url=/v2/manager/utls/secret, Reason:",
  "source_ip": " ",
  "domain_id": " ",
  "trace_type": "ApiCall",
  "service_type": "SWR",
  "event_type": "system",
  "project_id": " ",
  "response": "",
  "resource_id": "",
  "tracker_name": "system",
  "time": "Nov 16, 2023 10:54:04 GMT+08:00",
  "resource_name": "dockerlogincmd",
  "user": {
    "domain": {
      "name": " ",
      "id": " "
```

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.

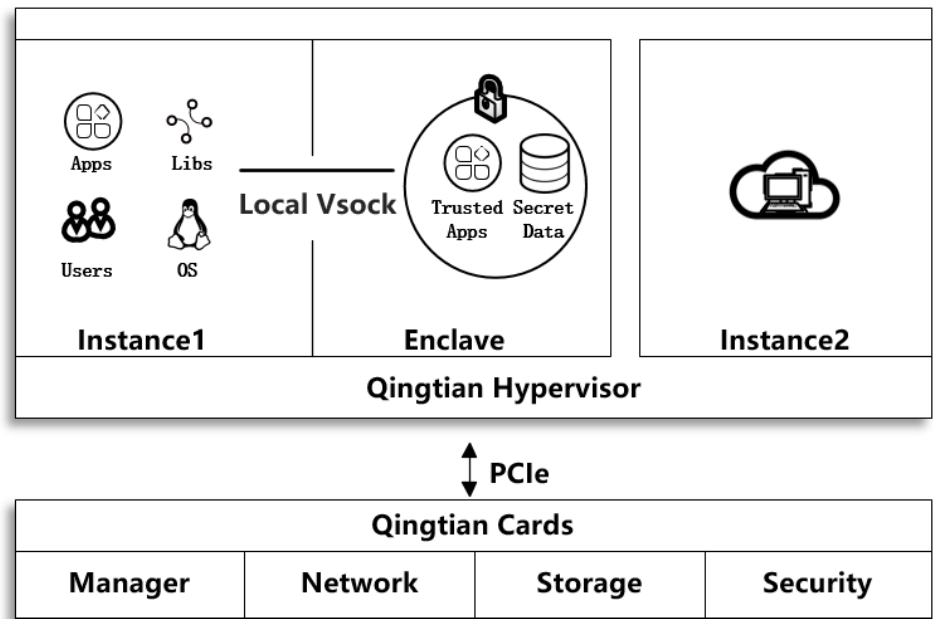
19 QingTian Enclave Management

19.1 QingTian Enclave Overview

19.1.1 What Is QingTian Enclave?

Introduction

- 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) | <div>1. At least two vCPUs and 512 MiB of memory are required.</div> <div>2. Only the Linux OS is supported.</div> <div>3. The parent instance must be a C7t and kC2 instance.</div> <div>4. QingTian Enclave is available in the following regions: CN North-Beijing4, CN East-Shanghai1, CN South-Guangzhou, AP-Singapore, and TR-Istanbul.</div> |

| Name | Constraints |
|--|---|
| QingTian Enclave instances (secondary VMs) | <ol style="list-style-type: none">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. If your services running in the QingTian Enclave instances are terminated unexpectedly, you need to manually run the services again.
2. By default, 1 GiB hugepages are configured for QingTian Enclave instances, with 1 GiB of memory and 2 vCPUs.
3. A QingTian Enclave parent instance must have at least 8 vCPUs and 16 GiB of memory.

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.

19.1.2 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.
- **qproxy**
The qproxy provides the capability of converting TCP and vsock network packets. Compared with qt-proxy, the network code can run in QingTian Enclave without modification. For details, see [QingTian Enclave Network Proxy](#).
- **qlog**
The qlog allows you to export specified files from QingTian Enclave to the parent instance. This helps you export service logs for subsequent O&M.
- **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).

19.2 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. Purchase an ECS and use it as the parent instance of QingTian Enclave. Select a Linux image for the OS and select **Enclave**.

For details, see [Purchasing an ECS in Custom Config Mode](#). Huawei Cloud EulerOS 2.0 is recommended.

2. Connect to the parent instance of QingTian Enclave. For details, see [Login Overview \(Linux\)](#).

3. Configure the parent instance of QingTian Enclave.

- If you select the Huawei Cloud EulerOS 2.0 image:

- i. Install the **qt CLI** tool and required peripheral components in the parent instance.

Install python modules required by qt CLI. For details, see [Installation of the qt CLI](#).

```
yum install qt-enclave-bootstrap  
yum install virtio-qtbox  
yum install qingtian-tool
```

The parameters are described as follows:

- **virtio-qtbox**: indicates the enclave-related driver.
 - **qt-enclave-bootstrap**: contains the files required for creating an enclave image.
 - **qingtian-tools**: manages the enclave lifecycle.
- ii. In the **qt-enclave-env.conf** configuration file, configure the isolation parameters as required. For details, see [Introduction to qt-enclave-env](#).

You are advised to isolate a large amount of memory resources for QingTian Enclave at a time. This prevents system memory fragmentation caused by repeated isolation, which may lead to start failures of the isolation service.

