

Elastic Cloud Server

User Guide

Issue 01
Date 2024-11-15



Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei 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 Selecting an ECS Billing Mode.....	5
2.1.1 Yearly/Monthly Billing.....	5
2.1.2 Pay-per-Use Billing.....	6
2.1.3 Spot Pricing.....	6
2.1.3.1 Spot Pricing ECSs.....	6
2.1.4 Reserved Instances.....	10
2.1.4.1 Reserved Instance Overview.....	10
2.1.4.2 Enabling and Purchasing a Reserved Instance.....	15
2.1.4.3 Modifying RI Attributes.....	19
2.1.5 Changing Pay-per-Use to Yearly/Monthly.....	20
2.1.6 Changing Yearly/Monthly to Pay-per-Use.....	21
2.2 Purchasing an ECS.....	22
2.2.1 Purchasing a Spot ECS.....	22
2.2.2 Purchasing the Same ECS.....	25
2.3 Logging In to a Windows ECS.....	26
2.3.1 Login Overview (Windows).....	26
2.3.2 Logging In to a Windows ECS Using VNC.....	27
2.3.3 Logging In to a Windows ECS Using MSTSC.....	28
2.3.4 Logging In to a Windows ECS from a Linux Computer.....	35
2.3.5 Logging In to a Windows ECS from a macOS Server.....	37
2.3.6 Logging In to a Windows ECS from a Mobile Terminal.....	40
2.4 Logging In to a Linux ECS.....	45
2.4.1 Login Overview (Linux).....	45
2.4.2 Logging In to a Linux ECS Using VNC.....	46
2.4.3 Logging In to a Linux ECS Using an SSH Key Pair.....	48
2.4.4 Logging In to a Linux ECS Using an SSH Password.....	54
2.4.5 Logging In to a Linux ECS from a macOS Server.....	56
2.4.6 Logging In to a Linux ECS from a Mobile Terminal.....	57
2.5 Managing GPU Drivers of GPU-accelerated ECSs.....	69

2.5.1 GPU Driver.....	69
2.5.2 Obtaining a Tesla Driver and CUDA Toolkit.....	70
2.5.3 Manually Installing a GRID Driver on a GPU-accelerated ECS.....	71
2.5.4 Manually Installing a Tesla Driver on a GPU-accelerated ECS.....	80
2.6 Managing ECS Configurations.....	88
2.6.1 Changing the Time Zone for an ECS.....	88
2.6.2 Enabling or Disabling Hyper-Threading.....	90
2.6.3 Obtaining Metadata and Passing User Data.....	92
2.6.3.1 Obtaining Metadata.....	92
2.6.3.2 Passing User Data.....	101
2.6.4 Changing ECS Names.....	106
2.6.5 Managing ECS Groups.....	107
2.6.6 Automatically Recovering ECSs.....	109
2.6.7 Configuring Mapping Between Hostnames and IP Addresses in the Same VPC.....	110
2.7 Modifying ECS Specifications (vCPUs and Memory).....	111
2.7.1 Modifying Individual ECS Specifications.....	111
2.8 Reinstalling or Changing the OS.....	113
2.8.1 Reinstalling the OS.....	113
2.8.2 Changing the OS.....	114
2.9 Viewing ECS Information.....	117
2.9.1 Viewing ECS Creation Statuses.....	117
2.9.2 Viewing Failed Tasks.....	117
2.9.3 Viewing ECS Details (List View).....	118
2.9.4 Exporting ECS Information.....	119
2.9.5 Searching for ECSs.....	119
3 Images.....	127
3.1 Overview.....	127
3.2 Creating an Image.....	129
4 Disks.....	131
4.1 Overview.....	131
4.2 Adding a Disk to an ECS.....	131
4.3 Attaching a Disk to an ECS.....	132
4.4 Detaching an EVS Disk from a Running ECS.....	133
4.5 Expanding the Capacity of an EVS Disk.....	135
4.6 Expanding the Local Disks of a Disk-intensive ECS.....	135
5 Elastic Network Interfaces.....	137
5.1 Overview.....	137
5.2 Attaching a Network Interface.....	138
5.3 Detaching a Network Interface.....	140
5.4 Changing a VPC.....	140
5.5 Modifying a Private IP Address.....	142

5.6 Managing Virtual IP Addresses.....	142
5.7 Enabling NIC Multi-Queue.....	143
6 EIPs.....	148
6.1 Overview.....	148
6.2 Binding an EIP.....	149
6.3 Unbinding an EIP.....	150
6.4 Changing an EIP.....	150
6.5 Modifying an EIP Bandwidth.....	151
6.6 Enabling Internet Connectivity for an ECS Without an EIP.....	152
7 Security.....	156
7.1 Methods for Improving ECS Security.....	156
7.2 Security Groups.....	160
7.2.1 Overview.....	160
7.2.2 Default Security Groups and Rules.....	161
7.2.3 Security Group Configuration Examples.....	165
7.2.4 Configuring Security Group Rules.....	171
7.2.5 Changing a Security Group.....	176
7.3 HSS.....	177
7.4 Project and Enterprise Project.....	178
8 Backup Using CBR.....	180
8.1 Overview.....	180
8.2 Backing Up an ECS.....	188
9 Passwords and Key Pairs.....	191
9.1 Password Reset.....	191
9.1.1 Application Scenarios for Using Passwords.....	191
9.1.2 Resetting the Password for Logging In to an ECS in the OS.....	192
9.1.3 Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed.....	193
9.1.4 Resetting the Password for Logging In to a Linux ECS.....	195
9.2 Key Pairs.....	198
9.2.1 Application Scenarios for Using Key Pairs.....	198
9.2.2 (Recommended) Creating a Key Pair on the Management Console.....	199
9.2.3 Creating a Key Pair Using PuTTY Key Generator.....	200
9.2.4 Importing a Key Pair.....	203
9.2.5 Obtaining and Deleting the Password of a Windows ECS.....	204
9.2.5.1 Obtaining the Password for Logging In to a Windows ECS.....	204
9.2.5.2 Deleting the Initial Password for Logging In to a Windows ECS.....	205
10 Resources and Tags.....	207
10.1 Tag Management.....	207
10.1.1 Overview.....	207
10.1.2 Adding Tags.....	208

10.1.3 Searching for Resources by Tag.....	210
10.1.4 Deleting a Tag.....	211
10.2 Quota Adjustment.....	212
11 Monitoring Using Cloud Eye.....	214
11.1 Monitoring ECSs.....	214
11.2 Basic ECS Metrics.....	215
11.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed.....	223
11.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed.....	265
11.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed and Using Simplified Monitoring Metrics.....	269
11.6 Setting Alarm Rules.....	273
11.7 Viewing ECS Metrics.....	273
12 Audit Using CTS.....	275
12.1 ECS Operations Supported by CTS.....	275
12.2 Viewing Traces.....	276

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](#)).

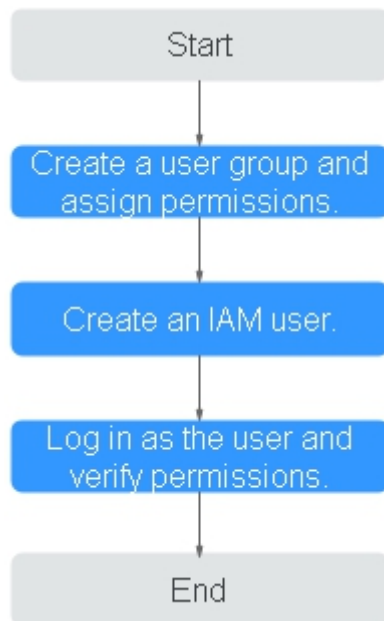
Prerequisites

Before assigning permissions to user groups, you should learn about system policies supported by ECS and select the policies based on service requirements.

For more information about system policies supported by ECS, see [ECS Permissions](#). For the permissions of other services, see [System-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 **1**.
3. Log in to the management console as the created user.
In the authorized region, perform the following operations:
 - Choose **Compute > Elastic Cloud Server** in the service list. On the ECS console, click **Create ECS**. If the purchase attempt failed, the **ECSReadOnlyAccess** policy has already taken effect.
 - Choose any service other than ECS in the service list. If a message appears indicating that you have insufficient permissions to access the service, the **ECSReadOnlyAccess** policy has already taken effect.

1.2 ECS Custom Policies

Custom policies can be created to supplement the system-defined policies of ECS. For the actions that can be added to custom policies, see "Permissions and Supported Actions" in [Elastic Cloud Server API Reference](#).

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see [Creating a Custom Policy](#). The following provides examples of common ECS custom policies.

