

Elastic Compute Service

User Guide

User Guide

Quick reference

This article provides a quick reference for you on the use of ECS on the ECS console.

Must read: ECS operation instructions

- ECS operation instructions

Log on to an instance

How to use a username and password to log on to a Linux instance?

How to use an SSH key pair to log on to a Linux instance?

How to log on to a Windows instance?

If you forget your instance logon password (not the Management Terminal password), you can reset the password.

Operate disks

- How to attach a data disk after you purchase a cloud disk?

Change the operating system

You can change the operating system, such as:

From Windows to Linux, or from Linux to Windows.

From one version to another, for example, from Windows Server 2008 to Windows Server

2012.

Change the image, for example, change to custom images, or shared images.

Use images and snapshots

How to copy images across different regions?

How to define automatic snapshot policies when you want an automatic update policy for configurations or applications?

Enable intranet communication

- How to use security groups to enable intranet communication?

For the answers to the preceding questions, click the corresponding node in the left-side navigation pane.

To guarantee proper operation of your ECS instance, you must take the considerations outlined in this section into account before use.

Prohibitions

Alibaba Cloud prohibits you from:

- Using your instance for flow-through services. Any violation results in punishment up to shutdown and lockout of instance, and termination of services.
- Activating SELinux.
- Uninstalling hardware related drivers.
- Arbitrarily modifying the MAC address of the network adapter.

General operating system considerations

- For an ECS with more than 4 GB RAM, we recommend that you use a 64-bit OS, because a 32-bit OS supports a maximum of 4 GB RAM. Currently available 64-bit systems include:
 - Aliyun Linux
 - CoreOS
 - CentOS
 - Debian

- FreeBSD
 - OpenSUSE
 - SUSE Linux
 - Ubuntu
 - Windows
- 32-bit Windows OS supports CPUs with up to 4 cores.
 - A minimum of 2 GB RAM is needed for building a website on a Windows instance, and an instance type with 1 vCPU core and 1 GB RAM cannot be used for MySQL service.
 - To guarantee service continuity and avoid service downtime, you must enable auto-start of service applications upon OS boot.
 - For I/O-optimized instances, do not stop the **aliyun-service** process.
 - We do not recommend that you update the kernel and the OS.

Linux restrictions

To guarantee stable system operation, **DO NOT**:

- Modify the content of the default `/etc/issue` file. Modifying this file renders management console buttons unusable.
- Modify directory permissions in partitions, particularly permissions for directories such as `/etc`, `/sbin`, `/bin`, `/boot`, `/dev`, `/usr`, and `/lib`. Improper modification of permissions may cause errors.
- Rename, delete, or disable the Linux **root** account.
- Compile or perform any other operations on the Linux kernel.
- Enable the **NetWorkManager** service. This service conflicts with the internal network service of the system and causes network errors.

Windows restrictions

- Do not close the built-in **shutdownmon.exe** process, which may delay the restart of your Windows server.
- Do not rename, delete, or disable the **Administrator** account.
- We do not recommend that you use virtual memory.

ECS does not support:

- Sound card applications.
- The installation of external hardware devices such as hardware dongles, USB drives, external hard drives, and the USB security keys issued by banks.
- SNAT and other IP packet address translation services. Achieve this using an external VPN or proxy.
- Multicast protocol. If multicasting services are required, unicast point-to-point method is recommended.

- Virtual application installation or subsequent virtualization such as when using VMware.
 Besides the preceding limits, using ECS has the following limits. Unless otherwise specified, all the resource limits listed in the table are for one region.

Restricted item	Parameter	To apply for an exception or unlock configuration rights
User creation of ECS resources	N/A	Users must undergo real-name registration
Default Pay-As-You-Go instance quota in all regions	30	Request increase by ticket
Disk quota within a single instance	17, including one system disk and 16 data disks	No higher configurations exist
Support for Pay-As-You-Go ephemeral disks	Supported	For APIs, open a ticket. Other types do not support this disk type
Capacity of a single ephemeral disk	20 to 1,024 GB	No higher configurations exist
Total ephemeral disk size for a single instance	2,048 GB	No higher configurations exist
Number of snapshots	(Number of disks)*64	No higher configurations exist
Capacity of a single basic cloud disk	5 to 2,000 GB	No higher configurations exist
Number of accounts to share a single custom image with	50	No higher configurations exist
List of available public images	List of images for sale on the official website	Common users cannot change; for others, open a ticket to add other images
Number of Elastic Network Interfaces (ENI)	100	No higher configurations exist
Available inbound bandwidth for Internet access	Up to 200 Mbps	No higher configurations exist
Available outbound bandwidth for Internet access	Up to 200 Mbps	No higher configurations exist
Number of instances allowed in a single security group	1,000	No higher configurations exist
Number of authorization rules for a single security group	100	No higher configurations exist
Security group quota	100	Request increase by ticket
Maximum number of security groups that a single instance can belong to	5	No higher configurations exist

Restrictions on image and instance types	Instances with 4 GB or more of memory cannot use 32-bit images	N/A
Adding new disks for ephemeral disk instances	Not supported	N/A
Changing the configuration of an instance with ephemeral disks	Bandwidth is adjustable	No exceptions (cloud disks can be used for attachment)
Relationship between the system disk and data disk	If the system disk is a cloud disk, all data disks must be cloud disks	N/A
Total number of Pay-As-You-Go cloud disks that can be purchased	ECS Instance Quota * 5	Open a ticket
Basic cloud disk capacity	5 to 2,000 GB	No higher configurations exist
User restrictions for the creation of Pay-As-You-Go cloud disks	Users must pass real-name registration (for buy only)	N/A
Range of system disk attaching points	/dev/xvda	N/A
Range of data disk attaching points	/dev/xvd[b-z]	N/A
Capacity of a single SSD Cloud Disk	20 to 2,048 GB	No higher configurations exist
Capacity of a single ultra Cloud Disk	20 to 2,048 GB	No higher configurations exist
Zones in which users may create instances	1 online zone	Request increase by ticket
Zones in which users may create disks	The zone combinations in which users may create instances and the zones where instances remain after overlap removal	N/A
Default Pay-As-You-Go instance types	ecs.t1.small (single-core 1 GB)	Request change by ticket
	ecs.s1.small (single-core 2 GB)	
	ecs.s1.medium (single-core 4 GB)	
	ecs.s2.small (dual-core 2 GB)	
	ecs.s2.large (dual-core 4 GB)	
	ecs.s2.xlarge (dual-core 8 GB)	
	ecs.s3.medium (quad-core 4 GB)	

	GB)	
	ecs.s3.large (quad-core 8 GB)	
	ecs.m1.medium (quad-core 16 GB)	

For limits of VPC, see **Limits** in *VPC Product Introduction*.

Connect

You can use the Management Terminal, also called VNC, to connect to an ECS instance, especially when the remote access software program that you are using, such as PuTTY, Xshell, or SecureCRT, cannot work.

Usage scenarios

The Management Terminal is used to:

- Check the status of an ECS instance if boot speed is slow.
- Reconfigure the firewall if a remote connection fails due to a software error within the ECS instance.
- End abnormal processes consuming excessive CPU usage or bandwidth.

Note:The Management Terminal can be used to connect to an instance even if purchased bandwidth is insufficient.

Procedure

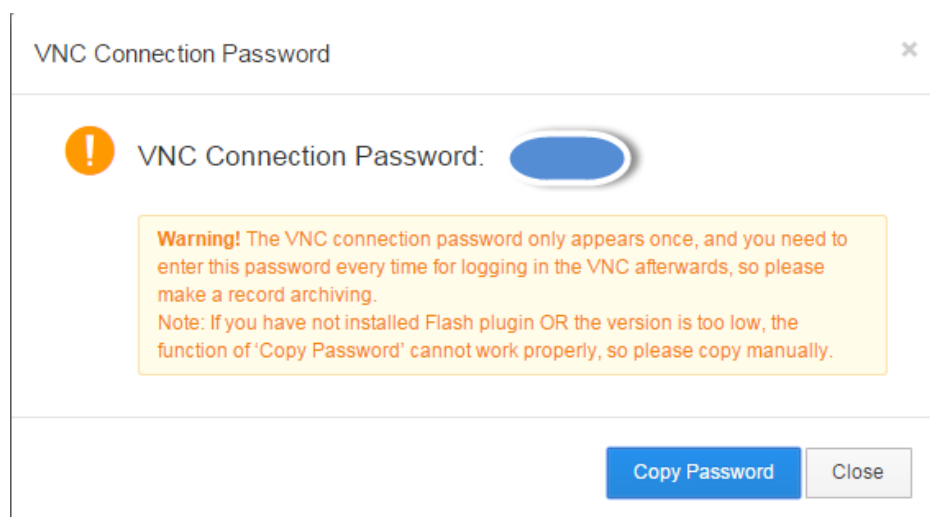
Log on to the ECS console.

Go to the ECS instance to connect to, and click **Connect**.

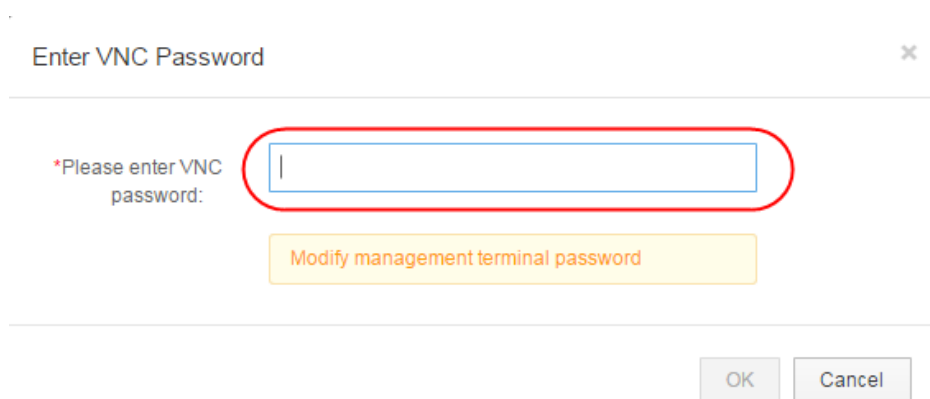
Follow the tips below to connect to the Management Terminal:

- If you connect the Management Terminal for the first time, follow the steps below:
 - a. On the **VNC Connection Password** dialog, copy the password. This dialog appears only once, but you need to enter the connection password each time you want to connect to the Management Terminal, so **write down**

the password.



- b. Click the **Close** button to close the **VNC Connection Password** dialog.
- c. On the **Enter VNC Password** dialog, paste the connection password that you copied, and then click the **OK** button to connect to the Management Terminal.



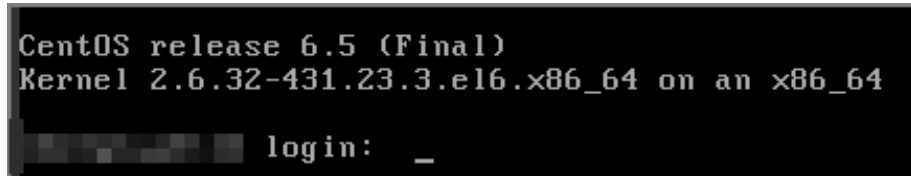
- If this is not your initial connection to the Management Terminal, the **Enter VNC Password** dialog appears, and you need to enter the connection password and click the **OK** button to connect to the Management Terminal.
- If you forget the password, you can follow the steps below to connect to the Management Terminal:
 - a. Change password.
 - b. On the upper left corner of the Management Terminal interface, click **Send remote command > Connect to management terminal**.
 - c. On the **Enter VNC Password** dialog, enter the new password to finish connection.

Follow the steps below to connect to an instance:

For a Linux instance, enter the user name ("root") and the password to connect

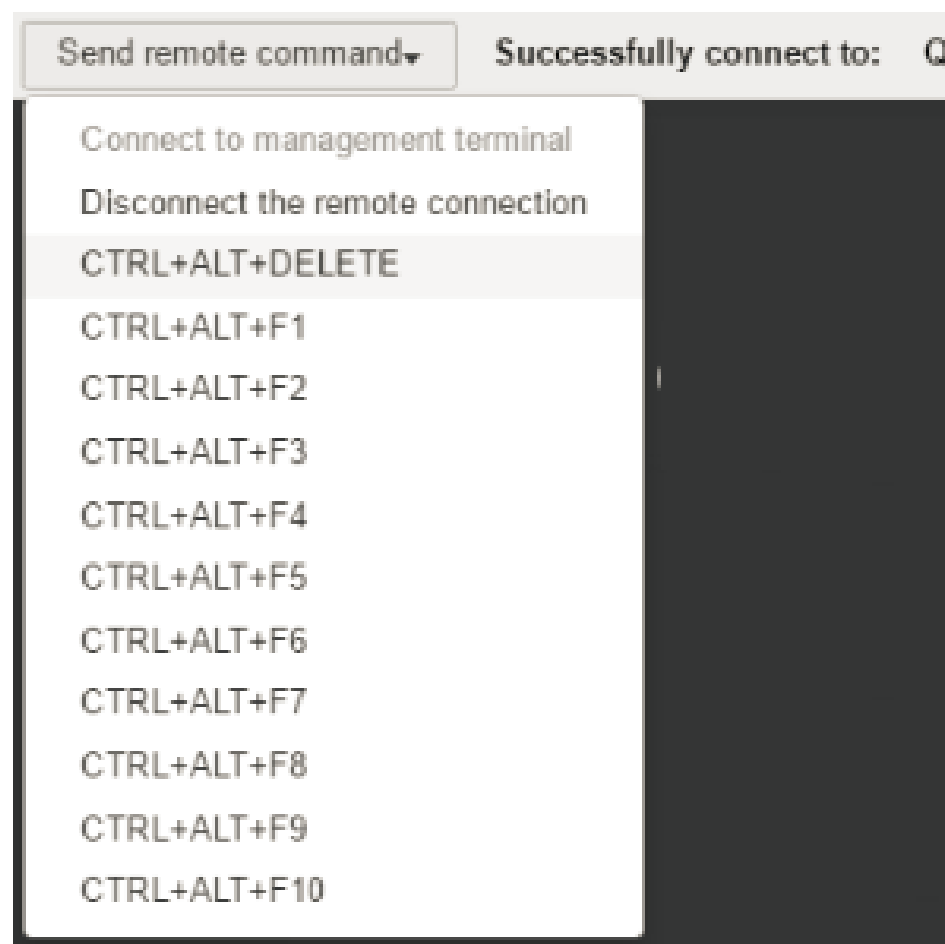
to it. Your screen may go black constantly, which occurs when the Linux instance is in sleep mode. Click the mouse or press any key to wake it up.

If you are operating several Linux instances, you can click **Send remote command > CTRL+ALT+Fx**, of which **Fx** can be any one from **F1** to **F10**, to switch management terminals.



```
CentOS release 6.5 (Final)
Kernel 2.6.32-431.23.3.el6.x86_64 on an x86_64
login: _
```

For a Windows instance, on the upper left corner of the Management Terminal interface, click **Send remote command > CTRL+ALT+DELETE** to get to the logon screen. Enter the user name, Administrator, and password to log on.



Change password

If you prefer a password that you are familiar with rather than the password displayed on the **VNC**


Connection Password dialog, or if you forget your password, you can change the connection password.

Note: If the instance that you are connecting to is not I/O optimized, you need restart your instance to make the new VNC connection password take effect after you change it. The restart operation will stop the work of your instance and interrupt your business. So you must be cautious to change the password.

Log on to the ECS console.

Go to the ECS instance to connect to, and click **VNC** on the right.

Close the **VNC Connection Password** dialog or **Enter VNC Password** dialog, and on the upper right corner of the Management Terminal interface, click the **Modify management terminal password** button to start changing the password.



Enter a new password, which must be of 6-character length, composed of uppercase letters, lowercase letters, digits, or a combination of them, but not special characters.

Make the new password take effect:

- If the instance that you are connecting to is I/O optimized, the new password takes effect immediately.
- If the instance that you are connecting to is not I/O optimized, **restart the instance** through the Management Console for the new password to take effect. Restarting within the instance does not work.

FAQ

Q: Can multiple users be simultaneously connected to the Management Terminal?

A: No. Only one user can be connected to the Management Terminal at any given time.

Q: Why can I not connect to an instance using the Management Terminal after changing the password?

A: Ensure you are entering the right **VNC connection password**. If the instance that you are connecting to is not I/O optimized, you must **restart the instance** through the Management Console for the new **VNC connection password** to take effect. Restarting directly within the instance will not make the password take effect.

Q: Why do I see a black screen after logging on to my instance?

A: A black screen indicates that the instance is in sleep mode.

- For a Linux instance, press any key to wake it up.
- For a Windows instance, click **Send remote command** > **Ctrl+ALT+DELETE** to bring back the logon interface.

Q: Why can I not access the Management Terminal?

A: To address login issues, open your browser and connect to the Management Terminal. Press the F12 key to open the developer tool. The Management Terminal information can then be analyzed to locate faults under the Console tab.