```
vim /etc/qingtian/enclave/qt-enclave-env.conf
```

Run the following commands with **memory_mib** set to **8192** and **cpu_count** to **2**:

```
# qingtian enclave configuration file.  
  
# Which hugepage size to reserve for qingtian enclave  
# can only configure it by 2 (2MB hugepage size) or 1024 (1GB hugepage size)  
hugepage_size: 1024  
  
# How much memory to allocate for qingtian enclave (in MiB).
```

```
memory_mib:8192
```

```
# User can use cpu_count to set how many CPUs need to be reserved.  
# Or cpu_list to set which CPUs need to be reserved.  
#  
# cpu_count and cpu_list conflict with each other. Only use exactly one of them.  
#  
# How many CPUs to reserve for qingtian enclave.  
cpu_count:2  
  
# Which CPUs to reserve for qingtian enclave. You can configure it like below:  
# 2,3 means reserving CPUs 2, 3, and you can also use 2-5 to reserve 2 through 5.  
# cpu_list:2,3
```

iii. Start the resource isolation service.

```
systemctl start qt-enclave-env
```

If the service isolation fails and no sufficient memory can be allocated, memory fragments may exist. You are advised to restart the VM and then isolate the service.

– If you select the Ubuntu 22.04 image:

i. Download the **QingTian Enclave** code.

ii. Visit <https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian.git> to see the branch that supports the Ubuntu 22.04 image.

```
cd /home  
git clone https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian.git  
cd /home/huawei-qingtian
```

iii. Install the required dependencies.

```
apt-get update -y  
apt install libgnutls28-dev libcjson-dev libglib2.0-dev -y
```

iv. Compile **virtio-qtbox**.

```
cd /home/huawei-qingtian/virtio-qtbox  
make  
make install
```

v. Compile the **qingtian-tools** tool package.

```
cd /home/huawei-qingtian/qingtian-tools  
make  
make install
```

vi. Compile the **qt-proxy** package.

```
cd /home/huawei-qingtian/qingtian-tools/qt-proxy  
make  
make install
```

vii. Install the **qt-enclave-bootstrap** package.

Kernel compiling is needed. If there are no special requirements, you can decompress the **qt-enclave-bootstrap** package of Huawei Cloud EulerOS to the corresponding path.

```
apt install alien -y  
cd /home  
wget https://repo.huaweicloud.com/hce/2.0/updates/x86_64/Packages/qt-enclave-bootstrap-1.0-34.hce2.x86_64.rpm  
alien -d qt-enclave-bootstrap-1.0-34.hce2.x86_64.rpm  
dpkg -i qt-enclave-bootstrap_1.0-35_amd64.deb
```

4. Install the required Python packages on the QingTian Enclave parent instance.
`pip3 install docker knack`

Figure 19-1 Example installation result

```

[root@ ~]# pip3 install docker knack
WARNING: When the --extra-index-url option is used, the installation package is downloaded from an unverified URL. Exercise caution when using this option.
WARNING: Running pip install with root privileges is generally not a good idea. Try `pip3 install --user` instead.
Collecting docker
  Downloading docker-7.1.0-py3-none-any.whl (147 kB)
    | 147 kB 9.0 MB/s
Collecting knack
  Downloading knack-0.12.0-py3-none-any.whl (60 kB)
    | 60 kB 14.9 MB/s
Requirement already satisfied: requests>=2.26.0 in /usr/local/lib/python3.9/site-packages (from docker) (2.32.3)
Requirement already satisfied: urllib3>=1.26.0 in /usr/local/lib/python3.9/site-packages (from docker) (2.2.3)
Collecting packaging
  Downloading packaging-25.0-py3-none-any.whl (66 kB)
    | 66 kB 12.4 MB/s
Collecting tabulate
  Downloading tabulate-0.9.0-py3-none-any.whl (35 kB)
Collecting argcomplete
  Downloading argcomplete-3.6.2-py3-none-any.whl (43 kB)
    | 43 kB 5.5 MB/s
Collecting jmespath
  Downloading jmespath-1.0.1-py3-none-any.whl (20 kB)
Collecting pygments
  Downloading pygments-2.19.1-py3-none-any.whl (1.2 MB)
    | 1.2 MB 79.4 MB/s
Requirement already satisfied: pyyaml in /usr/local/lib64/python3.9/site-packages (from knack) (6.0.2)
Requirement already satisfied: idna<4,>=2.5 in /usr/local/lib/python3.9/site-packages (from requests>=2.26.0->docker) (3.10)
Requirement already satisfied: certifi>=2017.4.17 in /usr/local/lib/python3.9/site-packages (from requests>=2.26.0->docker) (2024.8.30)
Requirement already satisfied: charset-normalizer<4,>=2 in /usr/local/lib64/python3.9/site-packages (from requests>=2.26.0->docker) (3.4.0)
Installing collected packages: tabulate, pygments, packaging, jmespath, argcomplete, knack, docker
Successfully installed argcomplete-3.6.2 docker-7.1.0 jmespath-1.0.1 knack-0.12.0 packaging-25.0 pygments-2.19.1 tabulate-0.9.0
WARNING: Running pip as the 'root' user can result in broken permissions and conflicting behaviour with the system package manager. It is recommended to use a virtual environment instead: https://pip.pypa.io/warnings/venv

```

5. Install Docker on the parent instance. Binary mode is recommended. You can download the required Docker version from Docker's official website.
 - a. Download Docker, for example, Docker 27.0.1.
`wget https://download.docker.com/linux/static/stable/x86_64/docker-27.0.1.tgz`
 - b. Decompress the downloaded package.
`tar xzf docker-27.0.1.tgz`
 - c. After the decompression is complete, copy all the files in the Docker directory to the **/usr/bin** directory.
`cp docker/* /usr/bin`
 - d. Start the Docker service and set the log level to error.
`dockerd -l error &`
 - e. Verify the Docker version.
`docker version`
 - f. Run the **hello-world** container to check whether Docker is installed.
`docker run hello-world`
6. Build a QingTian Enclave image file.

- a. 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"]