Example Custom Policies

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

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

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

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

- Example 3: Only allowing VNC login

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

- Example 4: Denying ECS deletion

A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

The following method can be used if you need to assign permissions of the **ECSFullAccess** policy to a user but you want to prevent the user from deleting ECSs. Create a custom policy for denying ECS deletion, and attach both policies to the group which the user belongs to. Then, the user can perform all operations on ECSs except deleting ECSs. The following is an example of a deny policy:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "ecs:cloudServers:delete"
      ]
    }
  ]
}
```

2 Instances

2.1 Selecting an ECS Billing Mode

2.1.1 Yearly/Monthly Billing

Concept

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

For more billing information, see [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 any more, unsubscribe it. To do so, switch to the **Elastic Cloud Server** page, locate the target ECS, and choose **More** > **Unsubscribe** in the **Operation** column.
2. A detached system disk can be used as a data disk for any ECSs, but can only be used as a system disk for the ECS where it was attached before.
3. A detached data disk that is purchased together with an ECS can only be used as a data disk for this ECS.

Resources Supporting Yearly/Monthly Billing

Resources billed in yearly/monthly mode include:

- ECSs (vCPUs and memory)
 - Images, including prepaid 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 bandwidths.

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

For details about ECS prices, see [Price Calculator](#).

2.1.2 Pay-per-Use Billing

Concept

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

For more billing information, see [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 of each type of ECS can be obtained by dividing their hourly price by 3600. Obtain the hourly price on the [Product Pricing Details](#) page.

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

- If you use the ECS for 30 minutes, you need to pay for \$0.34 USD ($0.68/3600 \times 30 \times 60$).
- If you use the ECS for 1 hour and 30 minutes, you need to pay for \$1.02 USD ($0.68/3600 \times 90 \times 60$).

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](#).

2.1.3 Spot Pricing

2.1.3.1 Spot Pricing ECSs

Concept

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

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

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

Working Rules

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

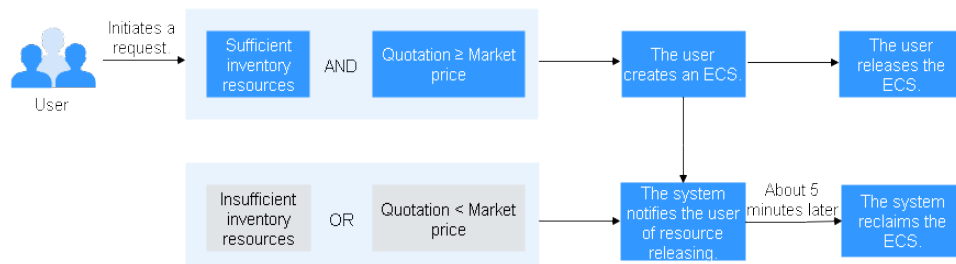
When purchasing a spot ECS, you are required to set the maximum price you are willing to pay for a specified flavor. A higher price ensures a greater success rate for you to purchase such an ECS.

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

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

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

Figure 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.
- 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. 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 price ECS and bills it based on the market price. Example:**

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

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

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

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

This ECS is billed in three billing periods.

During 08:30-09:00, the ECS had been running for 30 minutes and it is billed by the second: $0.02/3600 \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 ECS. If during that 5 minutes, the spot pricing hour is exceeded, any time in excess of that hour will be billed at the new market price.
- During the running of a spot ECS, its price is updated once an hour. After a spot ECS is restarted, or it is stopped and then started, it is billed at the market price when the ECS starts.

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

FAQs

See [FAQs About Spot ECSs](#).

2.1.4 Reserved Instances

2.1.4.1 Reserved Instance Overview

Concept

A reserved instance (RI) is not an actual instance, but a billing discount that can be applied to the use of pay-per-use ECSs in your account. When the attributes of your pay-per-use ECSs **match** those of an RI, the RI billing benefit automatically applies to your ECSs. The combination of RIs and pay-per-use billing fully utilizes the flexibility of pay-per-use resources at lower costs.

 NOTE

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

Table 2-1 ECS billing modes

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

Billing Mode	What It Is	How to Use
Yearly/ Monthly	ECS billed based on the service duration. This mode is ideal when the duration of ECS usage is predictable.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.
Spot pricing	ECS billed in spot pricing billing mode.	A basic computing unit that consists of vCPUs, memory, OS, and EVS disks. After purchasing such an ECS, you can use it on the cloud.

- For instructions about how to purchase an RI, see [Enabling and Purchasing a Reserved Instance](#).
- For instructions about how to modify an RI, see [Modifying RI Attributes](#).

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

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

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

Application Scenarios

If your ECSs are to be used in a short term, it is a good practice to use 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 2-2](#) lists RI attributes. You can purchase your desired RIs based on these attributes.

Table 2-2 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 ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

Zonal RIs

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

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

Regional RIs

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

- AZ flexibility:** The RI discount applies to pay-per-use ECS usage in any AZ within a region.
- ECS size flexibility:** The RI discount applies to pay-per-use ECS usage when the ECS OS, ECS series, and vCPU/memory ratio of the target ECS are the same as

those specified in the regional RI. ECS size flexibility is determined based on the normalization factor of the ECS size. ECS size flexibility does not apply to zonal RIs.

ECS size flexibility is applied from the smallest to the largest ECS size within the ECS series based on the normalization factor. [Table 2-3](#) describes ECS size within an ECS type and corresponding normalization factor per hour.

 **NOTE**

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

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


Table 2-3 Normalization factors

ECS Size	Normalization Factor
small	1
medium	1
large	2
xlarge	4
2xlarge	8
4xlarge	16
6xlarge	24
7xlarge	28
8xlarge	32
9xlarge	36
12xlarge	48
14xlarge	56
15xlarge	60
16xlarge	64
26xlarge	104
52xlarge	208
nxlarge	n x 4

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

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

Figure 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 2-4 Comparison between regional and zonal RIs

RI Type	AZ Flexibility	ECS Size Flexibility	Capacity Reservation
Regional RI	Supported A regional RI applies to any AZ in the region.	Supported A regional RI applies when the ECS series and vCPU/memory ratio of the target ECS are the same as those specified in the RI.	Not supported Resources are not reserved so ECS creation may fail when resources are insufficient.

RI Type	AZ Flexibility	ECS Size Flexibility	Capacity Reservation
Zonal RI	Not supported A zonal RI applies only in a specified AZ.	Not supported A zonal RI applies only when the flavor of the target ECS is the same as that specified in the RI.	Supported Desired resources can be reserved for creating a pay-per-use ECS.

Examples

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

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

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

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

The RI benefits are applied as follows:

- The discount of the five s3.large.2 zonal RIs is used by the five s3.large.2 ECSs because the attributes (AZ, OS, and ECS type) between the RIs and ECSs match.
- The m3.large.2 regional RIs offer AZ flexibility and ECS size flexibility.
An m3.large.2 RI is equivalent to two normalization factors. The six m3.large.2 regional RIs are equal to 12 normalization factors (6 x 2). 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 ECS size flexibility and can be applied to c3.xlarge.2 ECSs.
A c3.large.2 RI is equivalent to two normalization factors (1 x 2). A c3.xlarge.2 ECS requires an RI with four normalization factors (1 x 4). Therefore, the c3.large.2 RI billing discount applies to 50% of c3.xlarge.2 usage. The remaining c3.xlarge.2 usage is billed at the pay-per-use rate.

2.1.4.2 Enabling and Purchasing a Reserved Instance

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

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

- For more information about RIs, see [Reserved Instance Overview](#).
- For instructions about how to modify an RI, see [Modifying RI Attributes](#).



Notes and Constraints

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

Enabling RIs

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

Purchasing an RI

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Reserved Instance**.
5. Click **Buy RI**.
The **Buy RI** page is displayed.
6. Confirm the region.
If the RIs in the selected region do not meet your requirements, select another region.
7. (Optional) Select **Show offerings that reserve capacity** to view the AZs that support capacity reservations.
 - Zonal RIs offer capacity reservation.
 - Regional RIs offer capacity reservation.
8. (Optional) Select an AZ to purchase a zonal RI for capacity reservation.
Perform this operation only when you purchase RIs for a specified AZ.
9. Select an RI type.
The cloud platform provides various RI types for you to choose from based on your application scenarios.
10. Filter for RI specifications.
Set flavor, OS, term, offering class, and payment option to search for the target RI specifications.

[Table 2-5](#) shows specifications parameters.

Table 2-5 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 ECS size can be modified during the term. However, the ECS type cannot be changed.
Payment Option	No upfront

11. Select specifications.

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

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

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

Hourly Rate: amortized hourly costs of the RI, which is equivalent to the difference between the total cost of the RI and the upfront payment divided by the total number of hours over the entire term (Hourly rate = Total cost of the RI – Upfront payment/Entire term of the RI)

12. Specify an RI name.

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

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

For details, click **Pricing details**.

14. Click **Next**.
To learn more about the price, click **Pricing details**.
15. On the page for you to confirm RI specifications, view details and submit the request.
After verifying the configurations and price, click **Submit** and pay for the order as prompted.
16. Return to the RI list as prompted and view the purchased RI.

Follow-up Operations

- **Purchase a pay-per-use ECS that matches an RI.**
Locate the target RI and click **Buy ECS** in the **Operation** column. The system automatically switches to the page for purchasing ECSs, and the specifications of the ECSs selected by default are the same as those specified in the RI.

NOTE

- If the OS of the target ECS does not match the OS specified in the RI, or the target ECS is not billed on a pay-per-use basis, the RI cannot be used. When the attributes of the ECS match those of the RI, including the ECS series and vCPU/memory ratio, the ECS automatically benefits from the billing discount offered by the RI.
- **Check the usage of RIs.**
 - 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"/> ri-7afl 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: ri-7afl

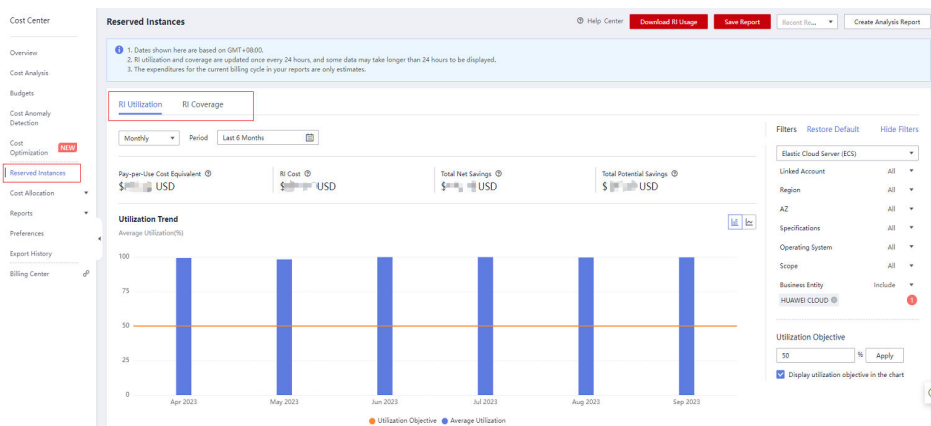
- c. Expand the RI details.

Figure 2-5 RI details



- 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



2.1.4.3 Modifying RI Attributes

Scenarios

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

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

- For more information about RIs, see [Reserved Instance Overview](#).
- For instructions about how to purchase an RI, see [Enabling and Purchasing a Reserved Instance](#).

Constraints

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

Procedure



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

Table 2-6 Common operations for modifying an RI

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

NOTICE

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

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

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

2.1.5 Changing Pay-per-Use to Yearly/Monthly

Scenarios



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

If you need to use an ECS for a long time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs by referring to the content in this section.

Prerequisites

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. Above the ECS list, choose **More > Change Billing Mode**.
6. Confirm the ECS details, specify the usage duration, and pay for the order.

2.1.6 Changing Yearly/Monthly to Pay-per-Use

Scenarios

Yearly/Monthly is a prepaid billing mode in which your ECS will be billed based on service duration. This cost-effective mode is ideal when the duration of ECS usage is predictable.

If you require a more flexible billing mode, in which your ECS will be billed based on usage frequency and duration, you can change the billing mode from yearly/monthly to pay-per-use.

NOTE

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

Constraints

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

Procedure

1. Log in to the management console.
2. On the top navigation bar, choose **Billing > Renewal**.
The **Renewals** page is displayed.
3. Customize search criteria.

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

2.2 Purchasing an ECS

2.2.1 Purchasing a Spot ECS

Scenarios

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

Compared with pay-per-use and yearly/monthly ECSs, spot ECSs offer the same level of performance while at lower costs.

For more information about spot ECSs, see [Spot Pricing \(for Spot Instances\)](#).

Notes and Constraints

- Only KVM ECSs support spot pricing payments. For details about supported ECS flavors, see the information displayed on the management console.
- The market prices of the ECSs of the same flavor may vary depending on AZs.
- Spot ECSs do not support OS change.
- Spot ECSs do not support automatic recovery.
- Spot ECSs do not support specifications modification.
- Spot ECSs cannot be created using a KooGallery image.
- Spot ECSs cannot be switched to yearly/monthly ECSs.
- Spot ECSs do not support system disk detachment.
- When a spot ECS is being reclaimed:
 - It cannot be used to create system disk images and full-ECS images. However, data disks of the ECS can be used to create data disk images.
 - It cannot be deleted.

- By default, the data disks and EIP of a spot ECS will not be released after it is reclaimed. If you want to be notified when a spot ECS is reclaimed so that you can determine whether to manually release data disks and EIP, set a reclaim notification.

Purchasing a Spot ECS

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

When purchasing a spot ECS:

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

NOTE

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

(Optional) Enabling Reclaim Notifications

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

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

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

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

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

Step 2 Configure reclaim notifications.

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


1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Management & Deployment**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Key Event Notifications**.
5. Click **Create Key Event Notification** in the upper right corner of the page and set parameters listed in [Table 2-7](#).

Table 2-7 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 & Deployment**, click **Cloud Trace Service**.
2. In the navigation pane on the left, choose **Trace List**.
3. Specify filter criteria listed in [Table 2-8](#) and search for traces as needed.

Table 2-8 Setting filter criteria to search for reclaimed ECSs

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

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

----End

2.2.2 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 go to the [ECS console](#).
2. Locate the row containing the target ECS and choose **More > Buy Same ECS** in the **Operation** column.
The system switches to the **Buy ECS** page and automatically copies the parameter settings of the selected ECS.
3. Adjust the settings of the new ECSs as needed.
For details about parameter settings, see [Purchasing an ECS](#).
4. Confirm the ECS details, select the agreement, and click **Create**.

 NOTE

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

- Manually add data disks if the quantity of data disks needed exceeds 10.
- Manually add NICs if the quantity of NICs needed exceeds 5.
- Manually add security groups if the quantity of security groups needed exceeds 5.
- Select a new data disk image if the disks of the source ECS are created using a data disk image.
- If the source ECS is created from a full-ECS image, only the disks included in this image are displayed. Add disks if necessary.
- Select **Encryption** if the disks of the source ECS have been encrypted.
- Configure the functions in **Advanced Options**.

2.3 Logging In to a Windows ECS

2.3.1 Login Overview (Windows)

Constraints

- Only a running ECS can be logged in.
- The username for logging in to a Windows ECS is **Administrator**.
- If an ECS uses key pair authentication, use the password obtaining function available on the management console to decrypt the private key used during ECS creation to obtain a password.
- Certain G series of ECSs do not support remote login provided by the cloud platform. If you need to remotely log in to the ECSs, install the VNC server on them. For details, see [GPU-accelerated ECSs](#). You are suggested to log in to the ECSs using MSTSC.
- If you log in to a GPU-accelerated ECS using MSTSC, GPU acceleration will fail. This is because MSTSC replaces the WDDM GPU driver with a non-accelerated remote desktop display driver. In such a case, you must log in to the ECS using other methods, such as VNC. If the remote login function available on the management console fails to meet your service requirements, you must install a suitable remote login tool, such as TightVNC, on the ECS. To download TightVNC, log in at <https://www.tightvnc.com/download.php>.

Login Modes

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

Table 2-9 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 has had an EIP bound. (If you log in to an ECS through an intranet, for example, through VPN or Direct Connect, the ECS does not require an EIP.)
	Linux	Install a remote connection tool, for example, rdesktop. For details, see Logging In to a Windows ECS from a Linux Computer .	
	macOS	Install a remote connection tool, for example, Microsoft Remote Desktop on the macOS. For details, see Logging In to a Windows ECS from a 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 .	

Helpful Links

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

2.3.2 Logging In to a Windows ECS Using VNC

Scenarios

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

Prerequisites

If an ECS uses key pair authentication, make sure that the key file has been used to resolve the login password before logging in to the ECS.

Logging In to a Windows ECS




1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Obtain the password for logging in to the ECS.
Before logging in to the ECS, you must have the login password.
 - If your ECS uses password authentication, log in to the ECS using the password you configured during the ECS creation.
5. In the **Operation** column of the target ECS, click **Remote Login**.
6. In the **Logging In to a Windows ECS** dialog box, click **Log In** in the **Other Login Modes** 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 2-7 Ctrl+Alt+Del



8. Enter the ECS password as prompted.

Helpful Links

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

2.3.3 Logging In to a Windows ECS Using MSTSC

Scenarios

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

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.
- Access to port 3389 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).

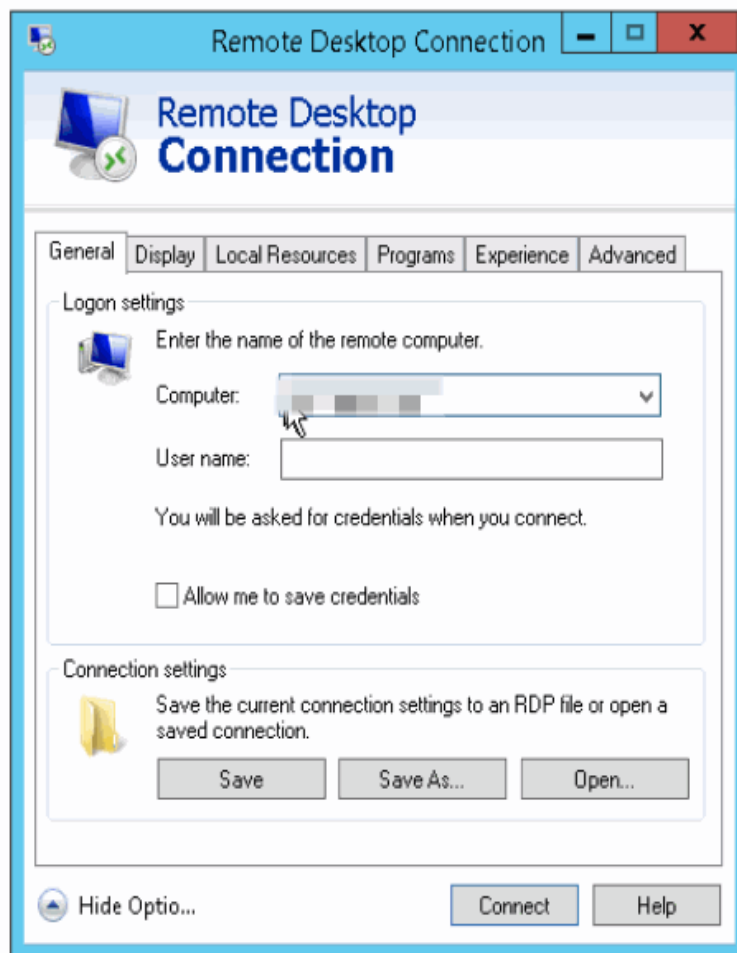
- The network connection between the login tool and the target ECS is normal. For example, the default port 3389 is not blocked by the firewall.
- Remote Desktop Protocol (RDP) needs to be enabled on the target ECS. For ECSs created using public images, RDP has been enabled by default. For instructions about how to enable RDP, see [Enabling RDP](#).

Logging In to a Windows ECS Using MSTSC

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

The following uses Windows Server 2012 ECS as an example.


Figure 2-8 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 2-9 Show Options

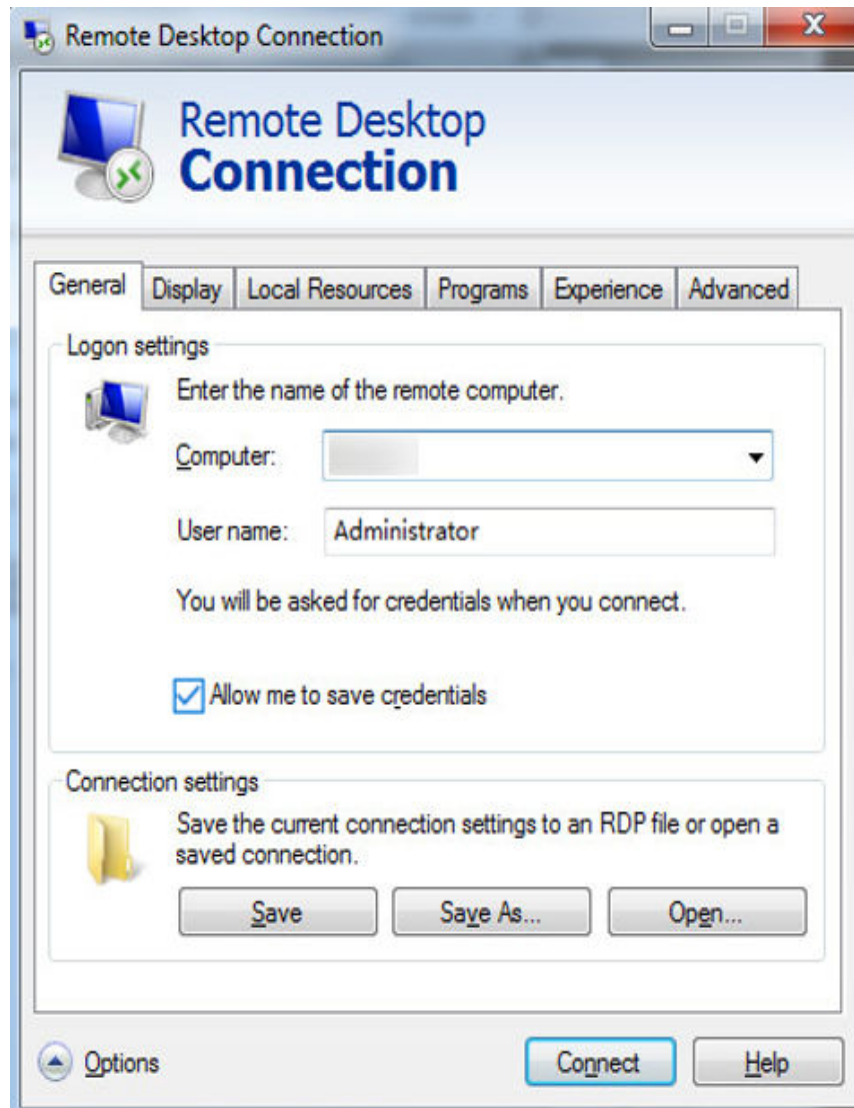


4. Enter the EIP and username (**Administrator** by default) of the target ECS.

 **NOTE**


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

Figure 2-10 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 2-11 Clipboard

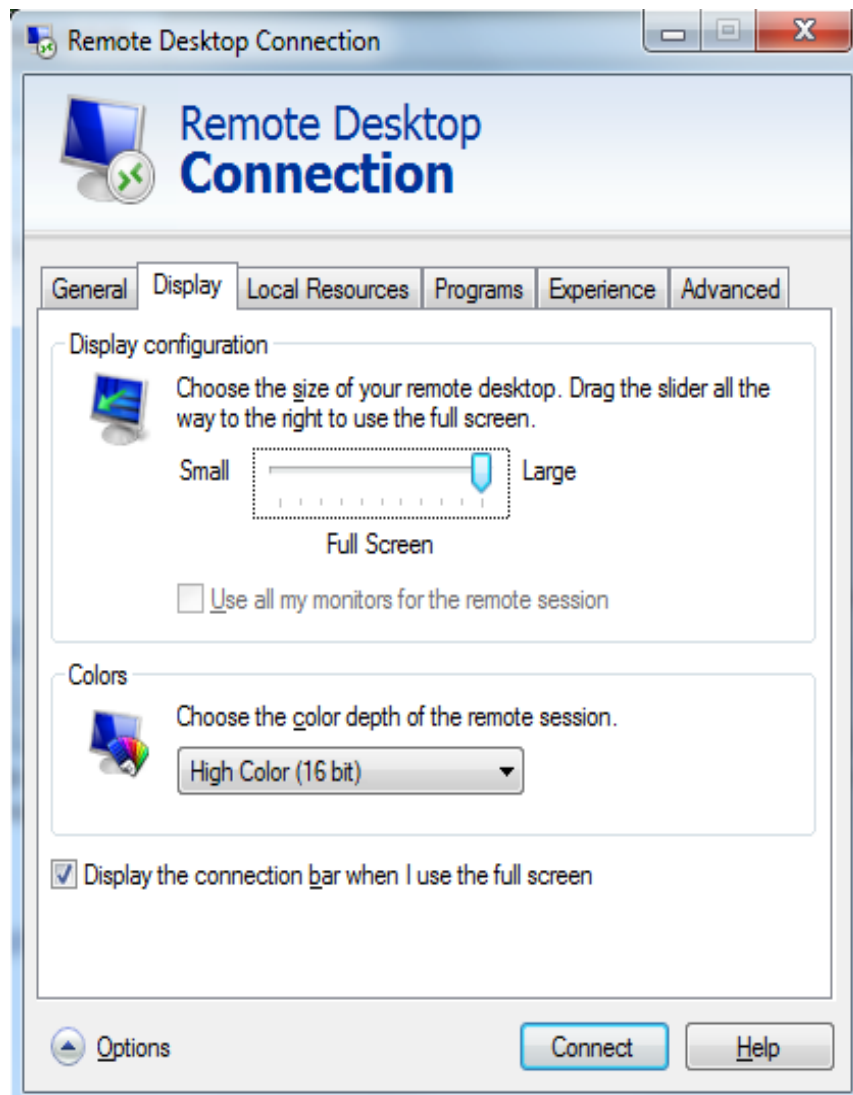


To copy files from the local server to your ECS, click **More** and select your desired disks.

Figure 2-12 Drives



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

Figure 2-13 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

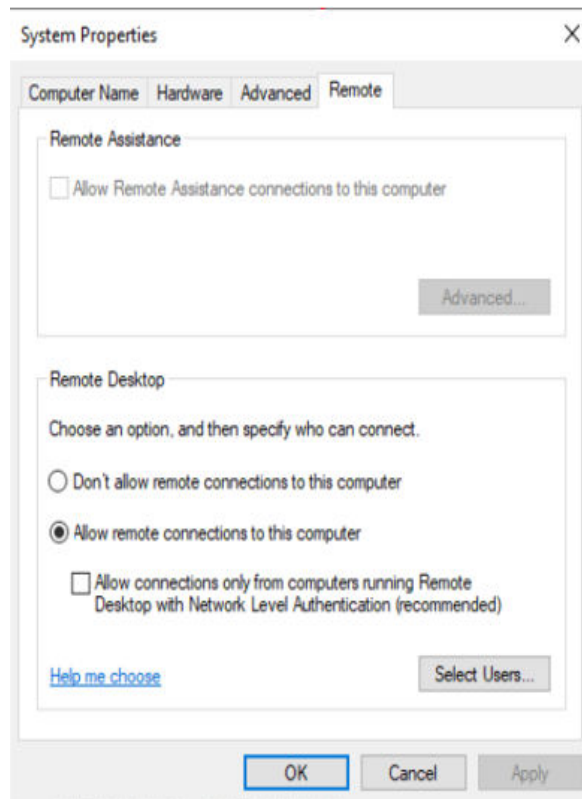
For your first login, use VNC to log in and enable RDP for your ECS. Then, use MSTSC to log in.

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 2-14 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](#)

2.3.4 Logging In to a Windows ECS from a Linux Computer

Scenarios

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

Prerequisites

- The target ECS is running.
- The ECS must have an EIP bound.
An EIP is not required if you log in to an ECS through an intranet using MSTSC, for example, through VPN or Direct Connect.

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

Procedure

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

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

rdesktop

If the message "command not found" is displayed, rdesktop is not installed. In such a case, obtain the rdesktop installation package at the [official rdesktop website](#).

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

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

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

Table 2-10 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 2-15 System Properties

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

2.3.5 Logging In to a Windows ECS from a macOS Server

Scenarios

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

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 3389 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).

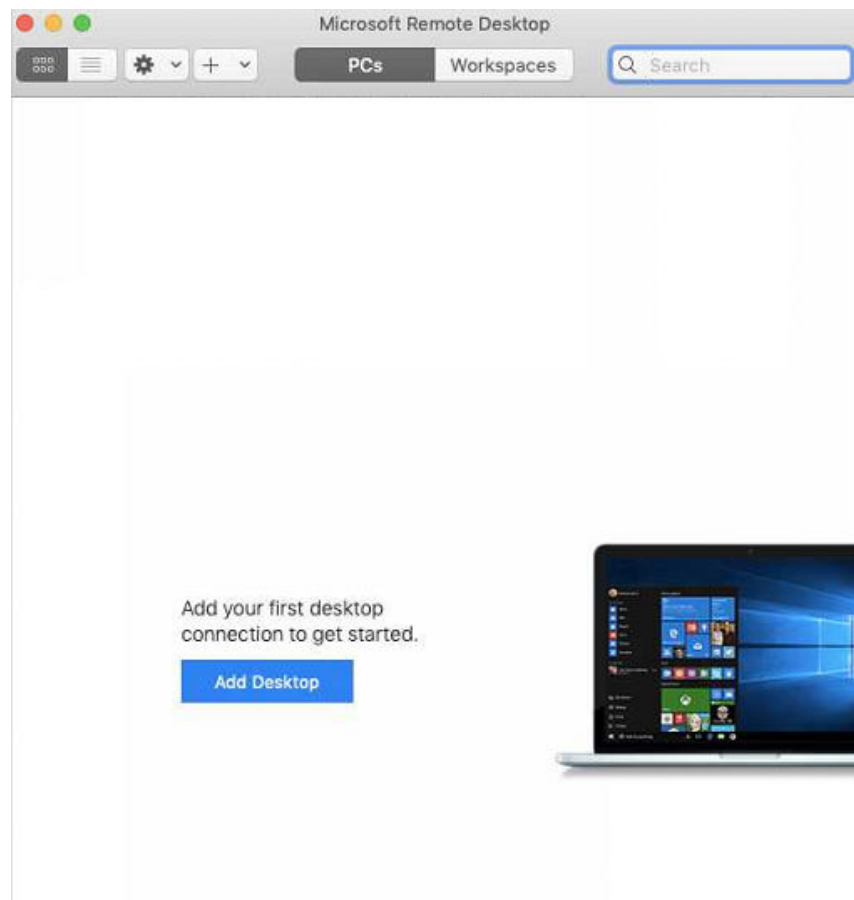
- The remote access tool supported by Mac, such as Microsoft Remote Desktop for Mac has been installed.

Microsoft stopped providing the link for downloading the Remote Desktop client. You can download the beta version by visiting [Microsoft Remote Desktop Beta](#).

Procedure

1. Start Microsoft Remote Desktop.
2. Click **Add Desktop**.

Figure 2-16 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 2-17 Add user account

Add a User Account

Username:

Password:

Show password

Friendly name:

Figure 2-18 Add PC

Add PC

PC name:

User account:

General | Display | Devices & Audio | Folders

Friendly name:

Group:

Gateway:

Bypass for local addresses

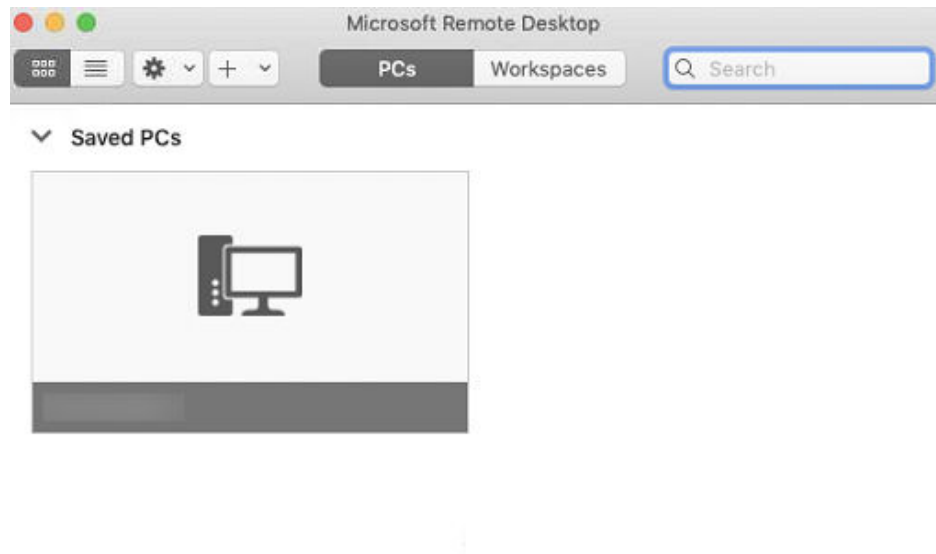
Reconnect if the connection is dropped

Connect to an admin session

Swap mouse buttons


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

Figure 2-19 Double-click for login



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

Figure 2-20 Successful login



2.3.6 Logging In to a Windows ECS from a Mobile Terminal

Scenarios

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

Prerequisites

- The target ECS is running.
- You have obtained the username and password for logging in to the ECS.
- 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 the Microsoft Remote Desktop client.
2. In the upper right corner of the **Remote Desktop** page, tap  and select **Desktop**.

Figure 2-21 Remote Desktop



3. On the **Add desktop** page, set login information and tap **SAVE**.
 - **PC name:** Enter the EIP bound to the target Windows ECS.
 - Perform the following operations to set **User name:**
 - i. Tap **User name** and select **Add user account** from the drop-down list.
The **Add user account** dialog box is displayed.
 - ii. Enter the username **administrator** and password for logging in to the Windows ECS and click **SAVE**.

Figure 2-22 Setting the login information





Figure 2-23 Saving the settings




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

Figure 2-24 Logging in to the Windows ECS




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

Figure 2-25 CONNECT



You have logged in to the Windows ECS.

Figure 2-26 Successful login



2.4 Logging In to a Linux ECS

2.4.1 Login Overview (Linux)

Constraints

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

Login Modes

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

Table 2-11 Linux ECS login modes

ECS OS	Local OS	Connection Method	Requirement
Linux	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 	The target ECS has an EIP bound. (If you log in to an ECS through an intranet, for example, through VPN or Direct Connect, the ECS does not require an EIP.)

ECS OS	Local OS	Connection Method	Requirement
	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?](#)

2.4.2 Logging In to a Linux ECS Using VNC

Scenarios

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

For instructions about how to copy and paste data on VNC pages after the ECS login, see [Follow-up Procedure](#).

NOTE

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

Constraints



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

Logging In to a Linux ECS Using an SSH Key Pair and Logging In to a Linux ECS Using an SSH Password.

Prerequisites

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

Procedure


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

NOTE

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

6. Enter the ECS password as prompted.

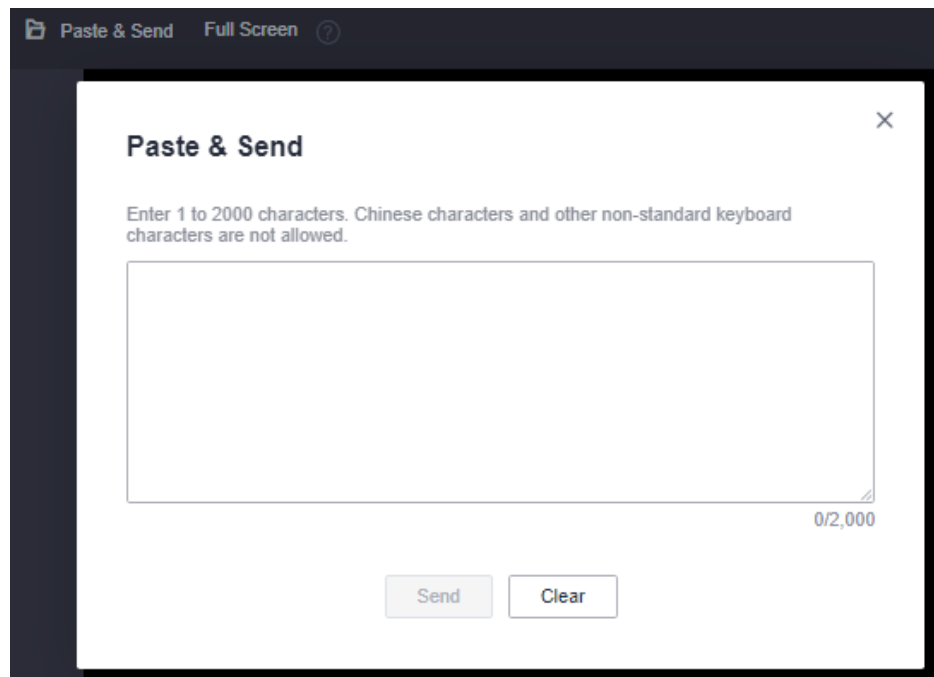
Figure 2-27 Username (root as an example) and password



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 2-28 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?](#)

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

Scenarios

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

Prerequisites

- You have obtained the private key file used for creating the ECS. For details about how to create a key pair, see [\(Recommended\) Creating a Key Pair on the Management Console](#).
- You have bound an EIP to the ECS. For details, see [Viewing ECS Details \(List View\)](#).

- You have configured the inbound rules of the security group. For details, see [Configuring Security Group Rules](#).
- The network connection between the login tool (PuTTY) and the target ECS is normal. For example, the default port 22 is not blocked by the firewall.

Logging In to a Linux ECS from a Local Windows Server

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

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

The following 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>

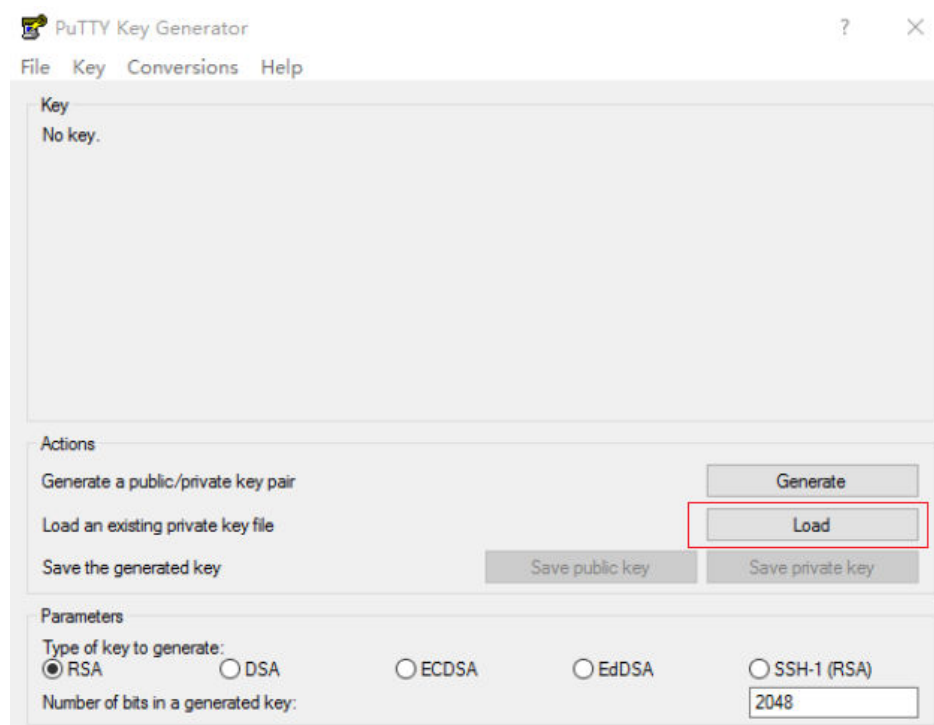
NOTE

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 2-29 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 2-30 Configuring the EIP



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

Figure 2-31 Entering the username



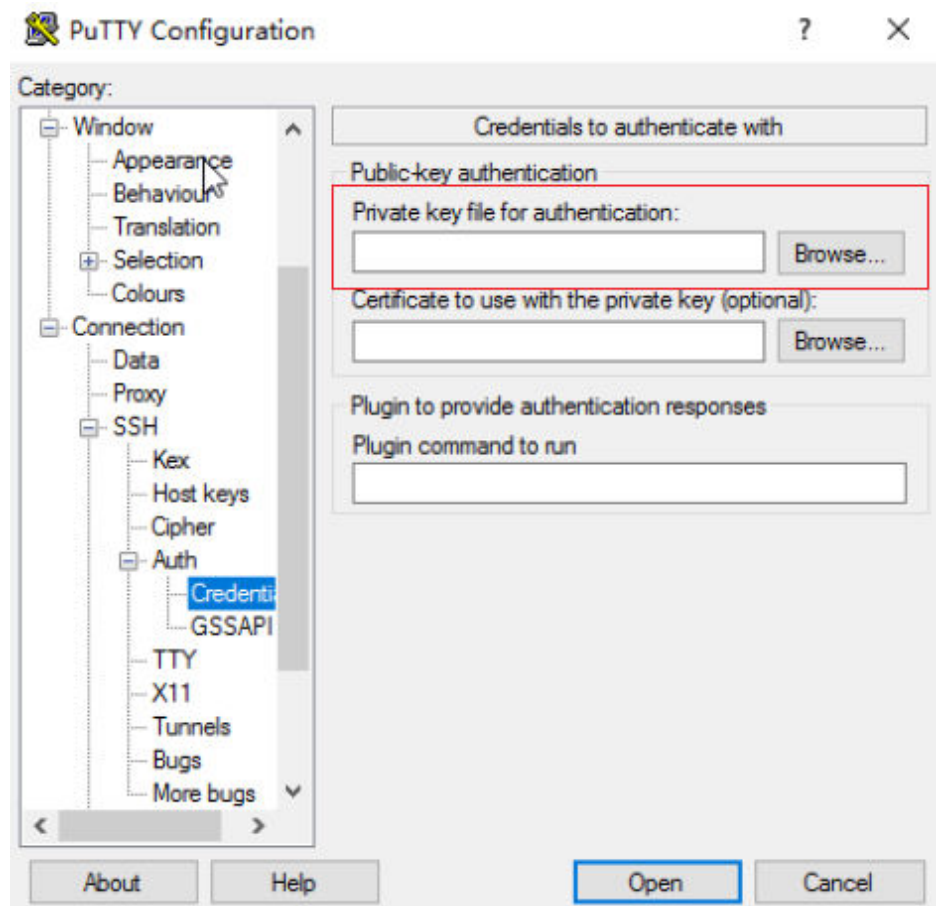
NOTE

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 2-32 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 2-33 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?](#)

2.4.4 Logging In to a Linux ECS Using an SSH Password

Scenarios

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

Prerequisites

- The target ECS is running.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).
- The network connection between the login tool (PuTTY) and the target ECS is normal. For example, the default port 22 is not blocked by the firewall.


Logging In to a Linux ECS from a Local Windows Server

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

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

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

Figure 2-34 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

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.  
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?](#)

2.4.5 Logging In to a Linux ECS from a macOS Server

Scenario

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

Prerequisites

- The target ECS is running.
- 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

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

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

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

 NOTE

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

Follow-up Procedure

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

2.4.6 Logging In to a Linux ECS from a Mobile Terminal

Scenarios

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

- For instructions about how to log in to a Linux ECS from an iOS terminal through Termius, see [Logging In to a Linux ECS from an iOS Terminal](#).
- For instructions about how to log in to a Linux ECS from an Android terminal through JuiceSSH, see [Logging In to a Linux ECS from an Android Terminal](#).

Prerequisites


- The target ECS is running.
- You have bound an EIP to the ECS. For details, see [Binding an EIP](#).
- Access to port 22 is allowed in the inbound direction of the security group which the ECS belongs to. For details, see [Configuring Security Group Rules](#).

Logging In to a Linux ECS from an iOS Terminal

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

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

Figure 2-35 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 2-36 Setting parameters

Cancel New Host Save

1 Alias

2 Hostname

Group >

Tags >

Backspace as CTRL+H

SSH / MOSH

3 Use SSH

Use Mosh (Beta)

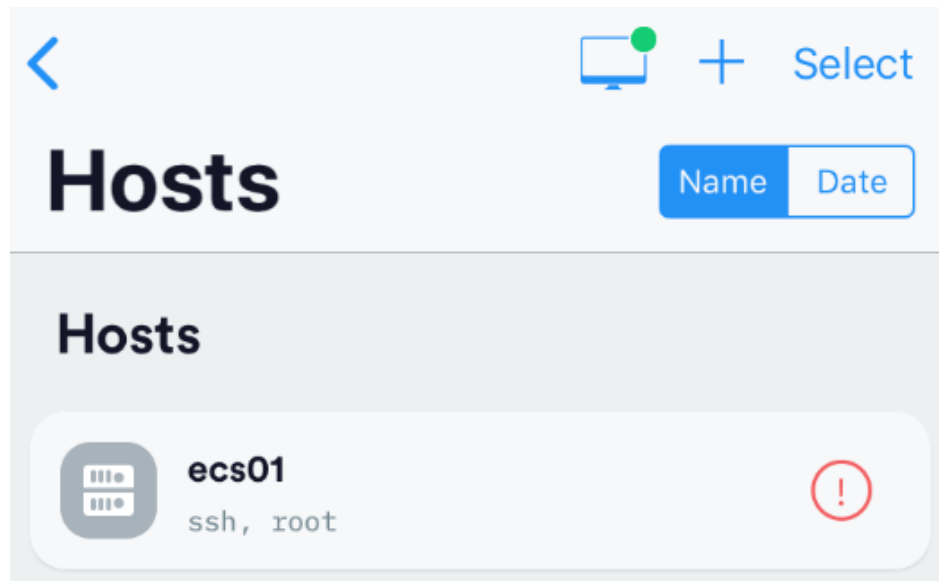
4 Port 22
Default

5 Username root

6 Password ●●●●●●●●

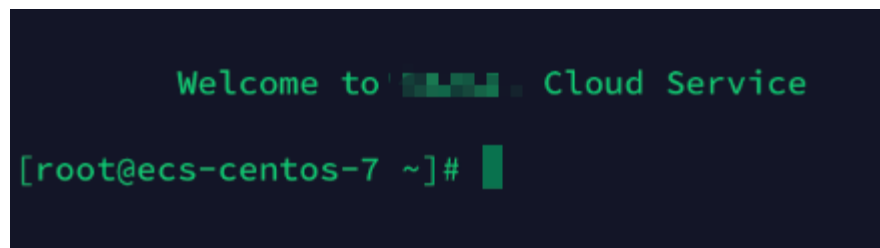
3. Tap **Save** in the upper right corner of the page to save the login settings. On the **Hosts** page, tap the name of the connection.

Figure 2-37 Login information



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

Figure 2-38 Connected

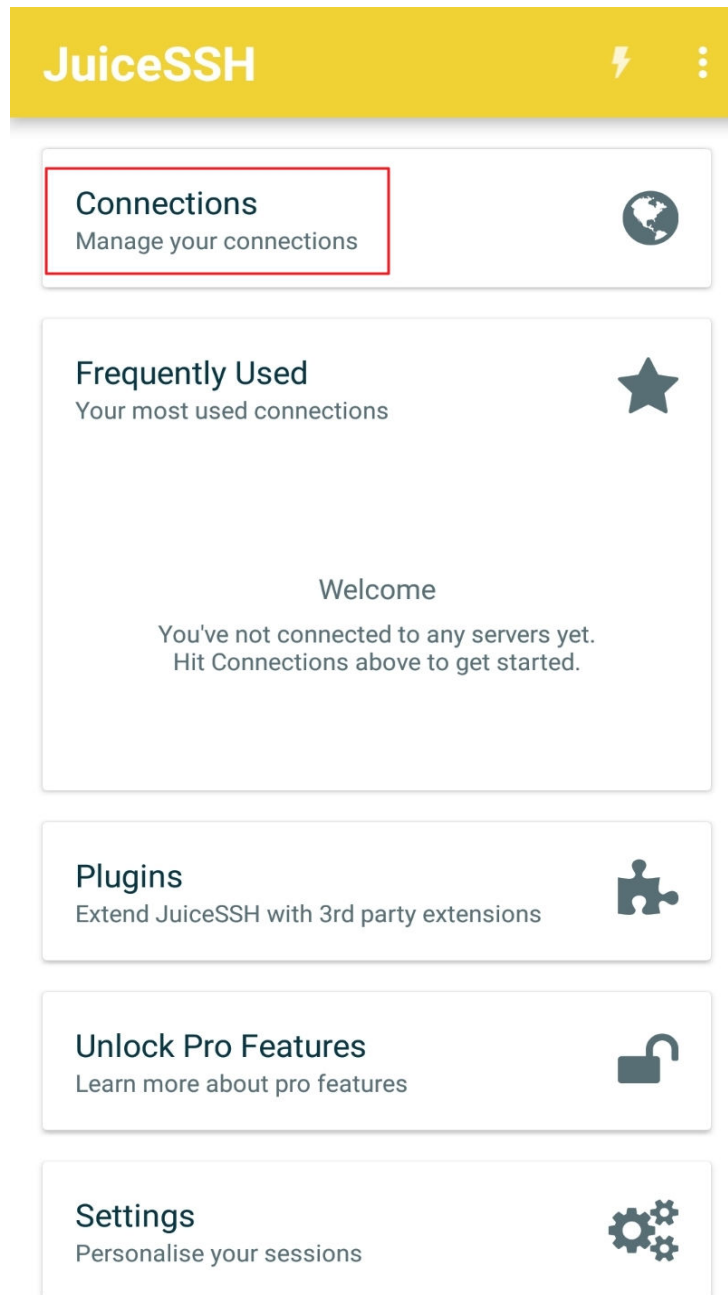


Logging In to a Linux ECS from an Android Terminal

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 2-39 Starting JuiceSSH




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

Figure 2-40 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 2-41 New Identity

← **New Identity** ✓

IDENTITY

Nickname: linux_test

Username: root

Password: **SET (OPTIONAL)**

Private Key: **SET (OPTIONAL)**

SNIPPET

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

GENERATE SNIPPET

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

Figure 2-42 Port

← **New Connection** ✓

BASIC SETTINGS

Nickname:

Type: ▼

Address:

Identity: ▼

ADVANCED SETTINGS

Port:

Connect Via: ▼

Run Snippet: ▼

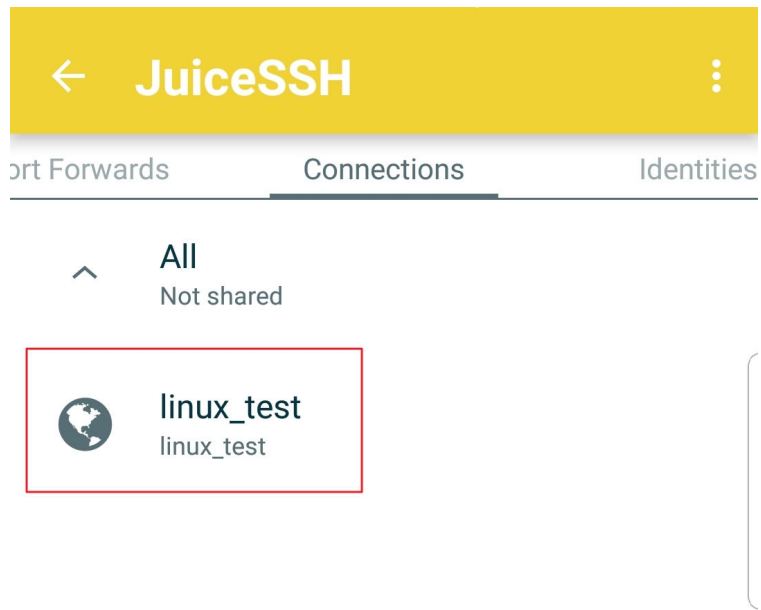
Backspace: ▼

GROUPS

ADD TO GROUP


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

Figure 2-43 Connections



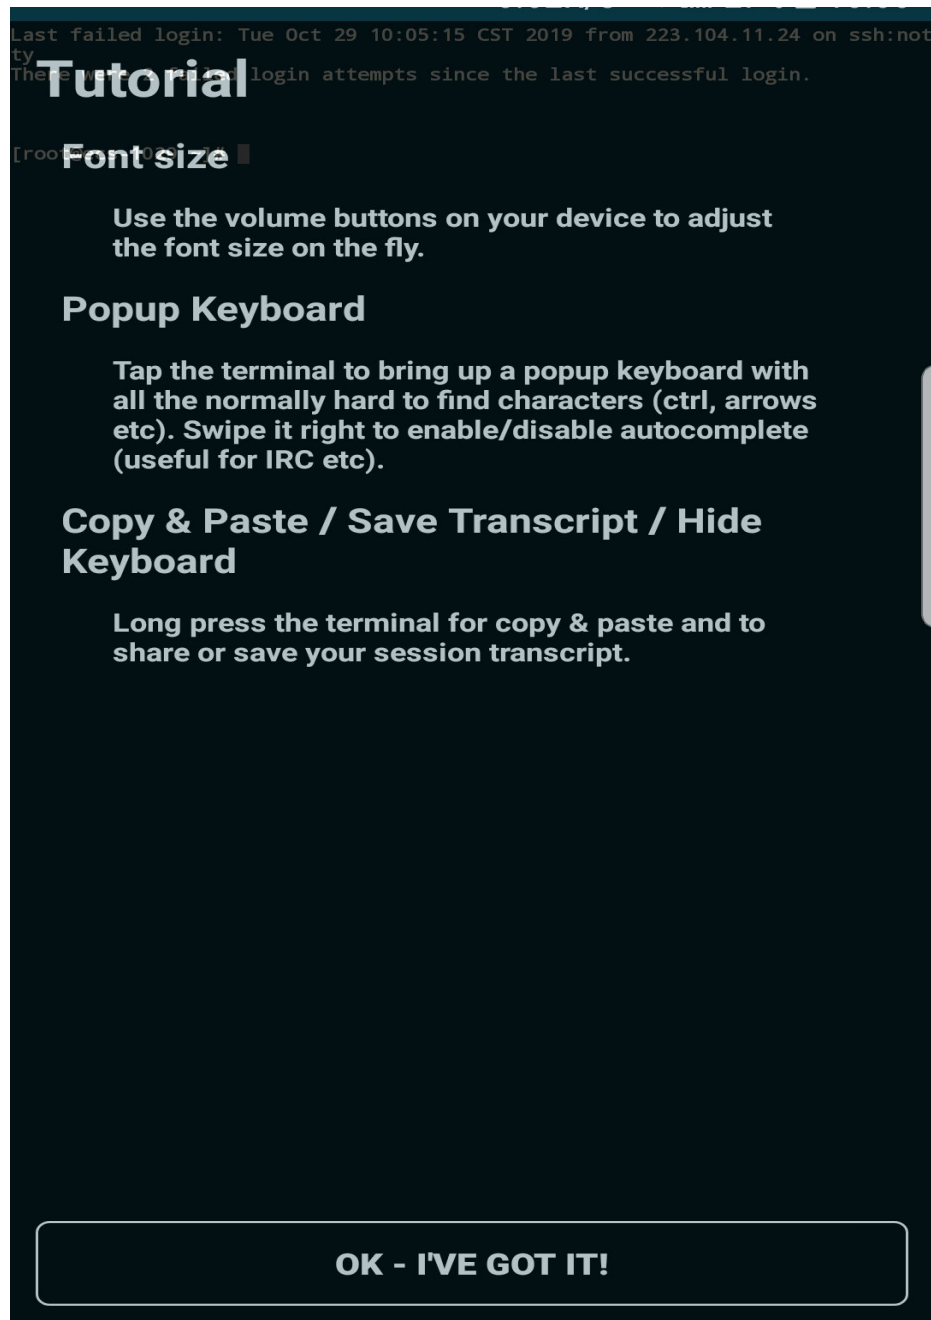
5. Confirm the information that is displayed and tap **ACCEPT**.

Figure 2-44 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 2-45 Tutorial



You have logged in to the Linux ECS.

Figure 2-46 Successful login



2.5 Managing GPU Drivers of GPU-accelerated ECSs

2.5.1 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 2-12 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.

2.5.2 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 2-13 Mapping between Tesla drivers and ECS types

ECS Type	Driver	Product Series	Product
Pi2	Tesla	T	T4

Downloading a CUDA Toolkit

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

NOTE

[NVIDIA Driver Downloads](#) provides the mapping between the driver version and CUDA Toolkit. If the versions do not match, the driver may be unavailable.

The following uses Tesla T4 as an example to describe how to download the driver package and CUDA Toolkit.

1. Select the Linux operating system and the CUDA Toolkit 11.6 version.

Figure 2-47 Selecting the CUDA Toolkit version

NVIDIA Driver Downloads

Select from the dropdown list below to identify the appropriate driver for your NVIDIA product. [Help](#)

Product Type: Data Center / Tesla

Product Series: T-Series

Product: Tesla T4

Operating System: Linux 64-bit

CUDA Toolkit: 11.6

Language: English (US)

Search

2. Select a CUDA Toolkit 11.6 package to download.

Figure 2-48 Downloading a CUDA Toolkit 11.6 package

Archived Releases

[CUDA Toolkit 11.7.1 \(August 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.7.0 \(May 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.6.2 \(March 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.6.1 \(February 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.6.0 \(January 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.5.2 \(February 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.5.1 \(November 2021\), Versioned Online Documentation](#)
[CUDA Toolkit 11.5.0 \(October 2021\), Versioned Online Documentation](#)
[CUDA Toolkit 11.4.4 \(February 2022\), Versioned Online Documentation](#)
[CUDA Toolkit 11.4.3 \(November 2021\), Versioned Online Documentation](#)
[CUDA Toolkit 11.4.2 \(September 2021\), Versioned Online Documentation](#)
[CUDA Toolkit 11.4.1 \(August 2021\), Versioned Online Documentation](#)
[CUDA Toolkit 11.4.0 \(June 2021\), Versioned Online Documentation](#)

2.5.3 Manually Installing a GRID Driver on a GPU-accelerated ECS

Scenarios

To use graphics acceleration, such as OpenGL, DirectX, or Vulkan, install a GRID driver and separately purchase and configure a GRID license. The GRID driver with a vDWS license also supports CUDA for both computing and graphics acceleration.

- A graphics-accelerated (G series) ECS created using a public image has had a GRID driver of a specified version installed by default, but the GRID license must be purchased and configured separately.
- If a GPU-accelerated ECS is created using a private image, install a GRID driver and separately purchase and configure a GRID license.

This section describes how to install a GRID driver, purchase or apply for a GRID license, and configure the license server.

Process of installing a GRID driver:

1. [Purchasing a GRID License](#)
2. [Downloading GRID Driver and Software License Packages](#)
3. [Deploying and Configuring the License Server](#)

 NOTE

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

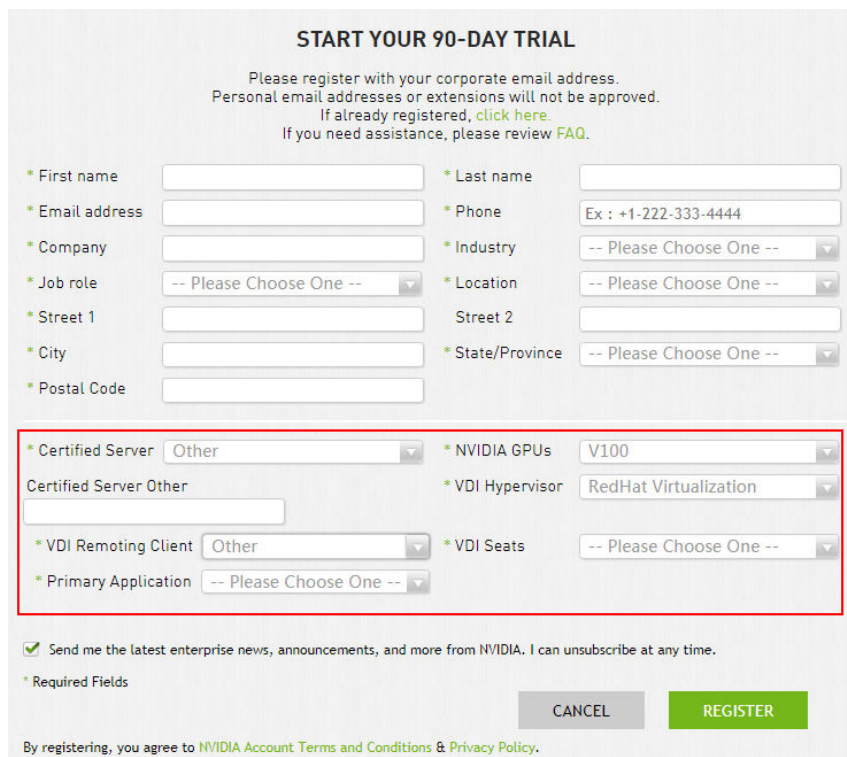
Purchasing a GRID License

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

 NOTE

The method of using a trial license is the same as that of using an official license. You can use an official license to activate an account with a trial license to prevent repetitive registration. The trial license has a validity period of 90 days. After the trial license expires, it cannot be used anymore. Purchase an official license then.

Figure 2-49 Applying for a trial license



START YOUR 90-DAY TRIAL

Please register with your corporate email address.
Personal email addresses or extensions will not be approved.
If already registered, [click here](#).
If you need assistance, please review [FAQ](#).

* First name * Last name
* Email address * Phone
* Company * Industry
* Job role * Location
* Street 1 Street 2
* City * State/Province
* Postal Code

* Certified Server * NVIDIA GPUs
Certified Server Other * VDI Hypervisor
* VDI Remoting Client * VDI Seats
* Primary Application

Send me the latest enterprise news, announcements, and more from NVIDIA. I can unsubscribe at any time.

* Required Fields

By registering, you agree to [NVIDIA Account Terms and Conditions & Privacy Policy](#).

Downloading GRID Driver and Software License Packages

1. Obtain the driver installation package required for an OS. For details, see [Table 2-14](#).

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.

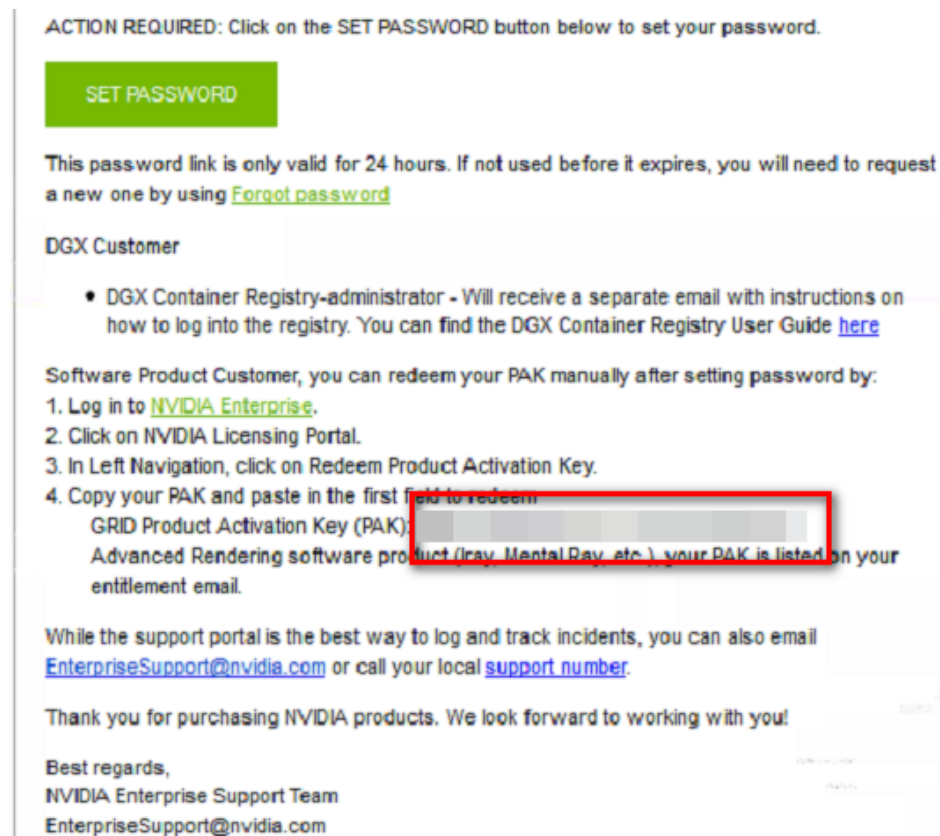
For a GPU virtualization ECS, select a driver version based on the following table.

Table 2-14 GRID driver versions supported by GPU-accelerated ECSs

ECS Type	GPU Attachment	OS	Driver Version	CPU Architecture
PI2	GPU passthrough	<ul style="list-style-type: none">• CentOS 8.2 64bit• CentOS 7.6 64bit• Ubuntu 20.04 Server 64bit• Ubuntu 18.04 Server 64bit	Select a version as needed.	x86_64


2. After the registration, log in at the [official NVIDIA website](#) and enter the account.
3. Check whether NVIDIA is used for the first time.
 - a. If yes, go to step [4](#).
 - b. If no, go to step [6](#).
4. Refer to [Figure 2-50](#) to obtain the Product Activation Key (PAK) from the email indicating successful registration with NVIDIA.

Figure 2-50 PAK



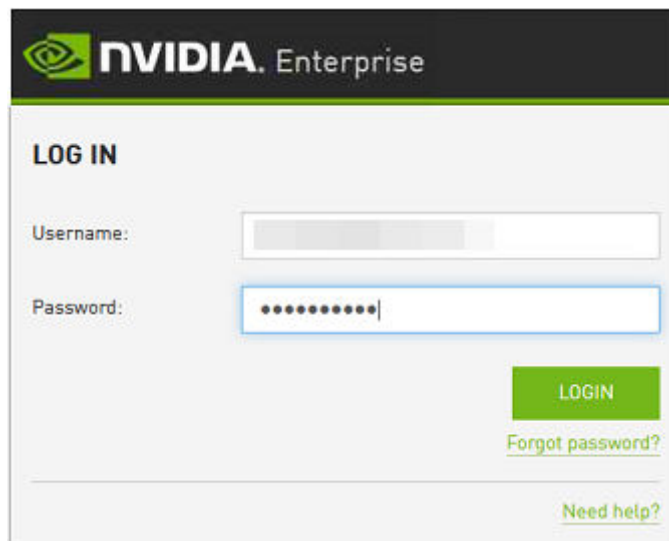
5. Enter the PAK obtained in step 4 on the **Redeem Product Activation Keys** page and click **Redeem**.

Figure 2-51 Redeem Product Activation Keys



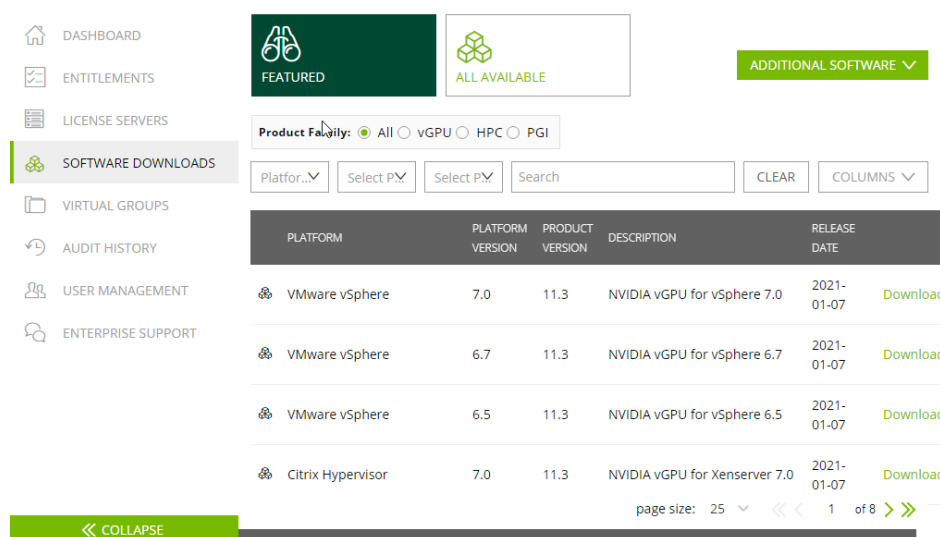
6. Specify **Username** and **Password** and click **LOGIN**.

Figure 2-52 Logging in to the official NVIDIA website




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

Figure 2-53 SOFTWARE DOWNLOADS page



8. Download the GRID driver of the required version. For details, see [Table 2-14](#).
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 2-54 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 [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
-----

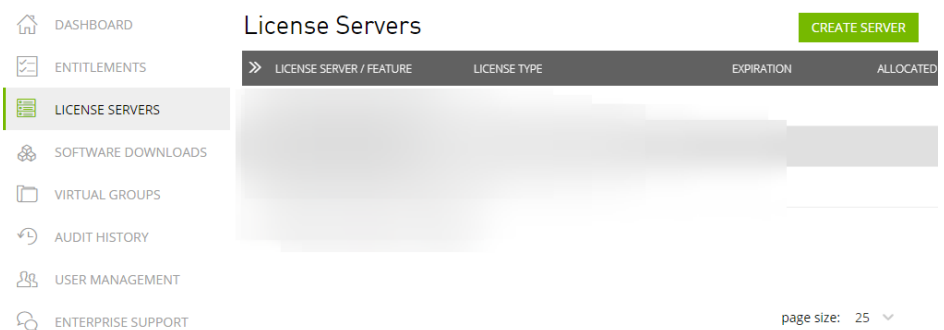
License Server has been successfully installed to:

  /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 2-55 LICENSE SERVERS



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


Figure 2-56 Create License Server

Table 2-15 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 2-57 Downloading the license file



4. In the web browser, access the homepage of the license server management page using the link configured during the installation.
Default URL: `http://IP address of the EIP:8080/licserver`
5. In the navigation pane on the left, click **License Server > License Management**.
6. Select the .bin license file to be uploaded and click **Upload**.

Figure 2-58 Uploading a license file



2.5.4 Manually Installing a Tesla Driver on a GPU-accelerated ECS

Scenarios

Before using a GPU-accelerated ECS, make sure that the desired Tesla driver and CUDA toolkit have been installed on the ECS for computing acceleration.

- A computing-accelerated (P series) ECS created using a public image has had a Tesla driver of a specified version installed by default.
- After a GPU-accelerated ECS is created using a private image, it must have a Tesla driver installed. Otherwise, computing acceleration will not take effect.

This section describes how to install a Tesla driver and CUDA toolkit on a GPU-accelerated ECS.

Notes

- The ECS must have an EIP bound.
- Check whether the CUDA toolkit and Tesla driver have been installed on the ECS.

NOTE

- If the CUDA toolkit has not been installed, download it from the official NVIDIA website and install it. A Tesla driver matching the CUDA version will be automatically installed then. However, if there are specific requirements or dependencies on the Tesla driver version, download the matching Tesla driver from the official NVIDIA website first and then install the driver before installing the CUDA toolkit.
- If a Tesla driver has been installed on the ECS, check the driver version. Before installing a new driver version, uninstall the original Tesla driver to prevent an installation failure due to driver conflicts.

Installation process:

- [Obtaining a Tesla Driver and CUDA Toolkit](#)
- Installing a Tesla Driver

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

Installing a Tesla Driver on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the Tesla driver matching CUDA 10.1 on a GPU-accelerated ECS.

NOTE

The Linux kernel version is compatible with the driver version. If installing the driver failed, check the driver installation log, which is generally stored in `/var/log/nvidia-installer.log`. If the log shows that the failure was caused by a driver compilation error, for example, the `get_user_pages` parameter setting is incorrect, the kernel version is incompatible with the driver version. In such a case, select the desired kernel version and driver version and reinstall them. It is recommended that the release time of the kernel version and driver version be the same.

1. Log in to the ECS.
2. Update the system software based on the OS.
 - Ubuntu
Update the software installation source: `apt-get -y update`
Install necessary programs: `apt-get install gcc g++ make`
 - CentOS
Update the software installation source: `yum -y update --exclude=kernel* --exclude=centos-release* --exclude=initscripts*`
Install the desired program: `yum install -y kernel-devel-`uname -r` gcc gcc-c++`

3. Download the NVIDIA driver package.

Select a driver version at [NVIDIA Driver Downloads](#) based on the ECS type. Click **SEARCH**.

Figure 2-59 Selecting a NVIDIA driver version

NVIDIA Driver Downloads


Advanced Driver Search

Product Type:	Operating System:
<input type="text" value="Tesla"/>	<input type="text" value="Linux 64-bit"/>
Product Series:	CUDA Toolkit:
<input type="text"/>	<input type="text" value="10.1"/>
Product:	Language:
<input type="text"/>	<input type="text" value="English (US)"/>
	Recommended/Beta:
	<input type="text" value="All"/> ?

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

Figure 2-60 Selecting a driver version

NVIDIA Driver Downloads



- Click the driver to be downloaded. On the **TESLA DRIVER FOR LINUX X64** page that is displayed, click **DOWNLOAD**.
- Copy the download link.

Figure 2-61 Copying the download link

Download


By clicking the "Agree & Download" button below, you are confirming that you have read and agree to be bound by the [License For Customer Use of NVIDIA Software](#) for use of the driver. The driver will begin downloading immediately after clicking on the "Agree & Download" button below. NVIDIA recommends users update to the latest driver version. Please review [NVIDIA Product Security](#) for more information.



- Run the following command on the ECS to download the driver:
wget Copied link

For example, **wget http://us.download.nvidia.com/tesla/418.67/NVIDIA-Linux-x86_64-418.67.run**

Figure 2-62 Obtaining the installation package




- Run the following command to install the driver:

```
sh NVIDIA-Linux-x86_64-418.67.run
```


9. (Optional) If the following information is displayed after the command for installing the driver is executed, disable the Nouveau driver.

Figure 2-63 Disabling the Nouveau driver



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

```
blacklist nouveau
options nouveau modeset=0
```
 - c. Run the following command to back up and create an initramfs application:
 - Ubuntu
sudo update-initramfs -u
 - CentOS:
mv /boot/initramfs-\$(uname -r).img /boot/initramfs-\$(uname -r).img.bak
dracut -v /boot/initramfs-\$(uname -r).img \$(uname -r)
 - d. Restart the ECS:
reboot
10. Select **OK** for three consecutive times as prompted to complete the driver installation.

Figure 2-64 Completing the NVIDIA driver installation

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

Figure 2-65 Viewing the NVIDIA driver version

```
root@ecs-474b:~# nvidia-smi
Thu Mar 26 20:05:17 2020

+-----+
| NVIDIA-SMI 418.67      Driver Version: 418.67      CUDA Version: 10.1   |
+-----+
| GPU   Name           Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf    Pwr:Usage/Cap|  Memory-Usage | GPU-Util  Compute M. |
+-----+-----+
|    0   Tesla          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 the CUDA Toolkit on a Linux ECS

The following uses Ubuntu 16.04 64bit as an example to describe how to install the CUDA 10.1 toolkit on a GPU-accelerated ECS.

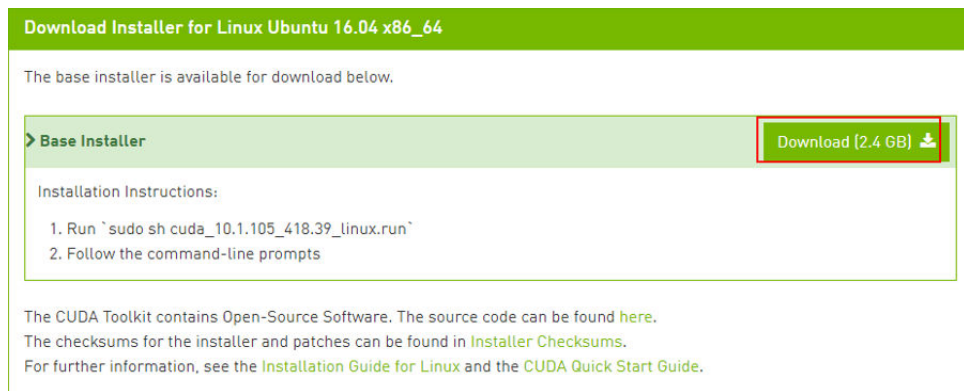
1. Log in to the ECS.
2. Update the system software based on the OS.
 - Ubuntu
Update the software installation source: **apt-get -y update**
Install necessary programs: **apt-get install gcc g++ make**
 - CentOS
Update the software installation source: **yum -y update --exclude=kernel* --exclude=centos-release* --exclude=initscripts***
Install the desired program: **yum install -y kernel-devel`uname -r` gcc gcc-c++**
3. On the CUDA download page, set parameters according to the information shown in [Obtaining a Tesla Driver and CUDA Toolkit](#).

Figure 2-66 Selecting a CUDA version



- 4. Find the link for downloading CUDA 10.1 and copy the link.

Figure 2-67 Copying the link for downloading CUDA




- 5. Run the following command on the ECS to download CUDA:

wget Copied link


For example, **wget https://developer.nvidia.com/compute/cuda/10.1/Prod/local_installers/cuda_10.1.105_418.39_linux.run**

Figure 2-68 Downloading CUDA



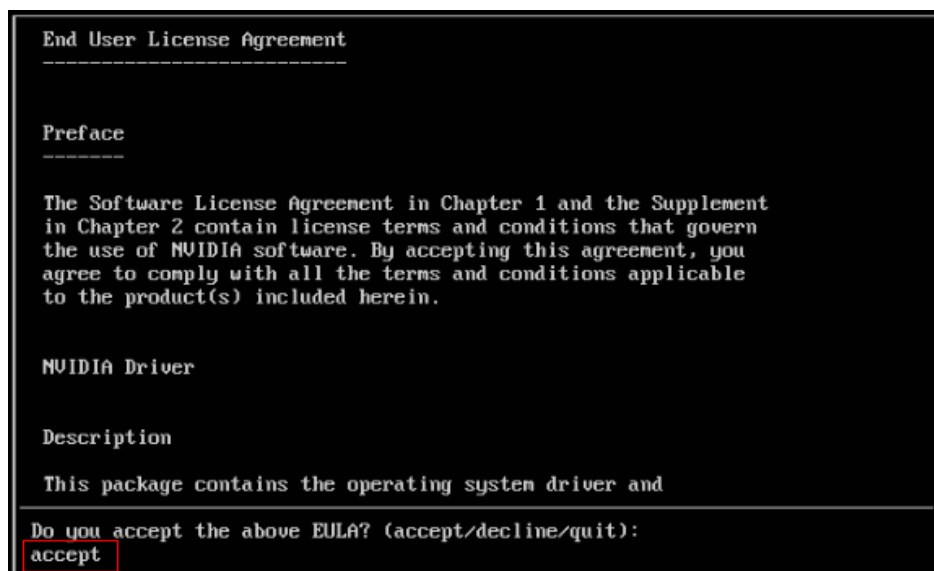
- 6. Install CUDA.
Follow the instructions provided on the official NVIDIA website.

Figure 2-69 Installing CUDA



7. Run the following command to install CUDA:
sh cuda_10.1.243_418.87.00_linux.run
8. Select **accept** on the installation page and press **Enter**.