Q: I cannot use IE8.0 or Firefox to open the Management Terminal. How can I resolve this?

A: Only IE10 or higher is supported, and only certain versions of Firefox are supported.

To resolve this issue, update or change your browser to a compatible version.

Note: Google Chrome offers the best support for the Management Terminal function, and is recommended for use when you connect to the Management Terminal.

How to use a key pair to log on to a Linux instance depends on the local operating system.

- Windows OS
- Linux OS or other systems supporting SSH commands

Windows OS

In this section, it is demonstrated how to use a key pair to log on to a Linux instance on a Windows system, using the popular SSH tools PuTTY and PuTTYgen as an example.

Prerequisites

- PuTTY and PuTTYgen must have been installed. You can download them at:
 - PuTTY
 - PuTTYgen
- You must have a Linux instance that has been bound to an instance. You can allocate an SSH key pair when creating an instance or bind an SSH key pair to an instance.
- Add the following rule in the security group to enable the access to the TCP Port 22 of the instance. For more information, see [Add a security group rule](#).

Network Type	NIC	Rule Direction	Authorization Policy	Protocol Type	Port Range	Authorization Type	Authorization Object	Priority
VPC	N/A	Inbound	Allow	SSH(22)	22/22	Address Field Access	0.0.0.0/0	1
Classic	Internet							


Procedure

To connect to a Linux instance by using an SSH key pair on a Windows system, follow these steps:

(Optional) If you are using a key pair generated by Alibaba Cloud, of which the private key is a .pem file, you must convert it to a .ppk file. If your private key is a .ppk file, you can skip this step.

- i. Start PuTTYgen. In this example, we use PuTTYgen version 0.68.

Under the **Type of key to generate** option, select **RSA**. The value of **Number of bits in a generated key** can be left as is. The software automatically update the value based on the imported private key information.



Click **Load**. By default, PuTTYgen only displays files with an extension of .ppk. To find your .pem file, select to display **All Files (*.*)**.



Select the downloaded private key file from Alibaba Cloud, or the ready private key file, and click **Open**.

Click **OK** to close the confirmation dialog box.

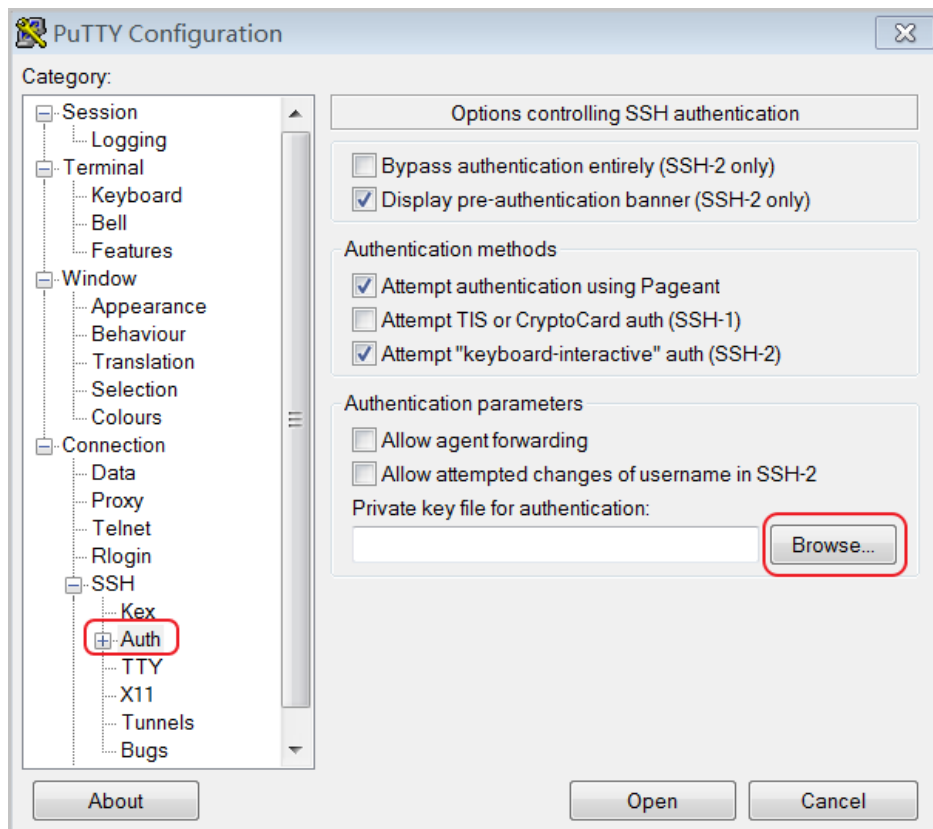
Click **Save private key**. PuTTYgen displays warning about saving the key without a password. Click **Yes**.

Specify the same name for the private key with the key pair, and save the settings. PuTTY automatically adds the .ppk file.

Start PuTTY.


Click **Connection > SSH > Auth**.

Click **Browse...** and select the .ppk file generated in Step 1.



Click **Session**.

- In **Host Name (or IP address)**, enter your account and the public IP address of the instance to be connected to. The format is root@IP address.
- In **Port**, enter the port number **22**.
- For **Connection type**, select **SSH**.



Click **Open** to start accessing your Linux instance.

When the window shows *Connection established*, it indicates you have successfully logged on to the instance using the key pair.

Linux OS or other systems supporting SSH commands

In this section, it is demonstrated how to use a key pair to log on to a Linux instance on a Linux system or a system supporting SSH commands, such as MobaXterm for Windows.

Prerequisites

You must have a Linux instance that has been bound to an instance. You can allocate an SSH key pair when creating an instance or bind an SSH key pair to an instance.

Procedure

To connect to a Linux instance by using an SSH key pair on Linux or Unix-like system, follow these steps:

Locate directory of your private key, for example, `/root/xxx.pem`.

xxx.pem is the private key file.

To modify the attributes of the private key, run the command: `chmod 400 [directory of the private key file]`. For example, `chmod 400 /root/xxx.pem`.

To connect to the instance, run the command `ssh -i [directory of the private key file] root@Internet IP address`. For example, `ssh -i /root/xxx.pem root@10.10.10.100`.

Connect to a Linux instance

The utilities used to remotely connect to Linux ECS instances vary based on the local OS as follows:

- For a Linux OS, use Secure Shell (SSH) Command Line.
- For a Windows OS, use either Management Terminal or use SSH Command Line through PuTTY or other SSH clients.
- For a Mac OS, use Management Terminal or SSH Command Line.
- For iPhone, use SSH Control Lite.
- For Android, use SSH Control Lite.

Connect to a Linux instance using Windows OS

On a Windows system, you can connect to a Linux instance using either of the following methods:

Remote access software

This method is available only if you purchase bandwidth when creating your instance. Before using this method, make sure the instance can be accessed through the Internet.

Management Terminal (VNC)

Connection through the Management Terminal (VNC) can be completed disregarding whether bandwidth has been purchased.

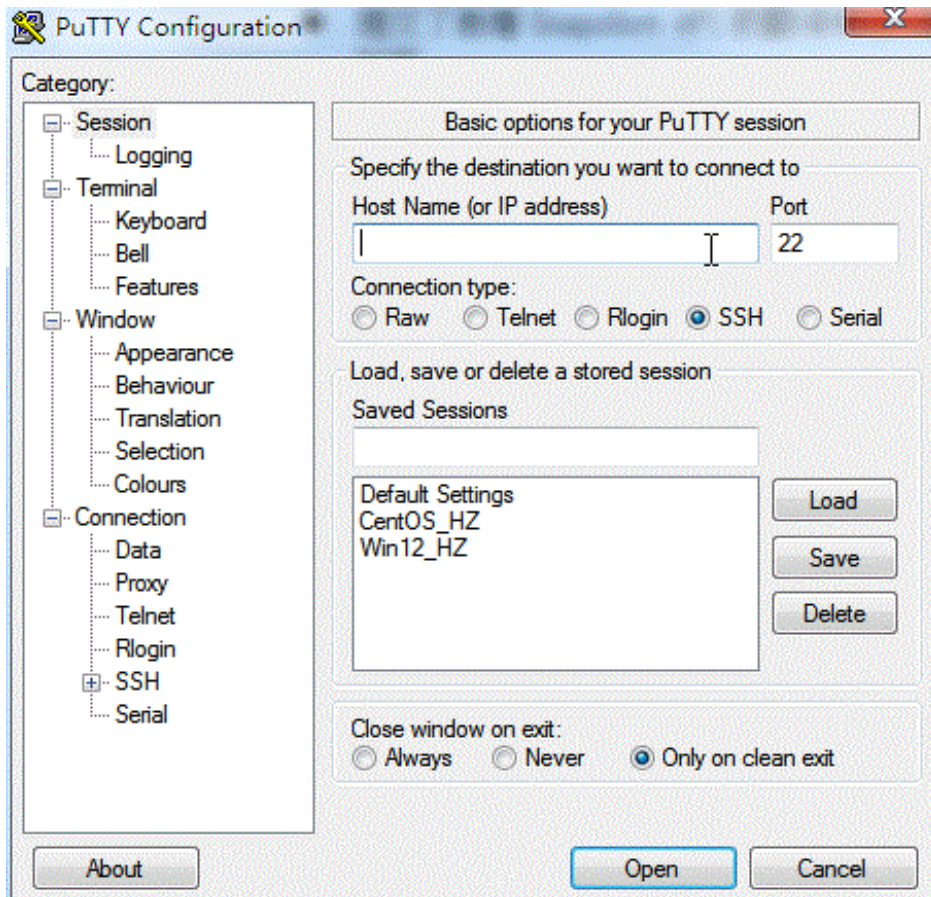
Use a remote connection Application

This section uses PuTTY as an example. PuTTY can be downloaded at <http://www.chiark.greenend.org.uk/~sgtatham/putty/>.

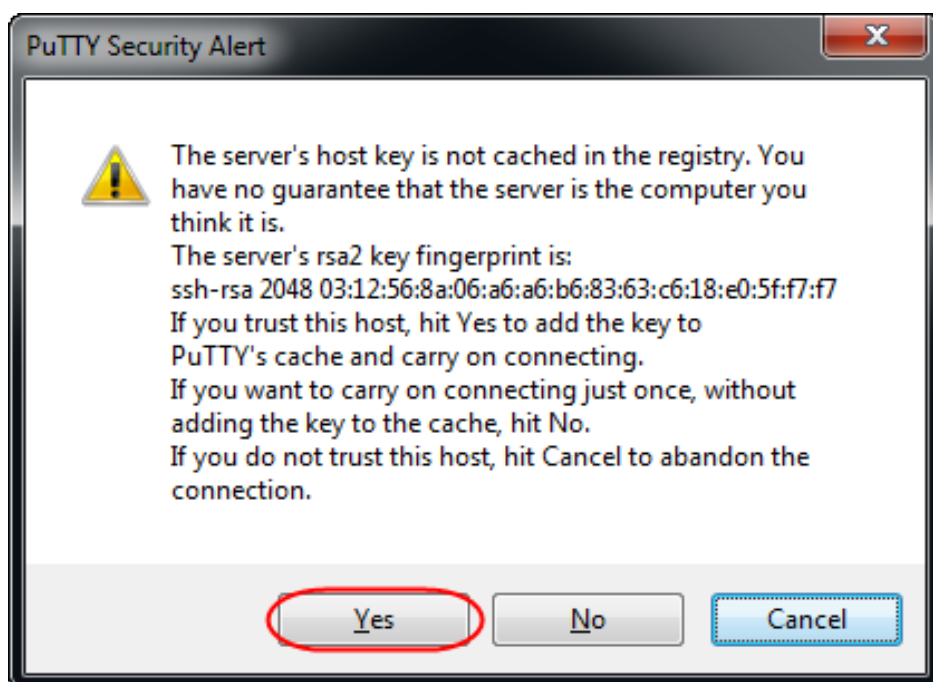
You can connect to a Linux instance by PuTTY as follows:

1. Start Putty.exe.


2. Enter the public IP address of the instance in **Host Name (or IP address)**.
3. Use the default port **22**.
4. Select **SSH** as **Connection Type**.
5. Type a session name in **Saved Sessions**, and then click **Save**. In later logins, you may directly load the session without re-entering the IP address.
6. Click **Open** to connect.



7. Upon first connection, the following dialog box is displayed. Click **Yes**.



8. As prompted, enter the username and password for the Linux ECS instance. The password is not displayed on-screen. Press the **Enter** key to complete connection to the instance.



When you connect your computer to the Linux instance successfully, you can operate the instance from your computer.

Use Management Terminal to connect to an ECS instance

See [Use Management Terminal \(VNC\) to connect to an ECS instance](#).

Connect to a Linux instance using Linux OS or Mac OS X

1. Connect to the instance using SSH commands. For example, `ssh root@Instance'` s public IP address.
2. Enter the root user password.

Connect to a Linux instance using an mobile app

You can connect to an instance by a remote desktop application installed on your smart phone. For

example, iPhone users can download **SSH Control Lite** from the App Store and use it to connect to Linux instances.

What if I forget my logon password?

If you forget your instance logon password (not the Management Terminal password), reset the logon password. For more information, see [Reset an instance password](#).

Connect to a Windows instance

The utilities used to remotely connect to Linux ECS instances vary based on local OS as follows:

- For a Linux OS, use rdesktop.
- For a Windows OS, use Management Terminal or Microsoft Terminal Services Client (MSTSC).
- For a Mac OS, use Management Terminal or MSTSC.
- For iPhone, use the Microsoft Remote Desktop app.
- For Android, use the Microsoft Remote Desktop app.

Connect to a Windows instance using Windows OS

Using a local Windows OS, connect to a Windows instance using one of the following:

Microsoft Terminal Services Client (MSTSC)

This method is available only if you purchased bandwidth when creating your instance. Prior to using this method, ensure the instance can be accessed through the Internet.

Management Terminal (VNC)


Connection through the Management Terminal can be completed disregarding whether bandwidth has been purchased.

Use MSTSC

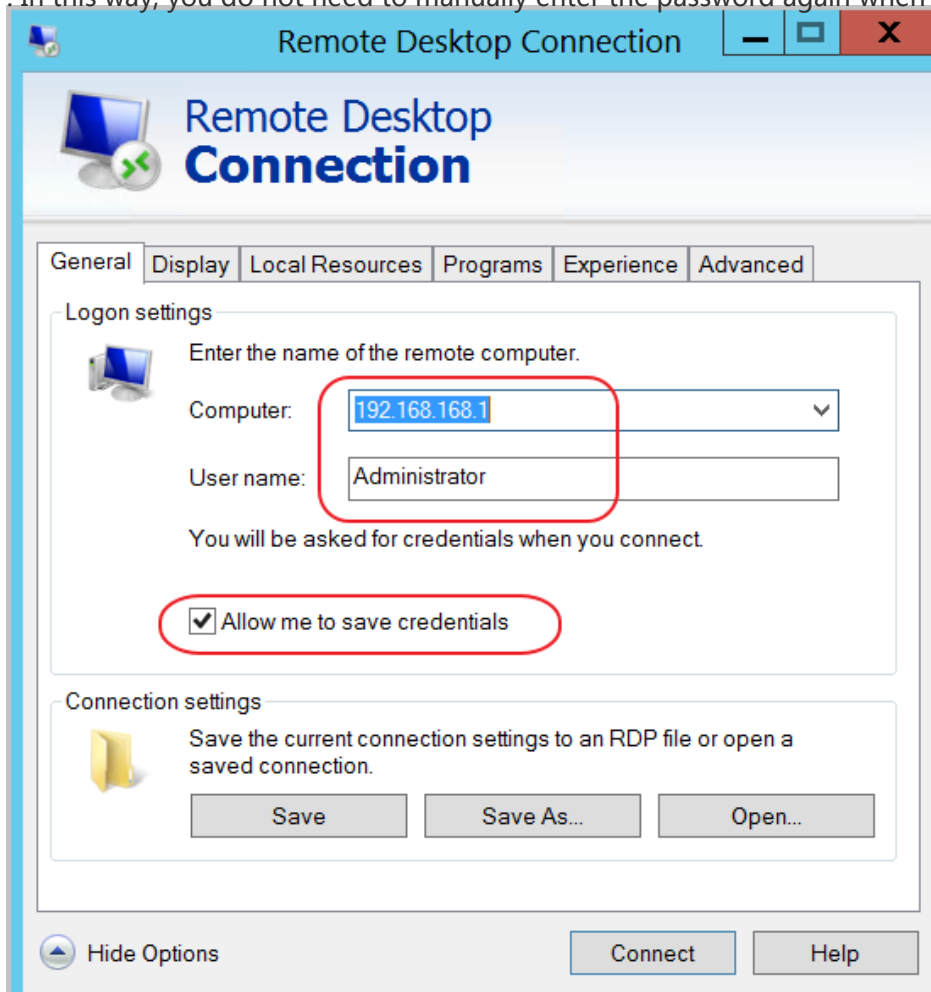
Perform one of the following to start Remote Desktop Connection:

- Click **Start > Remote Desktop Connection**.
- Click the **Start** icon and enter **mstsc** in the search box.
- Press the shortcut key **Windows Logo + R** to open the **Run** window, enter **mstsc**, and then press the **Enter** key.