```

The preceding example is used to directly print service information to the virtual terminal or serial port device. It only applies to the debug mode. During routine O&M, you are advised to redirect service information to a file and then forward the file to the parent instance to ensure system security. For details, see [QingTian Enclave Log Forwarding Tool](#).

For example, to export **/tmp/hello_enclave_output** in QingTian Enclave to a specified path on the parent instance, run the following command:

```

CMD ["/root/hello_enclave.sh > /tmp/hello_enclave_output"]

```

- b. Check that the script has execution permissions.

```
chmod +x hello_enclave.sh
```

- c. Build a Docker image named **hello-enclave**.

```
docker build -f Dockerfile -t hello-enclave .
```

- d. 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":
"63bf78ece7d2388ff773d0cad2ebc9a3070359db46d567ba271ff8adfb8b0b091be4ff4d5dda3f1c83
109096e3656f3b",
  "PCR8":
"000000000000000000000000000000000000000000000000000000000000000000000000000000000
000000000000000000"
}
```

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).

- e. Generate a certificate pair.

```
openssl ecparam -out private-key.pem -name secp384r1 -genkey
openssl req -new -key private-key.pem -out ssl.csr
openssl x509 -req -days 365 -in ssl.csr -signkey private-key.pem -out server.pem
```

- f. Generate the EIF file using the generated certificate and key.

```
qt enclave make-img --docker-uri hello-enclave --eif hello-enclave.eif --private-key private-
key.pem --signing-certificate server.pem
```

7. Run the 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.

8. 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.

9. 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
}
```

19.3 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.

19.3.1 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**. In the specified directory, download the huawei-qingtian-enclave source code.

```
cd /home
git clone https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian.git
```

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). For details, see [Creating a Key](#).
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).

On the IAM console, use an account with administrator permissions to create a custom identity policy. For details, see [Creating a Custom Policy](#).

The following is an example custom identity policy:

```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms::generateRandom",
        "kms:cmk:createDataKey",
        "kms:cmk:decryptData",
        "kms:cmk:decryptDataKey"
      ],
      "Condition": {
        "StringEqualsIgnoreCase": {
          "kms:RecipientAttestation/PCR0": [
            "8f2cbfb3930e59c6de5c4caff0a3f4c0457e8956bfb4556a7ca1f5f4614a741eeee39ae10447eb5baee48d49e6c1cb6c",
            "ff7ba807a385b49fc1c3346bb47215aef503dee6df22d32f733e22b90a9bc4b22424ca7de1a3537ac9608d7ebe461d67",
            "a28e765550d6ad1188860d30167b1fdb9e29c8da825543861bc76ef1e8427fac6b444ec6a1847fc2c22deae8170c2e67"
          ],
          "kms:RecipientAttestation/PCR8": [
            "a9add94b0ecbbd992bade2176370ecf3bfed2cb39b2ec547512b5174279799f2036fa0b8577bdaf503836178bd11ee2"
          ]
        }
      }
    }
  ]
}
```

```
    ]
  }
}
},
{
  "Effect": "Allow",
  "Action": [
    "kms:cmk:encryptData",
    "kms:cmk:encryptDataKey"
  ]
}
]
```

In the example identity policy, the parameters are described as follows:

- **kms::generateRandom**: Generate a random number.
- **kms::cmk:createDataKey**: Create data keys.
- **kms:cmk:decryptData**: Decrypt data.
- **kms:cmk:decryptDataKey**: Decrypt data keys.
- **kms:cmk:encryptDataKey**: Encrypt data keys.
- **kms:cmk:encryptData**: Encrypt data.

kms:RecipientAttestation/PCR0 and **kms:RecipientAttestation/PCR8** are condition keys determined during QingTian Enclave image creation. Multiple values are supported.

6. The parent instance administrator boots the parent instance, starts the qt-proxy service to forward the QingTian Enclave KMS network requests, and then boots the QingTian Enclave instance using the QingTian Enclave image file.

You can obtain the KMS endpoints in different regions by referring to [Regions and Endpoints](#).

Remote Attestation and Data Decryption

The following describes the execution process of a QingTian Enclave application.

1. With the qproxy 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.

For details about the sample code, see <https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian/tree/master/enclave/qtsm-sdk-c/samples>.

19.3.2 Creating a QingTian Enclave Image

After a QingTian Enclave application is developed, you need to create 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 create 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 image.