Figure 2-70 Installing CUDA_1



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

Figure 2-71 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 2-72 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. 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
```

11. Run the **make** command to automatically compile the deviceQuery program.
12. Run the following command to check whether CUDA has been installed:
./deviceQuery

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

Figure 2-73 deviceQuery common output

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

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

Detected 1 CUDA Capable device(s)

Device 0: "Tesla "
  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_Uutilities/deviceQuery#
```

13. Check the CUDA version.

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

Figure 2-74 Checking the CUDA version

```
root@ecs-474b:/usr/local/cuda-10.1/samples/1_Uutilities/deviceQuery# ./deviceQuery
[root@ecs-474b deviceQuery]# /usr/local/cuda/bin/nvcc -V
nvcc: NVIDIA (R) Cuda compiler driver
Copyright (c) 2005-2019 NVIDIA Corporation
Built on Fri Feb  8 19:08:17 PST 2019
Cuda compilation tools, release 10.1, V10.1.105
[root@ecs-474b deviceQuery]#
```

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

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

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

2.6 Managing ECS Configurations

2.6.1 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.
4. Set the target time zone.
 - a. Run the following command to open the **/etc/sysconfig/clock** file:
vim /etc/sysconfig/clock
 - b. Locate the **ZONE** entry and change its value to the name of the desired time zone file.
5. Press **Esc**. Then, run the following command to save and exit the **/etc/sysconfig/clock** file:
:wq
6. Run the following command to check whether the **/etc/localtime** file is available on the ECS:
ls /etc/localtime
 - If the file is available, go to step 7.
 - If the file is not available, go to step 8.
7. Run the following command to delete the existing **/etc/localtime** file:
rm /etc/localtime
8. Run the following command to create a symbolic link between **/etc/localtime** and your time zone file so that the ECS can find this time zone file when it references the local time:
ln -sf /usr/share/zoneinfo/Asia/city1 /etc/localtime
9. Run the following command to restart the ECS so that all services and applications running on the ECS use the new time zone:
reboot
10. Log in to the ECS again and run the following command as user **root** to check whether the time zone has been changed:
ls -lh /etc/localtime

The following information is displayed:

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

2.6.2 Enabling or Disabling Hyper-Threading

Scenarios

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

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

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

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

Background

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

A flavor defines the number of vCPUs. You can query the number of vCPUs that an x86 ECS has by referring to [Querying the Number of vCPUs of an ECS](#).

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

Notes and Constraints

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

Enabling or Disabling Hyper-Threading (During ECS Purchase)

1. Log in to the management console and access the [Buy ECS](#) page. Configure basic, network, and advanced settings for ECSs based on service requirements. For details, see [Purchasing an ECS](#).



2. Select **Configure now** to configure advanced options.
3. Select **Specify CPU options**.
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.
6. Click **Next**.
7. Confirm the settings, read and select the agreement, and then click **Submit Application**.

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

Querying the Number of vCPUs of an ECS

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

- For Linux ECSs:
 - a. [Log in to a Linux ECS](#).
 - b. Run the following command to view the number of logical cores of the ECS:

```
lscpu
```


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

Figure 2-75 Viewing the result

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

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

Figure 2-76 Viewing the result



2.6.3 Obtaining Metadata and Passing User Data

2.6.3.1 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 2-16](#). The following describes the URI and methods of using the supported ECS metadata.

Notes

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

Perform the following configuration on the firewall:

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

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

ECS Metadata Types

Table 2-16 does not contain the following metadata items: ami-id, ami-launch-index, ami-manifest-path, block-device-mapping/, instance-action, instance-id, reservation-id, ramdisk-id, and kernel-id. These metadata items are meaningless and are not recommended.

Table 2-16 ECS metadata types

Metadata Type	Metadata Item	Description
OpenStack	/meta_data.json	Displays ECS metadata. For the key fields in the ECS metadata, see Table 2-17 .
OpenStack	/password	Displays the password for logging in to an ECS.
OpenStack	/user_data	Displays ECS user data. This metadata allows you to specify scripts and configuration files for initializing ECSs. For details, see Passing User Data . For password-authenticated Linux ECSs, this metadata is used to save password injection scripts.
OpenStack	/network_data.json	Displays ECS network information.
OpenStack	/securitykey	Obtains temporary AKs and SKs. Before enabling an ECS to obtain a temporary AK and SK, authorize agency permissions to the op_svc_ecs account and ECSs in IAM. NOTE You can determine what permissions are granted to the agency based on the principal of least privilege (PoLP). ECSs will not use agencies to perform operations on resources.
OpenStack	/spot/instance-action	Queries the prompt of stopping a spot ECS.

Metadata Type	Metadata Item	Description
EC2-compatible	/meta-data/ hostname	Displays the name of the host accommodating an ECS. To remove the suffix .novalocal from an ECS, see: Is an ECS Hostname with Suffix .novalocal Normal?
EC2-compatible	/meta-data/ local-hostname	The meaning of this field is the same as that of hostname.
EC2-compatible	/meta-data/ public-hostname	The meaning of this field is the same as that of hostname.
EC2-compatible	/meta-data/ instance-type	Displays an ECS flavor.
EC2-compatible	/meta-data/ local-ipv4	Displays the fixed IP address of an ECS. If there are multiple NICs, only the IP address of the primary NIC is displayed.
EC2-compatible	/meta-data/ placement/ availability-zone	Displays the AZ accommodating an ECS.
EC2-compatible	/meta-data/ public-ipv4	Displays the EIP bound to the ECS. If there are multiple NICs, only the EIP of the primary NIC is displayed.
EC2-compatible	/meta-data/ public-keys/0/ openssh-key	Displays the public key of an ECS.
EC2-compatible	/user-data	Displays ECS user data.
EC2-compatible	/meta-data/ security-groups	Displays the security group of an ECS.

Table 2-17 Metadata key fields

Parameter	Type	Description
uuid	String	Specifies an ECS ID.
availability_zone	String	Specifies the AZ where an ECS locates.
meta	Dict	Specifies the metadata information, including the image name, image ID, and VPC ID.

Parameter	Type	Description
hostname	String	Specifies the name of the host accommodating an ECS. To remove the suffix .novalocal from an ECS, see: Is an ECS Hostname with Suffix .novalocal Normal?
enterprise_project_id	String	Specifies the ID of the enterprise project accommodating an ECS.

Prerequisites

- The target ECS has been logged in.
- Security group rules in the outbound direction meet the following requirements:
 - **Protocol: TCP**
 - **Port: 80**
 - **Destination: 169.254.0.0/16**

NOTE

If you use the default security group rules for the outbound direction, the metadata can be accessed because the default rules meet the preceding requirements. For details about the default security group rules for the outbound direction, see [Default Security Groups and Rules](#).

Metadata (OpenStack Metadata API)

This API is used to query ECS metadata.

- URI
`/169.254.169.254/openstack/latest/meta_data.json`
- Usage method
Supports GET requests.

- Example

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

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

```
{
  "random_seed": "rEocCViRS+dNwlydGlxJHUp+00poeUsAdBFkbPbYQTmpNwpoEb43k9z+96TyrekNKS
+iLYDdRny4kKGoNPEVBCC05Hg1TcDbIAPfJwgJS1okqEtlcofUhKmL3K0fto
+5KXEDU3GNuGwyZXjdVb9HQWU+E1jztAjjsahnU+g/tawABTVySLBKlAT8fMGax1mTGGArucn/
WzDcy19DGioKPE7F8lLtSQ4Ww3VClK5VYB/h0x+4r7IVHrPmYX/
bi1Yhm3Dc4rRYNaTjdOV5gUOsbo3oAeQkmKwQ/
NO0N8qw5Ya4l8ZUW4tMav4mOsRySOOB35v0bvaJc6p
+50DTbWNeX5A2MLiEhTP3vsPrmvk4LRF7CLz2J2TGIM14OoVBw7LARwmv9cz532zHki/c8tlhRzLmOTXh/
wL36zFW10DeuReUGmXth7IGNmRMQKV6+mil78jm/KMPpgAdK3vwYF/
GcelOFJD2HghMUUCeMbwYnvijLTejuBpwhJMNIHA/NvlEsxDxqBCoss/Jfe+yCmUFyxovJ
+L8oNkTzkmCNzw3Ra0hiKchGhqK3BleToV/kVx5DdF081xrEA
+qyoM6CVyftEoz1zIRRyoo9bJ65Eg6Jd8dj1UCVsDqRY1pljgzE/
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

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLkI0IGZlZWxzIGFuIGltcHVsc2lvbi4uLnRoaXMgaXMgdGhllH
BsYWNlIHVlGdviG5vdy4gQnV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IHV0IH
cyBiZWphbmQgYW5kIG5sb3VkcycwYm5kIGhvdSB3aWxsIGtub3csIHV0IHV0IHV0IHV0IHV0IHV0IHV0IH
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYm5kIGhvcml6b25kLlINCg0KLlVpY2hhcmQgQmFjaA=
```

NOTE

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

Figure 2-77 404 Not Found

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

Network Data (OpenStack Metadata API)

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

- URI
/openstack/latest/network_data.json
- Usage method
Supports GET requests.
- Example

NOTE

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

Linux:

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

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

Security Key (OpenStack Metadata API)

This API is used to obtain temporary AKs and SKs.

 NOTE

- If an ECS needs to obtain a temporary AK and SK, go to the ECS details page, and configure **Agency** for the ECS in the **Management Information** area so that the ECS is authorized on IAM.
For details, see [Cloud Service Delegation](#).
- The validity period of a temporary AK and SK is one hour. The temporary AK and SK are updated 10 minutes ahead of the expiration time. During the 10 minutes, both the new and old temporary AKs and SKs can be used.
- When using temporary AKs and SKs, add '**X-Security-Token**':{**securitytoken**} in the message header. **securitytoken** is the value returned when a call is made to the API.
- URI
`/openstack/latest/securitykey`
- Usage method
Supports GET requests.
- Examples
Linux:
curl http://169.254.169.254/openstack/latest/securitykey

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
```

```
ICAgICAgDQoiQSBjbG91ZCBkb2VzIG5vdCBrbm93IHdoeSBpdCBtb3ZlcyBpbjBqdXN0IHN1Y2ggYSBkaXJlY  
3Rpb24gYW5kIGF0IHN1Y2ggYSBzcGVlZC4uLk0IGZlZWxzIGFuIGltcHVsc2lubi4uLnRoaXMgaXMgdGhllH  
BsYWNlIHhrvIGdvlG5vdy4gQnV0IHRobzSBza3kga25vd3MgdGhllHJlYXNvbnMgYW5kIHRobzSBwYXR0ZXJl  
cyBiZWwhpbmQgYWxslGNsb3VkcycgYW5kIHlvdSB3aWxslGtub3csIHRobybygd2hlbiB5b3UgbGlmdCB5b3  
Vyc2VsZiBoaWdoIGVub3VnaCB0byBzZWUgYmV5b25kIGhvcml6b25zLiINCg0KLVJpY2hhcmQgQmFjaA=  
=
```

Hostname (EC2 Compatible API)

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

- URI
/169.254.169.254/latest/meta-data/hostname
- Usage method
Supports GET requests.
- Example

Linux:

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

```
vm-test.novalocal
```

Instance Type (EC2 Compatible API)

This API is used to query an ECS flavor.

- URI
/169.254.169.254/latest/meta-data/instance-type
- Usage method
Supports GET requests.
- Example

Linux:

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

```
s3.medium.2
```

Local IPv4 (EC2 Compatible API)

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

- URI
/169.254.169.254/latest/meta-data/local-ipv4
- Usage method
Supports GET requests.
- Example

Linux:

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

192.1.1.2

Availability Zone (EC2 Compatible API)

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

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

Public IPv4 (EC2 Compatible API)

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

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

Public Keys (EC2 Compatible API)

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

- URI
`/169.254.169.254/latest/meta-data/public-keys/0/openssh-key`
- Usage method
Supports GET requests.
- Example
Linux:
curl http://169.254.169.254/latest/meta-data/public-keys/0/openssh-key
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQADIA5Fw5k8Fgzajn1zJwLoV3+wMP+6CyvsSilc/
hioggSnYu/AD0Yqm8vVO0kWlun1rFbdO+QUZKyVr/OPUjQSw4SRh4qsTKf/+eFoWTjplFvd1WCBZzS/
WRenxlwR00KkczHSJro763+wYcwKieb4eKRxaQoQvoFgVjLBULXajH4eKoKTVNtMXAvPP9aMy2SLgsJNt
Mb9ArfziAiblQynq7UifLnN3VclzPeiWrqtzjyOp6CPUXnL0lVPTvble8sUteBsJZwLL6K4i
+Y0lf3ryqnmQgC21yW4Dzu+kwk8FVT2MgWkCwiZd8gQ/+uJzrJFyMfUOBiklOBfuJENIJUuhAB
Generated-by-Nova

Helpful Links

[Why Can't My Linux ECS Obtain Metadata?](#)

2.6.3.2 Passing User Data

Scenarios

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

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

Use Restrictions

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

Passing User Data

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

NOTE

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

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

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

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

User Data Scripts of Linux ECSs

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

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

 **NOTE**

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

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

Table 2-18 Linux ECS script types

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

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

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

- Script usage examples

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

Example 1: Inject a user-data script.

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

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

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

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

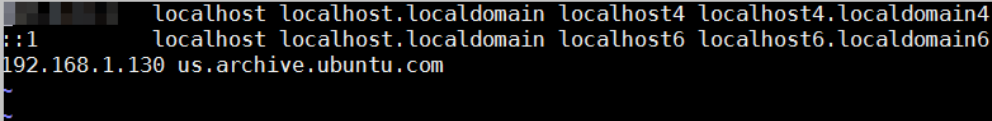
Example 2: Inject a Cloud-Config data script.

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

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

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

Figure 2-78 Viewing operating results



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

Case 1

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

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

User data example:

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

Case 2

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

NOTE

The new password must meet the password complexity requirements listed in [Table 2-19](#).

Table 2-19 Password complexity requirements

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

User data example:

Using a ciphertext password (recommended)

```
#!/bin/bash  
echo 'root:$6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlgi' | chpasswd -e;
```

In the preceding command output, **\$6\$V6azyeLwcD3CHlpY\$BN3VVq18fmCkj66B4zdHLWevqcxlgi** is the ciphertext password, which can be generated as follows:

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

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

The following information is displayed:

```
$6$V6azyeLwcD3CHlpY
```

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

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

The following information is displayed:

```
$6$V6azyeLwcD3CHlpY$BN3VVq18fmCkj66B4zdHLWevqcxlgi
```

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

Case 3

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

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

NOTE

The new password must meet the password complexity requirements listed in [Table 2-20](#).

Table 2-20 Password complexity requirements

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

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

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

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

Case 5

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

User data example:

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

Case 6

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

User data example:

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


Helpful Links

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

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




2.6.4 Changing ECS Names

Scenarios

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

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



Changing the Name of a Single ECS

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Click the name of the target ECS.
5. On the ECS details page, click  next to the ECS name and change 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 **OK**.

Changing the Names of Multiple ECSs in a Batch

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Select the target ECSs.
5. Click **More** above the ECS list and select **Change ECS Name** from the drop-down list.
6. Enter a new name.
7. Click **OK**.

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

2.6.5 Managing ECS Groups

Scenarios

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

Currently, only the anti-affinity policy is supported.

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

You can perform the following operations on an ECS group:



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

Constraints

- ECS groups support the anti-affinity policy only.
- ECSs are deployed on different physical hosts.
- If the maximum number of ECS groups is reached, you need to contact customer service to increase the quota.
- The maximum number of ECSs that can be added to an ECS group varies depending on the region. You can view the quota on the [ECS Group](#) page.

Creating an ECS Group

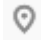

Create an ECS group and associate the same policy to all group members. ECS groups are independent from each other.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. On the **ECS Group** page, click **Create ECS Group**.
6. Enter the name of an ECS group.
7. Select a policy for the ECS group.
8. Click **OK**.

Adding an ECS to an ECS Group

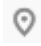

To improve service reliability, you can add ECSs to an ECS group so that these ECSs in this group can run on different hosts.

 **NOTE**

- ECSs of specific types must be stopped before being added to an ECS group. Stop these ECSs as prompted when adding them to an ECS group.
 - After an ECS is added to an ECS group, the system reallocates a host to run this ECS to ensure that ECSs in this group are running on different hosts. When the ECS is being restarted, the startup may fail due to insufficient resources. In such a case, remove the ECS from the ECS group and try to restart the ECS again.
 - ECSs that have local disks attached can be added to an ECS group only during the creation process. Once created, they can no longer be added to any ECS groups.
1. Log in to the management console.
 2. Click  in the upper left corner and select your region and project.
 3. Click . Under **Compute**, choose **Elastic Cloud Server**.
 4. In the navigation pane on the left, choose **ECS Group**.
 5. Locate the row that contains the target ECS group and click **Add ECS** in the **Operation** column.
 6. On the **Add ECS** page, select an ECS to be added.
 7. Click **OK**. The ECS is added to the ECS group.

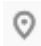

Removing an ECS from an ECS Group

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

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. Expand the ECS group information and view the ECSs in the ECS group.
6. Locate the ECS to be removed and click **Remove** in the **Operation** column.
7. In the displayed dialog box, click **OK**.
The ECS is removed from the ECS group.

Deleting an ECS Group

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

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **ECS Group**.
5. Locate the ECS group to be deleted and click **Delete** in the **Operation** column.

6. In the displayed dialog box, click **OK**.

2.6.6 Automatically Recovering ECSs

Scenarios


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

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

Notes

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. Set **Auto Recovery** to **Enable** or **Disable**.
Automatic recovery is enabled by default.

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

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

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

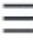
Constraints

This method applies only to Linux ECSs.

Procedure

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

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

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

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

ecs-01: 192.168.0.1

ecs-02: 192.168.0.2

Step 2 Obtain the hostnames for the two ECSs.

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

```
sudo hostname
```

For example, the obtained hostnames are as follows:

ecs-01: hostname01

ecs-02: hostname02

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

1. Log in to ecs-01.
2. Run the following command to switch to user **root**:

```
sudo su -
```

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

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

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