In the **Remote Desktop Connection** dialog box, enter the public IP address of the instance. Click **Show Options**.

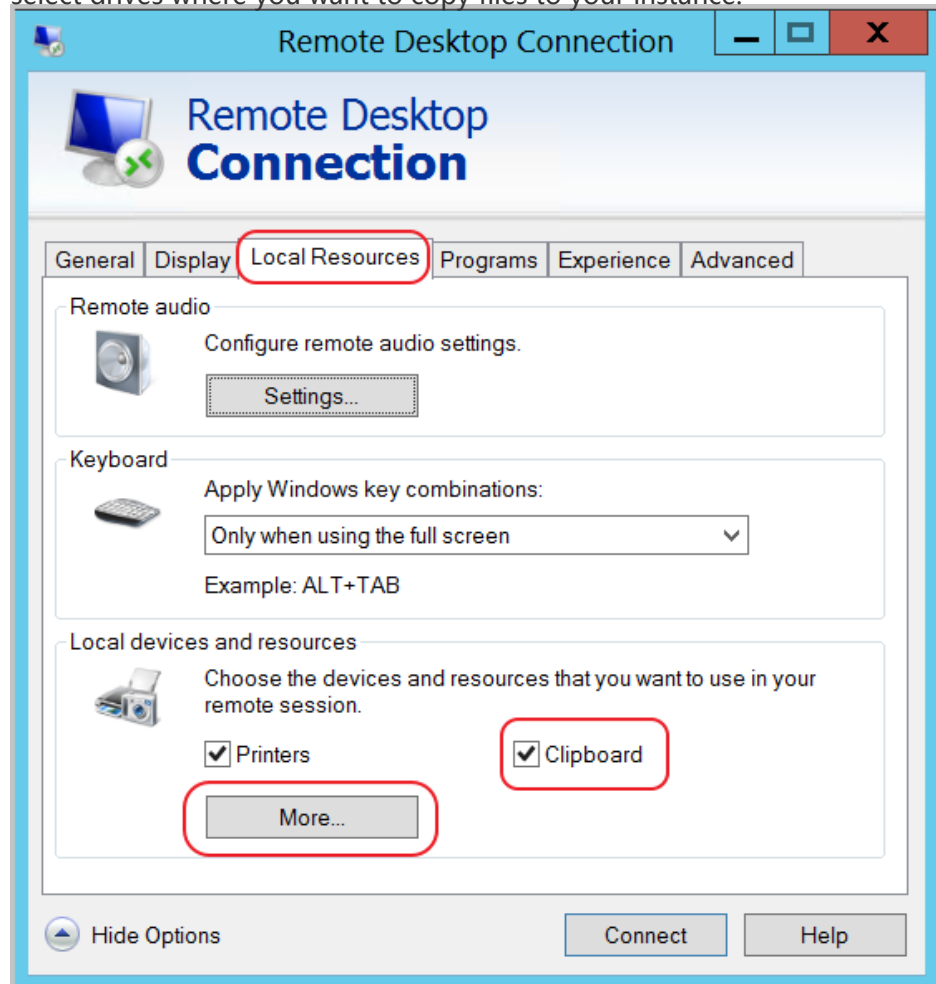


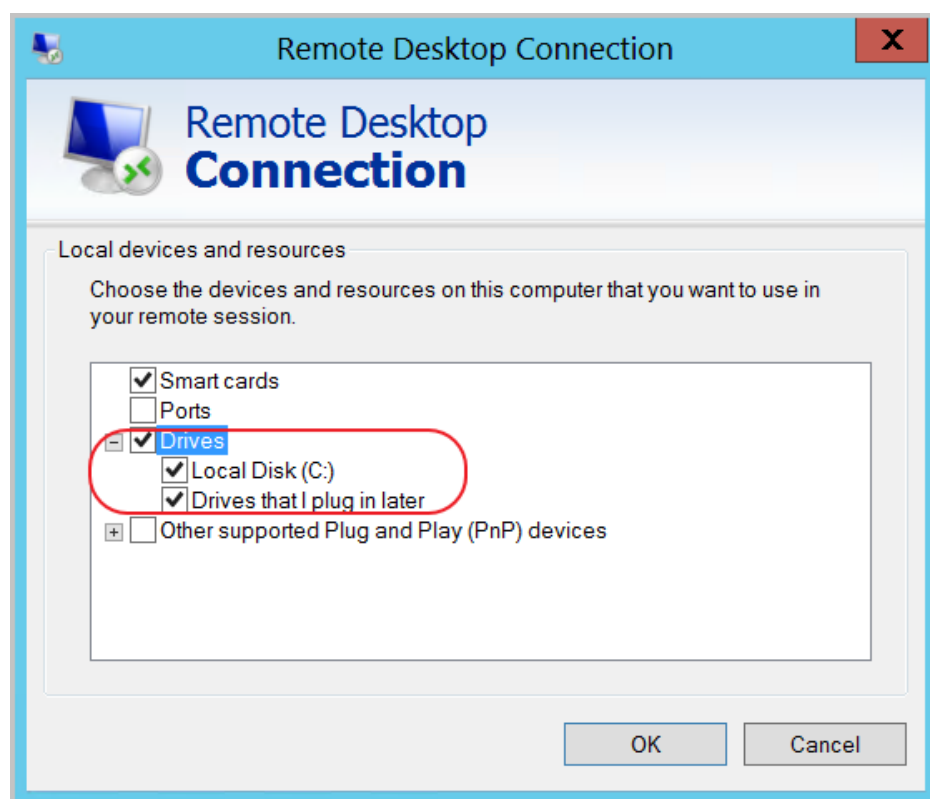
Enter the user name. The default value is **Administrator**. Check **Allow me to save credentials**. In this way, you do not need to manually enter the password again when you log on later.



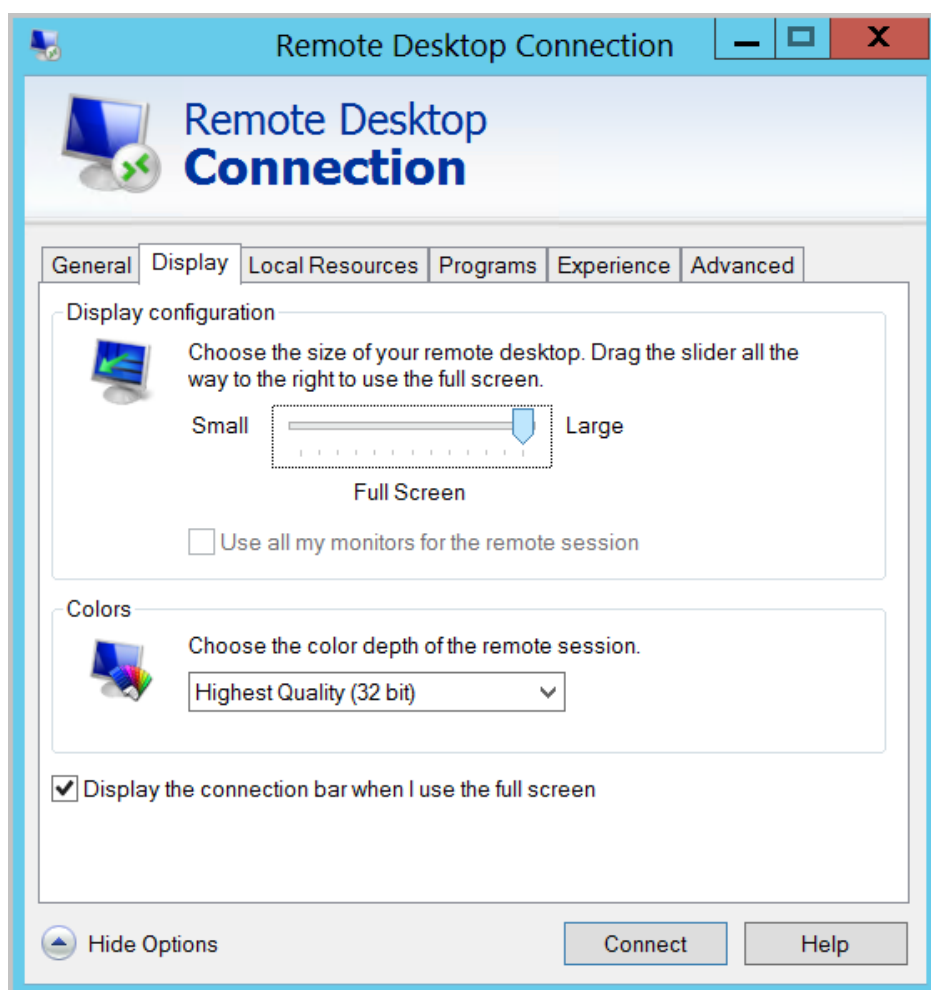
(Optional) If you want to copy text or files from your computer to the instance, click the **Local Resources** tab to see options for sharing local computer resources.

- If you need to copy text only, check **Clipboard**.
- If you need to copy files from your computer as well, click **More > Drives**, and select drives where you want to copy files to your instance.





(Optional) Open the **Display** tab to resize the remote desktop window. **Full Screen** is recommended.



Click the **Connect** button to complete connection to the instance.

When you connect your computer to the Windows instance successfully, you can operate the instance from your computer.

Use Management Terminal (VNC) to connect to an ECS instance

Refer to [Use Management Terminal \(VNC\) to connect to an ECS instance](#).

Connect to a Windows instance on a Linux OS

If you are connecting to a Windows Instance on a Linux OS, you can use either remote access software or the Management Terminal.

If no bandwidth has been purchased, you must log on to the Alibaba Cloud Console and use the Management Terminal to connect to the instance.

Use a remote connection application

This section uses the rdesktop application as an example. rdesktop can be downloaded at <http://www.rdesktop.org/>. To connect to a Windows instance using rdesktop, perform the following:

1. Start rdesktop.
2. Enter the following command (using your parameters):

```
rdesktop -u administrator -p password -f -g 1024*720 192.168.1.1 -r clipboard:PRIMARYCLIPBOARD -r disk:sunray=/home/yz16184
```

The following table defines the parameters used in the rdesktop command.

Parameter	Description
-u	Indicates username. For a Windows instance, the default username is Administrator.
-p	Indicates the Windows instance login password.
-f	Indicates full screen is the default view. Use the key combination Ctrl+Alt+Enter to exit full screen mode.
-g	Indicates the resolution. If the connector "*" is omitted, the default resolution will be the native resolution in full screen.
IPADDRESS	Enter the IP address of your Windows instance
-d	Indicates the domain name. For example, INC domains will use -d INC.
-r	Indicates a multimedia redirection. For example: To enable sound, use -r sound. To use the local sound card, use -r sound : local. To enable a Udisk, use -r disk:usb=/mnt/usbdevice
-r clipboard:PRIMARYCLIPBOARD	This parameter can be used to directly copy and paste text between the local Linux system and the remote Windows instance. Chinese characters are also supported.
-r disk:sunray=/home/yz16184	This indicates that a directory on the local Linux system is mapped onto the Windows hard disk. This allows you to transfer files without relying on Samba or FTP.

Use the Management Terminal (VNC) to connect to an ECS instance

The operation procedure is the same as that from a local Windows OS.

Connect to a Windows instance using Mac OS X

Download and install the remote desktop client for Mac OS X.

Follow the in-app directions to complete login.

Connect to a Windows instance using a mobile app

Download and install **Microsoft Remote Desktop** from the iTunes App Store. Follow the in-app directions to complete login.

What if I forget my logon password?

If you forget your instance logon password (not the VNC connection password), see [Reset the password](#).

This document describes how to connect to an ECS instance on a mobile device. The procedure varies with the operating system of your instance.

Connect to a Linux instance: We take SSH Control Lite as an example to describe how to connect to a Linux instance on an iOS device, and JuiceSSH to describe how to connect to a Linux instance on an Android device.

Connect to a Windows instance: We take Microsoft Remote Desktop as an example to describe how to connect to a Windows instance on an iOS or Android device.

Connect to a Linux instance

Prerequisites

Confirm the following before connecting to your instance:

- The instance is **Running**.
- The instance has a public IP address and is accessible from public network.
- You have set the logon password for the instance. If the password is lost, you must **reset the instance password**.

- The security group of the instance has the following security group rules:

Network type	NIC	Rule Direction	Authorization Policy	Protocol Type	Port Range	Authorization Type	Authorization Object	Priority
VPC	No configuration required	Inbound	Allow	SSH(22)	22/22	Address Field Access	0.0.0.0/0	1
Classic	Internet							

- You have downloaded and installed the appropriate app:

- The iOS device has SSH Control Lite installed.
- The Android device has JuiceSSH installed.

Procedure

For iOS devices, see [Use SSH Control Lite to connect to a Linux instance](#). In this example, user name and password are used for authentication.

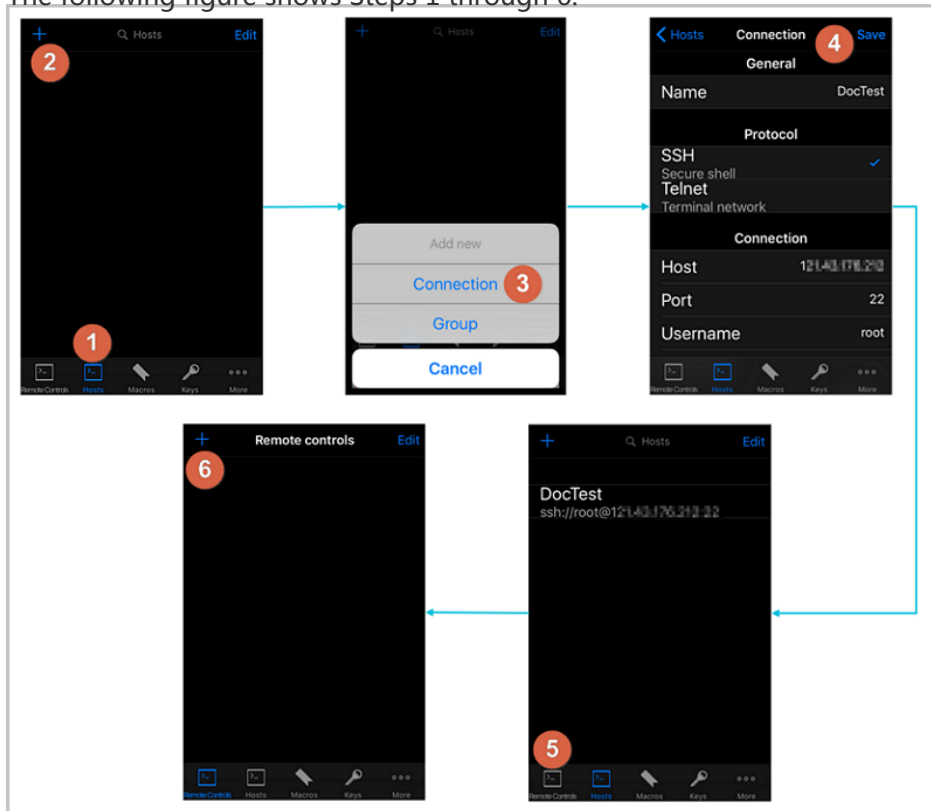
For Android devices, see [Use JuiceSSH to connect to a Linux instance](#). In this example, user name and password are used for the authentication.