Convert the Docker image to a QingTian Enclave image.

- a. (Optional) Create a private key (private-key.pem) and certificate (server.pem) using OpenSSL or other tools.

```
openssl ecparam -out private-key.pem -name secp384r1 -genkey
openssl req -new -key private-key.pem -out ssl.csr
openssl x509 -req -days 365 -in ssl.csr -signkey private-key.pem -out server.pem
```

NOTE

For security purposes, only the elliptic curve cryptography and secp384r1 elliptic curve are supported.

- b. Convert the Docker image to a QingTian Enclave image.

In the **qt make-img** command, the parameters for the Docker source image and the generated target QingTian Enclave image are mandatory.

```
# 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":
    "b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c036
4a9b88dd58017a62"
  "PCR8":
    "dbf4a7f9fab7f18619b5899c407081981ad6762fb9a809da78548821b5021965423181584acd7b20
1703376f1133a546"
}
```

Then you have obtained a QingTian Enclave image. 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](#).

19.3.3 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.

19.4 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.

19.4.1 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 |

| PCR | Measurement | Remarks |
|------|---|--|
| 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. |
| 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:
0d1ae7330f437ee563178df30a7c7b7634125d31cac14f6784933db5e90080008438b38fdbb39c886ffe0586ab099b56
EXTEND_PCR res: data:
b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c0364a9b88dd58017a62
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:

c5b3e075e00c261e7fc364f1541067b2a42d4b793225ab10e5cfb8eaca31b3d598af9dd2e491828c2569a9953401abcb

EXTEND_PCR res: data:

4f8b066ce5ac24150612ba9a55bbb9211f626152ada40ede160f4d7ecbfa214c2a549181f6611a3d16a12ec88a577a01

19.4.2 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.

19.4.3 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(
  {1: -35},           // COSE_Sign1 CBOR tag
  {},                 // 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 ]
```

19.4.4 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": [  
            "c5158cb6ee9dbb0ead648c3dc80e472c85e0d67f19fb53fbd3fb94c3371aec63cdb93b80d727a7084248873b1d8e8b41"  
          ],  
          "kms:RecipientAttestation/PCR8": [  
            "705afb1012d27f4e07a25e674e6a17dec57305e29cd412184b7bcb78d9e67f16a0cc26d8706a4fab418a5da5788bc949"  
          ]  
        }  
      }  
    ]  
  }  
}
```

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](#).

19.5 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.

19.5.1 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. A typical example is provided to show how to call KMS APIs for decryption in QingTian Enclave.

Table 19-1 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. |

| Type | API | Description |
|----------|--|--|
| | 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 Example

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:22.04
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
```

KMS API Calling and Permission Configuration Example

The following uses open-source sample code to describe how to call KMS APIs and configure permissions in a QingTian Enclave application. The sample application can only run on Linux instances.

Before the configuration, purchase a C7t ECS and use it as the QingTian Enclave parent instance. Configure the ECS by referring to [Getting Started with QingTian Enclave](#).

1. Download the huawei-qingtian-enclave code from the Gitee repository.

The following uses qtsm-java-sdk as an example.

- a. Download the huawei-qingtian-enclave code.

```
cd /home
```

```
git clone https://gitee.com/heathjay/huawei-qingtian.git
```

Switch to the newjay-java-sdk branch.

```
cd /home/huawei-qingtian
```

```
git checkout newjay-java-sdk
```

2. Obtain the required parameters according to the test case.

For details, see <https://gitee.com/HuaweiCloudDeveloper/huawei-qingtian/pulls/21/files>.

Obtain the parameters listed in [Table 19-2](#) and enter them in the test file.

Table 19-2 KMS API parameters

| Parameter | Meaning | How to Obtain |
|------------|-------------------|--|
| ak | Access Key ID | For details, see AK/SK Signing and Authentication Guide . |
| sk | Secret Access Key | |
| project_id | Project ID | For details about how to obtain the project ID, see Obtaining a Project ID . |
| key_id | Key ID in KMS | None |
| host | KMS endpoint | None |

3. Create custom identity policies in IAM.

Use an account with administrator permissions to create a custom identity policy named **enclave-test-kms-api**. For details, see [Creating a Custom Policy](#).

The following is an example custom identity policy:

```
{
  "Version": "5.0",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "kms::generateRandom",

```