```
ping Hostname
```

```
----End
```

2.7 Modifying ECS Specifications (vCPUs and Memory)

2.7.1 Modifying Individual ECS Specifications

Scenarios

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

Notes

- The ECS needs to be stopped during the 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 are unavailable for selection.
- Downgrading ECS specifications (vCPU or memory) will reduce performance.
- Certain ECS types do not support specifications modification currently. For details about available ECS types and functions, see [ECS Types](#). For details about restrictions on using different types of ECSs, see their notes.
- When the disk status is **Expanding**, you are not allowed to modify the specifications of the ECS where the disk is attached.

- Before modifying specifications, make sure that the ECS has been stopped.
- For yearly/monthly ECSs that use paid images, the instance specifications cannot be downgraded. This means you cannot change the specifications to lower-cost ones.

Preparations

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

NOTE



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

For more information about NIC flapping, see [What Should I Do If NIC Flapping Occurs After My ECS Specifications Are Modified?](#)

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

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

Step 1: Modify Specifications

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, view the status of the target ECS.
If the ECS is not in **Stopped** state, choose **More > Stop** in the **Operation** column.
5. Choose **More > Modify Specifications** in the **Operation** column.
The **Modify ECS Specifications** page is displayed.
6. Select the new ECS type, vCPUs, and memory as prompted.
7. Click **Next**.
8. Confirm the modified configurations and click **Submit**.
9. Check whether the specifications have been modified.
After modifying the specifications, you can check whether the specifications have been modified in **Failures**.
 - a. Check whether **Failures** is displayed on the management console. For details, see [Viewing Failed Tasks](#).
 - If yes, go to step **9.b**.
 - If no, the specifications have been modified.
 - b. Click **Failures**. Then, in the **Failures** dialog box, click **Operation Failures** and check whether the task is contained in the list by **Name/ID**, **Operated At**, or **Task**.

- If yes, the specifications modification failed. See [Follow-up Procedure](#) for failure causes.
- If no, the specifications have been modified.

Step 2: Check Disk Attachment

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

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

Follow-up Procedure

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

1. Log in to the management console.
2. Under **Management & Deployment**, choose **Cloud Trace Service**.
3. In the navigation pane on the left, choose **Trace List**.
4. In the **Trace Name** column, locate the **resizeServer** event by resource ID.
The resource ID is the ID of the ECS on which the specifications modification failed.
5. Click **View Trace** in the **Operation** column to view the failure cause.
If the fault cannot be rectified based on logs, contact customer service.

2.8 Reinstalling or Changing the OS

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



Constraints

- The EVS disk quotas must be greater than 0.
- If the target ECS is created using a private image, ensure that the private image is available.
- H2 ECSs do not support OS reinstallation.

Prerequisites

- The target ECS has a system disk attached.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Reinstall OS** in the **Operation** column.
Before reinstalling the OS, stop the ECS or select **Stop the ECS** in the **Reinstall OS** dialog box.
5. Select the login mode.
If the target ECS uses key pair authentication, you can replace the original key pair.
6. Click **OK**.
7. In the **Reinstall OS** dialog box, confirm the settings, read and select the agreement or disclaimer, and click **OK**.

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

NOTE

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

Follow-up Procedure

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

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

2.8.2 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 10 to 20 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. Otherwise, the login will fail.
- The target ECS must have a system disk attached.
- The EVS disk quota must be greater than 0.
- The system disk type cannot be changed.
- An ISO image created from an ISO file cannot be used to change the OS of an ECS. You need to install an OS and drivers on the ECS and use the ECS to create a system disk image first.
- The boot mode (BIOS or UEFI) cannot be changed.

Notes



- After the OS is changed, the original OS is not retained, and the original system disk is deleted, including the data in all partitions of the system disk.
- 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.
- It takes about 10 to 20 minutes to change the OS. During this process, the ECS is in **Changing OS** state.

Prerequisites

- The data is backed up.
For details, see [Cloud Backup and Recovery](#).
- If you want to change the login authentication mode from password to key pair during the OS change, create a key file in advance.
For details, see [\(Recommended\) Creating a Key Pair on the Management Console](#).
- If you plan to use a private image to change the OS, ensure that a private image is available. For details about how to create a private image, see [Image Management Service > User Guide](#).

- If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
- If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
- If a private image from another region is required, make sure that the image has been copied.
- If a private image from another user account is required, make sure that the image has been shared with you.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image/Disk > Change OS** in the **Operation** column.
Before changing the OS, stop the ECS first or select **Stop the ECS** in the **Change OS** dialog box.
5. Select the target image.
6. Configure the login mode.
If the target ECS uses key pair authentication, you can replace the original key pair.
7. Click **OK**.
8. In the **Change OS** dialog box, confirm the specifications, and click **OK**.
After the application is submitted, the status **Changing OS** is displayed. When this status disappears, the OS change is complete.

NOTE

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

Follow-up Procedure

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

- If the OS change is unsuccessful, perform steps **3** to **8** again to retry the OS change.
- If the second OS change attempt is unsuccessful, contact customer service for manual recovery at the backend.



2.9 Viewing ECS Information

2.9.1 Viewing ECS Creation Statuses

Scenarios

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. After creating an ECS, view the creation status above the ECS list beside the common operations (**Start**, **Stop**, **Restart**, and **More**).
5. Click the number displayed above **Creating** and view task details.

NOTE

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

2.9.2 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 2-21](#) lists the types of failures that can be recorded in the **Failures** area.

Table 2-21 Failure types

Failure Type	Description
Creation failures	A task failed. For a failed task, the system rolls back the task and displays an error code, for example, Ecs.0013 Insufficient EIP quota .
Operation failures	<ul style="list-style-type: none">• Modifying ECS specifications If an ECS specifications modification failed, this operation is recorded in Failures.• Automatic recovery enabled during ECS creation Automatic recovery is enabled during ECS creation. After the ECS is created, if the system fails to enable automatic recovery, this operation is recorded in Failures.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. View **Failures** on the right side of common operations.

NOTE

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

- The ECS specifications are modified.
 - Automatic recovery is enabled during ECS creation.
5. Click the number displayed in the **Failures** area to view task details.
 - **Creation Failures**: show the failed ECS creation tasks.
 - **Operation Failures**: show the tasks with failed operations and error codes, which help you troubleshoot the faults.

2.9.3 Viewing ECS Details (List View)



Scenarios

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

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

Procedure

1. Log in to the management console.




2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
The **Elastic Cloud Server** page is displayed. On this page, you can view your ECSs and the basic information about the ECSs, such as their specifications, images, and IP addresses.
4. In the search box above the ECS list, select a filter (such as ECS name, ID, or private IP address), enter the corresponding information, and 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](#).

2.9.4 Exporting ECS Information

Scenarios

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the upper right corner above the ECS list, click .
The system will automatically export all ECSs in the current region under your account to a local directory.

NOTE

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

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

2.9.5 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 2-22](#).

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 2-22 Search syntax

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

Search Type	Supported Property	Format	Example	Description
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 -6bd9-40 2a-934b- 26c8969f 7169,bf6c 0281- f749-42d 7- b732-23a c69d80eb e Flavor Name: s2,s3	Select a property and enter or choose multiple values. The values are in OR relationship. The ID only supports exact search. You need to enter a complete ECS ID for search.
	Status Billing mode	Property: Value 1 Property: Value 2	Status: Running Status: Stopped	Select a property and choose multiple values. The values are in OR relationship.

Search Type	Supported Property	Format	Example	Description
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">Multiple properties are in AND relationship.Multiple values of the same property are in OR relationship. The ID only supports exact search. You need to enter a complete ECS ID for search. You can use vertical bars () to separate values of the following properties: status and billing mode. Alternatively, you can directly select the properties and corresponding values.

Procedure


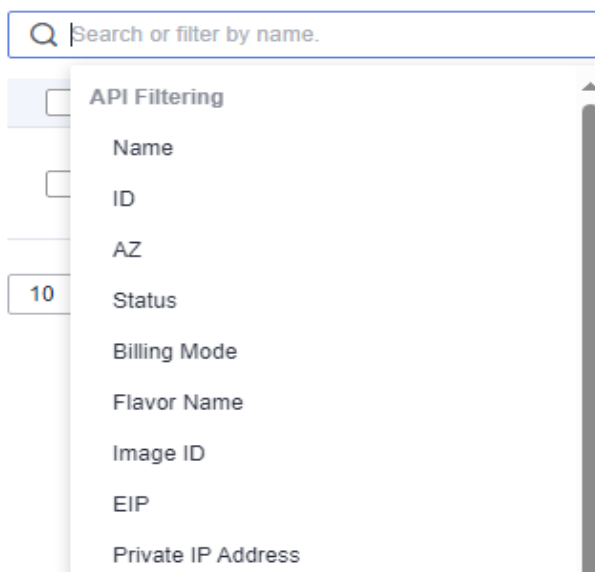
1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
The ECS console is displayed.
3. In the search box, specify search criteria.
You can use either of the following methods:
 - Directly enter a property value for search.
 - Select a property first and specify the property value for search.
 - i. Click the search box and select a property.



Figure 2-79 Selecting a property

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

Example 1: Searching by Property Value

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

- Searching by a single value
Enter a complete ECS ID and press **Enter** to search.

Figure 2-80 Entering a complete ECS ID**Figure 2-81** Automatic property matching

- Searching by multiple values
Enter multiple complete flavor names and press **Enter** to search.



Figure 2-82 Entering multiple complete flavor names

Figure 2-83 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
 - a. Select **Private IP Address** in the search box.
 - b. Enter a private IP address and press **Enter** to search. Private IP addresses can be used for fuzzy search. For example, you can enter **192.168.0** to search for all ECSs that use the 192.168.0 IP address.

Figure 2-84 Searching by private IP address



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

Figure 2-85 Searching for an ECS by ID



Example 3: Searching by Multiple Properties

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

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

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

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

Figure 2-86 Searching by name and private IP address

Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address	Billing Mode	Enterprise Project
ecs-9f64 f984337e-ccb0-4e3c-8...			Running	AZ2	1 vCPU 1 GiB s6.small.1 CentOS 7.5 64bit	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 2-87 Searching by flavor name



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

Figure 2-88 Searching by private IP address

A screenshot of a table displaying search results for private IP addresses. The search filter at the top is "Private IP Address: 10.0.0.146,10.0.0.235". The table has columns for Name/ID, Monitoring, Security, Status, AZ, Specifications/Image, and IP Address. Two rows are visible, both with status "Running" and AZ "AZ2". The IP Address column for the first row shows "10.0.0.146 (Private ...)" and for the second row "10.0.0.235 (Private ...)".

Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
ecs-976d f964337e-cc8b-4e3c-8...			Running	AZ2	1 vCPU 1 GiB s6.small.1 CentOS 7.5 64bit	10.0.0.146 (Private ...)
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 2-89 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 2-90 Searching by status and ID

Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address
ecs-976d f984337e-ccb-4e3c-8752-3bf13e7d111d			Stopped	AZ2	1 vCPU 1 GiB s6_small.1 CentOS 7.5 64bit	10.0.0.146 (Pr...
ecs-98c8 2ef2311b-c5b3-4273-b0b4-a3f527193393			Running	AZ2	1 vCPU 2 GiB s6_medium.2 CentOS 7.5 64bit	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 a tag value, and press **Enter**.

Figure 2-91 Searching by tag

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

Figure 2-92 Searching by tag

Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	IP Address	Billing Mode	Enterprise Project	Tag
ecs-976d f984337e-ccb-4e3c-8...			Stopped	AZ2	1 vCPU 1 GiB s6_small.1 CentOS 7.5 64bit	10.0.0.146 (Private ...)	Pay-per-usage Created on Oct 18, 202...	default	_type:Empty value: b7b

3 Images

3.1 Overview

Image

An image is an ECS or BMS template that contains an OS or service data. It may also contain proprietary software and application software, such as database software. Images are classified into public, private, and shared images.

Image Management Service (IMS) allows you to easily create and manage images. You can create an ECS using a public image, private image, or shared image. You can also use an existing ECS or external image file to create a private image.

Public Image

A public image is a standard, widely used image that contains a common OS, such as Ubuntu, CentOS, or Debian, and preinstalled public applications. This image is available to all users. Select your desired public image. Alternatively, create a private image based on a public image to copy an existing ECS or rapidly create ECSs in a batch. You can customize a public image by configuring the application environment or software.

For more information about public images, see [Overview](#).

Private Image

A private image contains an OS or service data, preinstalled public applications, and private applications. It is available only to the user who created it.

Table 3-1 Private image types

Image Type	Description
System disk image	Contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud.
Data disk image	Contains only service data. You can use a data disk image to create EVS disks and migrate your service data to the cloud.
Full-ECS image	Contains an OS, application software, and data for running services. A full-ECS image contains the system disk and all data disks attached to it.
ISO image	Created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.

If you plan to use a private image to change the OS, ensure that the private image is available. For instructions about how to create a private image, see [Image Management Service User Guide](#).

- If the image of a specified ECS is required, make sure that a private image has been created using this ECS.
- If a local image file is required, make sure that the image file has been imported to the cloud platform and registered as a private image.
- If a private image from another region is required, make sure that the image has been copied.
- If a private image from another user account is required, make sure that the image has been shared with you.

Shared Image

A shared image is a private image shared by another user and can be used as your own private image. For details, see [Sharing Images](#).

- Only the private images that have not been published in KooGallery can be shared.
- Images can be shared within a region only.
- Each image can be shared to a maximum of 128 tenants.
- You can stop sharing images anytime without notifying the recipient.
- You can delete shared image anytime without notifying the recipient.
- Encrypted images cannot be shared.
- Only the full-ECS images created using CBR can be shared.

KooGallery Image

A KooGallery image is a third-party image that has an OS, application environment, and software preinstalled. You can use the images to deploy websites and application development environments with a few clicks. No additional configuration is required.

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

Helpful Links

- [Creating a Private Image](#)

3.2 Creating an Image

Scenarios

You can use an existing ECS to create a system disk image, data disk image, and full-ECS image.

- **System disk image:** contains an OS and application software for running services. You can use a system disk image to create ECSs and migrate your services to the cloud.
- **Data disk image:** contains only service data. You can create a data disk image from an ECS data disk. You can also use a data disk image to create EVS disks and migrate your service data to the cloud.
- **Full-ECS image:** contains all the data of an ECS, including the data on the data disks attached to the ECS. A full-ECS image can be used to rapidly create ECSs with service data.
- **ISO image:** is created from an external ISO image file. It is a special image that can only be used to create temporary ECSs.

You can use a private image to change the OS. For instructions about how to create a private image, see [Image Management Service User Guide](#).

Prerequisites

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

For details, see [How Do I Configure an ECS, BMS, or Image File Before I Use It to Create an Image?](#)

Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, choose **More > Manage Image/Disk/Backup > Create Image** in the **Operation** column.
5. Configure the following information:
[Table 3-2](#) and [Table 3-3](#) list the parameters in the **Image Type and Source** and **Image Information** areas, respectively.

Table 3-2 Image type and source

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

Table 3-3 Image information

Parameter	Description
Encryption	This parameter specifies whether the image will be encrypted. The value is provided by the system and cannot be changed. <ul style="list-style-type: none">• Only an unencrypted private image can be created from an unencrypted ECS.• Only an encrypted private image can be created from an encrypted ECS.
Name	Set a name for the image.
Enterprise Project	Select an enterprise project from the drop-down list. This parameter is available only if you have enabled enterprise projects or your account is an enterprise account. To enable this function, contact your customer manager. An enterprise project provides central management of cloud resources on a project.
Tag	(Optional) Set a tag key and a tag value for the image to make identification and management of your images easier.
Description	(Optional) Enter a description of the image.

6. Click **Next** and submit the request.

4 Disks

4.1 Overview

What Is Elastic Volume Service?

Elastic Volume Service (EVS) offers scalable block storage for ECSs. With high reliability, high performance, and rich specifications, EVS disks can be used for distributed file systems, development and test environments, data warehouses, and high-performance computing (HPC) scenarios to meet diverse service requirements.

Disk Types

EVS disk types differ in performance. Choose a disk type based on your requirements.

For more information about EVS disk specifications and performance, see [Elastic Volume Service User Guide](#).

Helpful Links

- [Attaching an EVS Disk to an ECS](#)
- [Initializing an EVS Data Disk](#)
- [How Can I Adjust System Disk Partitions?](#)
- [Can I Attach Multiple Disks to an ECS?](#)
- [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)

4.2 Adding a Disk to an ECS



Scenarios

The disks attached to an ECS include one system disk and one or more data disks. The system disk is automatically created and attached when the ECS is created. You do not need to purchase it again. The data disks can be added in either of the following ways:

- If you add data disks when purchasing an ECS, the data disks will be automatically attached to the ECS.
- If you purchase data disks after an ECS is created, the data disks need to be manually attached to the ECS.

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Locate the row containing the target ECS and choose **More > Manage Image/Disk/Backup > Add Disk** in the **Operation** column.

The page for adding a disk is displayed.

5. Set parameters for the new EVS disk as prompted.

For instructions about how to set EVS disk parameters, see [Purchasing an EVS Disk](#).

NOTE

- By default, the billing mode of the new disk is the same as that of the ECS.
 - By default, the new disk is in the same region as the ECS.
 - By default, the new disk is in the same AZ as the ECS, and the AZ of the disk cannot be changed.
 - After the new disk is created, it is attached to the ECS by default.
6. Click **Next** to confirm the order and click **Submit** to complete the payment.
The system automatically switches back to the **Disks** tab on the ECS management console. Then, you can view the information of the new disk.

Follow-up Procedure

The system automatically attaches the new disk to the ECS, but the disk can be used only after it is initialized. To do so, log in to the ECS and initialize the disk.

For details about how to initialize a data disk, see [Initializing an EVS Data Disk](#).

4.3 Attaching a Disk to an ECS

Scenarios



If the existing disks of an ECS fail to meet service requirements, for example, due to insufficient disk space or poor disk performance, you can attach more available EVS disks to the ECS, or purchase more disks (under **Storage > Elastic Volume Service**) and attach them to the ECS.

Prerequisites

- EVS disks are available.

For instructions about how to purchase an EVS disk, see [Purchasing an EVS Disk](#).

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID for search.
5. Click the name of the target ECS.
The page providing details about the ECS is displayed.
6. Click the **Disks** tab. Then, click **Attach Disk**.
The **Attach Disk** dialog box is displayed.
7. Select the target disk and specify the disk as the system disk or data disk
 - For KVM ECSs, you can specify a disk as a system disk or data disk but cannot specify a device name for the disk.
 - For Xen ECSs, you can specify the device name of a disk, such as **/dev/vdb**.

NOTE

- If no EVS disks are available, click **Create Disk** in the lower part of the list.
 - For the restrictions on attaching disks, see [What Are the Requirements for Attaching an EVS Disk to an ECS?](#)
8. Click **OK**.
After the disk is attached, you can view the information about it on the **Disks** tab.

Follow-up Procedure

If the attached disk is newly created, the disk can be used only after it is initialized.

For details about how to initialize a data disk, see [Initializing an EVS Data Disk](#).

4.4 Detaching an EVS Disk from a Running ECS

Scenarios

You can detach EVS disks from an ECS.

- System disks (mounted to **/dev/sda** or **/dev/vda**) can only be detached offline. They must be stopped before being detached.
- Data disks (mounted to points other than **dev/sda**) can be detached online if the attached ECS is running certain OSs. You can detach these data disks without stopping the ECS.

This section describes how to detach a disk from a running ECS.

Constraints

- The EVS disk to be detached must be mounted to a point other than `/dev/sda` or `/dev/vda`.
EVS disks mounted to `/dev/sda` or `/dev/vda` are system disks and cannot be detached from running ECSs.
- Before detaching an EVS disk from a running Linux ECS, you must log in to the ECS and run the `umount` command to cancel the association between the disk and the file system. In addition, ensure that no programs are reading data from or writing data to the disk. Otherwise, detaching the disk will fail.

Notes

- Do not detach an EVS disk from an ECS that is being started, stopped, or restarted.
- Do not detach an EVS disk from a running ECS whose OS does not support this feature. OSs supporting EVS disk detachment from a running ECS are listed in [OSs Supporting EVS Disk Detachment from a Running ECS](#).
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and then attached to it again. This is a normal case due to the drive letter allocation mechanism of the Linux system.
- For a running Linux ECS, the drive letter may be changed after an EVS disk is detached from it and the ECS is restarted. This is a normal case due to the drive letter allocation mechanism of the Linux system.

OSs Supporting EVS Disk Detachment from a Running ECS

OSs supporting EVS disk detachment from a running ECS include two parts:

- For the first part, see [External Image File Formats and Supported OSs](#).
- [Table 4-1](#) lists the second part of supported OSs.

Table 4-1 OSs supporting EVS disk detachment from a running ECS

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

 NOTE

Online detachment is not supported by the ECSs running OSs not listed in the preceding table. For such ECSs, stop the ECSs before detaching disks from them to prevent any possible problems from occurring.

Procedure

1. On the **Elastic Cloud Server** page, click the name of the ECS from which the EVS disk is to be detached. The page providing details about the ECS is displayed.
2. Click the **Disks** tab. Locate the row containing the EVS disk to be detached and click **Detach**.

4.5 Expanding the Capacity of an EVS Disk

Scenarios

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

Procedure

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

- Create an EVS disk and attach it to an ECS.
- Expand the capacity of an existing EVS disk. The capacities of both system disks and data disks can be expanded.

For details about how to expand the capacity of an EVS disk, see [Disk Capacity Expansion](#).

For details, see [Expansion Overview](#).

 NOTE

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

4.6 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 [7](#) to the local disks of the new disk-intensive ECS.

5 Elastic Network Interfaces

5.1 Overview

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

Network Interface Types

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

Application Scenarios

- Flexible migration
You can detach a network interface from an instance and then attach it to another instance. The network interface retains its private IP address, EIP, and security group rules. In this way, service traffic on the faulty instance can be quickly migrated to the standby instance, implementing quick service recovery.
- Traffic management
You can attach multiple network interfaces that belong to different subnets in a VPC to the same instance, and configure the network interfaces to carry the private network traffic, public network traffic, and management network traffic of the instance. You can configure access control policies and routing policies for each subnet, and configure security group rules for each network interface to isolate networks and service traffic.

5.2 Attaching a Network Interface

Scenarios

If your ECS requires multiple network interfaces, you can attach them to your ECS.

Procedure



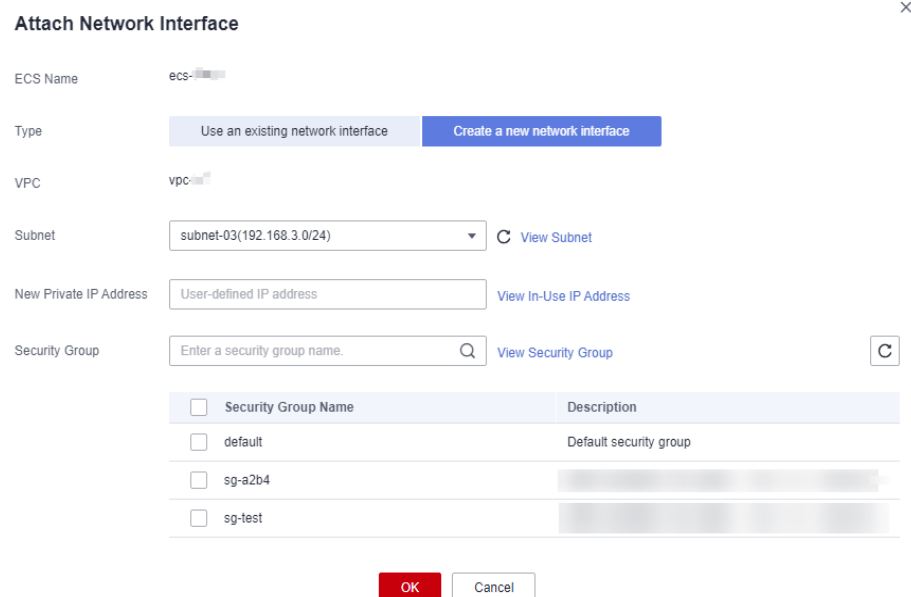

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Click the name of the target ECS.
The page providing details about the ECS is displayed.
5. On the **Network Interfaces** tab, click **Attach Network Interface**.
6. Select either of the following methods to attach the network interface.
 - Use an existing network interface.
 - i. (Optional) Search for the network interface by name, ID, or private IP address.
 - ii. In the network interface list, select the target one.
 - Create a new network interface.
Set the subnet and security group for the network interface to be attached.


Figure 5-1 Configuring the subnet and security group



Attach Network Interface ×

ECS Name: ecs-



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

VPC: vpc-

Subnet: [View Subnet](#)

New Private IP Address: [View In-Use IP Address](#)

Security Group: [View Security Group](#) C

<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.
 - **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 network interfaces can communicate with external networks, you need to configure policy-based routes for the network interfaces on the ECS after they are attached.
- For details, see [How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?](#)

Follow-up Procedure

Some OSs cannot identify newly added network interfaces. In this case, you must manually activate the network interfaces. Ubuntu is used as an example in the following network interface activation procedure. Required operations may vary among systems. For additional information, see the documentation for your OS.

1. Locate the row containing the target ECS and click **Remote Login** in the **Operation** column.
Log in to the ECS.
2. Run the following command to view the network interface name:
ifconfig -a
In this example, the network interface name is **eth2**.
3. Run the following command to switch to the target directory:
cd /etc/network
4. Run the following command to open the **interfaces** file:
vi interfaces
5. Add the following information to the **interfaces** file:
auto eth2
iface eth2 inet dhcp
6. Run the following command to save and exit the **interfaces** file:
:wq
7. Run either the **ifup eth2** command or the **/etc/init.d/networking restart** command to make the newly added network interface take effect.
X in the preceding command indicates the network interface name and SN, for example, **ifup eth2**.
8. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:
ifconfig
For example, check whether **eth2** is displayed in the command output.


- If yes, the newly added network interface has been activated, and no further action is required.
 - If no, the newly added network interface failed to be activated. Go to step 9.
9. Log in to the management console. Locate the row containing the target ECS, click **More** in the **Operation** column, and select **Restart**.
 10. Run the following command to check whether the network interface name obtained in step 2 is displayed in the command output:
 - If yes, no further action is required.
 - If no, contact customer service.

5.3 Detaching a Network Interface

Scenarios

An ECS can have up to 12 network interface, including one primary network interface that cannot be detached. This section describes how to detach an extension network interface.

Procedure

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
4. On the **Network Interfaces** tab, locate the target network interface and click **Detach**.

NOTE

You are not allowed to delete the primary ECS network interface. By default, the primary ECS network interface is the first network interface displayed in the network interface list.

5. In the displayed dialog box, click **OK**.

NOTE

Certain ECSs do not support network interface detachment when they are running. For details, see the GUI display. To detach a network interface from such an ECS, stop the ECS first.

5.4 Changing a VPC

Scenarios

This section describes how to change a VPC.

Constraints

- Only running or stopped ECSs support VPC change.
- A VPC can be changed for an ECS only if the ECS has one NIC.
- If you have reinstalled or changed the OS of an ECS before changing the VPC, log in to the ECS and check whether the password or key pair configured during the reinstallation or change is successfully injected.
 - If the login is successful, the password or key pair is injected. Perform operations as required.
 - Otherwise, the system is injecting the password or key pair. During this period, do not perform any operations on the ECS.
- During the VPC switchover, do not bind, unbind, or replace the EIP. Otherwise, a message indicating insufficient permissions will be displayed, but you do not need to take any action.
- If an ECS NIC has an IPv6 address, the VPC of the ECS cannot be changed.

Notes

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

NOTE

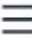
If you intend to change the VPC for a running ECS, the VPC change may fail when traffic is routed to the ECS NIC. In this case, you are advised to try again later or stop the ECS first and then try to change the VPC.

- After the VPC is changed, the subnet, private IP address, MAC address, and OS NIC name of the ECS will change.
- After the VPC is changed, the source/destination check and virtual IP address must be configured again.
- After the VPC is changed, you are required to reconfigure network-related application software and services, such as ELB, VPN, NAT, and DNS.

Prerequisites

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

Procedure

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network > Change VPC**.
The **Change VPC** dialog box is displayed.
4. Select an available VPC and subnet from the drop-down lists, and set the private IP address and security group as prompted.
You can select multiple security groups. In such a case, the access rules of all the selected security groups will apply to the ECS.

 NOTE

Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click **OK**.

5.5 Modifying a Private IP Address


Scenarios

You can modify the private IP address of the primary NIC. If you want to modify the private IP address of an extension NIC, delete the NIC and attach a new NIC.

Constraints

- The ECS must be stopped.
- If a virtual IP address or DNAT rule has been configured for the NIC, cancel the configuration before modifying the private IP address.
- If the NIC has an IPv6 address, its private IP address (IPv4 or IPv6 address) cannot be modified.
- To change the private IP address for a backend server of a load balancer, remove the backend server from the backend server group first.

Procedure

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. Click the name of the target ECS.
The ECS details page is displayed.
4. Click the **Network Interfaces** tab. Locate the row containing the primary network interface and click **Modify Private IP**.
The **Modify Private IP** dialog box is displayed.
5. Change the subnet and private IP address of the primary NIC as required.

 NOTE

Subnets can be changed only within the same VPC.


If the target private IP address is not specified, the system will automatically assign one to the primary NIC.

5.6 Managing Virtual IP Addresses

Scenarios

A virtual IP address provides the second IP address for one or more ECS NICs, improving high availability between the ECSs.

Binding a Virtual IP Address

1. Log in to the management console.
2. Click  . Under **Compute**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
4. On the **Network Interfaces** tab, locate the target virtual IP address and click **Manage Virtual IP Address**.
5. On the **IP Addresses** tab of the displayed page, locate the row containing the target virtual IP address and select **Bind to EIP** or **Bind to Server** in the **Operation** column.
Multiple ECSs deployed to work in active/standby mode can be bound with a virtual IP address to improve DR performance.
6. Click **OK**.

Configuring a Virtual IP Address for an ECS

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

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

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

5.7 Enabling NIC Multi-Queue

Scenarios

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

The ECS described in this section is assumed to comply with the requirements on specifications and virtualization type.

- If the ECS was created using a public image listed in [Support of NIC Multi-Queue](#), NIC multi-queue has been enabled on the ECS by default. Therefore, you do not need to perform the operations described in this section.
- If the ECS was created using a private image and the OS of the external image file is listed in [Support of NIC Multi-Queue](#), perform the following operations to enable NIC multi-queue:
 - a. [Importing the External Image File to the IMS Console](#)
 - b. [Setting NIC Multi-Queue for the Image](#)
 - c. [Creating an ECS Using a Private Image](#)
 - d. [Running the Script for Configuring NIC Multi-Queue](#)

 NOTE

After NIC multi-queue is enabled on an ECS, you need to enable this function on the ECS again after you add or delete a NIC or change the VPC for the ECS. For details, see [Running the Script for Configuring NIC Multi-Queue](#).

Support of NIC Multi-Queue

NIC multi-queue can be enabled on an ECS only when the ECS specifications, virtualization type, and image OS meet the requirements described in this section.

- For details about the ECS specifications that support NIC multi-queue, see [ECS Types](#).

 NOTE

If the number of NIC queues is greater than 1, NIC multi-queue is supported.

- The virtualization type must be KVM.
- The Linux public images listed in [Table 5-1](#) support NIC multi-queue.

 NOTE

- It is a good practice to upgrade the kernel version of the Linux ECS to 2.6.35 or later. Otherwise, NIC multi-queue is not supported.

Run the `uname -r` command to obtain the kernel version. If the kernel version is earlier than 2.6.35, contact customer service to upgrade the kernel.

Table 5-1 Support of NIC multi-queue for Linux ECSs

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

Importing the External Image File to the IMS Console


For details, see "Registering an Image File as a Private Image" in *Image Management Service User Guide*. After the image file is imported, view the value of **NIC Multi-Queue** on the page providing details about the image.

- If the value is **Supported**, go to [Creating an ECS Using a Private Image](#).
- If the value is **Not supported**, go to [Setting NIC Multi-Queue for the Image](#).


Setting NIC Multi-Queue for the Image

Use one of the following methods to set the NIC multi-queue attribute:

Method 1:

1. Log in to the management console.
2. Click  . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab, locate the row containing the target image, click **Modify** in the **Operation** column.
4. Set the NIC multi-queue attribute of the image.

Method 2:

1. Log in to the management console.
2. Click  . Under **Compute**, click **Image Management Service**.
3. Click the **Private Images** tab. In the image list, click the name of the target image to switch to the page providing details about the image.
4. Click **Modify** in the upper right corner. In the displayed **Modify Image** dialog box, set the NIC multi-queue attribute.

Method 3: Add `hw_vif_multiqueue_enabled` to an image through the API.

1. For instructions about how to obtain the token, see [Authentication](#).
2. For instructions about how to call an API to update image information, see [Updating Image Information \(Native OpenStack API\)](#).
3. Add **X-Auth-Token** to the request header.
The value of **X-Auth-Token** is the token obtained in step 1.
4. Add **Content-Type** to the request header.

The value of **Content-Type** is **application/openstack-images-v2.1-json-patch**.

The request URI is in the following format:

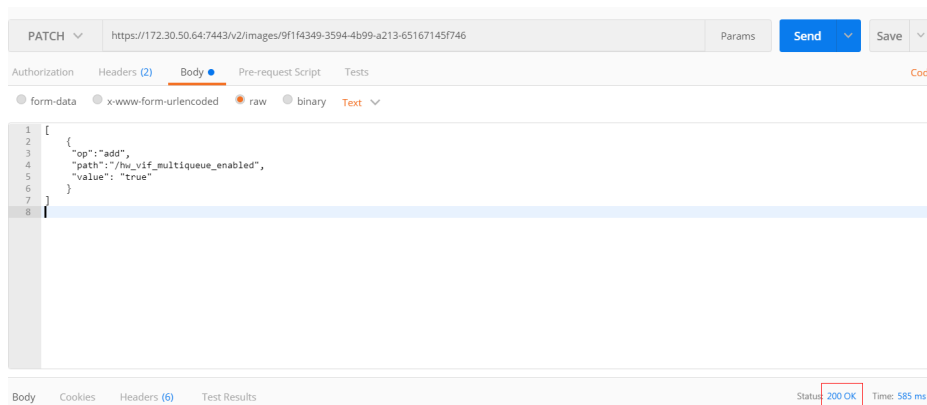
PATCH /v2/images/{image_id}

The request body is as follows:

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

Figure 5-2 shows an example request body for modifying the NIC multi-queue attribute.

Figure 5-2 Example request body



Creating an ECS Using a Private Image

Create an ECS using a registered private image. Note the following when setting the parameters:

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

Running the Script for Configuring NIC Multi-Queue

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 N/C

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 Queues of the NIC

NIC multi-queue has been enabled.

1. Log in to the ECS.
2. Run the following command to obtain the number of queues supported by the NIC and the number of queues with NIC multi-queue enabled:

ethtool -l NIC

Example:

```
[root@localhost ~]# ethtool -l eth0 #View the number of queues used by NIC eth0.
Channel parameters for eth0:
Pre-set maximums:
RX:          0
TX:          0
Other:       0
Combined: 4 #Indicates that a maximum of four queues can be enabled for the NIC.
Current hardware settings:
RX:          0
TX:          0
Other:       0
Combined: 1 #Indicates that four queues have been enabled.
```

6 EIPs


6.1 Overview

EIP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways or load balancers. Various billing modes are provided to meet different service requirements.

Each EIP can be used by only one cloud resource at a time.

Figure 6-1 Accessing the Internet using an EIP



Helpful Links



- [Binding an EIP](#)
- [Unbinding an EIP](#)
- [Modifying an EIP Bandwidth](#)

6.2 Binding an EIP

Scenarios

You can assign an EIP and bind it to an ECS to enable the ECS to access the Internet.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, select the ECS to which an EIP is to be bound, and choose **More > Manage Network > Bind EIP** in the **Operation** column.

5. In the displayed dialog box, select an EIP

 **NOTE**

If no EIP is available in the current region, the EIP list is empty. In such a case, purchase an EIP and then bind it.

6. Click **OK**.

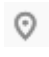

After the EIP is bound, view it in the ECS list on the **Elastic Cloud Server** page.

6.3 Unbinding an EIP

Scenarios

This section describes how to unbind an EIP from an ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Unbind EIP**.
5. Confirm the EIP to be unbound and click **OK**.

 **NOTE**

Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

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

6.4 Changing an EIP

Scenarios

You can change the EIP bound to your ECS as needed.

 **NOTE**

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, purchase one first.

Restrictions

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

An EIP has been assigned.

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


Unbinding an EIP

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Unbind EIP**.
4. Confirm the displayed information and click **Yes**.

NOTE

Unreleased EIPs will continue to be billed. To stop the EIPs from being billed, release them.

Binding a New EIP

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Locate the row containing the target ECS. Click **More** in the **Operation** column and select **Manage Network > Bind EIP**.
4. Select the desired EIP and click **OK**.

NOTE



If no EIP is available in the current region, the EIP list is empty. In such a case, purchase an EIP and then bind it.

6.5 Modifying an EIP Bandwidth

Scenarios

If an EIP has been bound to the ECS, the ECS can access the Internet using the bandwidth associated with the EIP. This section describes how to adjust the bandwidth of an ECS.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the ECS list, locate the row containing the target ECS and choose **More > Manage Network > Modify Bandwidth** in the **Operation** column.

5. Change the bandwidth name, billing mode, or bandwidth size as prompted.

6.6 Enabling Internet Connectivity for an ECS Without an EIP

Scenarios

To ensure platform security and conserve EIPs, EIPs are assigned only to specified ECSs. ECSs without EIPs cannot access the Internet directly. If these ECSs need to access the Internet (for example, to perform a software upgrade or install a patch), you can select an ECS with an EIP bound to function as a proxy ECS, providing an access channel for these ECSs.

NOTE




NAT Gateway is recommended, which provides both the SNAT and DNAT functions for your ECSs in a VPC and allows the ECSs to access or provide services accessible from the Internet. For details, see [NAT Gateway](#).

Prerequisites

- A proxy ECS with an EIP bound is available.
- The IP address of the proxy ECS is in the same network and same security group as the ECSs that need to access the Internet.

Linux Proxy ECS

The following uses a proxy ECS running CentOS 6.5 as an example. For details about other OSs and versions, see the official help documentation.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the proxy ECS name for search.
5. Click the name of the proxy ECS. The page providing details about the ECS is displayed.
6. On the **Network Interfaces** tab, click . Then, disable **Source/Destination Check**.

By default, the source/destination check function is enabled. When this function is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. However, this mechanism prevents the packet sender from receiving returned packets. Therefore, disable the source/destination check.
7. Log in to the proxy ECS.

For more details, see [Login Overview \(Linux\)](#).

8. Check whether the proxy ECS can access the Internet.

ping www.huaweicloud.com

The proxy ECS can access the Internet if information similar to the following is displayed:

Figure 6-2 Checking connectivity

```
root@ecs-f4f8 ~]# 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. Install iptables.

yum install iptables-services -y

10. Set the automatic startup of iptables.

systemctl start iptables

systemctl enable iptables

11. Disable the firewall.

systemctl stop firewalld

systemctl disable firewalld

12. Check whether IP forwarding is enabled on the proxy ECS.

cat /proc/sys/net/ipv4/ip_forward

– If **0** (disabled) is displayed, go to **13**.

– If **1** (enabled) is displayed, go to **18**.

13. Open the IP forwarding configuration file in the vi editor.

vi /etc/sysctl.conf

14. Press **i** to enter editing mode.

15. Set the **net.ipv4.ip_forward** value to **1**.

Set the **net.ipv4.ip_forward** value to **1**.

NOTE

If the **sysctl.conf** file does not contain the **net.ipv4.ip_forward** parameter, run the following command to add it:

```
echo net.ipv4.ip_forward=1 >> /etc/sysctl.conf
```

16. Press **Esc**, type **:wq**, and press **Enter**.

The system saves the configurations and exits the vi editor.

17. Apply the modification.

sysctl -p /etc/sysctl.conf

18. Configure default iptables rules.

iptables -P INPUT ACCEPT

iptables -P OUTPUT ACCEPT

iptables -P FORWARD ACCEPT

CAUTION

Running **iptables -P INPUT ACCEPT** will set default INPUT policy to ACCEPT, which poses security risks. You are advised to set security group rules to restrict inbound access.

19. Configure source network address translation (SNAT) to enable ECSs in the same network segment to access the Internet through the proxy ECS.

```
iptables -t nat -A POSTROUTING -o eth0 -s subnet/netmask-bits -j SNAT --to nat-instance-ip
```

For example, if the proxy ECS is in network 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

```
iptables -t nat -A POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4
```

NOTE

To retain the preceding configuration even after the ECS is restarted, run the **vi /etc/rc.local** command to edit the **rc.local** file. Specifically, copy the rule described in step 19 into **rc.local**, press **Esc** to exit Insert mode, and enter **:wq** to save the settings and exit.

20. Save the iptables configuration and set the automatic startup of iptables.

```
service iptables save
```

```
chkconfig iptables on
```

21. Check whether SNAT has been configured.

```
iptables -t nat --list
```


SNAT has been configured if information similar to [Figure 6-3](#) is displayed.

Figure 6-3 Successful SNAT configuration

```
[root@host- ~]# iptables -t nat --list
Chain PREROUTING (policy ACCEPT)
target     prot opt source                destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source                destination
SNAT      all  --  192.168.125.0/24      anywhere             to:192.168.125.4
SNAT      all  --  anywhere              anywhere             to:192.168.125.4
```

22. Add a route.

- a. Log in to the management console.
- b. Click  in the upper left corner and select your region and project.
- c. Under **Network**, click **Virtual Private Cloud**.
- d. Choose **Route Tables** in the left navigation pane. In the route table list, click a target route table. On the displayed page, click **Add Route**.
- e. Set route information on the displayed page.
 - **Destination:** indicates the destination network segment. The default value is **0.0.0.0/0**.