Use SSH Control Lite to connect to a Linux instance

To connect to a Linux instance by using SSH Control Lite, follow these steps:

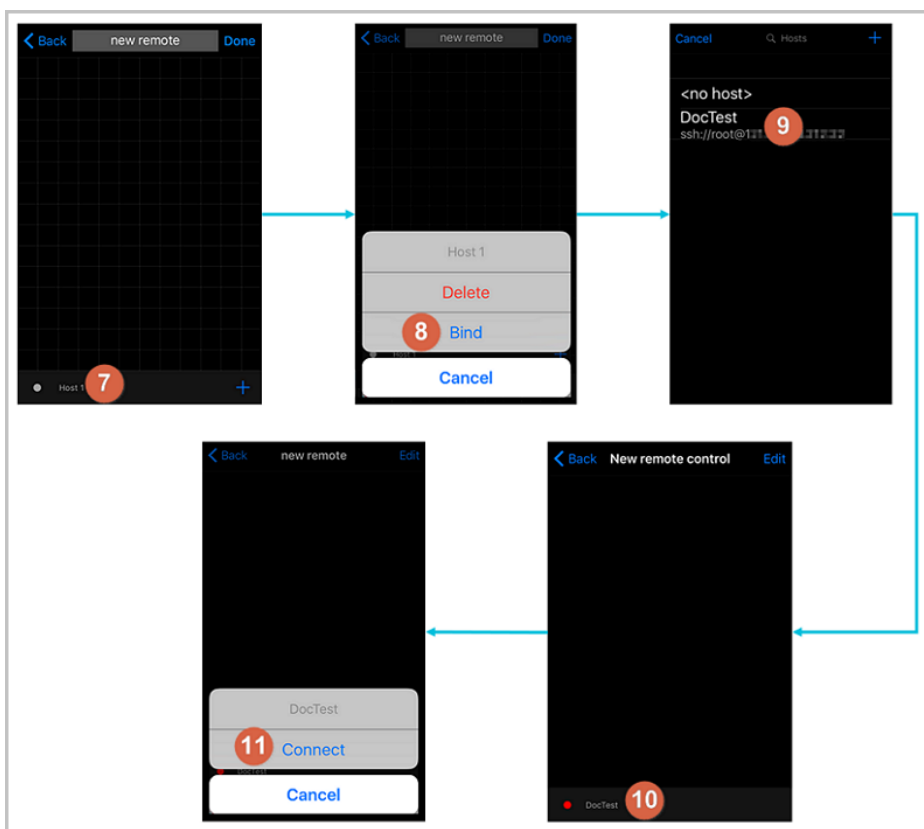
1. Start **SSH Control Lite**, and tap **Hosts**.
2. Tap the + icon in the upper left corner of the **Hosts** page.
3. In the action sheet, tap **Connection**.
4. On the **Connection** page, set the connection information and tap **Save**. The following connection information is required:
 - **Name**: Specify the Host name. *DocTest* is used in this example.
 - **Protocol**: Use the default value **SSH**.
 - **Host**: Type the public IP address of the Linux instance to connect to.
 - **Port**: Type the port number for SSH protocol. **22** is used in this example.
 - **Username**: Type **root** for the user name.
 - **Password**: Type the logon password of the instance.
5. In the tool bar, tap **Remote Controls**.
6. On the **Remote Controls** page, tap the + icon in the upper left corner to create a remote connection session. *New remote* is used in this example.

The following figure shows Steps 1 through 6.

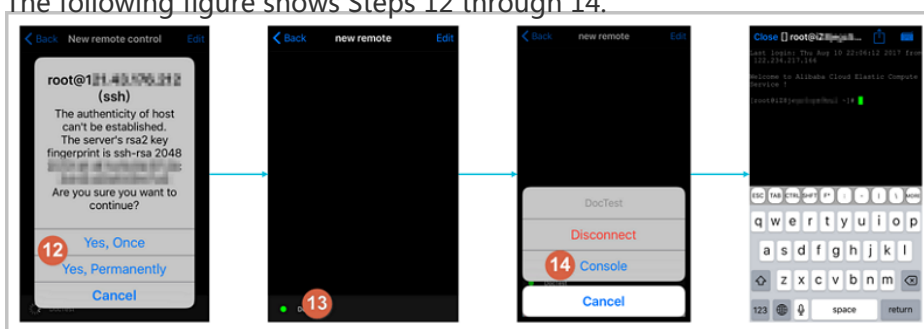


7. On the **New remote** page, tap **Host1**.
8. In the action sheet, tap **Bind**.
9. Select the new Linux instance. In this example, select *DocTest*.
10. On the **New remote** page, tap **Done** to switch it to the **Edit** mode, and then tap **DocTest**.
11. In the action sheet, tap **Connect**.

The following figure shows Steps 7 through 11.



12. In the action sheet, select **Yes, Once** or **Yes, Permanently**. Once the connection is successful, the indicator in front of *DocTest* turns green.
 13. On the **New remote** page, tap *DocTest*.
 14. In the action sheet, tap **Console** to open Linux instance console.
- The following figure shows Steps 12 through 14.

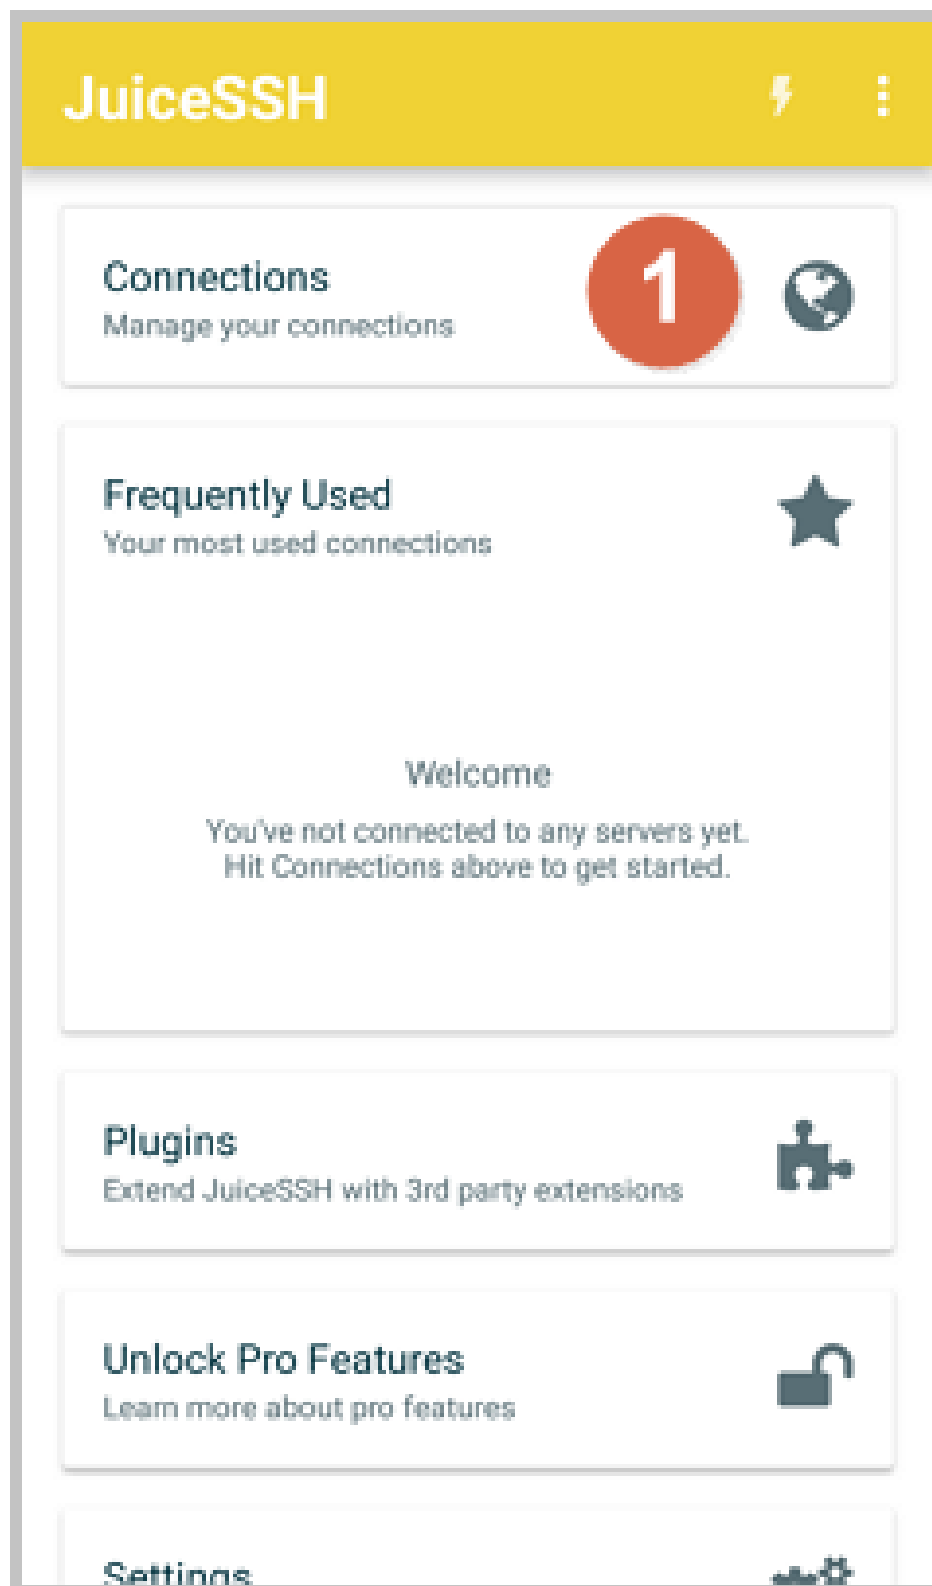


Now, you are connected to the Linux instance.

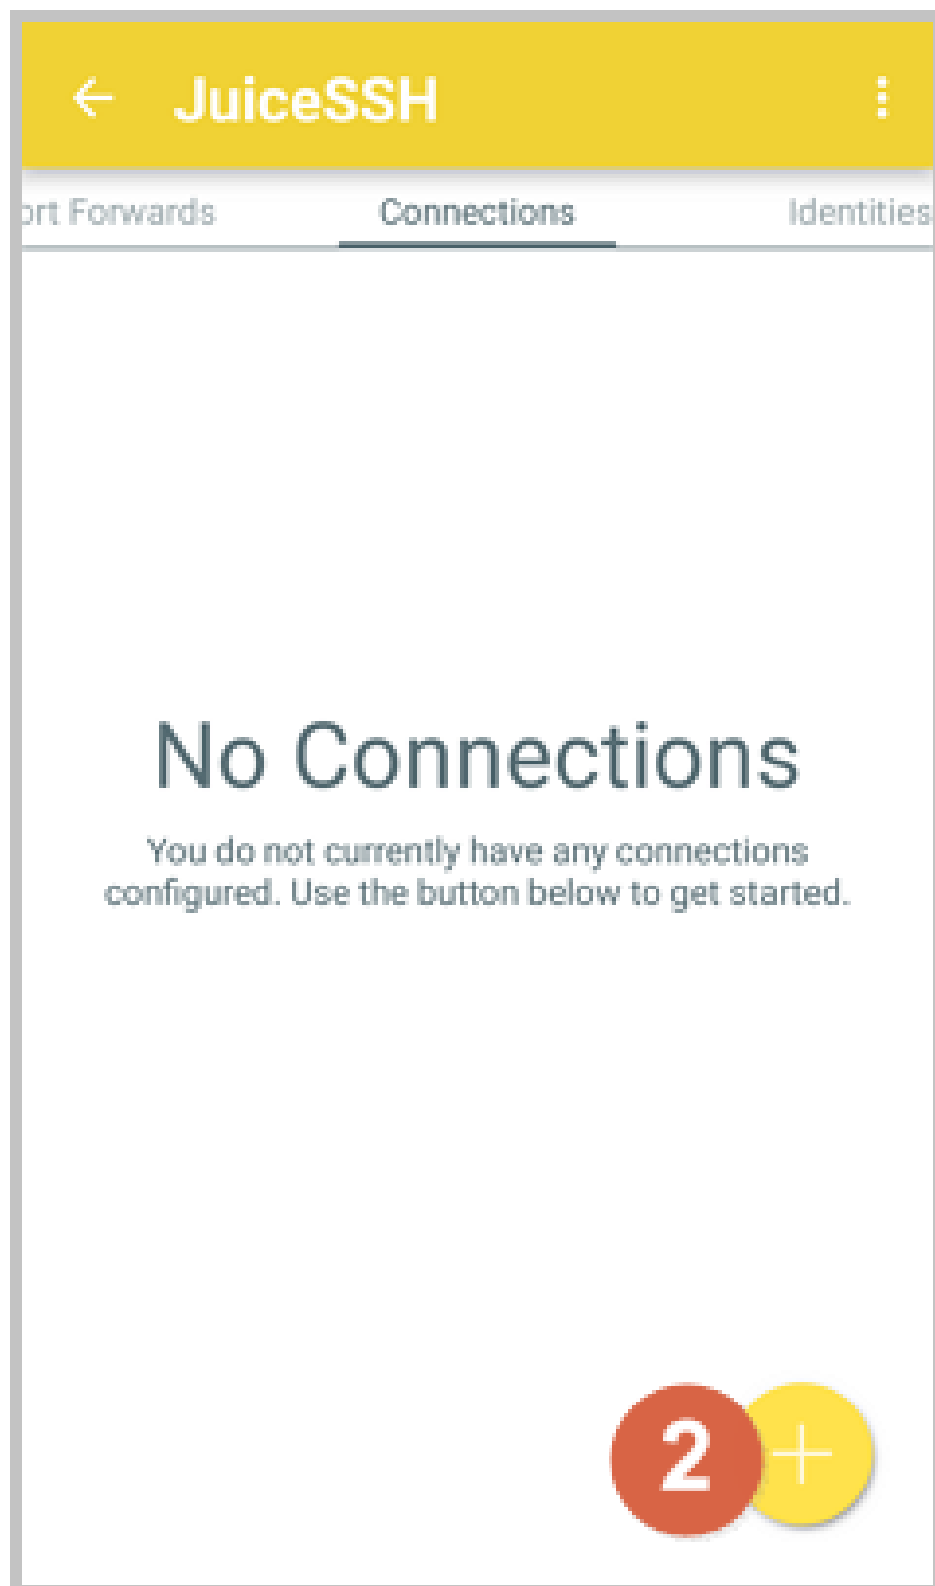
Use JuiceSSH to connect to a Linux instance

To connect to a Linux instance by using JuiceSSH, follow these steps:


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



2. Under the **Connections** tab, tap the + icon.




3. On the **New Connection** page, add the connection information and tap the




icon. The


following connection information is required:

- **Name:** Specify the name of the connection session. *DocTest* is used in this example.
- **Type:** Use the default value **SSH**.
- **Address:** Type the public IP address of the Linux instance to connect to.
- To set **Identity**, follow these steps:
 - a. Tap **Identity**, and tap **New** in the drop-down list.
 - b. On the **New Identity** page, add the following information and tap the



- **NickName:** Optional. You may set a nickname to ease management. *DocTest* is used in this example.
- **Username:** Type **root** for the user name.
- **Password:** Tap **SET(OPTIONAL)**, and type the logon password of the instance.

 **New Identity**



IDENTITY

Nickname:

DocTest

Username:

root

Password:

UPDATE / CLEAR

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:** Type the port number for SSH protocol. In this example, 22 is used.