```
"kms:cmk:createDataKey",
"kms:cmk:decryptData",
"kms:cmk:decryptDataKey"
],
"Condition": {
  "StringEqualsIgnoreCase": {
    "kms:RecipientAttestation/PCR0": [
      "8f2cbfb3930e59c6de5c4caff0a3f4c0457e8956bfb4556a7ca1f5f4614a741eeee39ae10447eb5baee48d49e6c1cb6c",
      "ff7ba807a385b49fc1c3346bb47215aef503dee6df22d32f733e22b90a9bc4b22424ca7de1a3537ac9608d7ebe461d67",
      "a28e765550d6ad1188860d30167b1fdb9e29c8da825543861bc76ef1e8427fac6b444ec6a1847fc2c22dae8170c2e67"
    ]
  },
  "kms:RecipientAttestation/PCR8": [
    "a9add94b0ecbbd992bade2176370ecf3bfed2cb39b2ec547512b5174279799f2036fa0b8577bdaf503836178bd11ee2"
  ]
}
},
{
  "Effect": "Allow",
  "Action": [
    "kms:cmk:encryptData",
    "kms:cmk:encryptDataKey"
  ]
}
]
```

The values of **PCR0** and **PCR8** need to be configured after the QingTian Enclave image is created.

4. Create an IAM user and grant permissions to the user.
 - a. In the upper right corner of the IAM console, click **Go to New Console**.
 - b. Choose **Users** from the left navigation pane. On the **Users** page, create an IAM user named **enclave-test-kms-uers** and grant permissions defined in the **enclave-test-kms-api** identity policy to the user.

Figure 19-2 Creating an IAM user (1)




Figure 19-3 Creating an IAM user (2)

Identity and Access Management

Users / Create User

1 Set Details — 2 Set Permissions — 3 Finish

★ Username

Description

Management Console Access ☐ Allow user access to the console.

Figure 19-4 Granting permissions to the user

Identity and Access Management

Users / Create User

1 Set Details — 2 Set Permissions — 3 Finish

Assign By: ☐ User group ☒ Identity policy

Identity Policy: Identity policies (1/326)

| Name | Type | Description |
|----------------------|--------|-------------|
| enclave-test-kms-ops | Custom | Custom |

- c. Obtain the AK and SK of the **enclave-test-kms-users** user.
- Click the username to go to the user details page. On the **Security Settings** tab, create and download access keys. For details, see [Creating an Access Key](#).
- Obtain the AK and SK of the user from the downloaded access key file.
- d. Set the corresponding parameters of the test program to the AK and SK.
- Test program: /home/huawei-qingtian/enclave/qt-sm-sdk-java/kms-cms-java/com/huawei/src/test/TestKmsCmsProxy.java

Figure 19-5 Entering the AK and SK

```
9
10 public class TestKmsCmsProxy {
11     public static final String TAG = "TestCmsKmsProxy";
12     public static void main(String[] args) {
13         /* cid is a fixed number of 3, port should be consistent with the port in the qt-
14         proxy.conf file,
15         * and it is the configuration for the vsoc channel, which facilitates communica-
16         tion between the
17         * enclave and the host.
18         */
19         int vsocCid = 3;
20         int vsocPort = 8000;
21         /* The security server needs to create a new user with permissions of encryption,
22         please refer to
23         * the user manual on the website of the Huawei Cloud to get this user's ak sk.
24         */
25         // String ak = System.getenv("HUAWEICLOUD_SDK_AK");
26         // String sk = System.getenv("HUAWEICLOUD_SDK_SK");
27         String ak = "xxxxxxxxxxxx";
28         String sk = "xxxxxxxxxxxx";
29         /* KMS tenant-side domain name */
30         // String host = System.getenv("HUAWEICLOUD_SDK_HOST");
31         String host = "kms.xxx.myhuaweicloud.com";
```

5. Obtain the KMS endpoint.
 - a. For details about the KMS endpoints in different regions, see [Regions and Endpoints](#).
 - b. Set the **host** parameter in the **huawei-qingtian/enclave/qt-sm-sdk-java/kms-cms-java/com/huawei/src/test/TestKmsCmsProxy.java** file to the KMS endpoint.

Figure 19-6 Entering the endpoint

```
17     int vsockCid = 3;
18     int vsockPort = 8000;
19     /* The security server needs to create a new user with permissions of encryption,
    please refer to
20     * the user manual on the website of the Huawei Cloud to get this user's ak sk.
21     */
22     // String ak = System.getenv("HUAWEICLOUD_SDK_AK");
23     // String ak = "xxxxxxxxxxxx";
24     String ak = "xxxxxxxxxxxx";
25     // String sk = System.getenv("HUAWEICLOUD_SDK_SK");
26     //String sk = "xxxxxxxxxxxx";
27     String sk = "xxxxxxxxxxxx";
28     /* KMS tenant-side domain name */
29     // String host = System.getenv("HUAWEICLOUD_SDK_HOST");
30     //String host = "kms.xxx.mvhuaweicloud.com";
31     String host = "kms.xxx.mvhuaweicloud.com";
32     /* The security officer needs to purchase an elastic cloud server with the enclav
```

6. Obtain the project ID.
 - a. Obtain the project ID on the console. For details, see [Obtaining a Project ID](#).
 - b. Set the **uriPrefix** parameter in the **/home/huawei-qingtian/enclave/qt-sm-sdk-java/kms-cms-java/com/huawei/src/test/TestKmsCmsProxy.java** file to the obtained project ID.

Figure 19-7 Entering the project ID