- **Next Hop:** indicates the private IP address of the proxy ECS.
You can obtain the private IP address of the ECS on the **Elastic Cloud Server** page.

23. Delete the added iptables rules as needed.

```
iptables -t nat -D POSTROUTING -o eth0 -s subnet/netmask-bits -j SNAT --to nat-instance-ip
```

For example, if the proxy ECS is in network segment 192.168.125.0, the subnet mask has 24 bits, and the private IP address is 192.168.125.4, run the following command:

```
iptables -t nat -D POSTROUTING -o eth0 -s 192.168.125.0/24 -j SNAT --to 192.168.125.4
```

7 Security

7.1 Methods for Improving ECS Security

Scenarios

If ECSs are not protected, they may be attacked by viruses, resulting in data leakage or data loss.

You can use the methods introduced below to protect your ECSs from viruses or attacks.

Protection Types

ECS can be protected externally and internally.

Table 7-1 Methods for improving ECS security

Type	Description	Protection Method
External security	DDoS attacks and Trojan horses or other viruses are common external security issues. To address these issues, you can choose services such as Host Security Service (HSS) and cloud-native anti-DDoS based on your service requirements:	<ul style="list-style-type: none">• Installing an Agent on Linux• Enabling the Basic/Enterprise/Premium Edition• Monitoring ECSs• Enabling Anti-DDoS• Backing Up Data Periodically

Type	Description	Protection Method
Internal security	Incorrect ports opening and weak passwords may cause internal security issues. Improving the internal security is the key to improving the ECS security. If the internal security is not improved, external security solutions cannot effectively intercept and block various external attacks.	<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 so that your ECSs are protected by the HSS cloud protection center. You will be able to check the security statuses and risks (if any) of all ECSs in a region on the HSS console.

We provide different methods for you to install the HSS agent depending on whether your ECSs are to be created or already exist.

- **An ECS is already created and HSS is not configured for it.**
For an existing ECS without HSS configured, you can manually install an Agent on it.
For details, see [Installing the Agent on Huawei Cloud Servers](#) and [Enabling Protection](#).

Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

Server monitoring includes basic monitoring, OS monitoring, and process monitoring for servers.

- **Basic monitoring**
Basic monitoring does not require the agent to be installed and automatically reports ECS metrics to Cloud Eye. Basic monitoring for KVM ECSs is performed every 5 minutes.
- **OS monitoring**
By installing the Agent on an ECS, OS monitoring provides system-wide, active, and fine-grained monitoring. OS monitoring for KVM ECSs is performed every minute.

To enable OS monitoring when purchasing an ECS:

Select **Enable Detailed Monitoring** when purchasing an ECS. After this option is selected, the cloud platform automatically installs the agent required for OS monitoring.

 **NOTE**

Currently, you can enable OS monitoring only when you purchase ECSs running specific OSs in specific regions.

To enable OS monitoring for a created ECS:

You need to manually install the agent if **Enable Detailed Monitoring** is not selected during the creation.

For instructions about how to install and configure the Agent, see [Agent Installation and Configuration](#).

- Process monitoring

Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. Processes are monitored at an interval of 1 minute (for KVM ECSs).

After server monitoring is enabled, you can set ECS alarm rules to customize the monitored objects and notification policies and learn about the ECS running status at any time.

Enabling Anti-DDoS

To defend against DDoS attacks, Huawei Cloud provides multiple security solutions. You can select an appropriate one based on your service requirements. Anti-DDoS Service on Huawei Cloud provides three sub-services: Cloud Native Anti-DDoS (CNAD) Basic (also known as Anti-DDoS), CNAD Pro, and Advanced Anti-DDoS (AAD).

Anti-DDoS is free while CNAD Pro and AAD are paid services.

For details about CNAD Pro and AAD, see [What Is Anti-DDoS?](#)

If you choose to purchase an EIP when purchasing an ECS, the console will display a message indicating that you have enabled free-of-charge Anti-DDoS protection.

Anti-DDoS defends ECSs against DDoS attacks and sends alarms immediately when detecting an attack. In addition, Anti-DDoS improves the bandwidth utilization to further safeguard user services.

Anti-DDoS monitors the service traffic from the Internet to public IP addresses and detects attack traffic in real time. It then scrubs attack traffic based on user-configured defense policies without service interruptions. It also generates monitoring reports that provide visibility into the security of network traffic.

Backing Up Data Periodically

Data backup is a process of storing all or part of data in different ways to prevent data loss. The following uses Cloud Backup and Recovery (CBR) as an example. For more backup methods, see [Overview](#).

CBR enables you to back up ECSs and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any

point in the past when the data was backed up. CBR protects your services by ensuring the security and consistency of your data.

You can use the cloud server backup and cloud disk backup to [back up your ECS data](#).

- Cloud server backup (recommended): Use this backup method if you want to back up the data of all EVS disks (system and data disks) attached to an ECS. This prevents data inconsistency caused by the time difference in creating a backup.
- Cloud disk backup: Use this backup method if you want to back up the data of one or more EVS disks (system or data disk) attached to an ECS. This minimizes backup costs on the basis of data security.

Enhancing the Login Password Strength

Key pair authentication is recommended because it is more secure than password-based authentication.

If you select **Password**, ensure that the password meets complexity requirements to prevent malicious attacks. For details, see [Application Scenarios for Using Passwords](#).

The system does not periodically change the ECS password. It is recommended that you change your password regularly for security.

The password must conform to the following rules:

- The password must consist of at least 10 characters.
- Do not use easily guessed passwords (for example, passwords in common rainbow tables or passwords with adjacent keyboard characters). The password must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special characters.
- Do not include accounts in passwords, such as administrator, test, root, oracle, and mysql.
- Change the password at least every 90 days.
- Do not reuse the latest five passwords.
- Set different passwords for different applications. Do not use the same password for multiple applications.

Improving the Port Security

You can use security groups to protect the network security of your ECSs. A security group controls inbound and outbound traffic for your ECSs. Inbound traffic originates from the outside to the ECS, while outbound traffic originates from the ECS to the outside.

You can configure security group rules to grant access to or from specific ports. You are advised to disable high-risk ports and only enable necessary ports.

[Table 7-2](#) lists common high-risk ports. You are advised to change these ports to non-high-risk ports. For details, see [Common Ports Used by ECSs](#).

Table 7-2 Common high-risk ports

Protocol	Port
TCP	42, 135, 137, 138, 139, 444, 445, 593, 1025, 1068, 1434, 3127, 3128, 3129, 3130, 4444, 4789, 5554, 5800, 5900, and 9996
UDP	135 to 139, 1026, 1027, 1028, 1068, 1433, 1434, 4789, 5554, and 9996

Periodically Upgrading the Operating System

After ECSs are created, you need to maintain and periodically upgrade the operating system. The officially released vulnerabilities will be released in Security Notices.

7.2 Security Groups

7.2.1 Overview

Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and that are mutually trusted. After a security group is created, you can create various access rules for the security group, these rules will apply to all ECSs added to this security group.

You can also customize a security group or use the default one. The system provides a default security group for you, which permits all outbound traffic and denies inbound traffic. ECSs in a security group are accessible to each other. For details about the default security group, see [Default Security Groups and Rules](#).

NOTE

If two ECSs are in the same security group but in different VPCs, the security group does not take effect. You can use a VPC peering connection to connect the two VPCs first. For details, see [VPC Connectivity Options](#).

Security Group Rules

After a security group is created, you can add rules to the security group. A rule applies either to inbound traffic (ingress) or outbound traffic (egress). After ECSs are added to the security group, they are protected by the rules of that group.

Each security group has default rules. For details, see [Default Security Groups and Rules](#). You can also customize security group rules. For details, see [Configuring Security Group Rules](#).

Notes and Constraints

- For better network performance, you are advised to associate an instance with no more than five security groups.

- A security group can have no more than 6,000 instances associated, or its performance will deteriorate.
- For inbound security group rules, the sum of the rules with **Source** set to **Security group**, of the rules with **Source** set to **IP address group**, and of the rules with inconsecutive ports, cannot exceed 120. If there are both IPv4 and IPv6 security group rules, up to 120 rules can be added for each type.

The limits on outbound security group rules are the same as those on inbound rules.

For example, to add inbound IPv4 rules to a security group (Sg-A), you can refer to [Table 7-3](#) for rules that meet the restrictions. Of these rules, rule A02 uses inconsecutive ports (TCP: 22,25,27) and security group Sg-B as the source. In this case, only one quota is occupied.

Table 7-3 Inbound security group rules

Rule No.	Action	Type	Protocol & Port	Source
Rule A01	Allow	IPv4	All	Current security group: Sg-A
Rule A02	Allow	IPv4	TCP: 22,25,27	Another security group: Sg-B
Rule A03	Allow	IPv4	TCP: 80-82	IP address group: ipGroup-A
Rule A04	Allow	IPv4	TCP: 22-24,25	IP address: 192.168.0.0/16

- Traffic from load balancers is not restricted by network ACL and security group rules if:

Transfer Client IP Address is enabled for the listener 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.


7.2.2 Default Security Groups and Rules

Note the following when using default security group rules:

- Inbound rules control incoming traffic to instances in the default security group. The instances can only communicate with each other but cannot be accessed from external networks.
- Outbound rules allow all traffic from the instances in the default security group to external networks.

[Figure 7-1](#) shows the default security group.

Figure 7-1 Default security group



NOTE

- Both default and custom security groups are free of charge. The name of a default security group is **default**.
- You cannot delete the default security group, but you can modify existing rules or add rules to the group.
- The default security group is automatically created to simplify the process of creating an instance for the first time. The default security group denies all external requests. To log in to an instance, add a security group rule by referring to [Remotely Logging In to an ECS from a Local Server](#).

Table 7-4 describes the rules in the default security group.

Table 7-4 Default security group rules

Direction	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	Allow	IPv4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inbound	Allow	IPv6	All	Source: default security group (default)	Allows IPv6 instances in the security group to communicate with each other using any protocol over any port.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows all traffic from the instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: :/0	Allows all traffic from the instances in the security group to any IPv6 address over any port.

When you create an ECS for the first time, the system automatically creates a VPC **vpc-default** and:

- Add the **Sys-WebServer** security group.
- Add the **Sys-FullAccess** security group.
- Add security group rules to the default security group **default**.

Table 7-5 Default security group rules

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Windows ECSs through the default Windows remote desktop.
Inbound	Allow	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows all IPv4 addresses to access Linux ECSs over SSH.
Inbound	Allow	IPv4	All	Source: Default security group (default)	Allows instances in the security group to communicate with each other over IPv4 protocols.
Inbound	Allow	IPv6	All	Source: Default security group (default)	Allows instances in the security group to communicate with each other over IPv6 protocols.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows access from instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows access from instances in the security group to any IPv6 address over any port.

Table 7-6 Sys-WebServer security group rules

Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	ICMP: All	Source: 0.0.0.0/0	Allows the use of the ping command to test the network connectivity over IPv4 protocols.

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 7-7 Sys-FullAccess security group rules

Direction	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	Allow	IPv4	All	Source: current security group (Sys-FullAccess)	Allows instances in the security group to communicate with each other over IPv4 protocols.

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.

7.2.3 Security Group Configuration Examples

When you create instances, such as cloud servers, containers, and databases, in a VPC subnet, you can use the default security group or create a security group. You can add inbound and outbound rules to the default or your security group to control traffic from and to the instances in the security group. Here are some common security group configuration examples:

- [Remotely Logging In to an ECS from a Local Server](#)
- [Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files](#)
- [Setting Up a Website on an ECS to Provide Services Externally](#)
- [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 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. You need to add inbound rules to allow specific traffic to the instances in the security group.
- By default, outbound security group rules allow all requests from the instances in the security group to access external resources. If outbound rules are deleted, the instances in the security group cannot communicate with external resources. To allow outbound traffic, you need to add outbound rules by referring to [Table 7-8](#).

Table 7-8 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

Remotely Logging In to an ECS from a Local Server

A security group denies all external requests by default. To remotely log in to an ECS in a security group from a local server, add an inbound rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable port 22. For details, see [Table 7-9](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 7-10](#).

Table 7-9 Remotely logging in to a Linux ECS using SSH

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 22	IP address: 0.0.0.0/0

Table 7-10 Remotely logging in to a Windows ECS using RDP

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3389	IP address: 0.0.0.0/0

NOTICE

If the source is set to 0.0.0.0/0, all external IP addresses are allowed to remotely log in to the ECS. To ensure network security and prevent service interruptions caused by network intrusions, set the source to a known IP address. For details, see [Table 7-11](#).

Table 7-11 Remotely logging in to an ECS using a known 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 FTP Files

By default, a security group denies all external requests. If you need to remotely connect to an ECS from a local server to upload or download files, you need to enable FTP ports 20 and 21.

Table 7-12 Remotely connecting to an ECS from any server to upload or download files

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 0.0.0.0/0

NOTICE

- 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 known IP address. For details, see [Table 7-13](#).
- You must first install the FTP server program on the ECSs and check whether ports 20 and 21 are working properly.

Table 7-13 Remotely connecting to an ECS from a known 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 Services Externally

A security group denies all external requests by default. If you have set up a website on an ECS that can be accessed externally, you need to add an inbound rule to the ECS security group to allow access over specific ports, such as HTTP (80) and HTTPS (443).

Table 7-14 Setting up a website on an ECS to provide services externally

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 80	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv4	TCP: 443	IP address: 0.0.0.0/0

Using ping Command to Verify Network Connectivity

Ping works by sending an Internet Control Message Protocol (ICMP) Echo Request. To ping an ECS from your PC to verify the network connectivity, you need to add an inbound rule to the security group of the ECS to allow ICMP traffic.

Table 7-15 Using ping command to verify network connectivity

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	ICMP: All	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv6	ICMP: All	IP address: ::/0

Enabling Communications Between Instances in Different Security Groups

Instances in the same VPC but associated with different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** to access MySQL databases in security group **sg-B**, you need to add an inbound rule to security group **sg-B** to allow access from ECSs in security group **sg-A**.

Table 7-16 Enabling communications between instances in different security groups

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A

NOTICE

In the example in "Allowing Traffic from a Virtual IP Address" in [How Security Groups Are Used](#), two ECSs in Subnet-A and Subnet-B are connected by a virtual IP address. If you set the source of inbound rules to the security groups associated with the ECSs, the ECSs in the two security groups cannot communicate with each other. To enable communications between them, set the source to the private IP address or subnet CIDR block of the virtual IP address.

Allowing External Instances to Access the Database Deployed on an ECS

A security group denies all external requests by default. If you have deployed a database on an ECS and want the database to be accessed from external instances on a private network, you need to add an inbound rule to the security group of the ECS to allow access over corresponding ports. Here are some common ports for databases:

- MySQL: port 3306
- Oracle: port 1521
- MS SQL: port 1433
- PostgreSQL: port 5432
- Redis: port 6379

Table 7-17 Allowing external instances to access the database deployed on an ECS

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A	Allows the ECSs in security group sg-A to access the MySQL database service.

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 service.
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 service.
Inbound	1	Allow	IPv4	TCP: 5432	IP address: 192.168.0.0/24	This rule allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database service.
Inbound	1	Allow	IPv4	TCP: 6379	IP address group: ipGroup-A	Allows ECSs whose private IP addresses are in IP address group ipGroup-A to access the Redis database service.

NOTICE

In this example, the source is for reference only. Set the source address based on your requirements.

Allowing ECSs to Access Specific External Websites

By default, a security group allows all outbound traffic. [Table 7-19](#) lists the default rules. If you want to allow ECSs to access specific websites, configure the security group as follows:

1. Add outbound rules to allow traffic over specific ports to specific IP addresses.

Table 7-18 Allowing ECSs to access 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 original outbound rules that allow all traffic.

Table 7-19 Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

7.2.4 Configuring Security Group Rules



Scenarios

Similar to firewall, a security group is a logical group used to control network access. You can define access rules for a security group to protect the ECSs that are added to this security group.

- Inbound: Inbound rules allow external network traffic to be sent to the ECSs in the security group.
- Outbound: Outbound rules allow network traffic from the ECSs in the security group to be sent out of the security group.

For details about the default security group rules, see [Default Security Groups and Security Group Rules](#). For details about configuration examples for security group rules, see [Security Group Configuration Examples](#).

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
5. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
6. Click the security group ID.
The system automatically switches to the security group details page.
7. Configure required parameters.


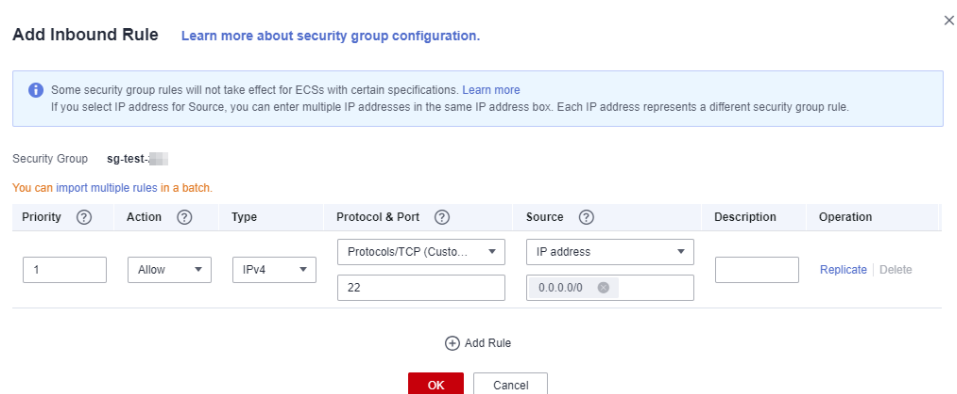
You can click  to add more inbound rules.

Figure 7-2 Add Inbound Rule



Add Inbound Rule [Learn more about security group configuration.](#) ×

i Some security group rules will not take effect for ECSs with certain specifications. [Learn more](#)
If you select IP address for Source, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.

Security Group **sg-test**

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1	Allow	IPv4	Protocols/TCP (Custo... 22	IP address 0.0.0.0/0		Replicate Delete

+ Add Rule

OK Cancel

Table 7-20 Inbound rule parameter description

Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1

Parameter	Description	Example Value
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 the source is allowed to ECSs in the security group over specified ports.If the Action is set to Deny, access from the source is denied to ECSs in the security group over specified ports. <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>Source IP address version. You can select:</p> <ul style="list-style-type: none">IPv4IPv6	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
	<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>Specify one of the following:</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.All ports: Leave it empty or enter 1-65535.	22, or 22-30
Source	<p>Source of the security group rule. The value can be an IP address or a security group, to allow access from IP addresses or instances in the security group.</p> <ul style="list-style-type: none">IP address:<ul style="list-style-type: none">Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) <p>If the source is a security group, this rule will apply to all instances associated with the selected security group.</p>	192.168.0.0/24

Parameter	Description	Example Value
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

8. Configure required parameters.

You can click  to add more outbound rules.

Figure 7-3 Add Outbound Rule

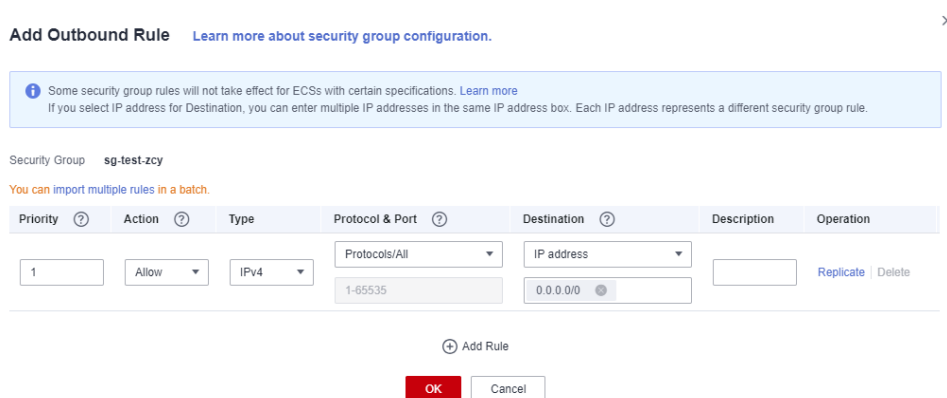


Table 7-21 Outbound rule parameter description

Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The value can be Allow or Deny . <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. 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

Parameter	Description	Example Value
Type	Destination IP address version. You can select: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Outbound rules control outgoing traffic over specific ports from instances in the security group. Specify one of the following: <ul style="list-style-type: none">• Individual port: Enter a port, such as 22.• Consecutive ports: Enter a port range, such as 22-30.• All ports: Leave it empty or enter 1-65535.	22, or 22-30
Destination	Destination of the security group rule. The value can be an IP address or a security group, to allow access to IP addresses or instances in the security group. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)- All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	0.0.0.0/0
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

9. Click **OK** to complete the security rule configuration.

Verifying Security Group Rules

After allowing traffic over a port in a security group rule, you need to ensure that the port used by the instance is also opened.

For example, if you have deployed a website on an ECS and want users to access your website through HTTP (80), you need to add an inbound rule to the ECS security group to allow access over the port. [Table 7-22](#) shows the rule.

Table 7-22 Security group rule

Direction	Protocol & Port	Source
Inbound	TCP: 80	IP address: 0.0.0.0/0

After adding the security group rule, perform the following operations to check whether the ECS port is opened and whether the rule is applied:

1. Log in to the ECS and check whether the ECS port is opened.
 - **Checking the port of a Linux server**
Run the following command to check whether TCP port 80 is being listened on:
netstat -an | grep 80
If the following figure is displayed, TCP port 80 is enabled.

Figure 7-4 Command output for the Linux ECS

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

- **Checking the port of a Windows server**
 - 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.
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.

Figure 7-5 Command output for the Windows ECS

```
TCP      0.0.0.0:80          0.0.0.0:0          LISTENING
```

7.2.5 Changing a Security Group


Scenarios

To change the security group associated with an ECS network interface, perform the operations described in this section.

Constraints

- Changing the security group will overwrite the original security group settings.
- Using multiple security groups may deteriorate ECS network performance. You are advised to select no more than five security groups.

Procedure

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. In the ECS list, locate the row that contains the target ECS. Click **More** in the **Operation** column and select **Manage Network > Change Security Group**. The **Change Security Group** dialog box is displayed.
4. Select the target NIC and security groups.

You can select multiple security groups. In such a case, the rules of all the selected security groups will apply to the ECS.

To create a security group, click **Create Security Group**.

NOTE

Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click **OK**.

7.3 HSS

What Is HSS?

Host Security Service (HSS) helps you identify and manage the assets on your servers, eliminate risks, and defend against intrusions and web page tampering. There are also advanced protection and security operations functions available to help you easily detect and handle threats.

After installing the HSS agent on your ECSs, you will be able to check the ECS security status and risks in a region on the HSS console.

How Do I Use HSS?

Before using the HSS service, install the agent on your ECS.

- **An ECS is already created and HSS is not configured for it.**
For an existing ECS without HSS configured, you can manually install an Agent on it.
For details, see [Installing the Agent on Huawei Cloud Servers](#) and [Enabling Protection](#).

How Do I Check Host Security Statuses?

On the **Server** tab, you can view the ECS security statuses in the current region.


1. Log in to the management console.
2. Click  and choose **Security > Host Security Service**.
3. Go to the **Servers** page to view the protection status of the target servers.

Table 7-23 Statuses

Parameter	Description
Agent Status	<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?](#)

7.4 Project and Enterprise Project

Creating a Project and Assigning Permissions

- **Creating a project**

Log in to the management console, click the username in the upper right corner, and select **Identity and Access Management** from the drop-down list box. In the navigation pane on the left, choose **Projects**. In the right pane, click **Create Project**. On the displayed **Create Project** page, select a region and enter a project name.

- **Assigning permissions**

You can assign permissions (of resources and operations) to user groups to associate projects with user groups. You can add users to a user group to control projects that users can access and the resources on which users can perform operations. To do so, perform the following operations:

- a. On the **User Groups** page of the IAM console, locate the target user group and click **Authorize** in the **Operation** column.
- b. Select policies or roles from the list.
- c. Click **Next** and select **Region-specific projects**.
- d. In the displayed regional project list, select one or more projects and click **OK**.
- e. On the **Users** page, locate the target user and click **Authorize** in the **Operation** column.
- f. Select **Inherit permissions from user groups** and select the user group authorized in **a**.

- g. Click **OK**.

Creating an Enterprise Project and Assigning Permissions

- **Creating an enterprise project**

On the management console, choose **Enterprise > Project Management** in the upper right corner. On the **Enterprise Project Management** console, click **Create Enterprise Project**.

 **NOTE**

Enterprise is available on the management console only if you have enabled the enterprise project, or your account is the primary account. To enable this function, contact customer service.

- **Assigning permissions**

You can add a user group to an enterprise project and configure a policy to associate the enterprise project with the user group. You can add users to a user group to control projects users can access and the resources on which users can perform operations. To do so, perform the following operations:

- a. On the **Enterprise Management** console, click the name of an enterprise project to go to the enterprise project details page.
- b. On the **Permissions** tab, click **Authorize User Group** to go to the **User Groups** page on the IAM console. Associate the enterprise project with a user group and assign permissions to the group.

For details, see [Creating a User Group and Assigning Permissions](#).

- **Associating ECSs with enterprise projects**

You can use enterprise projects to manage cloud resources.

- Select enterprise projects when purchasing ECSs.

On the page for buying an ECS, select an enterprise project from the **Enterprise Project** drop-down list.

- Add ECSs to an enterprise project.

On the **Enterprise Project Management** page, you can add existing ECSs to an enterprise project.

Value **default** indicates the default enterprise project. Resources that are not allocated to any enterprise projects under your account are displayed in the default enterprise project.

For more details, see [Enterprise Management User Guide](#).

8 Backup Using CBR

8.1 Overview

What Is CBR?

Cloud Backup and Recovery (CBR) enables you to back up cloud servers and disks with ease. In case of a virus attack, accidental deletion, or software or hardware fault, you can restore data to any point in the past when the data was backed up.

CBR protects your services by ensuring the security and consistency of your data.

What Are the Differences Between Backup, Snapshot, and Image?

You can use the cloud server backup function to create ECSs and the cloud disk backup function to create EVS disks.

An image can be a system disk image, data disk image, or full-ECS image.

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Cloud server backup	All disks (system and data disks) on an ECS	<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?

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Cloud disk backup	One or more specified disks (system or data disks)	<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. 	<p>Backup data is stored in OBS, instead of disks. This ensures data restoration upon disk data loss or corruption.</p> <p>Backup cost is reduced without compromising data security.</p>	<p>Creating a Cloud Disk Backup</p>	<ul style="list-style-type: none"> • Restoring Data Using a Cloud Disk Backup • Using a Backup to Create a Disk

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Snapshot	One or more specified disks (system or data disks)	<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 backup and restoration. ● You can create snapshots to rapidly save disk data as it was at specified points in time. You can also use snapshots to create new disks so that the created disks will contain the snapshot data in the beginning. 	Creating a Snapshot	Rolling Back Data from a Snapshot

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	Creating a Data Disk Using a Data Disk Image

Backup Type	Backup Object	Application Scenario	Differences and Advantages	Backup Method	Restoration Method
Full-ECS image	All disks (system and data disks) on an ECS	<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 8-1 One-off backup and periodic backup

Item	One-Off Backup	Periodic Backup
Backup policy	Not required	Required
Number of backup tasks	One manual backup task	Periodic tasks driven by a backup policy

Item	One-Off Backup	Periodic Backup
Backup name	User-defined backup name, which is manualbk_XXXX by default	System-assigned backup name, which is autobk_XXXX by default
Backup mode	Full backup for the first time and incremental backup subsequently, by default	Full backup for the first time and incremental backup subsequently, by default
Application scenario	Executed before patching or upgrading the OS or upgrading an application on a resource. A one-off backup can be used to restore the resource to the original state if the patching or upgrading fails.	Executed for routine maintenance of a resource. The latest backup can be used for restoration if an unexpected failure or data loss occurs.

8.2 Backing Up an ECS

Scenarios



Cloud Backup and Recovery (CBR) enhances data integrity and service continuity. For example, if an ECS or EVS disk is faulty or a misoperation causes data loss, you can use data backups to quickly restore data. This section describes how to back up ECSs and EVS disks.

For more information, see [CBR Architecture](#), [Backup Mechanism](#), and [Backup Options](#).

You can back up ECS data using Cloud Server Backup or Cloud Disk Backup.



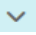
- Cloud Server Backup (recommended): Use this backup function if you want to back up the data of all EVS disks (system and data disks) on an ECS. This prevents data inconsistency caused by time difference in creating a backup.
- Cloud Disk Backup: Use this backup function if you want to back up the data of one or more EVS disks (system or data disk) on an ECS. This minimizes backup costs on the basis of data security.

ECS Backup Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the ECS list, locate the target ECS and choose **More > Manage Image/Disk/Backup > Create Server Backup**.
 - If the ECS has been associated with a vault, configure the backup information as prompted.

- **Server List:** The ECS to be backed up is selected by default.
 - **Name:** Customize your backup name.
 - **Description:** Supplementary information about the backup.
 - **Full Backup:** If this option is selected, the system will perform full backup for the ECS to be associated. The storage capacity used by the backup increases accordingly.
- If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.
- For details, see [Purchasing a Server Backup Vault](#).
5. Click **OK**. The system automatically creates a backup for the ECS.
- On the **Backups** tab page, if the status of the backup is **Available**, the backup task is successful.
- The ECS can be restarted if the backup progress of an ECS exceeds 10%. However, to ensure data integrity, restart it after the backup is complete.
- After the backup is complete, you can restore server data or create images on the **Backups** tab page. For details, see [Restoring Data Using a Cloud Server Backup](#) and [Using a Backup to Create an Image](#).