 **New Connection** 3 ✓

BASIC SETTINGS

Nickname: DocTest

Type: SSH ▼

Address: 121.43.176.212

Identity: DocTest ▼

ADVANCED SETTINGS

Port: 22

Connect Via: (Optional) ▼

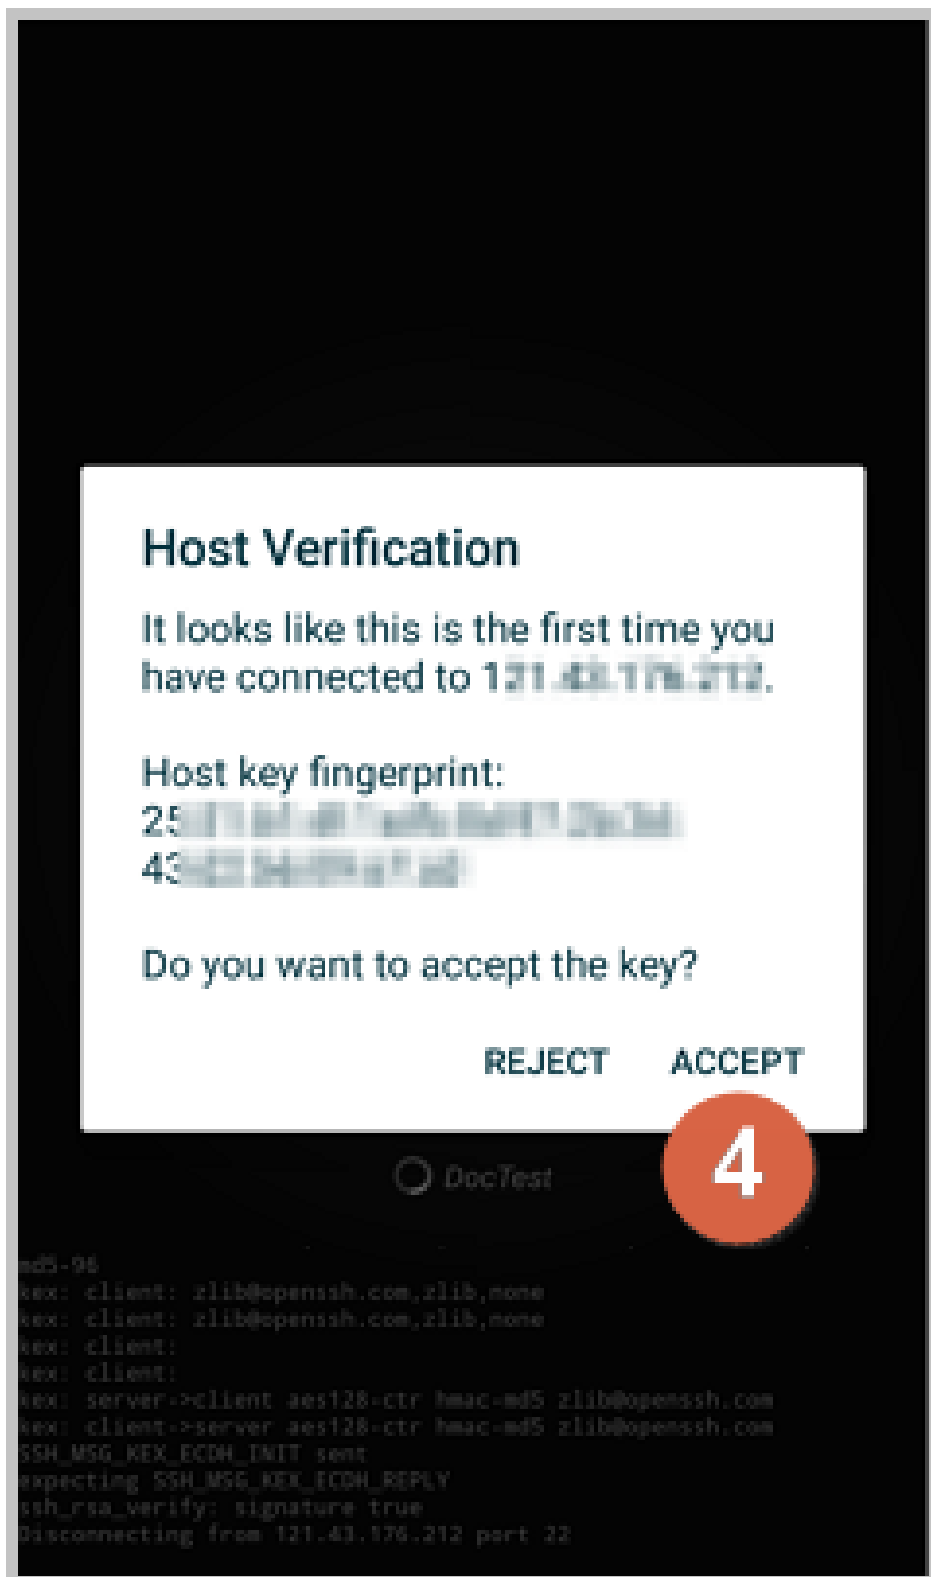
Run Snippet: (Optional) ▼

Backspace: Default (sends DEL) ▼

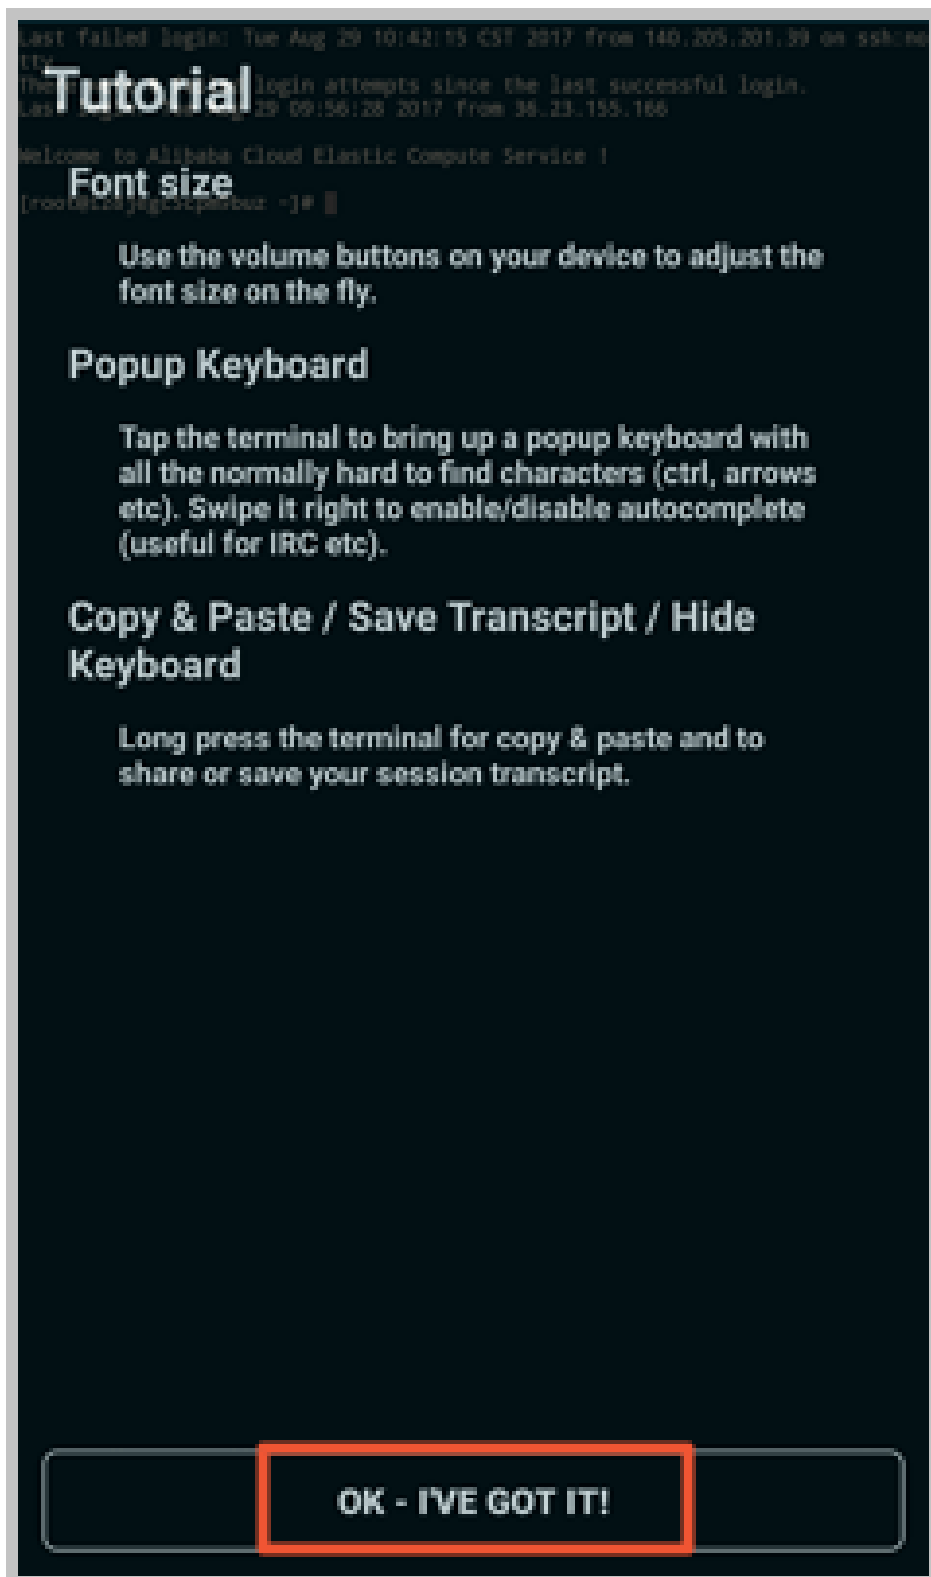
GROUPS

ADD TO GROUP

4. Confirm the message, and tap **ACCEPT**.



5. (Optional) For the first connection, the app would offer you some tips about font setting and the like. Confirm the message, and tap OK - I' VE GOT IT!



Now, you are connected to the Linux instance.

```
Last failed login: Tue Aug 29 10:42:15 CST 2017 from 188.166.166.88 on ssh:node
tty
There were 8 failed login attempts since the last successful login.
Last login: Tue Aug 29 09:56:18 2017 from 36.23.133.166

Welcome to Alibaba Cloud Elastic Compute Service !

[root@t-l4-jpgk-4gw8um ~]#
```

Connect to Windows instances

In this section, we take Microsoft Remote Desktop as an example to describe how to use an app to connect to a Windows instance on a mobile device.

Prerequisites

Confirm the following before connecting to your instance:

- The instance is **Running**.
- The instance has a public IP address and is accessible from public network.
- You have set the logon password for the instance. If the password is lost, you must reset the instance password.
- The security group of the instance has the following security group rules:

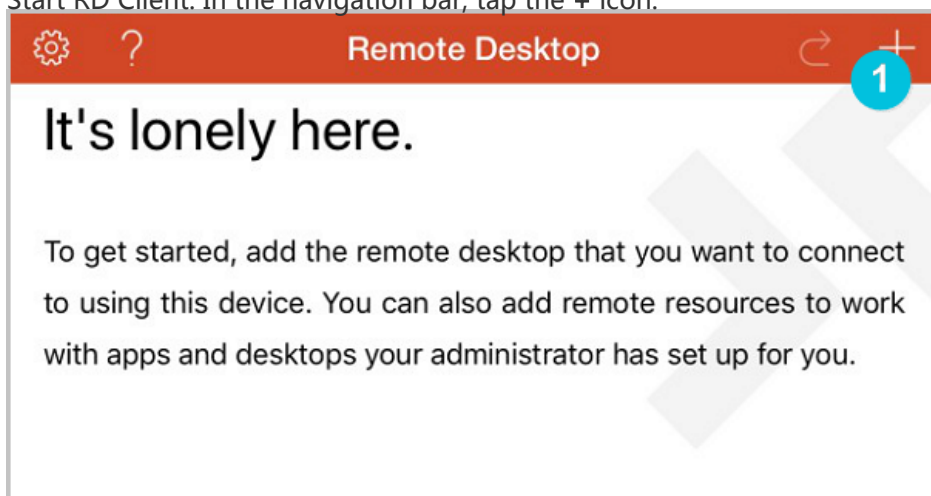
Network Type	NIC	Rule Direction	Authorization Policy	Protocol Type	Port Range	Authorization Type	Authorization Object	Priority
VPC	No configuration required	Inbound	Allow	RDP(3389)	3389/3389	Address Field Access	0.0.0.0/0	1
Classic	Internet							

- You have downloaded and installed Microsoft Remote Desktop.
 - For iOS devices, download the app from iTunes.
 - For Android devices, download the app from Google Play.

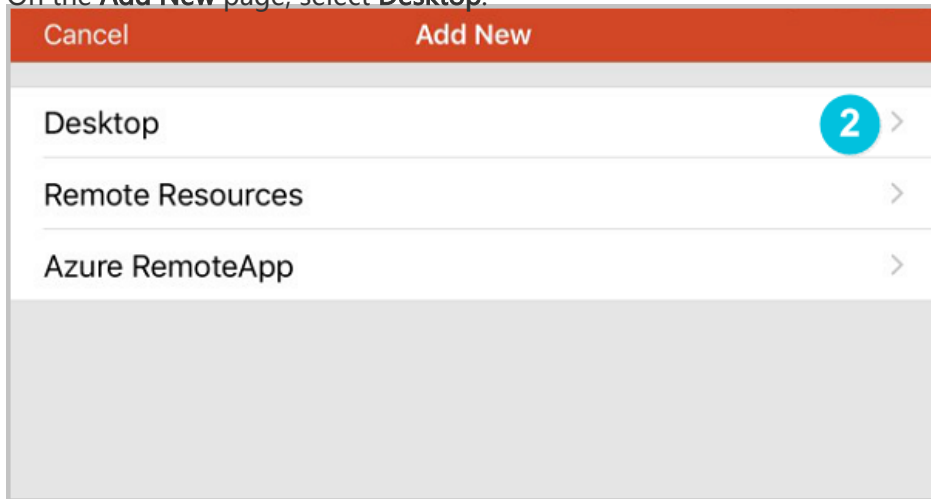
Procedure

To connect to a Windows instance by using Microsoft Remote Desktop, follow these steps:

Start RD Client. In the navigation bar, tap the + icon.

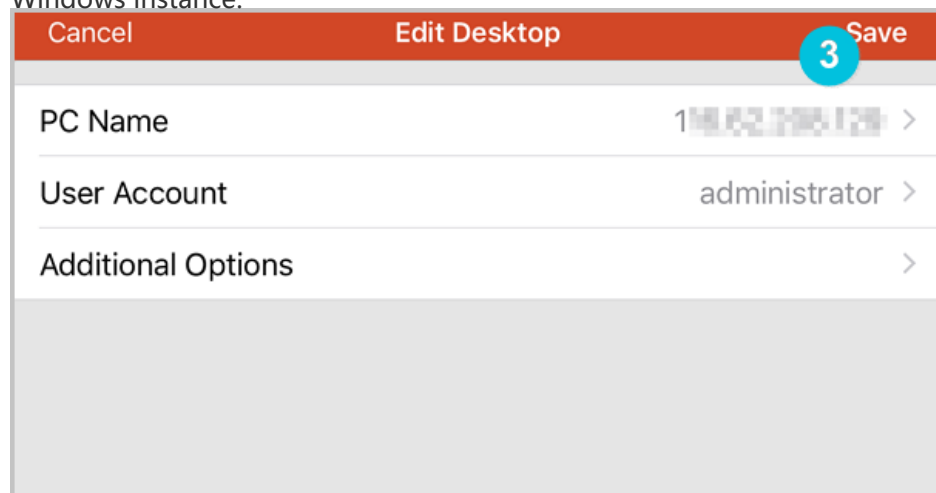


On the **Add New** page, select **Desktop**.

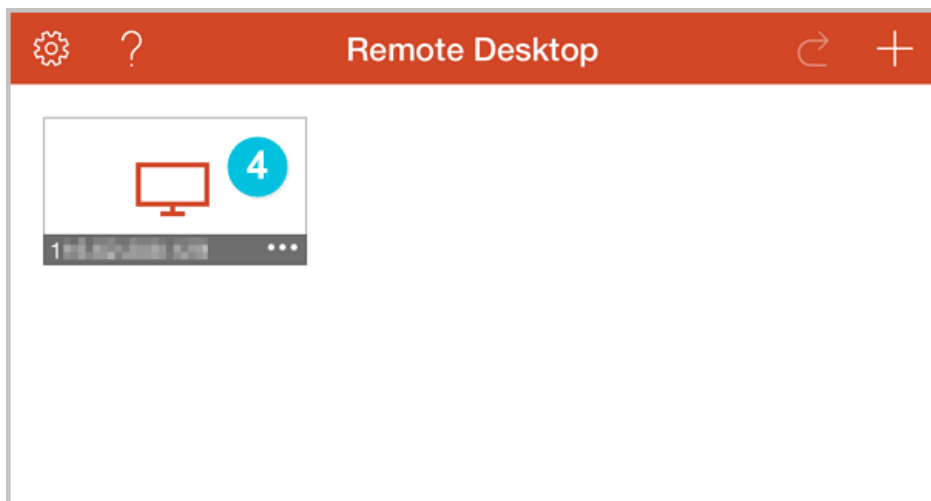


On the **Edit Desktop** page, type the connection information and tap **Save**. The following connection information is required:

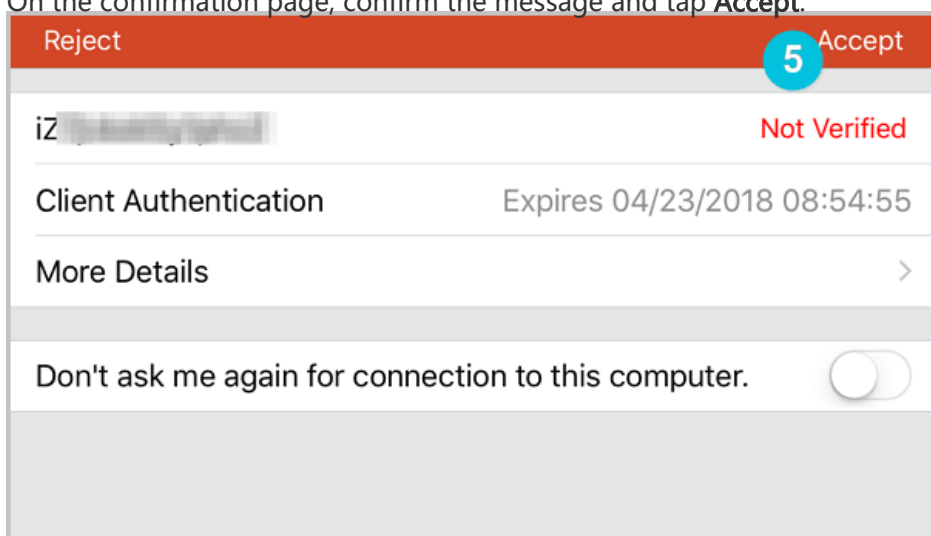
- **PC Name:** Type the public IP address of the Windows instance to connect to.
- **User Account:** Type the account name *administrator* and the logon password of the Windows instance.



On the **Remote Desktop** page, tap the icon of a Windows instance.



On the confirmation page, confirm the message and tap **Accept**.



Now, you are connected to the Windows instance.

Instances

Create an instance

Create an instance

You can create instances running Linux, Windows, or a custom system (based on a public image). Detailed information for each is provided as follows:

- To create an instance, refer to [Quick Start](#).
- If you want to clone the operating system, installed applications, and data of an existing instance, refer to [Create an instance using a custom image](#).