```
32     /* The security officer needs to purchase an elastic cloud server with the enclav
    e service.
33     * The security officer needs to go to the unified identity authentication servic
    e center of Huawei Cloud
34     * to create a new project., to get project ID, then uriPrefix is /v1.0/{project I
    D}/kms/
35     */
36     // String uriPrefix = System.getenv("HUAWEICLOUD_SDK_URI");
37     //String uriPrefix = "/v1.0/{project ID}/kms/";
38     String uriPrefix = "/v1.0/{project ID}/kms/";
```

7. Obtain the key ID.
 - a. Create a key and obtain the key ID on the DEW console. For details, see [Creating a Custom Key](#).
 - b. Set the **uriPrefix** parameter in the **/home/huawei-qingtian/enclave/qt-sm-sdk-java/kms-cms-java/com/huawei/src/test/TestKmsCmsProxy.java** file to the obtained key ID.

Figure 19-8 Entering the key ID

```
50     System.out.println("Failed to get environment variable.");
51     System.exit(1);
52     }
53     SigParams params = new SigParams(ak, sk, host, uriPrefix);
54     /* Create a new key in the kms service for data encryption and obtain the k
    y ID
55     */
56     //String keyid = System.getenv("HUAWEICLOUD_SDK_KEY");
57     String keyid = "xxxxxxxxxxxx";
58     KeyIdHandle keyIdHandle = new KeyIdHandle(keyid, keyid.length());
59     int rndLen = 256;
60     int datakey_length = 512; // plainkey_len = 512 / 8
```


Figure 19-11 Entering the enclave

```
[root@ecs-6b79 scripts]# /root/huawei-qingtian/nc-vsock/nc-vsock 4 9999
sh: 0: can't access tty; job control turned off
#
```

- c. Run the following commands in the QingTian Enclave instance:

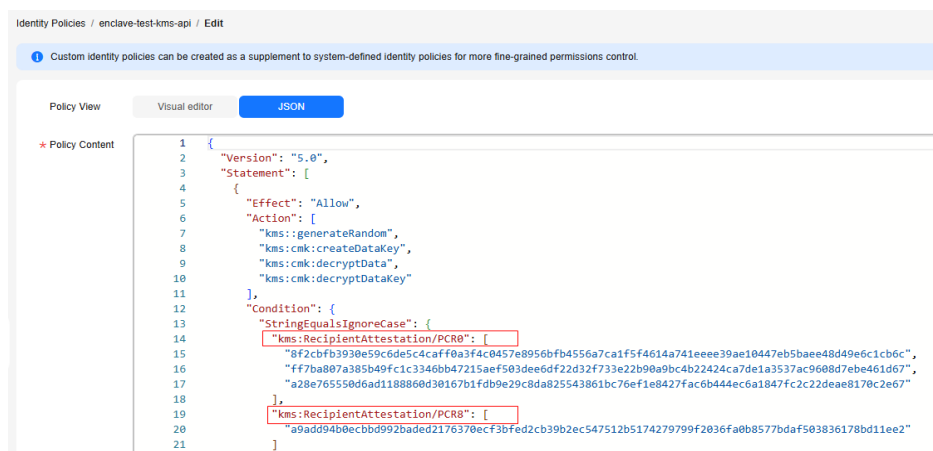
```
cd /home/builder/enclave/qtsm-sdk-java/kms-cms-java/target
# We can perform it manually
java -cp ../lib/lombok-1.18.26.jar:/lib/junit-4.13.1.jar -Djava.library.path=../lib
com.huawei.src.test.TestKmsCmsProxy
```

Figure 19-12 Error message

```
* Connection #0 to host kms.ap-southeast-3.myhuaweicloud.com left intact
HTTPS status: 403
KMS response content:
{"error":{"error_msg":"The user role has no permission to access the interface. User: iam::c9
0494f3fa3a44cf95607da97fae75b1:user:enclave-test-kms-uers is not authorized to perform: kms:c
mk:createDataKey on resource: kms:ap-southeast-3:c90494f3fa3a44cf95607da97fae75b1:KeyId:a64c0
d3c-4a53-4a4c-b0f6-b88f3f51f32a because no identity-based policy allows the kms:cmk:createDat
aKey action.", "encoded_authorization_message": "Eg5hcC1zb3V0aGVhcnQ3dXMxogYzkwNDk0ZjNmYTZhNDRjZj
k1NjA3ZGE5N2ZlZTc1YjEgogYyQDI/oTM/wafknJjKmHg==.ew2cLq8brghe0E4WuQa3x0bC5IFFyA7ctAAq0fVIDyaECd
wy9Q0nMPdJC8dDTh7VPpoBbQq3DuJZpe0X7IdQbs2VyKZKRnQNFx4bZ+l3oX05vMPaexmECMtdDjIc0LE8gFi0Uy/Au1p
4zEhJhr27tDKkwgXlfstIFJBV0BCLLT4Ce/7LVmJpK35L/E5DaadhLjLBf8hLVDOE0pjiUJC4roVHS0cg0fWJsh3cQfIT
1xAxBudvrJGQ1HxEb9IiLnR4CmjDksnnw7bR9WY9Ujz9TMGytqPxWIZgHGANGC+Dur/irDGtjsMoWCRBTBj1mIfqufA2h
ymBgC HdIkMe1Ro9bqHH+uAK8U9WfV6QhNaAODNosKYt/0SHemoVXRgFD0TumHPKV90B1r80uArR23rX4z9WG+h3g4lLn
CA68cc6R4Y62/6P/gKobyLQWVIS05bpc4R/H4K19Jd7uho2wmxmNyJJP8hNHRMNLG0cm4qkXzN+ojgU4qQUrA0V3E8GmL
```

- If the command execution failed, the system displays a message indicating that the user does not have sufficient permissions, as shown in [Figure 19-12](#). Modify the custom identity policy for the user based on PCR0 and PCR8 of the Enclave image. For details, see [step 11.d](#).
 - If the command execution is successful, go to [step 11.e](#).
- d. (Optional) Modify the custom identity policy. The modification will be applied in about 30 to 90 seconds.

After the command is executed, if the user does not have sufficient permissions, information shown in [Figure 19-12](#) is displayed. Modify the custom identity policy based on PCR0 and PCR8 of the Enclave image.

Figure 19-13 Modifying a custom identity policy

- e. Run the following commands in the QingTian Enclave instance to debug the KMS API:
- ```
cd /home/builder/enclave/qtsm-sdk-java/kms-cms-java/target
We can perform it manually
```