EVS Disk Backup Procedure

1. Log in to the management console.
 2. Click  in the upper left corner and select your region and project.
 3. Click . Under **Compute**, choose **Elastic Cloud Server**.
 4. In the ECS list, locate the target ECS and choose **More > Manage Image/Disk/Backup > Create Disk Backup**.
 - If the ECS has been associated with a vault, configure the backup information as prompted.
 - **Server List:** The ECS to be backed up is selected by default. Click  to view the disks attached to the ECSs. Select the disks to be backed up.
 - **Name:** Customize your backup name.
 - **Description:** Supplementary information about the backup.
 - **Full Backup:** If this option is selected, the system will perform full backup for the disks to be associated. The storage capacity used by the backup increases accordingly.
 - If the ECS is not associated with a vault, buy a vault first and then configure the backup information as prompted.

For details, see [Purchasing a Disk Backup Vault](#).
 5. Click **OK**. The system automatically creates a backup for the disk.
- On the **Backups** tab of the CBR console, if the status of the backup is **Available**, the backup task is successful.

If some files are deleted from the disk during the backup, the deleted files may fail to be backed up. Therefore, to ensure data integrity, delete the target data after the backup is complete.

After the backup is complete, you can restore disk data on the **Backups** tab page. For details, see [Restoring Data Using a Cloud Disk Backup](#).

9 Passwords and Key Pairs

9.1 Password Reset

9.1.1 Application Scenarios for Using Passwords

The password for logging in to your ECS is important and please keep it secure. You can reset the password if it is forgotten or expires.

[Table 9-1](#) provides guidance on how to reset your password in different scenarios.

Table 9-1 Resetting a password

Reference	Prerequisites
Resetting the Password for Logging In to an ECS in the OS	N/A
Resetting the Password for Logging In to a Linux ECS	The password reset plug-in has not been installed.

Background

[Table 9-2](#) shows the ECS password complexity requirements.

Table 9-2 Password complexity requirements

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

9.1.2 Resetting the Password for Logging In to an ECS in the OS

Scenarios

This section describes how to reset the password for logging in to an ECS in the OS when the password is about to expire, the password is forgotten, or you are logging in to the ECS for the first time. It is a good practice to change the initial password upon the first login.

Prerequisites

The ECS can be logged in.

Background

[Table 9-3](#) shows the ECS password complexity requirements.

Table 9-3 Password complexity requirements

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

Linux

1. Use the existing key file to log in to the ECS as user **root** through SSH.
For details, see [Login Using an SSH Key](#).
2. Run the following command to reset the password of user **root**:
passwd
To reset the password of another user, replace **passwd** with **passwd username**.
3. Enter a new password that meets the requirements listed in [Table 9-3](#) as prompted.
New password:
Retype new password:
If the following information is displayed, the password has been changed:
passwd: all authentication tokens updates successfully

9.1.3 Resetting the Password for Logging In to a Windows ECS Without the Password Reset Plug-in Installed

Scenarios

You can reset your ECS password if:

- The password is forgotten.
- The password has expired.

The method described in this section can only be used to change the password of a local Windows account, but not the password of a domain account.



Prerequisites

- A temporary Linux ECS running Ubuntu 14.04 or later is available. It is located in the same AZ and has the same CPU architecture as the target ECS.

NOTE

- Currently, this operation can be performed only for Ubuntu 16.04 and Ubuntu 18.04 public images.
- You can select an existing ECS or purchase a temporary ECS.
After the password of the purchased ECS is reset, you are advised to delete the ECS to avoid additional billing.
- You have bound an EIP to the temporary ECS and configured the apt-get source.
- You have used either of the following methods to install **ntfs-3g** and **chntpw** software packages on the temporary ECS:
Method 1:
Run the following command to install the **ntfs-3g** and **chntpw** software packages:
sudo apt-get install ntfs-3g chntpw
Method 2:
Download the ntfs-3g and chntpw software packages of the version required by the temporary ECS OS.

Procedure

1. Stop the original ECS and detach the system disk.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select your region and project.
 - c. Click . Under **Compute**, click **Elastic Cloud Server**.
 - d. Stop the original Windows ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

NOTE

Do not forcibly stop the Windows ECS. Otherwise, password reset may fail.

- e. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
2. Attach the system disk to the temporary ECS.
 - a. On the temporary ECS details page, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step 1.e and attach it to the temporary ECS.
 - c. Remotely log in to the temporary ECS.
 - d. Run the following command to view the directory of the system disk detached from the original Windows ECS now attached to the temporary ECS:
fdisk -l
 - e. Run the following command to mount the file system of the detached system disk to the temporary ECS:
mount -t ntfs-3g /dev/*Result obtained in step 2.d* /mnt/
For example, if the result obtained in step 2.d is **xvde2**, run the following command:
mount -t ntfs-3g /dev/xvde2 /mnt/
If the following error information is displayed after the preceding command is executed, the NTFS file systems may be inconsistent. In such a case, rectify the file system inconsistency.

```
The disk contains an unclean file system (0, 0).
Metadata kept in Windows cache, refused to mount.
Failed to mount '/dev/xvde2': Operation not permitted
The NTFS partition is in an unsafe state. Please resume and shutdown
Windows fully (no hibernation or fast restarting), or mount the volume
read-only with the 'ro' mount option.
```


Back up the disk data, run the following command to rectify the NTFS file system inconsistency, and attach the system disk:
ntfsfix /dev/*Result obtained in step 2.d*
For example, if the result obtained in step 2.d is **xvde2**, run the following command:
ntfsfix /dev/xvde2
 3. Change the password of the specified user and clear the original password.
 - a. Run the following command to back up the SAM file:

```
cp /mnt/Windows/System32/config/SAM /mnt/Windows/System32/
config/SAM.bak
```

- b. Run the following command to change the password of the specified user:

```
chntpw -u Administrator /mnt/Windows/System32/config/SAM
```

- c. Enter **1**, **q**, and **y** as prompted, and press **Enter**.

The password has been reset if the following information is displayed:

```
Select: [q] > 1
Password cleared!
Select: [q] > q
Hives that have changed:
#Name
0<SAM>
Write hive files? (y/n) [n] : y
0<SAM> - OK
```

4. Stop the temporary ECS, detach the system disk, and attach the system disk to the original Windows ECS.
 - a. Stop the temporary ECS, go to the ECS details page, and click the **Disks** tab.
 - b. Click **Detach** to detach the data disk temporarily attached in step **2.b**.
 - c. On the original Windows ECS details page, click the **Disks** tab.
 - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in step **4.b** and attach it to the original ECS as the system disk.
5. Start the original Windows ECS and set a new login password.
 - a. Click **Start** to start the original Windows ECS. After the status becomes **Running**, click **Remote Login** in the **Operation** column.
 - b. Click **Start**. Enter **CMD** in the search box and press **Enter**.
 - c. Run the following command to set a new password. The new password must meet the password complexity requirements described in [Application Scenarios for Using Passwords](#).
`net user Administrator New password`

9.1.4 Resetting the Password for Logging In to a Linux ECS

Scenarios

Keep your password secure. Reset the password if:

- The password is forgotten.
- The password has expired.

This section describes how to reset the password of user **root**. After resetting the password, you can log in to the ECS, and change the private key or reset the password of a non-**root** user.

Prerequisites

- A temporary Linux ECS is available. It is located in the same AZ and has the same CPU architecture as the target ECS.

 **NOTE**

You can select an existing ECS or purchase a temporary ECS.

After the password of the purchased ECS is reset, you are advised to delete the ECS to avoid additional billing.



- You have bound an EIP to the temporary ECS.

Procedure

1. Download the script for resetting the password and upload the script to the temporary ECS.

Download and decompress the password reset script. Use a connection tool, such as WinSCP, to upload the obtained **changepasswd.sh** script to the temporary ECS.

To download WinSCP, log in at <https://winscp.net/>.

2. Stop the original Linux ECS, detach the system disk from it, and attach the system disk to the temporary ECS.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select your region and project.
 - c. Click . Under **Compute**, click **Elastic Cloud Server**.
 - d. Stop the original ECS, switch to the page providing details about the ECS, and click the **Disks** tab.

 **NOTE**

Do not forcibly stop the original ECS. Otherwise, password reset may fail.

- e. Locate the row containing the system disk to be detached and click **Detach** to detach the system disk from the ECS.
3. Attach the system disk to the temporary ECS.
 - a. On the page providing details about the temporary ECS, click the **Disks** tab.
 - b. Click **Attach Disk**. In the displayed dialog box, select the system disk detached in step 2.e and attach it to the temporary ECS.
 4. Log in to the temporary ECS remotely and reset the password.
 - a. Locate the row containing the temporary ECS and click **Remote Login** in the **Operation** column.
 - b. Run the following command to view the directory of the system disk detached from the original Linux ECS now attached to the temporary ECS:
fdisk -l

Figure 9-1 Viewing the directory of the system disk

```
root@ecs-~:~# fdisk -l
Disk /dev/vda: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x43591807

Device            Boot Start          End  Sectors  Size Id Type
/dev/vda1         *            2048 83884031 83881984   40G 83 Linux

Disk /dev/vdb: 40 GiB, 42949672960 bytes, 83886080 sectors
Units: sectors of 1 * 512 = 512 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disklabel type: dos
Disk identifier: 0x5e9a7bb5

Device            Boot Start          End  Sectors  Size Id Type
/dev/vdb1         *            2048 83886079 83884032   40G 83 Linux
```

- c. Run the following commands in the directory where the **changepasswd.sh** script is stored to run the script for resetting the password:

```
chmod +x changepasswd.sh
```

```
./changepasswd.sh
```

When you run the password reset script, if the system displays a message indicating that there is no command related to logical volume manager (LVM), such as the message "no lvs command", install an LVM tool on the temporary ECS. The LVM2 tool is recommended, which can be installed by running the **yum install lvm2** command.

NOTE

If the original ECS and the temporary ECS both run CentOS 7, a mount failure may occur during script execution. To resolve this issue, replace **mount \$dev \$mountPath** with **mount -o nouuid \$dev \$mountPath** in the script.

- d. Enter the new password and the directory obtained in step 4.b as prompted.

If the following information is displayed, the password has been changed:
set password success.

5. (Optional) Enable remote root login for non-root users.

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

Modify the following settings:

- Change **PasswordAuthentication no** to **PasswordAuthentication yes**.
Alternatively, uncomment **PasswordAuthentication yes**.
- Change **PermitRootLogin no** to **PermitRootLogin yes**.
Alternatively, uncomment **PermitRootLogin yes**.
- Change the value of **AllowUsers** to **root**.

Search for **AllowUsers** in the file. If **AllowUsers** is missing, add **AllowUsers root** at the end of the file.

6. Stop the temporary ECS, detach the system disk, attach the system disk to the original Linux ECS, and restart the original Linux ECS.
 - a. Stop the temporary ECS, switch to the page providing details about the ECS, and click the **Disks** tab.
 - b. Click **Detach** to detach the data disk temporarily attached in step 3.
 - c. On the page providing details about the original Linux ECS, click the **Disks** tab.
 - d. Click **Attach Disk**. In the displayed dialog box, select the data disk detached in 6.b.
7. Restart the original Linux ECS.

9.2 Key Pairs

9.2.1 Application Scenarios for Using Key Pairs

Key Pairs

Key pairs are a set of security credentials for identity authentication when you remotely log in to ECSs.

A key pair consists of a public key and a private key. Key Pair Service (KPS) stores the public key and you store the private key. If you have imported a public key into a Linux ECS, you can use the corresponding private key to log in to the ECS without a password. You do not need to worry about password interception, cracking, or leakage.

You can use [Data Encryption Workshop \(DEW\)](#) to manage key pairs, including creating, importing, binding, viewing, resetting, replacing, unbinding, and deleting key pairs.

Scenarios

When purchasing an ECS, you are advised to select the key pair login mode.

- Logging in to a Linux ECS
You can directly use a key pair to log in a Linux ECS.
 - During the ECS creation, select the key pair login mode. For details, see "Set Login Mode" in [Step 3: Configure Advanced Settings](#).

Creating a Key Pair

You can create a key pair or use an existing one for remote login authentication.

- Creating a key pair
You can create a key pair using either of the following methods:
 - Follow the instructions in [\(Recommended\) Creating a Key Pair on the Management Console](#). The public key is automatically stored in the system, and the private key is stored locally.

- Follow the instructions in [Creating a Key Pair Using PuTTY Key Generator](#). Both the public and private keys are stored locally.
After the key pair is created, import the key pair following the instructions provided in [Importing a Key Pair](#) so that you can use it.
- Using an existing key pair
If an existing key pair (created using PuTTYgen, for example) is available, you can import the public key by referring to [Importing a Key Pair](#) on the management console to let the system maintain your public key.

NOTE

If the public key of the existing key pair is stored by clicking **Save public key** on 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?](#)

Notes and Constraints

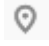

- Key pairs can be used to remotely log in to Linux ECSs only.
- SSH-2 key pairs created on the console support only the RSA-2048 cryptographic algorithms.
- Key pairs can be used only for ECSs in the same region.
- Imported key pairs support the following cryptographic algorithms:
 - RSA-1024
 - RSA-2048
 - RSA-4096
- Store your private key in a secure place because you need to use it to prove your identity when logging in to your ECS. The private key can be downloaded once only.

9.2.2 (Recommended) Creating a Key Pair on the Management Console

Scenarios

You can use the management console to create a key pair. ECS stores the public key and you store the private key.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.
5. On the displayed page, click **Create Key Pair**.
6. Enter a key pair name.
A key pair name consists of two parts: KeyPair and four random digits (KeyPair-xxxx).

7. Click **OK**.
8. Manually or automatically download a .pem private key file with the name that you specify as the key name. Store it in a secure place and click **OK**.

 **NOTE**

This is the only chance for you to save the private key file. Keep it secure. You'll need to provide the key pair name when you create an ECS, and the corresponding private key each time you connect to the ECS through SSH.

9.2.3 Creating a Key Pair Using PuTTY Key Generator

Scenarios

You can use puttygen.exe to create a key pair and store both the public key and private key locally.

 **NOTE**

Key pairs created using puttygen.exe must be imported by referring to [Importing a Key Pair](#) before they are used.

Procedure


1. Download and install PuTTY and PuTTYgen.
<https://www.chiark.greenend.org.uk/~sgtatham/putty/latest.html>

 **NOTE**

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


2. Obtain the public and private keys.
 - a. Double-click **puttygen.exe** to open **PuTTY Key Generator**.

Figure 9-2 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 9-3](#) is the public key.

Figure 9-3 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.

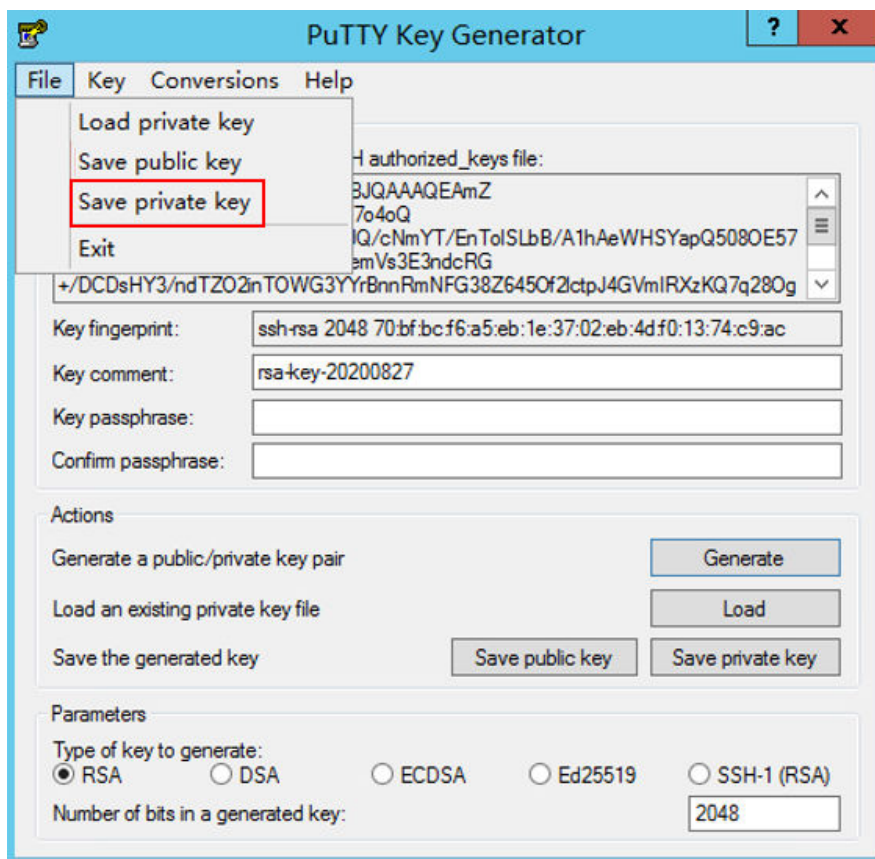
The format in which to save your private key file varies depending on application scenarios.

- 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 9-4 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:
5. After you have saved the key pair, import your public key to the ECS by referring to [Importing a Key Pair](#).

9.2.4 Importing a Key Pair

Scenarios

You need to import a key pair in either of the following scenarios:

- Create a key pair using PuTTYgen and import the public key to the ECS.
- Import the public key of an existing key pair to the ECS to let the system maintain your public key.




NOTE

If the public key of the existing key pair is stored by clicking **Save public key** on PuTTY Key Generator, the public key cannot be imported to the management console.

If you want to use this existing key pair for remote login, see [Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?](#)

Procedure

1. Log in to the management console.

2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the navigation pane on the left, choose **Key Pair**.
5. On the **Key Pair Service** page, click **Import Key Pair**.
6. Use either of the following methods to import the key pair:
 - Selecting a file
 - i. In the **Import Key Pair** dialog box of the management console, click **Select File** and select the locally stored public key file (for example, the .txt file saved in 3 in [Creating a Key Pair Using PuTTY Key Generator](#)).
 -  **NOTE**

Make sure that the file to be imported is a public key file.
 - ii. Click **OK**.

After the public key is imported, you can change its name.
- Copying the public key content
 - i. Copy the public key content from the locally stored .txt file into the **Public Key Content** text box.
 - ii. Click **OK**.

Helpful Links

- [What Should I Do If a Key Pair Cannot Be Imported?](#)
- [Why Does a Key Pair Created Using puttygen.exe Fail to Be Imported on the Management Console?](#)

9.2.5 Obtaining and Deleting the Password of a Windows ECS

9.2.5.1 Obtaining the Password for Logging In to a Windows ECS



Scenarios

Password authentication is required to log in to a Windows ECS. You must use the key file used when you created the ECS to obtain the administrator password generated during ECS creation. The administrator user is **Administrator** or the user configured using Cloudbase-Init. This password is randomly generated, offering high security.

You can obtain the initial password for logging in to a Windows ECS through the management console or APIs. For details, see this section.

Obtaining the Password Through the Management Console

1. Obtain the private key file (.pem file) used when you created the ECS.
2. Log in to the management console.

3. Click  in the upper left corner and select your region and project.
4. Click . Under **Compute**, click **Elastic Cloud Server**.
5. On the **Elastic Cloud Server** page, select the target ECS.
6. In the **Operation** column, click **More** and select **Get Password**.

NOTE

If **Get Password** is not displayed, the one-click password reset plug-in may not be installed.

In this case, you can reset the password by referring to .

7. Use either of the following methods to obtain the password through the key file:
 - Click **Select File** and upload the key file from a local directory.
 - Copy the key file content to the text field.
8. Click **Get Password** to obtain a random password.

Obtaining the Password Through APIs

1. Obtain the private key file (.pem file) used when you created the ECS.
2. Set up the API calling environment.
3. Call APIs. For details, see "Before You Start" in *Elastic Cloud Server API Reference*.
4. Obtain the ciphertext password.

Call the password obtaining APIs to obtain the ciphertext password of the public key encrypted using RSA. The API URI is in the format "GET /v2/{*project_id*}/servers/{*server_id*}/os-server-password".

NOTE

For details, see "Obtaining the Password for Logging In to an ECS" in the *ECS API Reference*.

5. Decrypt the ciphertext password.

Use the private key file used when you created the ECS to decrypt the ciphertext password obtained in step 4.

 - a. Run the following command to convert the ciphertext password format to ".key -nocrypt" using OpenSSL:
openssl pkcs8 -topk8 -inform PEM -outform DER -in *rsa_pem.key* -out *pkcs8_der.key* -nocrypt
 - b. Invoke the Java class library **org.bouncycastle.jce.provider.BouncyCastleProvider** and use the key file to edit the code decryption ciphertext.



9.2.5.2 Deleting the Initial Password for Logging In to a Windows ECS

Scenarios

After you obtain the initial password, it is a good practice to delete it to ensure system security.

Deleting the initial password does not affect ECS operation or login. Once deleted, the password cannot be retrieved. Before you delete a password, it is a good practice to record it.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. On the **Elastic Cloud Server** page, select the target ECS.
5. In the **Operation** column, click **More** and select **Delete Password**.
The system displays a message, asking you whether you want to delete the password.
6. Click **OK** to delete the password.

10 Resources and Tags

10.1 Tag Management

10.1.1 Overview

Scenarios

A tag identifies an ECS. Adding tags to an ECS facilitates ECS identification and management.

You can add a tag to an ECS during the ECS creation or after the ECS is created. You can add a maximum of 10 tags to each ECS.

NOTE

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

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

Basics of Tags

Tags are used to identify cloud resources. When you have many cloud resources of the same type, you can use tags to classify cloud resources by dimension (for example, usage, owner, or environment).

Figure 10-1 Example tags




Figure 10-1 shows how tags work. In this example, you assign two tags to each cloud resource. Each tag contains a key and a value that you define. The key of one tag is **Owner**, and the key of another tag is **Usage**. Each tag has a value.

You can quickly search for and filter specific cloud resources based on the tags added to them. For example, you can define a set of tags for cloud resources in an account to track the owner and usage of each cloud resource, making resource management easier.

Tag Naming Rules

- Each tag consists of a key-value pair.
- A maximum of 10 tags can be added to an ECS.
- For each resource, a tag key must be unique and can have only one tag value.
- A tag consists of a tag key and a tag value. [Table 10-1](#) lists the tag key and value requirements.

Table 10-1 Tag key and value requirements

Parameter	Requirement	Example Value
Key	<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

10.1.2 Adding Tags



Tags are used to identify cloud resources, such as ECSs, images, and disks. If you have multiple types of cloud resources which are associated with each other, you can add tags to the resources to classify and manage them easily. For more details, see [Overview](#).

You can add tags to an ECS in any of the following ways:

- [Adding Tags During ECS Creation](#)
- [Adding Tags on the ECS Details Page](#)
- [Adding Tags on the TMS Console](#)

For details about how to use predefined tags, see [Using Predefined Tags](#).



Adding Tags During ECS Creation

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Click **Buy ECS**.
5. Configure parameters for the ECS.
Select **Configure now** for **Advanced Options**. Then, add a tag key and tag value. For the tag key and tag value requirements, see [Table 10-1](#).

NOTE

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

Adding Tags on the ECS Details Page



1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the ECS list, click the name of the target ECS.
The ECS details page is displayed.
5. Click the **Tags** tab and then **Add Tag**. In the displayed dialog box, enter the tag key and tag value. For the tag key and tag value requirements, see [Table 10-1](#).
6. Click **OK**.
After tags are added, you can click **Edit** in the **Operation** column to modify them.

Adding Tags on the TMS Console

NOTE

This method is suitable for adding tags with the same tag key to multiple resources.

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the displayed **Resource Tags** page, select the region where the resource is located, select **ECS-ECS** for **Resource Type**, and click **Search**.
All ECSs matching the search criteria are displayed.

4. In the **Search Result** area, click **Create Key**. In the displayed dialog box, enter a key (for example **project**) and click **OK**.
After the tag is created, the tag key is added to the resource list. If the key is not displayed in the resource list, click  and select the created key from the drop-down list.
By default, the value of the tag key is **Not tagged**. You need to set a value for the tag of each resource to associate the tag with the resource.
5. Click **Edit** to make the resource list editable.
6. Locate the row containing the target ECS, click , and enter a value (for example **A**).
After a value is set for a tag key, the number of tags is incremented by 1. Repeat the preceding steps to add tag values for other ECSs.

Using Predefined Tags

If you want to add the same tag to multiple ECSs or other resources, you can create a predefined tag on the TMS console and then select the tag for the ECSs or resources. This frees you from having to repeatedly enter tag keys and values. To do so, perform the following operations:



1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. Choose **Predefined Tags** in the left navigation pane and click **Create Tag**. In the displayed dialog box, enter a key (for example, **project**) and a value (for example, **A**).
4. Choose **Compute** > **Elastic Cloud Server** from the service list and select the predefined tag by following the procedure for adding a tag.

10.1.3 Searching for Resources by Tag

After tags are added to resources, you can search for resources by tag using either of the following methods.

Searching for ECSs by Tag

On the **Elastic Cloud Server** page, search for ECSs by tag key or value.

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. Click **Search by Tag** above the upper right corner of the ECS list to expand the search area.
5. Enter the tag of the ECS to be queried.
Neither the tag key nor value can be empty. When the tag key or value is matched, the system automatically shows the target ECSs.
6. Add tags.

The system supports multiple tags and uses the intersection set of all tags to search for ECSs.

7. Click **Search**.

The system searches for ECSs based on tag keys and values.

Filtering Resources on the TMS Console

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria, including **Region**, **Resource Type**, and **Resource Tag**.
4. Click **Search**.



All the resources that meet the search criteria will be displayed in the **Search Result** area.

10.1.4 Deleting a Tag

If you no longer need a tag, delete it in any of the following ways:


- [Deleting a Tag on the ECS Details Page](#)
- [Deleting a Tag on the TMS Console](#)
- [Batch Deleting Tags on the TMS Console](#)



Deleting a Tag on the ECS Details Page

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, choose **Elastic Cloud Server**.
4. In the ECS list, click the name of the target ECS.
The ECS details page is displayed.
5. Click the **Tags** tab, locate the tag to be deleted, and click **Delete** in the **Operation** column.
6. Click **OK**.

Deleting a Tag on the TMS Console

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. In the **Search Result** area, click **Edit** to make the resource tag list editable.


If the key of a tag you want to delete is not contained in the list, click  and select the tag key from the drop-down list. It is a good practice to select at most 10 keys to display.

5. Locate the row containing the target ECS and click .
6. (Optional) Click  in the upper right of the **Search Result** area.
The resource list is refreshed and the refresh time is updated.

Batch Deleting Tags on the TMS Console

NOTICE

Exercise caution when deleting tags in a batch. After you delete the tags, they will be removed from all the associated ECSs and cannot be recovered.

1. Log in to the management console.
2. In the upper right corner of the page, click the username and select **Tag Management** from the drop-down list.
3. On the **Resource Tags** page, set the search criteria for ECSs and click **Search**.
4. Select the target ECSs.
5. Click **Manage Tag** in the upper left corner of the list.
6. In the displayed **Manage Tag** dialog box, click **Delete** in the **Operation** column. Click **OK**.
7. (Optional) Click  in the upper right of the **Search Result** area.
The resource list is refreshed and the refresh time is updated.


10.2 Quota Adjustment

What Is Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECS or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

11 Monitoring Using Cloud Eye

11.1 Monitoring ECSs

Monitoring is key for ensuring ECS performance, reliability, and availability. Using monitored data, you can determine ECS resource utilization. The cloud platform provides Cloud Eye to help you obtain the running statuses of your ECSs. You can use Cloud Eye to automatically monitor ECSs in real time and manage alarms and notifications to keep track of ECS performance metrics.

Server Monitoring includes **Basic Monitoring** and **OS Monitoring**.

- **Basic Monitoring** automatically reports ECS metrics to Cloud Eye.
- Using the agent installed on the target ECS, **OS Monitoring** provides system-wide, active, and fine-grained ECS monitoring.

For instructions about how to install and configure the agent, see **Server Monitoring** in *Cloud Eye User Guide*.

This section covers the following content:

- Viewing basic ECS metrics
- Viewing OS metrics (Agent installed on ECS)
- Viewing process monitoring metrics (Agent installed on ECS)
- Customizing ECS alarm rules
- Viewing ECS running statuses for routine monitoring

One-Click Monitoring

ECSs run on physical hosts. Although there are multiple mechanisms to ensure system reliability, fault tolerance, and high availability, host hardware might be damaged or power failures might occur. The cloud platform supports automatic recovery by default. If a physical host accommodating ECSs breaks down, the ECSs will automatically be migrated to a functional physical host to minimize the impact on your services. During the process, the ECSs will restart. For details, see [Can ECSs Automatically Recover After the Physical Host Accommodating the ECSs Becomes Faulty?](#)

You can enable one-click monitoring on the Cloud Eye console so that you will be notified if high availability occurs (if a physical host accommodating ECSs is faulty,

the ECSs will automatically be migrated to a functional physical host). For details, see [One-Click Monitoring](#).

11.2 Basic ECS Metrics

Description

This section describes basic monitoring metrics reported by ECS to Cloud Eye. You can use Cloud Eye to view these metrics and alarms generated for ECSs.

Namespace

SYS.ECS

Basic ECS Metrics

Basic ECS metrics vary depending on ECS OSs and types. For details, see [Table 11-1](#).