Introduction to gn4 type family

For detailed information about various types in gn4 type family, see [Instance generations and type families](#) in *Product Introduction* of Elastic Compute Service.

Create an instance of gn4 type family

Create an instance of gn4 type family by following the steps described in [Create an instance](#) in *Quick Start of Elastic Compute Service (ECS)*. When creating the instance, note the following items:

- **Network Type:** Select **VPC**.
- **Instance Type:** Choose **Generation III > GPU computing type gn4**.
- **Network Bandwidth:** Select the peak bandwidth as needed.

Note:

If the Windows 2008 R2 image is used and it is enabled, you must enable Internet access of the instance because you cannot connect to an instance of gn4 type family by using the **Management Terminal** in the ECS console. For remote connection to the instance, you can set the peak bandwidth to a non-zero value when creating an instance, or **bind an elastic IP address to an instance** after the instance is created.

Download and install the GPU driver

Before using an instance of gn4 type family, you need to install the GPU driver for the instance. Follow the steps to download and install the GPU driver.

Go to NVIDIA official website to download the corresponding driver for the operating system and P100 GPU. The download URL:
<http://www.nvidia.com/Download/index.aspx?lang=en-us>.

Manually find the driver for the instance. Set the parameters as follows:

- **Product Type:** Tesla
- **Product Series:** M-Class
- **Product:** M40
- **Operating System:** The corresponding version according to the instance image

If the operating system is not displayed in the drop-down list, click **Show all Operating Systems** at the bottom of the drop-down list.

NVIDIA Driver Downloads

Option 1: Manually find drivers for my NVIDIA products. [Help](#)

Product Type:

Product Series:

Product:

Operating System:

- Windows 10 64-bit
- Windows 7 64-bit
- Windows Vista 64-bit
- Windows XP
- Windows XP 64-bit
- Windows Server 2008
- Windows Server 2008 R2 64
- Windows Server 2012 R2 64
- Windows Server 2016
- Windows Server 2003
- Windows Vista 32-bit
- Windows Server 2003 x64
- VMware vSphere ESXi 5.1
- VMware vSphere ESXi 5.5
- Linux 64-bit
- Linux 64-bit RHEL6
- Linux 64-bit RHEL7
- Linux POWER8 RHEL
- Linux 64-bit Ubuntu 16.04
- Linux POWER8 Ubuntu
- Linux 64-bit Ubuntu 14.04
- Linux 64-bit Fedora 23
- Linux 64-bit SLES 12
- Linux 64-bit Opensuse 13.2

[Show less Operating Systems](#)

CUDA Toolkit:

Language:

SEARCH

Click **SEARCH**.

After confirming the information, click **DOWNLOAD**.

Install the GPU driver by following the **ADDITIONAL INFORMATION** on the download page.

Take Linux 64-bit Opensuse 13.2 as an example:

TESLA DRIVER FOR LINUX OPENSUSE 13.2

Version: 375.66
Release Date: 2017.5.9
Operating System: Linux 64-bit Opensuse 13.2
Language: English (US)
File Size: 133.05 MB

DOWNLOAD

RELEASE HIGHLIGHTS

SUPPORTED PRODUCTS

ADDITIONAL INFORMATION

Once you accept the download please follow the steps listed below

```
i) `rpm -i nvidia-diag-driver-local-repo-opensuse132-375.66-1.x86_64.rpm`  
ii) `zypper refresh`  
iii) `zypper install cuda-drivers`  
iv) `reboot`
```

Notes

For Windows 2008 R2 or earlier version, if you go to the **Management Terminal** by clicking **Connect** in the ECS console after the GPU driver is installed, the **Management Terminal** will be stuck at either a black screen or the startup interface. If the instance can access the Internet, you need to connect to the ECS instance remotely using other protocols, such as the Remote Desktop Protocol developed by Microsoft.

The RDP does not support DirectX, OpenGL, and other related applications. Therefore, you need to install VNC services and the client, or other protocols that supports these applications, such as PCoIP or XenDesktop HDX 3D.

Introduction to gn5 type family

For more information about various types in gn5 type family, see Instance generations and type families in *Product Introduction* of Elastic Compute Service.

Create an instance of gn5 type family

Create an instance of gn5 type family by following the steps described in *Create an instance* in *Quick Start* of Elastic Compute Service (ECS). When creating the instance, note the following items:

- **Region:** Currently, gn5 is available only in the following regions: China East 1, China East 2, China North 2, China South 1, US East 1 (Virginia), US West 1, Hong Kong, Asia Pacific SE 1, Asia Pacific SE 2, and Germany 1.
- **Network Type:** Select **VPC** because gn5 is available for virtual private cloud network.
- **Instance Type:** Select **GPU Compute Type gn5** under **Generation III**.

- **Network Bandwidth:** Select the peak bandwidth as needed.

If the Windows 2008 R2 image is used and you want to connect to the instance of gn5 type family, you must enable Internet access of the instance because you cannot connect to an instance of gn5 type family by using the **Management Terminal** in the ECS console. If you want to allocate an Internet IP address to the instance, you must not set the peak bandwidth to 0 Mbps.

- **Image:** Select the image as needed.

Download and install the GPU driver

Before using an instance of gn5 type family, you must install the GPU driver for the instance. Follow the steps to download and install the GPU driver.

Go to NVIDIA official website to download the corresponding driver for the operating system and P100 GPU. The download URL:
<http://www.nvidia.com/Download/index.aspx?lang=en-us>.

Manually find the driver for the instance. Set the parameters as follows:

- **Product Type:** Tesla
- **Product Series:** P-Series
- **Product:** Tesla P100
- **Operating System:** The corresponding version according to the instance image
 - If the operating system is not displayed in the drop-down list, click **Show all Operating Systems** at the bottom of the drop-down list.
 - If the instance uses a Linux image that is not in the list, select **Linux 64-bit**.

Click **SEARCH**.

After confirming the information, click **DOWNLOAD**.

Install the GPU driver by following the **ADDITIONAL INFORMATION** on the download page.

Take Linux 64-bit Opensuse 13.2 as an example:

Notes

For Windows 2008 R2 or earlier version, if you enter the **Management Terminal** by clicking **Connect**

in the ECS console after the GPU driver is installed, the **Management Terminal** is stuck at either a black screen or the startup interface. If the instance can access the Internet, you must connect to the ECS instance remotely using other protocols, such as the Remote Desktop Protocol developed by Microsoft.

To create an instance that has the same operating system, software applications, and data with those of your existing ECS instance or server, create a copy of the existing ECS instance or server as a custom image, and then use it to create an instance. This method improves the deployment efficiency.

Prerequisites

If the image and the instance are in the same region, you have created a custom image by using one of the following methods:

- Importing a physical image
- Creating a custom image by using an ECS instance
- Creating a custom image by using a snapshot of a system disk

If the custom image and instance are in different regions, you have copied the custom image to the target region.

Procedure

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

In the upper-right corner of the page, click **Create Instance**.

On the purchase page,

- Select the expected billing method, target region, instance type, network type, and other parameters. For more information, see the instance creation process in the **Quick Start**.
- Select a custom image.

Note: If the selected custom image contains more than one data disk snapshot, an equal number of data disks are automatically created. By default, the size of each data disk is equal to that of the source snapshot. You are only allowed to increase, but not decrease, the size of a data disk.

Click **Buy Now**.

You can create a spot instance in the ECS console following the steps in [Create an ECS instance](#). Confirm the following settings during the procedure:

- Billing method: Select **Spot Instance**.
- Bidding policy: Select **Set maximum bid**. Set the highest hourly price you are willing to pay. When your bid exceeds the current market price, the instance runs. The final price you pay is the market price. Therefore, you must enter the highest price you are willing to pay.
- Select or enter the number of instances to purchase.
- Click **Buy Now**. If your maximum bid price exceeds the current market price, the system creates an instance for you according to your settings.

After a spot instance is created, you can log on to the ECS console to view its information in the instance list. Spot instances are marked as **Pay-As-You-Go - Spot Instance**. After going to the instance details page, in the **Billing Information** area, you can find the **Bidding Policy** that you set for creating the instance.

This document describes how to create an f1 instance.

Note:

The f1 instance type has been released only for testing. You can [open a ticket](#) to submit your request for it.

Procedure

Follow the steps described in [Create an ECS instance](#) to create an f1 instance. During the procedure, confirm the following settings:

- Region and zone: Select **China East 1 (Hangzhou) > China East 1 Zone F**.
- Network type: Select **VPC**.
- Instance type: Select **Generation III > FPGA compute f1**.
- Image: Select **Shared Image**, and then select the specified image.

Note:

Now, the development environments of Intel FPGA is only available by sharing images. You can find quartus17.0, vcs2017.3, and dcp sdk in the opt directory.

After an f1 instance is created, log on to the instance, and run the following command to check whether the License is configured.

```
echo $LM_LICENSE_FILE #To check whether the variable is set.
```

Best practices

See best practices of f1 instances:

- Use OpenCL on an f1 instance
- Download the bitstream file to an FPGA chip
- Use RTC compiler on an f1 instance

This document describes how to create an f2 instance.

Note:

Now, the f2 instance type is released for test. You can [open a ticket](#) to apply for using an f2 instance.

Procedure

Create an f2 instance according to the procedure described in [Create an ECS instance](#). Consider the following configurations:

- Region and zone: Select **China East 1 (Hangzhou)** > **China East 1 Zone B**.
- Network type: Select **VPC**.
- Instance type: Select **Generation III** > **FPGA compute f2**.
- Image: Click **Shared Image**, and then select the specified image.

Note:

Now, the development environment of Xilinx is only available by sharing images.

Best practices

See best practices of f2 instances:

- Use RTL compiler on an f2 instance
- Use OpenCL on an f2 instance

Change the operating system

Use the management console to convert the instance OS to your preferred OS. For details, see [Change the operating system of an ECS instance](#).

Change the system disk to your custom image or [Change the system disk to a public image](#).

Note: Regions outside of mainland China do not currently support transition between Linux and Windows OSs. If your instance is hosted in one of these regions, you are not allowed to change the operating system between Windows and Linux. You can only change the version of Windows OS, or replace one Linux OS with another Linux OS.

Change configurations

You can change the specifications of an instance and its Internet bandwidth after it is created.

Upgrade or downgrade instance specifications

You can only upgrade or downgrade the specifications of vCPU and RAM simultaneously by changing instance types. The methods to change an instance type vary according to the billing method of the instance:

Subscription:

- Upgrade: Use the [Upgrade Configuration] feature. The new specification takes effect immediately.
- Downgrade: Use the [Renew for Configuration Downgrade](#) feature. You can downgrade the specification of an instance when you are renewing the instance. The new specification takes effect after you restart the instance in the ECS console within the first seven days of the new billing cycle.

Pay-As-You-Go: Use the [Change Instance Type](#) feature. You must stop the instance to use this feature.

Note:

Stopping an instance disrupts your business traffic. Proceed with caution.

Adjust Internet bandwidth

You can adjust the Internet bandwidth of an instance. The methods vary by your business needs and the billing method of the instance. The following table lists the methods.

Billing method of instances	Effective immediately	Available feature	Description
-----------------------------	-----------------------	-------------------	-------------

Subscription	Yes	Upgrade Configuration	Only applicable to VPC instances to which no EIP addresses are attached or instances of the Classic network type. After you upgrade your configurations, the Internet and the intranet IP addresses remain unchanged. If no Internet IP address is assigned to your instance when you create it, you can use this feature to assign the instance an Internet IP address.
Subscription	No. Effective in the new billing cycle	Renew for Configuration Downgrade	Adjust bandwidth in the new billing cycle. When the Internet bandwidth is set to 0 Mbit/s, the Internet IP address of a VPC instance is released in the new billing cycle, but that of an instance of the Classic network type is retained. If no Internet IP address is assigned to your instance when you create it, you can use this feature to assign the instance an Internet IP address.
Pay-As-You-Go or Subscription	Yes	Change Bandwidth	Only applicable to those VPC instances that EIP addresses are bound to. You can adjust the Internet bandwidth on an EIP address at any time.

If you find that the instance specifications exceed or are insufficient for your application requirements, you can change the instance type, that is, the specification of the memory and the CPU. Different operations are allowed based on the billing method of an instance:

For Pay-As-You-Go instances, see descriptions in this document.

For Subscription instances, see [Upgrade configurations](#).

Note:

You must restart your instance after you change its instance type, which may lead to interruptions in your service. We recommend that you perform this operation during non-peak hours.

Limits

Pay-As-You-Go instances are subject to the following limits for configuration change:

The interval between two configuration change operations must be more than five minutes.

You cannot change the instance type across instance generations. For example, instance types of Generation I are not allowed to be changed to those of Generation II or Generation III.

For instance types of Generation III, you cannot change the configuration within or between the following instance type families:

- GPU-based instance type families, including gn5, gn4, gn5i, and ga1.
- FPGA-based instance type families, including f1.
- Big data instance type families, including d1 and d1ne.
- Local SSD instance type families, including i1 and i2.

For instance types of Generation III, you can change the instance types according to the following table.

Instance type families	ecs.sn1ne	ecs.sn2ne	ecs.mn4	ecs.se1ne	ecs.cm4	ecs.c4	ecs.se1	ecs.ce4	ecs.xn4	ecs.e4	ecs.n4
ecs.sn1ne	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.sn2ne	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.mn4	-	-	Y	-	-	-	-	-	Y	Y	Y

ecs.se1ne	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.cm4	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.c4	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.se1	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.ce4	Y	Y	-	Y	Y	Y	Y	Y	-	-	-
ecs.xn4	-	-	Y	-	-	-	-	-	Y	Y	Y
ecs.e4	-	-	Y	-	-	-	-	-	Y	Y	Y
ecs.n4	-	-	Y	-	-	-	-	-	Y	Y	Y

- For instance types of Generation II, you can change the instance type according to the following table.

Instance type families	ecs.n2	ecs.e3	ecs.n1	ecs.sn2	ecs.sn1
ecs.n2	Y	Y	Y	-	-
ecs.e3	Y	Y	Y	-	-
ecs.n1	Y	Y	Y	-	-
ecs.sn2	-	-	-	Y	Y
ecs.sn1	-	-	-	Y	Y

- You can change the configuration of all instance types within Generation I.

Note:

In the preceding table, "Y" indicates that you can change the configuration between the instance type families, and "-" indicates that you are not allowed to change the configuration between the instance type families.

Prerequisites

You must stop the instance.

Procedure

To change the memory and CPU configurations of a Pay-As-You-Go instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

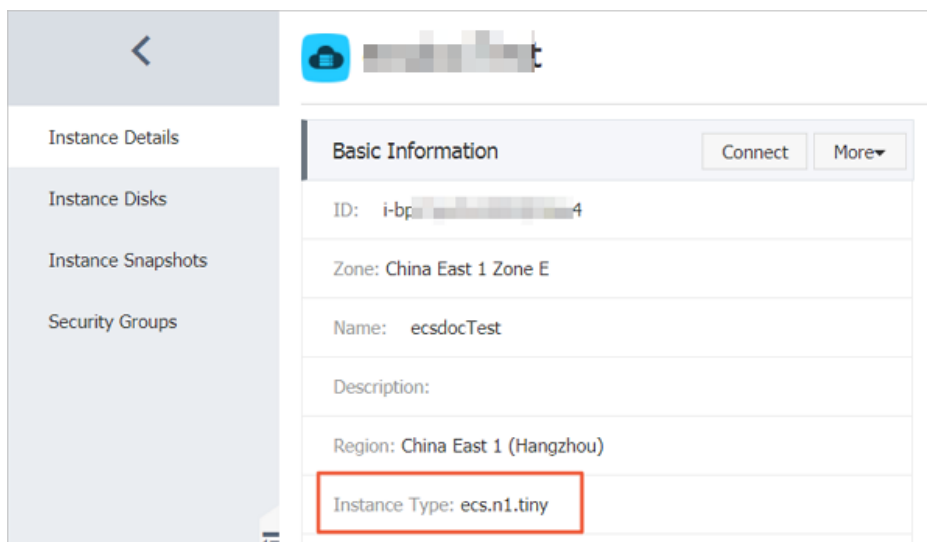
Find the Pay-As-You-Go instance you want to change the configuration, and in the **Actions** column, click **Change Instance Type**.

In the **Change Instance Type** dialog box, select an instance type and click **OK**.

Note:

You can enter the instance type information in the search box to filter instance types in real time.

The change is immediately effective after you complete the operation. You can view the instance type information in the **Basic Information** area of the **Instance Details** page.



- For Linux instances, the default username is **root**.
- For Windows instances, the default username is **Administrator**.

Note: You must restart an instance after its password is reset, which has an impact on your

service. To reduce the impact, we recommend you to reset the password when the related service is not busy.

To reset the instance password, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Select the target instance. You may select multiple instances with identical operating statuses.