```
java -cp ../lib/lombok-1.18.26.jar:../lib/junit-4.13.1.jar -Djava.library.path=../lib
com.huawei.src.test.TestKmsCmsProxy
```

- f. Verify that the KMS API can be called and check the debugging result.

Figure 19-14 Decryption API

```
TestCmsKmsProxy [testKmsDecryptData] after native kmsDecryptDataBlocking data_output length i
s 1000
hello world!
```

Figure 19-15 Data key API

```
key_id is:
a64c0d3c-4a53-4a4c-b0f6-b88f3f51f32a
len of plain_datakey is 512
plain_datakey is:
e6bcd1de9ba14ddac256838dc4c085a092316c8e06a6807f1c2beeb5f08d6f2f4fcbec54e644a375e2b26a0813587
18c2c4ca8b377397354e6082d52766408d1
len of cipher_datakey is 376
cipher_datakey is:
303530303834303037633535373265303437316436326432623132666338393430353939383339646434613535316
63430376461396262383323766656665343133653439316632663431396231373535666564646462666536393539
333837333262306335633836633864313035316537336363323231383430353530383037363161303664353966666
438386539636663346138626433323730636635373262663964333761646234646132343237393032626362633434
3832343934373631333633343633333036343333633326433343631333533332643334363133343633326436323
33036333363264363233383338363633333633333331363633333323631303030303030303031343133313231
313130306630653064306330623061303930383037303630353034303330323031363663336163376131633134346
638373265303638336435323639656639633065313935623032663062386635343462393261303033373839323737
66363532
unix socket listening...
```

Figure 19-16 Random number API

```
len of rnd_data is 32
random data is:
38 12 77 b5 ae 2b c0 47 2e c5 c4 7e d4 7b e7 b6 5b 9c 80 1f 01 d5 fe ac 6a f1 60 11 2c TestCm
sKmsProxy [testKmsGenRandom]the length of random data is 32
TestCmsKmsProxy [testKmsGenRandom]the randomData is:
381277b5ae2bc0472ec5c47ed47be7b65b9c801f01d5feac6af160112cd23389
d2 33 89
```

## 19.5.2 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.

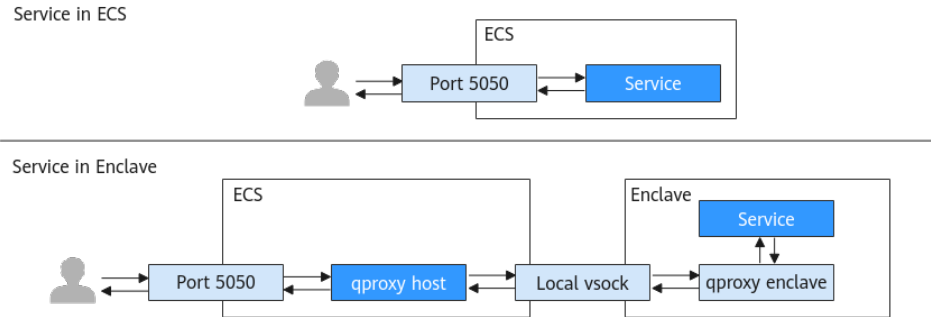
- On the parent instance, run **/path/to/qproxy host --config=/path/to/config\_qproxy.toml <cid>** to enable qproxy.
- On a QingTian Enclave instance, run **/path/to/qproxy host --config=/path/to/config\_qproxy.toml <cid>** to enable qproxy.

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 19-17 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 19-3 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**.

```
cargo build --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.

```
mkdir -p /home/workspace
cp ./qproxy /home/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:22.04
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 19-4 Configuration parameters

Bound	Variable	Type	Description
outbound_connections	hostname	String	Hostname used to forward traffic, for example, <b>api.myservice.com</b> . It also can be an IP address.
	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 <b>host_port</b> is set to <b>80</b> , 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 <b>enclave_port</b> is set to <b>80</b> , 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 19-5** Log file configuration

Variable	Type	Description
host_log	String	Log file name of qproxy on the host. For example, if <b>host_log</b> is set to <b>host.log</b> , the log file name of qproxy on the host is <b>host.log</b> .
enclave_log	String	Log file name of qproxy on the QingTian Enclave instance. For example, if <b>enclave_log</b> is set to <b>enclave.log</b> , the log file name of qproxy on the QingTian Enclave instance is <b>enclave.log</b> .
log_level	String	qproxy log level. For example, if <b>log_level</b> is set to <b>info</b> , 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 <b>/var/log/qproxy</b> .
enclave_log_dir	String	Log directory of qproxy on the QingTian Enclave instance. The default value is <b>/var/log/qproxy</b> .

### 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"
```