NOTE

- Certain ECS metrics require the installation of UVP VMTools on the image from which the ECS is created. For details about how to install UVP VMTools, see <https://github.com/UVP-Tools/UVP-Tools/>.
- Certain ECS metrics require the installation of the Agent on the ECS. After the Agent is installed, log in to the management console and choose **Cloud Eye** under **Management & Deployment**. On the Cloud Eye console, choose **Server Monitoring > Elastic Cloud Server** from the left navigation pane to view ECS metrics, such as **AGT. User Space CPU Usage**. For details, see [OS Monitoring Metrics Supported by ECSs with the Agent Installed](#).
- For details about how to install the Agent on a Linux ECS, see "Installing and Configuring the Agent (Linux)" in *Cloud Eye User Guide*.

Table 11-1 Basic ECS metrics

Metric	LinuxECS	
None	Xen	KVM
CPU Usage	Supported	Supported
Memory Usage	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Disk Usage	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Disk Read Bandwidth	Supported	Supported
Disk Write Bandwidth	Supported	Supported

Metric	LinuxECS	
Disk Read IOPS	Supported	Supported
Disk Write IOPS	Supported	Supported
Inband Incoming Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Inband Outgoing Rate	Supported (UVP VMTools must be included in the image. Otherwise, this metric is unavailable.)	Not supported
Outband Incoming Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband incoming rate.)	Supported
Outband Outgoing Rate	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
Inbound Bandwidth	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
Outbound Bandwidth	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
Inbound PPS	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported

Metric	LinuxECS	
Outbound PPS	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported
New Connections	Supported (If UVP VMTools is included in the image, this metric is unavailable. In such a case, use the inband outgoing rate.)	Supported

Table 11-2 describes these basic ECS metrics.

The monitoring intervals for the following ECSs with raw monitoring metrics are as follows:

- Xen ECSs: 4 minutes
- KVM ECSs: 5 minutes

Table 11-2 Basic metric description

Metric ID	Parameter	Description	Value Range	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 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>Unit: Percent</p> <p>Formula: CPU usage of an ECS/Number of vCPUs in the ECS</p>	≥ 0	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 VMTools installed.</p> <p>Unit: Percent</p> <p>Formula: Used memory of an ECS/ Total memory of the ECS</p> <p>NOTE The memory usage of QingTian ECSs cannot be monitored.</p>	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	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. Unit: Percent Formula: Used capacity of an ECS-attached disk/Total capacity of the ECS-attached disk	≥ 0	ECS	5 minutes
disk_read_bytes_rate	Disk Read Bandwidth	Number of bytes read from an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes read from an ECS-attached disk/ Monitoring interval $\text{byte_out} = (\text{rd_bytes} - \text{last_rd_bytes}) / \text{Time difference}$	≥ 0	ECS	5 minutes
disk_write_bytes_rate	Disk Write Bandwidth	Number of bytes written to an ECS-attached disk per second Unit: byte/s Formula: Total number of bytes written to an ECS-attached disk/ Monitoring interval	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	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 Unit: request/s Formula: Total number of read requests sent to an ECS-attached disk/ Monitoring interval $req_out = (rd_req - last_rd_req) / \text{Time difference}$	≥ 0	ECS	5 minutes
disk_write_requests_rate	Disk Write IOPS	Number of write requests sent to an ECS-attached disk per second Unit: request/s Formula: Total number of write requests sent to an ECS-attached disk/ Monitoring interval $req_in = (wr_req - last_wr_req) / \text{Time difference}$	≥ 0	ECS	5 minutes
network_incoming_bytes_rate_inband	Inband Incoming Rate	Number of incoming bytes on an ECS per second Unit: byte/s Formula: Total number of inband incoming bytes on an ECS/ Monitoring interval	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_outgoing_bytes_rate_inband	Inband Outgoing Rate	Number of outgoing bytes on an ECS per second Unit: byte/s Formula: Total number of inband outgoing bytes on an ECS/Monitoring interval	≥ 0	ECS	5 minutes
network_incoming_bytes_aggregate_rate	Outband Incoming Rate	Number of incoming bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband incoming bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	ECS	5 minutes
network_outgoing_bytes_aggregate_rate	Outband Outgoing Rate	Number of outgoing bytes on an ECS per second on the hypervisor Unit: byte/s Formula: Total number of outband outgoing bytes on an ECS/Monitoring interval This metric is unavailable if SR-IOV is enabled.	≥ 0	ECS	5 minutes

Metric ID	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Interval (Raw Metrics and KVM Only)
network_vm_connections	Network Connections	Total number of TCP and UDP connections to an ECS Unit: count NOTE This metric is collected out-of-band and its value may be greater than the number of network connections queried in the OS.	≥ 0	ECS	5 minutes
network_vm_bandwidth_in	Inbound Bandwidth	Number of public and private bits received by the ECS per second Unit: byte/s	≥ 0	ECS	5 minutes
network_vm_bandwidth_out	Outbound Bandwidth	Number of public and private bits sent by the ECS per second Unit: byte/s	≥ 0	ECS	5 minutes
network_vm_pps_in	Inbound PPS	Number of public and private packets received by the ECS per second Unit: packet/s	≥ 0	ECS	5 minutes
network_vm_pps_out	Outbound PPS	Number of public and private packets sent by the ECS per second Unit: packet/s	≥ 0	ECS	5 minutes
network_vm_newconnections	New Connections	Number of new connections (including TCP, UDP, and ICMP) created on the ECS Unit: count	≥ 0	ECS	5 minutes

Dimensions

Key	Value
instance_id	Specifies the ECS ID.

11.3 OS Monitoring Metrics Supported by ECSs with the Agent Installed

Description

OS monitoring provides system-level, proactive, and fine-grained monitoring. It requires the Agent to be installed on the ECSs to be monitored. This section describes OS monitoring metrics reported to Cloud Eye.

OS monitoring supports metrics about the CPU, CPU load, memory, disk, disk I/O, file system, GPU, NIC, NTP, and TCP.

After the Agent is installed, you can view monitoring metrics of ECSs running different OSs. Monitoring data is collected every 1 minute.

Namespace

AGT.ECS

OS Metrics: CPU

Table 11-3 CPU metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object Unit: percent <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	ECS	1 minute
cpu_usage_idle	(Agent) Idle CPU Usage	Percentage of time that CPU is idle Unit: percent <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_user	(Agent) User Space CPU Usage	Percentage of time that the CPU is used by user space Unit: percent <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	ECS	1 minute
cpu_usage_system	(Agent) Kernel Space CPU Usage	Percentage of time that the CPU is used by kernel space Unit: percent <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_other	(Agent) Other Process CPU Usage	Percentage of time that the CPU is used by other processes Unit: percent <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	ECS	1 minute
cpu_usage_nice	(Agent) Nice Process CPU Usage	Percentage of time that the CPU is in user mode with low-priority processes which can easily be interrupted by higher-priority processes Unit: percent <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
cpu_usage_iowait	(Agent) iowait Process CPU Usage	Percentage of time that the CPU is waiting for I/O operations to complete Unit: percent <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	ECS	1 minute
cpu_usage_irq	(Agent) CPU Interrupt Time	Percentage of time that the CPU is servicing interrupts Unit: percent <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	ECS	1 minute
cpu_usage_softirq	(Agent) CPU Software Interrupt Time	Percentage of time that the CPU is servicing software interrupts Unit: percent <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) si value.Windows is not supported currently.	0-100	ECS	1 minute

OS Metric: CPU Load

Table 11-4 CPU load metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
load_aver1	(Agent) 1-Minute Load Average	CPU load averaged from the last 1 minute Linux: Obtain the metric value from the number of logic CPUs in load1/ in file /proc/loadavg . Run the top command to check the load1 value.	≥ 0	ECS	1 minute
load_aver5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes Linux: Obtain the metric value from the number of logic CPUs in load5/ in file /proc/loadavg . Run the top command to check the load5 value.	≥ 0	ECS	1 minute
load_aver15	(Agent) 15-Minute Load Average	CPU load averaged from the last 15 minutes Linux: Obtain the metric value from the number of logic CPUs in load15/ in file /proc/loadavg . Run the top command to check the load15 value.	≥ 0	ECS	1 minute

OS Metric: Memory

Table 11-5 Memory metrics

Metric	Parameter	Description	Value Range	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> <p>Unit: GB</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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_usedPercent	(Agent) Memory Usage	<p>Memory usage of the monitored object</p> <p>Unit: percent</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	ECS	1 minute
mem_free	(Agent) Idle Memory	<p>Amount of memory that is not being used</p> <p>Unit: GB</p> <ul style="list-style-type: none">Linux: Obtain the metric value from /proc/meminfo.Windows is not supported currently.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
mem_buffers	(Agent) Buffer	Amount of memory that is being used for buffers Unit: GB <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	ECS	1 minute
mem_cached	(Agent) Cache	Amount of memory that is being used for file caches Unit: GB <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	ECS	1 minute
total_open_files	(Agent) Total File Handles	Total handles used by all processes Unit: count <ul style="list-style-type: none">Linux: Use the /proc/{pid}/fd file to summarize the handles used by all processes.Windows is not supported currently.	≥ 0	ECS	1 minute

OS Metric: Disk

NOTE

- Currently, only physical disks are monitored. The NFS-attached disks cannot be monitored.
- By default, Docker-related mount points are shielded. The prefix of the mount point is as follows:
`/var/lib/docker;/mnt/paas/kubernetes;/var/lib/mesos`

Table 11-6 Disk metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_free	(Agent) Available Disk Space	Free space on the disks Unit: GB <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	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_total	(Agent) Disk Storage Capacity	<p>Total space on the disks, including used and free Unit: GB</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	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_used	(Agent) Used Disk Space	Used space on the disks Unit: GB <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 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	ECS - Mount point	1 minute

Metric	Parameter	Description	Value Range	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> <p>Unit: percent</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	ECS - Mount point	1 minute

OS Metric: Disk I/O

Table 11-7 Disk I/O metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_read_bytes_rate	(Agent) Disks Read Rate	<p>Number of bytes read from the monitored disk per second Unit: byte/s</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"> Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). 	≥ 0 bytes/s	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		<ul style="list-style-type: none"> - When the CPU usage is high, monitoring data obtaining timeout may occur and result in the failure of obtaining monitoring data. 			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_read_requests_rate	(Agent) Disks Read Requests	<p>Number of read requests sent to the monitored disk per second</p> <p>Unit: request/s</p> <ul style="list-style-type: none"> • Linux: <p>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.</p> <p>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 (~).</p> • Windows: <ul style="list-style-type: none"> - Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data. - The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). - When the CPU usage is high, monitoring data obtaining timeout 	≥ 0 requests/s	<ul style="list-style-type: none"> • EC S - Disk • EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		may occur and result in the failure of obtaining monitoring data.			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_write_bytes_rate	(Agent) Disks Write Rate	<p>Number of bytes written to the monitored disk per second</p> <p>Unit: byte/s</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. <p>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 (~).</p> <ul style="list-style-type: none">Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).When the CPU usage is high, monitoring data obtaining timeout	≥ 0 bytes/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		may occur and result in the failure of obtaining monitoring data.			

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_agt_write_requests_rate	(Agent) Disks Write Requests	<p>Number of write requests sent to the monitored disk per second</p> <p>Unit: request/s</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. <p>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 (~).</p> <ul style="list-style-type: none">Windows:<ul style="list-style-type: none">Use Win32_PerformanceData_PerfDisk_LogicalDisk object in the WMI to obtain disk I/O data.The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).When the CPU usage is high, monitoring data obtaining timeout	≥ 0 requests/s	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
		<p>may occur and result in the failure of obtaining monitoring data.</p>			
disk_readTime	(Agent) Average Read Request Time	<p>Average amount of time that read requests have waited on the disks Unit: ms/count</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	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	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> <p>Unit: ms/count</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	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_ioUtil	(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> <p>Unit: percent</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	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_queue_length	(Agent) Disk Queue Length	<p>Average number of read or write requests queued up for completion for the monitored disk in the monitoring period</p> <p>Unit: count</p> <ul style="list-style-type: none"> Linux: The average disk queue length is calculated based on the data changes in the fourteenth column of the corresponding device in file /proc/diskstats in a collection period. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~). Windows is not supported currently. 	≥ 0	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	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> <p>Unit: byte/op</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 / proc/diskstats in a collection period.<p>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 (~).</p>Windows is not supported currently.	≥ 0 bytes/op	<ul style="list-style-type: none">EC S - DiskEC S - Mount point	1 minute

Metric	Parameter	Description	Value Range	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> <p>Unit: byte/op</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 bytes/op	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	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> <p>Unit: ms/op</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	<ul style="list-style-type: none"> EC S - Disk EC S - Mount point 	1 minute

Metric	Parameter	Description	Value Range	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 ECSs do not support this metric. 	0-100	ECS - Disk	1 minute

OS Metric: File System

Table 11-8 File system metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
disk_fs_rwstate	(Agent) File System Read/Write Status	<p>Read and write status of the mounted file system of the monitored object Possible values are 0 (read and write) and 1 (read only).</p> <p>Linux: Check file system information in the fourth column in file /proc/mounts.</p>	<ul style="list-style-type: none"> 0: readable and writable 1: read-only 	ECS - Mount point	1

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
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	ECS - Mount point	1 minute
disk_inodesUsed	(Agent) Total inode Used	Number of used index nodes on the disk Linux: Run the df -i command to check the value in the IUsed column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	≥ 0	ECS - Mount point	1 minute
disk_inodesUsedPercent	(Agent) Percentage of Total inode Used	Number of used index nodes on the disk Unit: percent Linux: Run the df -i command to check the value in the IUse% column. The path of the mount point prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), periods (.), and swung dashes (~).	0-100	ECS - Mount point	1 minute

 NOTE

The Windows OS does not support the file system metrics.

OS Metric: NIC

Table 11-9 NIC metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_bitR ecv	(Agent) Outbound Bandwidth	Number of bits sent by this NIC per second Unit: bit/s <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	ECS	1 minute
net_bitS ent	(Agent) Inbound Bandwidth	Number of bits received by this NIC per second Unit: bit/s <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Unit: count/s <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 Counts/s	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Unit: count/s <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 Counts/s	ECS	1 minute
net_errin	(Agent) Receive Error Rate	Percentage of receive errors detected by this NIC per second Unit: percent <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_error	(Agent) Transmit Error Rate	Percentage of transmit errors detected by this NIC per second Unit: percent <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	ECS	1 minute
net_dropin	(Agent) Received Packet Drop Rate	Percentage of packets received by this NIC which were dropped per second Unit: percent <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	ECS	1 minute
net_dropout	(Agent) Transmitted Packet Drop Rate	Percentage of packets transmitted by this NIC which were dropped per second Unit: percent <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	ECS	1 minute

OS Metric: NTP

Table 11-10 NTP metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
ntp_offset	(Agent) NTP Offset	NTP offset of the monitored object Unit: ms Collection method for Linux ECSs: Run chronyc sources -v to obtain the offset.	≥ 0 ms	ECS	1 minute

OS Metric: TCP

Table 11-11 TCP metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_total	(Agent) TCP TOTAL	Total number of TCP connections in all states Unit: count <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_established	(Agent) TCP ESTABLISHED	<p>Number of TCP connections in ESTABLISHED state</p> <p>Unit: count</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	ECS	1 minute
net_tcp_sys_sent	(Agent) TCP SYS_SENT	<p>Number of TCP connections that are being requested by the client</p> <p>Unit: count</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	ECS	1 minute
net_tcp_sys_rcv	(Agent) TCP SYS_RECEIVED	<p>Number of pending TCP connections received by the server</p> <p>Unit: count</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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_fin_wait1	(Agent) TCP FIN_WAIT1	Number of TCP connections waiting for ACK packets when the connections are being actively closed by the client Unit: count <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	ECS	1 minute
net_tcp_fin_wait2	(Agent) TCP FIN_WAIT2	Number of TCP connections in the FIN_WAIT2 state Unit: count <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	ECS	1 minute
net_tcp_time_wait	(Agent) TCP TIME_WAIT	Number of TCP connections in TIME_WAIT state Unit: count <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_close	(Agent) TCP CLOSE	Number of closed TCP connections Unit: count <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	ECS	1 minute
net_tcp_close_wait	(Agent) TCP CLOSE_WAIT	Number of TCP connections in CLOSE_WAIT TCP state Unit: count <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	ECS	1 minute
net_tcp_last_ack	(Agent) TCP LAST_ACK	Number of TCP connections waiting for ACK packets when the connections are being passively closed by the client Unit: count <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	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
net_tcp_listen	(Agent) TCP LISTEN	Number of TCP connections in the LISTEN state Unit: count <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	ECS	1 minute
net_tcp_closing	(Agent) TCP CLOSING	Number of TCP connections to be automatically closed by the server and the client at the same time Unit: count <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	ECS	1 minute
net_tcp_retrans	(Agent) TCP Retransmission Rate	Percentage of packets that are resent Unit: percent <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	ECS	1 minute

OS Metric: GPU

Table 11-12 GPU metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_status	GPU Health Status	Overall measurement of the GPU health Unit: none <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.	<ul style="list-style-type: none">ECS - GPUECS - GPU	1 minute
gpu_usage_encoder	Encoding Usage	Encoding capability usage on the GPU Unit: percent <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	<ul style="list-style-type: none">ECS - GPUECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_usage_decoder	Decoding Usage	<p>Decoding capability usage on the GPU</p> <p>Unit: percent</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.	0-100	<ul style="list-style-type: none">ECSECS - GPU	1 minute
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> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
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> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - GPU	1 minute
gpu_aggregate_correctable	Aggregate Correctable ECC Errors	<p>Aggregate correctable ECC errors on the GPU</p> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - GPU	1 minute
gpu_aggregate_uncorrectable	Aggregate Uncorrectable ECC Errors	<p>Aggregate uncorrectable ECC Errors on the GPU</p> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_retired_page_single_bit	Retired Page Single Bit Errors	<p>Number of retired page single bit errors, which indicates the number of single-bit pages blocked by the graphics card</p> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - 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> <p>Unit: count</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	<ul style="list-style-type: none">ECSECS - GPU	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_performance_state	(Agent) Performance Status	<p>GPU performance of the monitored object</p> <p>Unit: none</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. 	<p>P0-P15, P32</p> <ul style="list-style-type: none"> P0: indicates the maximum performance status. P15: indicates the minimum performance status. P32: indicates the unknown performance status. 	<ul style="list-style-type: none"> ECS - GPU 	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
gpu_usage_memory	(Agent) GPU Memory Usage	GPU memory usage of the monitored object Unit: percent <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	<ul style="list-style-type: none"> ECS ECS - GPU 	1 minute
gpu_usage_gpu	(Agent) GPU Usage	GPU usage of the monitored object Unit: percent <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	<ul style="list-style-type: none"> ECS ECS - GPU 	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

11.4 Process Monitoring Metrics Supported by ECSs with the Agent Installed

Description

Process monitoring provides monitoring of active processes on ECSs and it requires the Agent to be installed on the ECSs to be monitored. By default, Cloud Eye

collects CPU usage, memory usage, and number of opened files of active processes.

This section describes process monitoring metrics reported to Cloud Eye.

Namespace

AGT.ECS

Process Metrics

After the agent is installed, you can view the default process metrics listed in the following table, regardless of ECS types and OSs.

Table 11-13 Process metrics

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_cpu	CPU Usage	CPU consumed by a process. pHashId (process name and process ID) is the value of md5 . Unit: percent <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/pid/stat.	0-1 x Number of vCPUs.	ECS	1 minute
proc_pHashId_mem	Memory Usage	Memory consumed by a process. pHashId (process name and process ID) is the value of md5 . Unit: percent <ul style="list-style-type: none">Linux: RSS*PAGESIZE/ MemTotal Obtain the RSS value by checking the second column of file / proc/pid/statm. Obtain the PAGESIZE value by running the getconf PAGESIZE command. Obtain the MemTotal value by checking file / proc/meminfo.	0-100%	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_pHashId_file	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.	≥0	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.	≥0	ECS	1 minute
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.	≥0	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.	≥0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object & Dimension	Monitoring Period (Raw Data)
proc_blocked_count	(Agent) Blocked Processes	Number of processes that are blocked <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.	≥0	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.	≥0	ECS	1 minute
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.	≥0	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.	≥0	<ul style="list-style-type: none">ECSECS - Process	1 minute

Dimensions

Dimension	Key	Value
ECS	instance_id	Specifies the ECS ID.

11.5 OS Monitoring Metrics Supported by ECSs with the Agent Installed and Using Simplified Monitoring Metrics

Description

This section describes the OS metrics supported by ECSs. In the following region, the agent of the latest version is used with simplified monitoring metrics:

EU-Dublin

After installing the agent on an ECS, you can view its OS monitoring metrics. Monitoring data is collected every 1 minute.

OS Monitoring Metrics

Table 11-14 OS monitoring metrics

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
cpu_usage	(Agent) CPU Usage	CPU usage of the monitored object Unit: percent <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.	0-100	ECS	1 minute
load_average5	(Agent) 5-Minute Load Average	CPU load averaged from the last 5 minutes <ul style="list-style-type: none">Linux: Obtain the metric value from the number of logic CPUs in load5/ in file /proc/loadavg. Run the top command to check the load5 value.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mem_usedPercent	(Agent) Memory Usage	Memory usage of the monitored object Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value from the /proc/meminfo file: (MemTotal - MemAvailable)/MemTotal	0-100	ECS	1 minute
mountPointPrefix_disk_free	(Agent) Available Disk Space	Free disk space Unit: GB <ul style="list-style-type: none">Linux: Run the df -h command to check the value in the Avail column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	≥ 0	ECS	1 minute
mountPointPrefix_disk_usedPercent	(Agent) Disk Usage	Percentage of total disk space that is used Unit: percent <ul style="list-style-type: none">Linux: Obtain the metric value using following formula: Disk Usage = Used Disk Space/Disk Storage Capacity. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
mountPointPrefix_disk_ioUtils and volumePrefix_disk_ioUtils	(Agent) Disk I/O Usage	<p>Percentage of the time that the disk has had I/O requests queued to the total disk operation time</p> <p>Unit: percent</p> <ul style="list-style-type: none">Linux: Obtain the metric value by calculating the data changes in the thirteenth column of the monitored object in file /proc/diskstats in a collection period. <p>The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).</p>	0-100	ECS	1 minute
mountPointPrefix_disk_inodeUsedPercent	(Agent) Percentage of Total Inode Used	<p>Number of used index nodes on the disk</p> <p>Unit: percent</p> <ul style="list-style-type: none">Linux: Run the df -i command to check the value in the IUse% column. The path of the device name prefix cannot exceed 64 characters. It must start with a letter, and contain only digits, letters, hyphens (-), dots (.), and swung dashes (~).	0-100	ECS	1 minute
net_bitSent	(Agent) Inbound Bandwidth	<p>Number of bits received by the monitored object per second</p> <p>Unit: bit/s</p> <ul style="list-style-type: none">Linux: Check metric value changes in file /proc/net/dev in a collection period.	≥ 0	ECS	1 minute

Metric	Parameter	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
net_bitRecv	(Agent) Outbound Bandwidth	Number of bits sent by the target NIC per second Unit: bit/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.	≥ 0	ECS	1 minute
net_packetRecv	(Agent) NIC Packet Receive Rate	Number of packets received by this NIC per second Unit: count/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.	≥ 0	ECS	1 minute
net_packetSent	(Agent) NIC Packet Send Rate	Number of packets sent by this NIC per second Unit: count/s <ul style="list-style-type: none">Linux: Check metric value changes in file / proc/net/dev in a collection period.	≥ 0	ECS	1 minute
net_tcp_total	(Agent) Total Number of TCP Connections	Total number of TCP connections of this NIC	≥ 0	ECS	1 minute
net_tcp_established	(Agent) Number of ESTABLISHED TCP Connections	Number of ESTABLISHED TCP connections of this NIC	≥ 0	ECS	1 minute

Dimensions

Key	Value
instance_id	Specifies the ECS ID.


11.6 Setting Alarm Rules

Scenarios

Setting ECS alarm rules allows you to customize the monitored objects and notification policies so that you can closely monitor your ECSs.

This section describes how to set ECS alarm rules, including alarm rule names, monitoring objects, monitoring metrics, alarm thresholds, monitoring intervals, and notifications.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Under **Management & Deployment**, choose **Cloud Eye**.
4. In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
5. On the **Alarm Rules** page, click **Create Alarm Rule** to create an alarm rule, or modify an existing alarm rule.

The following uses modifying an existing alarm rule as an example.

- a. Click the target alarm rule.
- b. Click **Modify** in the upper right corner of the page.
- c. On the **Modify Alarm Rule** page, set parameters as prompted.
- d. Click **Modify**.

After an alarm rule is modified, the system automatically notifies you of an alarm when the alarm complying with the alarm rule is generated.

NOTE

For more information about ECS alarm rules, see [Cloud Eye User Guide](#).

11.7 Viewing ECS Metrics

Scenarios

The cloud platform provides Cloud Eye, which monitors the running statuses of your ECSs. You can obtain the monitoring metrics of each ECS on the management console.

There a short time delay between transmission and display of monitoring data. The status of an ECS displayed on Cloud Eye is the status obtained 5 to 10

minutes before. If an ECS is just created, wait for 5 to 10 minutes to view the real-time monitoring data.

Prerequisites



- The ECS is running properly.
Cloud Eye does not display the monitoring data for a stopped, faulty, or deleted ECS. After such an ECS restarts or recovers, the monitoring data is available in Cloud Eye.

NOTE

Cloud Eye discontinues monitoring ECSs that remain in **Stopped** or **Faulty** state for 24 hours and removes them from the monitoring list. However, the alarm rules for such ECSs are not automatically deleted.

- Alarm rules have been configured in Cloud Eye for the target ECS.
The monitoring data is unavailable for the ECSs without alarm rules configured in Cloud Eye. For details, see [Setting Alarm Rules](#).
- The target ECS has been properly running for at least 10 minutes.
The monitoring data and graphics are available for a new ECS after the ECS runs for at least 10 minutes.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select your region and project.
3. Click . Under **Compute**, click **Elastic Cloud Server**.
4. In the search box above the upper right corner of the ECS list, enter the ECS name, IP address, or ID to search for the target ECS.
5. Click the name of the target ECS. The page providing details about the ECS is displayed.
6. Click the **Monitoring** tab to view the monitoring data.
7. In the ECS monitoring area, select a duration to view the monitoring data.
You can view the monitoring data of the ECS in the last 1 hour, last 3 hours, last 12 hours, last 1 day, or last 7 days.

12 Audit Using CTS

12.1 ECS Operations Supported by CTS

Scenarios

Cloud Trace Service (CTS) records user operations performed on ECSs and related resources for further query, auditing, and backtracking.

Prerequisites

CTS has been enabled.

Key ECS Operations Recorded by CTS

Table 12-1 ECS operations recorded by CTS

Operation	Resource Type	Trace
Creating an ECS	ecs	createServer createServerV2 createServerV21
Deleting an ECS	ecs	deleteServer deleteServerV2 deleteServerV21
Starting an ECS	ecs	startServer
Restarting an ECS	ecs	rebootServer
Stopping an ECS	ecs	stopServer
Adding an ECS NIC	ecs	addNic
Deleting an ECS NIC	ecs	deleteNic delNic

Operation	Resource Type	Trace
Attaching a disk	ecs	attachVolume attachVolumeV2
Attaching a disk (on the EVS console)	ecs	attachVolume2
Detaching a disk	ecs	detachVolume
Reinstalling an OS	ecs	reinstallOs
Changing an OS	ecs	changeOs
Modifying specifications	ecs	resizeServer
Enabling automatic recovery on an ECS	ecs	addAutoRecovery
Disabling automatic recovery on an ECS	ecs	deleteAutoRecovery

12.2 Viewing Traces


Scenarios



After you enable CTS and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, the system starts recording operations on data in OBS buckets. CTS stores operation records generated in the last seven days.

This section describes how to query and export operation records of the last seven days on the CTS console.

- [Viewing Real-Time Traces in the Trace List](#)

Viewing Real-Time Traces in the Trace List

1. Log in to the management console.
2. Click  in the upper left corner and choose **Management & Deployment > Cloud Trace Service**. The CTS console is displayed.
3. Choose **Trace List** in the navigation pane on the left.
4. 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: You can query traces generated during any time range in the last seven days.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5000 records.
5. Click **Query**.
 6. 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 5000 records.
 - Click  to view the latest information about traces.
 7. 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: [redacted]

code: 200

trace_name: createDockerConfig

resource_type: dockerlogincmd

trace_rating: normal

api_version: ""

message: createDockerConfig, Method: POST Url=/v2/management/secret, Reason:

source_ip: [redacted]

domain_id: [redacted]

trace_type: ApiCall

8. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace ×

```
{
  "request": "",
  "trace_id": "[redacted]",
  "code": "200",
  "trace_name": "createDockerConfig",
  "resource_type": "dockerlogincmd",
  "trace_rating": "normal",
  "api_version": "",
  "message": "createDockerConfig, Method: POST Url=/v2/management/secret, Reason:",
  "source_ip": "[redacted]",
  "domain_id": "[redacted]",
  "trace_type": "ApiCall",
  "service_type": "SWR",
  "event_type": "system",
  "project_id": "[redacted]",
  "response": "",
  "resource_id": "",
  "tracker_name": "system",
  "time": "Nov 16, 2023 10:54:04 GMT+08:00",
  "resource_name": "dockerlogincmd",
  "user": {
    "domain": {
      "name": "[redacted]",
      "id": "[redacted]"
    }
  }
}
```

9. For details about key fields in the trace structure, see section "Trace References" > "Trace Structure" and section "Trace References" > "Example Traces" in the *CTS User Guide*.