Click **More > Reset Password**.

Enter a new password in the displayed dialog box. Click **Submit**.

Click **OK**.

Select the instance on which the password was changed and click **Restart**.

Note: The new password only takes effect after the instance is **restarted** through the console. Restarting directly within the instance do not apply the new password.

Click **OK** in the displayed dialog box to restart the instance.

Start, view, or stop an instance

This section describes how to start, view, and stop an instance.

Start an instance

Note: You can only start instances in **Stopped** status.

To start an instance, perform the following:

1. Log on to the ECS console.
2. In the left-side navigation pane, click **Instances**.

3. Select a region.
4. Select the desired instance. You can select multiple instances, as long as they are all in **Stopped** status.
5. Click **Start** at the bottom of the page.

View an instance

Use the console to view the following instance information:

- Quantity and operating statuses of instances in each region
- Basic configuration, payment, and monitoring information
- Disks
- Snapshots
- Related security groups

To view instances, perform the following:

1. Log on to the **ECS console**.
2. On the overview page, you can view the operating statuses of ECS instances in all regions.
3. In the left-side navigation pane, click **Instances**.
4. Select a region, and click the name of the instance you want to view.
(Optional) You can also search for the name of the instance you want to view.

You can then view instance details, including its:

- Region
- Zone
- Configuration specification
- Payment status
- CPU
- Network usage

Note: You can also view and manage the instance's disks, snapshots, and security group information.

Stop an instance

Note:

- A stopped instance still incurs fees. To stop charging, you need to **release** the instance.
- Stopping an instance may disrupt your business traffic. Proceed with caution.
- Only instances in **Running** status can be stopped.

To stop an instance, perform the following:

1. Log on to the **ECS console**.

2. In the left-side navigation pane, click **Instances**.
3. Select a region.
4. Select the desired instance. You can select multiple instances, as long as they are all in **Running** status.
5. Click **Stop**.
6. Select **Stop** in the displayed dialog box. Click **OK**.

Restart an instance

Instances can be restarted from within or through the management console.

Note: Restarting an instance may disrupt your business traffic. Proceed with caution.

Only instances in the **Running** status can be restarted.

To restart an instance, perform the following:

Log on to the ECS console.

Click **Instances** in the left-side navigation pane.

Select your desired region.

Select the desired instance. You can select multiple instances, as long as they are all in the **Running** status.

Click **Restart**.

In the displayed dialog box, click **Restart**, and then click **OK**.

For a Pay-As-You-Go instance, in the event of payment failure within 15 days (T+15) after the due date (T), the instance is stopped due to overdue payment and becomes **Expired**. You must open a ticket to settle the payment and reactivate the instance within 30 days (T+30) after the due date (T). Otherwise, the instance is released and the data cannot be recovered.

Note:

If you fail to reactivate the ECS instance within 30 days (T+30) after the due date (T), the instance is automatically released 30 days after the due date and the data cannot be recovered.

Prerequisites

The Pay-As-You-Go instance is in the **Expired** status.

You have settled the payment by opening a ticket.

Procedure

To reactivate an instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Select the instance to be reactivated, and at the bottom of the instance list, select **More > Reactivate**.

Determine that you reactivate the instance immediately or later at a specified time.

If you choose to reactivate immediately, the selected instance returns to the **Running** status in about 10 minutes.

If you no longer require a Pay-As-You-Go instance, we recommend that you release it immediately. Charges continue when the instance is stopped but not released.

To release a Pay-As-You-Go instance, you have two options:

- **Release Now:** Immediately releases the Pay-As-You-Go instance.
- **Timed Release:** Makes a release plan for your Pay-As-You-Go instances by defining release schedules. These schedules are definable to the hour. Applying new time schedules overwrites previous ones.

Note:

If set to be released automatically, instances are released every half hour or every hour. However, the system stops billing according to your specified release time.

This document describes how to release an instance and how to enable and disable the release schedule for automatic release.

Release an instance now

To release an instance now, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Select an instance, and in the **Actions** column, select **More > Release Instance**.

In the dialog box, select **Release Now**.

Click **Next > OK**.

Enable automatic release

If you want your instance to be released at a specified time, you can use the **Timed Release** feature to enable the instance to be released automatically.

To enable auto release, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Select an instance, and in the **Actions** column, select **More > Release Instance**.

In the dialog box, select **Timed Release**, enable automatic release, and then specify the release date and time.

Click **Next**, and then click **OK**.

Disable automatic release

If you want to cancel the automatic release feature of a Pay-As-You-Go instance, you can disable the feature.

To disable the automatic release feature, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Select an instance, and in the **Actions** column, select **More > Release Instance**.

In the dialog box, disable automatic release.

Click **Next**, and then click **OK**.

Add an instance to a security group

You can add an instance to a security group using the ECS Management console. One ECS instance can be added to up to five security groups. After adding the instance to a security group, the security group rules will automatically be applied to the instance.

To add an instance to a security group, perform the following:

Log on to the ECS console.

Click **Instances** in the left-side navigation pane.

Select your desired region.

Select the desired instance. Click the instance name or corresponding **Manage** button.

Click **Security Groups** in the left-side navigation pane.

Click **Add Security Group**. In the displayed dialog box, select the appropriate security group.

Click **OK**.

Remove an instance from a security group

You can remove instances from security groups. Note that an instance must be in at least two security group for this action to be performed, and you have done enough test before this operation to avoid any intranet communication error between instances.

To remove an instance from a security group, perform the following:

Log on to the ECS console.

Click **Instances** in the left-side navigation pane.

Select your desired region.

Select the desired instance. Click the instance name or corresponding **Manage** button.

Click **Security Groups** in the left-side navigation pane.

Select the security group to remove from and click **Remove**.

Click **OK**.

For use cases of security groups, see [Usage scenarios](#).

Use RAM roles of an ECS instance

Instance RAM (Resource Access Management) roles grant role-based permissions to ECS instances.

You can assign a **RAM role** to an ECS instance to allow applications hosted on that instance to access other cloud services by using a temporary **STS (Security Token Service)** credential. This feature guarantees the security of your AccessKey and allows you to use the fine-grained access control in virtue of **RAM**.

Background

Typically, the applications within an ECS instance use the AccessKey of the **user account** or **RAM user account**, including AccessKeyId and AccessKeySecret, to access various cloud services on the Alibaba

Cloud platform.

However, if Alibaba Cloud writes the AccessKey into the instance for calling purposes, for example, writing in the configuration files, the exposed AccessKey leads to problems such as excessive permission, data breaches, and maintenance complexities. Thus, Alibaba Cloud has designed the instance RAM role to solve the complexities.

Benefits

The instance RAM role enables you to:

- Associate a RAM role to an ECS instance.

- Access other cloud services securely, such as OSS, SLB, or ApsaraDB for RDS, by using the STS credential from the applications within the ECS instance.

- Assign roles that have different policies for different ECS instances, and let these instances have restrictive access level to other cloud services to obtain fine-grained access control.

- Maintain the access permission of the ECS instances efficiently only by modifying the policy of the RAM role, without manually changing the AccessKey.

Free of charge

ECS does not charge additional fee for the instance RAM role feature.

Limits

The instance RAM role has the following limits:

- The instance RAM role is only applicable to VPC instances.

- One ECS instance can only be authorized to one instance RAM role.

How to use an instance RAM role

The instance RAM role can be used by any of the following methods:

- Use the instance RAM role on the console

Use the instance RAM role by calling APIs

References

For a list of cloud services that support the STS credential, see *RAM* document *Cloud services supporting RAM*.

For instruction on how to access other cloud services, see *Access other cloud products by using the instance RAM role*.

Limits

The instance RAM role has the following limits:

- The instance RAM role is only applicable to VPC instances.
- One instance RAM role can be bound to one instance at a time.
- After an instance RAM role is attached to an ECS instance, if you want to access other cloud services, such as OSS, SLB, or ApsaraDB for RDS, from the applications within the ECS instance, you must obtain the authorization credential of the instance RAM role by using metadata.
- Before using this feature, the RAM user must be authorized to use the instance RAM role.

Prerequisites

You must have activated the RAM service.

Usage instructions

1. Create an instance RAM role

1. Log on to the RAM console.
2. On the left-side navigation pane, click **Roles**.
3. Click **Create Role**.
4. In the dialog box:
 - i. Select **Service Role** for **Role Type**.
 - ii. Select **ECS (Elastic Compute Service)** for **Type**.
 - iii. Enter the role name and description, for example, `EcsRamRoleDocumentTesting`.

Create Role

1 : Select Role Type 2 : Enter Type 3 : Configure Basic 4 : Role created

Role Name : EcsRamRoleDocumentTesting

Names must be 1-64 characters long. They may only contain letters, numbers, and hyphens.

Description : EcsRamRoleDocumentTesting

Previous Create

iv. Click **Create** to create the instance RAM role.

2. Authorize the instance RAM role

1. Log on to the RAM console.
2. On the left-side navigation pane, click **Policy**.
3. Click **Create Authorization Policy**.
4. In the dialog box:
 - i. Select **Blank Template** for authorization policy template.
 - ii. Enter the **Authorization Policy Name** and **Policy Content**, for example, EcsRamRoleDocumentTestingPolicy.

Note:

For more information about how to write the authorization policy by using the JSON language, see [Syntax](#).

Create Authorization Policy

Step 1: Select an authorization policy Step 2: Edit permissions and submit. Policy creation complete.

Authorization Policy Name : EcsRamRoleDocumentTestingPolicy
Names must be 1-128 characters long. They may only contain the letters A-Z, numbers 0-9, and hyphens.

Description : EcsRamRoleDocumentTestingPolicy

Policy Content :

```

1 {
2   "Version": "1",
3   "Statement": [
4     {
5       "Effect": "Allow",
6       "Action": [
7         "oss:Get*",
8         "oss:List*"
9       ],
10      "Resource": "*"
11    }
12  ]
13 }

```

[Authorization Policy Format](#)
[Authorization Policy FAQ](#)

Previous Create Authorization Policy Cancel

iii. Click **Create Authorization Policy** to complete authorization.

5. On the left-side navigation pane, click **Roles**.

6. Select the created role, for example, EcsRamRoleDocumentTesting, and click **Authorize**.

i. Enter the authorization policy name and click it, for example, EcsRamRoleDocumentTestingPolicy.

ii. Click the icon > to select the policy name, and click **OK**.

Edit Role Authorization Policy

Roles added to this group have all the permissions of this group. A role cannot be added to the same group more than once.

Search and Attach Input and Attach

Available Authorization Policy Names	Type
Ecs	
AliyunECSFullAccess	Provides full acce...
AliyunECSReadOnlyAccess	Provides read-only...
EcsRamRoleDocumentTestingPolic...	EcsRamRoleDocument...

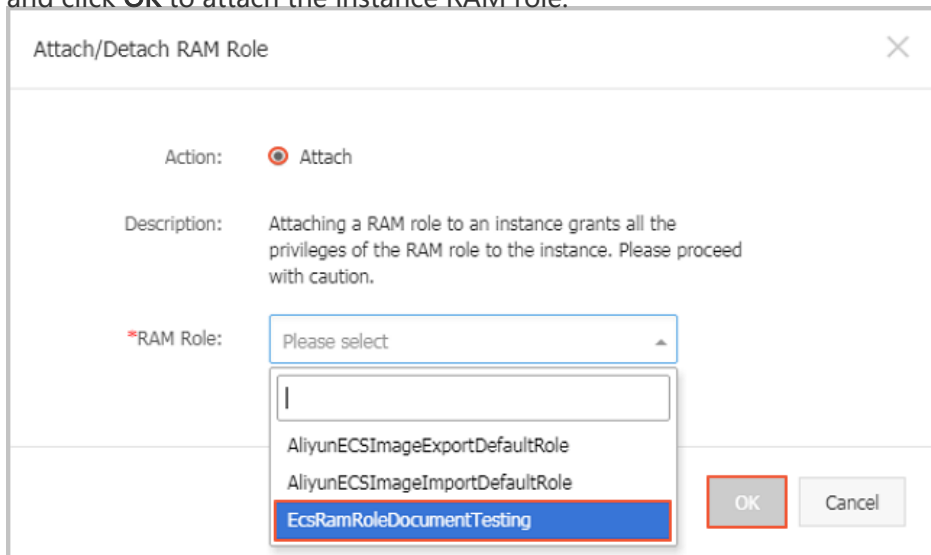
>

Selected Authorization Policy Name	Type
------------------------------------	------

OK Close

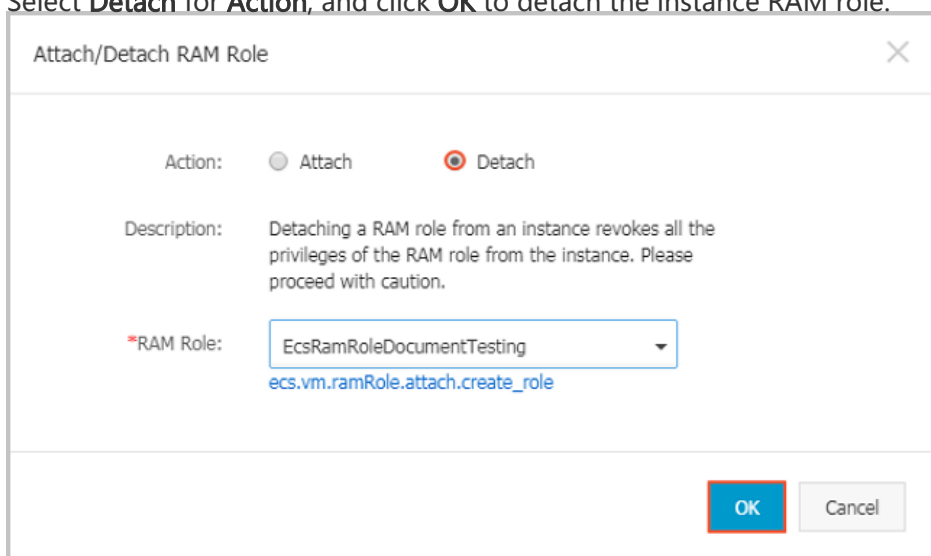
3. Attach an instance RAM role

1. Log on to the ECS console.
2. On the left-side navigation pane, click **Instances**.
3. Choose a region.
4. Find the target ECS instance and choose **More** > **Attach/Detach RAM Role**.
5. Select **Attach** for **Action**, select the created role, for example, EcsRamRoleDocumentTesting, and click **OK** to attach the instance RAM role.



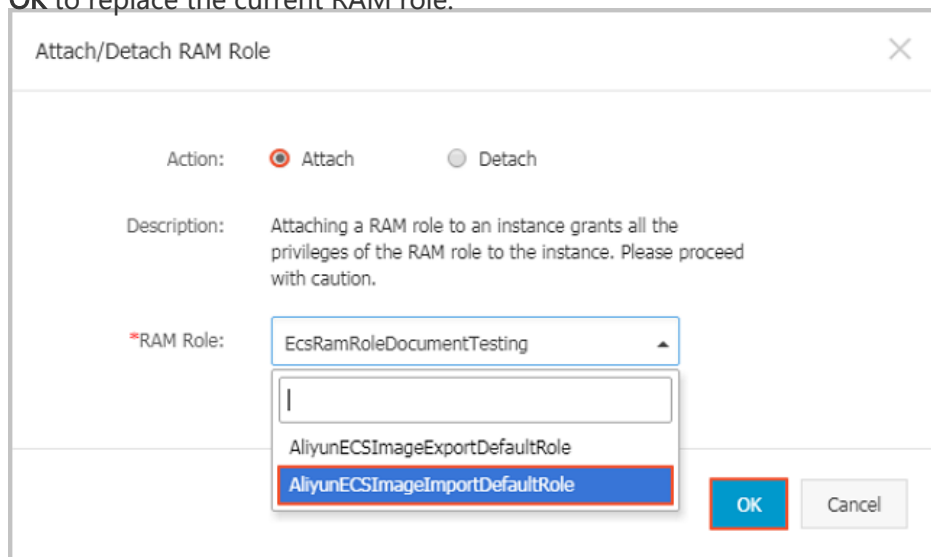
4. (Optional). Detach an instance RAM role

1. Log on to the ECS console.
2. On the navigation pane, click **Instances**.
3. Select a region.
4. Choose an ECS instance, and Select **More** > **Attach/Detach RAM Role**.
5. Select **Detach** for **Action**, and click **OK** to detach the instance RAM role.



5. (Optional). Replace an instance RAM role

1. Log on to the ECS console.
2. On the left-side navigation pane, click **Instances**.
3. Select a region.
4. Choose an ECS instance, and select **More** > **Attach/Detach RAM Role**.
5. Select **Attach** for **Action**, select another instance RAM role in the list of **RAM Role**, and click **OK** to replace the current RAM role.



6. (Optional). Obtain the authorization credential

For the internal application of an ECS instance, you can obtain the STS credential of the instance RAM role, which is a **metadata** of an instance, to access the role-authorized permissions and resources. The credential is updated periodically.