## 19.5.3 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.

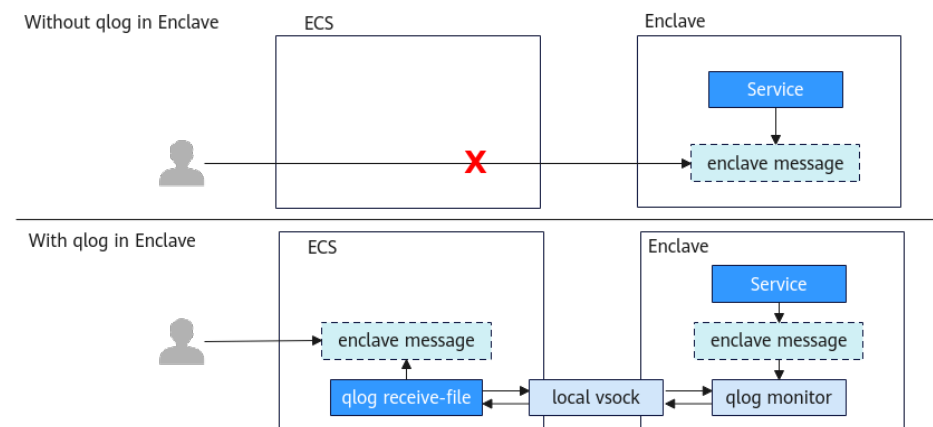
- On the parent instance, run **/path/to/qlog receive-file <cid>/path/to/config\_qlog.toml** to enable **qlog**.
- On a QingTian Enclave instance, run **/path/to/qlog monitor /path/to/config\_qlog.toml** to enable **qlog**.

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 19-18** 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 19-6** 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_interval: 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:22.04
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 19-7 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>

Variable	Type	Description
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 19-8** Monitoring item parameters

Variable	Type	Description
name	String	Name of the monitoring item.
monitor_type	String	Monitoring type. The value <b>file</b> indicates log files, and the value <b>resource</b> 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 <b>monitor_type</b> is <b>file</b> .

Variable	Type	Description
monitor_internel	u32	(Optional) Resource monitoring interval, in seconds. The default value is 15. This parameter can be specified only when <b>monitor_type</b> is <b>resource</b> .
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
```

## 19.6 QingTian CLI (qt CLI)

### 19.6.1 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



```
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": "b8c59692da8a5bcb739a83d15a0ceca670bd78da06cb2250ec70548f72254e674419e9888db9c0364a9b88dd58017a62",
 "PCR8": "dbf4a7f9fab7f18619b5899c407081981ad6762fb9a809da78548821b5021965423181584acd7b201703376f1133a546"
}
```

## 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.

## 19.7 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	loctl get sandbox capacity failure.	Failed to get sandbox capacity using loctl.	Contact Huawei Cloud technical support.

Error Code	Error Message	Description	Solution
05	loctl define sandbox failure.	Failed to define sandbox using loctl.	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	loctl add memory failure.	Failed to increase memory using loctl.	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.
11	loctl add cpu failure.	Failed to add vCPUs using loctl.	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.

Error Code	Error Message	Description	Solution
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 message is read from the QingTian Enclave instance.	Contact Huawei Cloud technical support.
19	loctl start an enclave failure.	Failed to start the QingTian Enclave instance using loctl.	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.
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.

Error Code	Error Message	Description	Solution
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 <b>qt enclave query</b> 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.
32	The required enclave is not running in the debug mode.	The requested QingTian Enclave instance is not running in debug mode.	Run the <b>qt enclave query</b> 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.

Error Code	Error Message	Description	Solution
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.
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.

## 19.8 FAQs About QingTian

## 19.8.1 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](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:
  - a. 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
```

- b. 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:

- a. 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.

- b. Apply the modification.

```
grub2-mkconfig -o /boot/efi/EFI/hce/grub.cfg
```

- c. Restart the VM and check the number of available huge pages again.

```
reboot
```

## 19.8.2 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.

## What Do I Need to Do If I Use an Ubuntu Image?

1. Use the [huawei-qingtian](#) source code to compile the **virtio-qtbox** and **qingtian-tools** packages.
2. Create the **qt-enclave-bootstrap** package required by the image.  
You can download the latest **qt-enclave-bootstrap** package and decompress it to the specified directory. The download URL is [https://repo.huaweicloud.com/hce/2.0/updates/x86\\_64/Packages/](https://repo.huaweicloud.com/hce/2.0/updates/x86_64/Packages/).  
**/usr/local/share/qingtian/enclave/init**  
**/usr/local/share/qingtian/enclave/qtsm.ko**  
**/usr/local/share/qingtian/enclave/vmlinux.bin**

**`/usr/local/share/qingtian/enclave/vmlinux.bin.bz2`**