1. Connect and log on to your ECS instance.
2. Obtain the STS credential of the instance RAM role, for example, `EcsRamRoleDocumentTesting`:
 - Linux instance: run `curl http://100.100.100.200/latest/meta-data/Ram/security-credentials/EcsRamRoleDocumentTesting`.
 - Windows instance: run `Invoke-RestMethod http://100.100.100.200/latest/meta-data/Ram/security-credentials/EcsRamRoleDocumentTesting` in PowerShell.
3. Get the credential. Return example:

```
{
  "AccessKeyId" : "XXXXXXXXXX",
  "AccessKeySecret" : "XXXXXXXXXX",
  "Expiration" : "2017-11-01T05:20:01Z",
  "SecurityToken" : "XXXXXXXXXX",
  "LastUpdated" : "2017-10-31T23:20:01Z",
  "Code" : "Success"
}
```

```
}
```

7. (Optional). Authorize a RAM user to use the instance RAM role

NOTE:

You must grant the RAM user with the PassRole permission to use the instance RAM role feature. Without the PassRole permission, a RAM user cannot carry out the permission of the authorization policy that is attached to the RAM user.

Log on to the RAM console and authorize a RAM user to complete the authorization, see the following code snippet as an authorization policy example:

```
{
  "Version": "2016-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs: [ECS RAM Action]",
        "ecs: CreateInstance",
        "ecs: AttachInstanceRamRole",
        "ecs: DetachInstanceRAMRole"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "ram:PassRole",
      "Resource": "*"
    }
  ]
}
```

The parameter [ECS RAM Action] indicates the action that a ram user can be authorized. For more information, see [Actions in RAM that can be authorized to an ECS instance](#).

References

- You can also Use the instance RAM role by calling APIs.
- You may want to Access other cloud products by using the instance RAM role.

Limits

The instance RAM role has the following limits:

- The instance RAM role is only applicable to VPC instances.
- One instance RAM role can be bound to one instance at a time.
- After an instance RAM role is attached to an ECS instance, if you want to access other cloud services, such as OSS, SLB, or ApsaraDB for RDS, from the applications within the ECS instance, you must obtain the authorization credential of the instance RAM role by using metadata.
- Before using this feature, the RAM user must be authorized to use the instance RAM role.

Prerequisites

You must have activated the RAM service.

Usage instructions

1. Create an instance RAM role

Call the `CreateRole` to create an instance RAM role.

- Set the parameter **RoleName**, for example, `EcsRamRoleDocumentTesting`.
- Set the **AssumeRolePolicyDocument** as follows:

```
{
  "Statement": [
    {
      "Action": "sts:AssumeRole",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "ecs.aliyuncs.com"
        ]
      }
    }
  ],
  "Version": "1"
}
```

2. Authorize the instance RAM role

1. Call the `CreatePolicy` to create an authorization policy.
 - Set the parameter **RoleName**, for example, set it to `EcsRamRoleDocumentTestingPolicy`.
 - Set the **PolicyDocument** as follows.

Note:

For more information about how to write the authorization policy by using the JSON language, see *RAM* document **Syntax**.

```
{
  "Statement": [
    {
      "Action": [
        "oss:Get*",
        "oss:List*"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ],
  "Version": "1"
}
```

2. Call the **AttachPolicyToRole** to authorize the role policy.
 - Set **PolicyType** to Custom.
 - Set the parameter **PolicyName**, for example, `EcsRamRoleDocumentTestingPolicy`.
 - Set the parameter **RoleName**, for example, `EcsRamRoleDocumentTesting`.

3. Attach the instance RAM role

Call the **AttachInstanceRamRole** to attach an instance RAM role to an ECS instance.

- Set the parameters **RegionId** and **InstanceIds** to specify an ECS instance.
- Set the parameter **RamRoleName**, for example, `EcsRamRoleDocumentTesting`.

4. (Optional). Detach an instance RAM role

Call the **DetachInstanceRamRole** to detach an instance RAM role.

- Set the parameters **RegionId** and **InstanceIds** to specify an ECS instance.
- Set the parameter **RamRoleName**, for example, `EcsRamRoleDocumentTesting`.

5. (Optional). Obtain the on-demand authorization credential

For the internal application of an ECS instance, you can obtain the STS credential of the instance RAM role, which is a **metadata** of an instance, to access the role-authorized permissions and resources. The credential is updated periodically.

1. Connect and log on to your ECS instance.
2. Obtain the STS credential of the instance RAM role, for example, `EcsRamRoleDocumentTesting`:
 - Linux instance: run `curl http://100.100.100.200/latest/meta-data/Ram/security-`

credentials/EcsRamRoleDocumentTesting.

- Windows instance: run Invoke-RestMethod http://100.100.100.200/latest/meta-data/Ram/security-credentials/EcsRamRoleDocumentTesting in PowerShell.

3. Get the credential. Return example:

```
{
  "AccessKeyId" : "XXXXXXXXXX",
  "AccessKeySecret" : "XXXXXXXXXX",
  "Expiration" : "2017-11-01T05:20:01Z",
  "SecurityToken" : "XXXXXXXXXX",
  "LastUpdated" : "2017-10-31T23:20:01Z",
  "Code" : "Success"
}
```

6. (Optional). Authorize a RAM user to use the instance RAM role

Note:

You must grant the RAM user with the PassRole permission to use the instance RAM role feature. Without the PassRole permission, a RAM user cannot carry out the permission of the authorization policy that is attached to the RAM user.

Log on to the RAM console and authorize a RAM user to complete the authorization, see the following code snippet as an authorization policy example:

```
{
  "Version": "2016-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ecs: [ECS RAM Action]",
        "ecs: CreateInstance",
        "ecs: AttachInstanceRamRole",
        "ecs: DetachInstanceRAMRole"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Allow",
      "Action": "ram:PassRole",
      "Resource": "*"
    }
  ]
}
```

The parameter [ECS RAM Action] indicates the action that a RAM user can be authorized. See [Actions in RAM](#) that can be authorized to an ECS instance for more information.

References

- You can also use the instance RAM role on the console.
- APIs related to the instance RAM role include:
 - **CreateRole**: Create an instance RAM role
 - **ListRoles**: Query the list of instance RAM roles
 - **CreatePolicy**: Create an instance RAM role policy
 - **AttachPolicyToRole**: Authorize an instance RAM role policy
 - **AttachInstanceRamRole**: Attach an instance RAM role
 - **DetachInstanceRamRole**: Detach an instance RAM role
 - **DescribeInstanceRamRole**: Query an instance RAM role

Renew instances

Introduction

Manual renewal only applies to instances in **Subscription** mode.

You can manually renew your instances in **Subscription** mode when they are in the **Expired** status or shut down. You can manually renew your instance for a month or a year. Therefore, if you want to modify the service duration of your subscription-mode instances, you can choose manual renewal.

Your instance will still work normally when the instance is in the **Expired** status. If the manual renewal is successfully completed within 15 days after expiration, your instance will go into the next billing cycle from the day of expiration.

For example, if your instance expired at 00:00:00 on April 25, 2016, but you successfully renewed it for one month on May 9, 2016, the billing cycle for this renewal was from April 25, 2016 to 00:00:00 on May 25, 2016.

If the instance fails to be renewed within 15 days after expiration, the instance will be shut down. Your instance will stop providing services, but Alibaba Cloud will still keep the data for you.

After the instance is shut down,

If the renewal is successful within 15 days, your instance will go into the new billing

cycle from the day of renewal.

For example, if your instance was shut down at 00:00:00 on May 10, 2016, but you successfully renewed it for one month at 08:09:35 on May 23, 2016, the billing cycle for this renewal is from 08:09:35 on May 23, 2016 to 00:00:00 on June 24, 2016.

If the renewal fails within 15 days, your instance will be automatically released on the 15th day and the data will not be restored.

Procedure

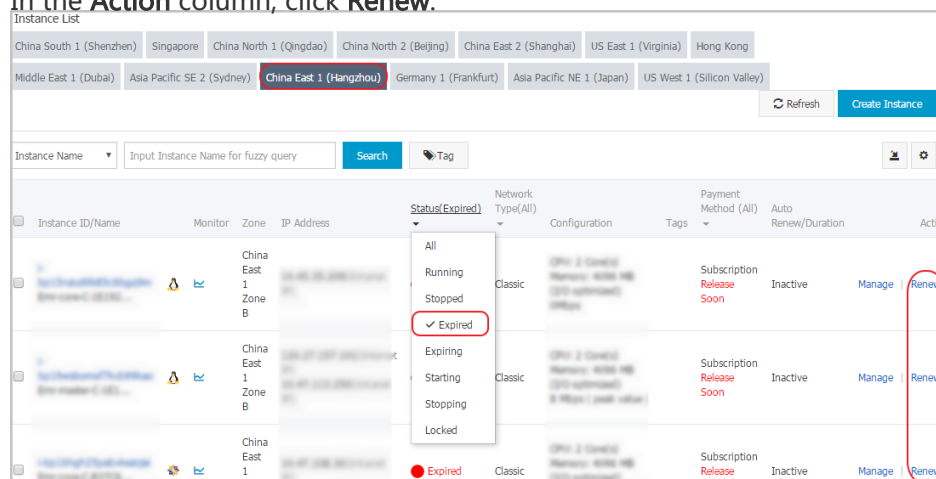
You can manually renew your instance with the following steps.

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

On the **Instance List** page, select the expected region and locate the ECS instance by the instance name, instance ID, or status (**Expired**).

In the **Action** column, click **Renew**.



On the renewal page,

- Confirm the instance configuration.
- Select the expected renewal length, **1 Month** or **1 Year**, and click **Place Order**.

On the **Pay** page, confirm the order information and click **Pay** to complete the renewal

operation.

Auto-renewal service only applies to the instances that use the Subscription billing method.

Introduction

If you have activated the auto-renewal service, Alibaba Cloud charges the subscription fee to your linked credit card or PayPal account when the instance expires.

The auto-renewal service can be activated any time after the ECS instance is purchased and before it expires. It cannot be activated after a Subscription instance expires. Features of auto-renewal service are as follows:

- The monthly subscription service automatically renews the instance on a monthly basis when a monthly subscription instance expires.

The annual subscription service automatically renews the instance on a yearly basis when a yearly subscription instance expires.

Note:

The auto-renewal service does not support switching between monthly subscription and annual subscription. If you want to change the service duration for an instance, you can choose the **Manual renewal** service.

After you activate the auto-renewal service,

You are notified of the expiration of your Subscription instances on the seventh day, third day, or one day before the expiration day (T).

Alibaba Cloud charges the subscription fee to your linked credit card or PayPal account on the expiration day (T). If the fee payment fails, Alibaba Cloud tries again on Day 7 (T+6) and Day 15 (T+14) until the payment is successful. If all the three payment attempts fail, the instance shuts down.

If the payment for the subscription is successful, your instance is no longer be in the **Expired** status and the next billing cycle starts from the expiration day.

For example, if your monthly subscription instance expired at 00:00:00 on April 25, 2016, but it was successfully renewed automatically on May 9, 2016, the billing cycle for this renewal is from 00:00:01 on April 25, 2016 to 00:00:00 on May 25, 2016.

If all the three payment attempts fail, the instance shuts down 15 days after its expiration day. And, if the instance shuts down, it stops providing services and you cannot log on or even remotely connect to the instance. At this point, you can only choose **Manual renewal**. If the instance is not renewed within 15 days after the expiration day, the instance is released and the data is lost.

If you manually renew the instance even before the auto-renewal is attempted, your instance gets renewed and no auto-renewal is attempted for the current billing cycle. And, the instance goes into the next billing cycle.


Alibaba Cloud sends a notification email to your linked email address for each failed auto-renewal attempt. Therefore, keep checking your mailbox frequently so that you do not miss any such notifications and can take necessary action to avoid further business impact.

Activate auto-renewal

To activate the auto-renewal service, follow these steps:

Log on to the ECS console.

At the top of the page, choose **Billing Management > Renew**.



In the left-side navigation pane, click **Elastic Compute Service**.

On the **Renew** page, select the **Manually Renew** tab.

Find an instance, and in the **Actions** column, click **Enable Auto-Renew**.

On the **Enable Auto-Renew** dialog box, click **Enable Auto-Renew**.


Then you can select the **Auto-Renew** tab and find the instance.

Deactivate auto-renewal

To deactivate the auto-renewal service for an instance, follow these steps:

Log on to the ECS console.

At the top of the page, choose **Billing Management > Renew**.



In the left-side navigation pane, click **Elastic Compute Service**.

On the **Renew** page, select the **Auto-Renew** tab.

Find an instance, and in the **Actions** column, click **Modify Auto-Renew**.

On the **Modify Auto-Renew** dialog box, select **Disable Auto-Renew** and click **OK**.

Then you can select the **Manually Renew** tab and find the instance.

User-defined data and metadata

Disks

Create a cloud disk

Cloud disks, also known as data disks, can be purchased from the ECS Management Console. Each user account can own up to 250 cloud disks simultaneously. Up to 16 data disks can be attached to any single ECS instance, with a maximum capacity of 32768 GB per data disk.

To purchase a cloud disk, perform the following:

Log on to the ECS console.

Click **Cloud Disks** in the left-side navigation pane.

Select your desired region, then click **Create Cloud Disk** in the top-right corner of the page.

Select a region and zone.

Note: A cloud disk can be attached to only a server in the same zone of the same region. Cloud disks do not support cross-regional functionality.

Select the cloud disk type, size, and quantity. Click **Buy Now** on the right side of the page.

Next step for Linux

For Linux instances, cloud disks must be attached, partitioned, and formatted before they can be displayed and used in the system.

- For details on attaching a data disk, see [Attach a data disk](#).
- For details on formatting partitions and mounting new partitions to an attached data disk, see [Format and mount a data disk](#).

Next step for Windows

For Windows instances, you must attach and format a cloud disk before using it.

- For details on attaching a data disk, see [Attach a data disk](#).
- For details on formatting an attached data disk, see [Format a data disk](#).

You can take a snapshot of an existing **system disk** or a **data disk**, and create a cloud disk from the snapshot. The new disk can be attached to any instance in **the same zone of the same region**.

Scenarios

If you need to access data from a snapshot, but do not want to roll back your disk, you can create a cloud disk from the snapshot to access the data you need.

For example, if your instance encounters a system disk failure, you can use an existing snapshot to create a cloud disk, and attach the disk to a healthy instance. In this way, you can restore the data of the impaired instance.

Disk Performance

SSD cloud disks and ultra cloud disks that are not created from snapshots can exhibit the maximum

performance of its capacity, no preconditioning is needed. However, for cloud disks created from snapshots, the initial performance can drop off because data has to be accessed from OSS before written into the disk.

We recommend that you write and read once to every data block before production use.

Prerequisites

You have created a snapshot for your instance.

Procedure

To create a cloud disk from a snapshot, follow these steps:

Log on to the ECS console.

On the navigation pane, select **Snapshots and Images** > **Snapshots**.

Find the expected snapshot and copy its ID.

On the navigation pane, select **Block Storage** > **Cloud Disks**.

Click **Create Cloud Disk**.

On the **Create Cloud Disk** page, do the following:

- i. Select a region and a zone. If you want to attach this disk to a instance, make sure that they are in the same region and zone.
- ii. For **Choose Storage**, click **Create from snapshot**, then select the snapshot you need.
- iii. For **Purchase Plan**, select an option.
- iv. Check **Overview**.
- v. Click **Buy Now**.
- vi. Confirm you order and make the payment.

You can go back to the **Cloud Disks** page, and click the refresh button next to **Create Cloud Disk**. You can find the cloud disk you just created with the **Disk Status** as **Available**.

Follow-up operations

You can attach the cloud disk to an instance, see [Attach a data disk](#).

Note:

- For Windows users, you can attach your new cloud disk to an instance by clicking **Attach** on the console.
- For Linux users, once you click **Attach** on the console, you must run the mount command.

Attach a data disk

The following disk types can be attached:

- Basic Cloud Disks
- Ultra Cloud Disks
- SSD Cloud Disks that serve as data disks.

Before attaching a cloud disk, an instance must meet the following requirements:

- The instance is in **Stopped** status.
- The security control marker is not Locked.
- The instance is not in payment arrears.

To attach a cloud disk, the following conditions must be adhered to:

- The cloud disk must be in available status.
- A single instance can have one system and up to 16 data disks attached. This includes disks of all types.
- A cloud disk can be attached to an instance of the same zone only.
- A cloud disk can be attached to only one instance at a time. Attachment to multiple instances is not supported.
- A cloud disk can be attached to any instance of the same region and zone. Both Subscription and Pay-As-You-Go instances are supported.
- When a cloud disk is acting as the system disk of an instance, it cannot be separately attached.

You can attach a disk through either of the following:

Instances Menu

Recommended if you require multiple disks to a single instance.

Disks Menu

Recommended if you require disks be attached to different instances.

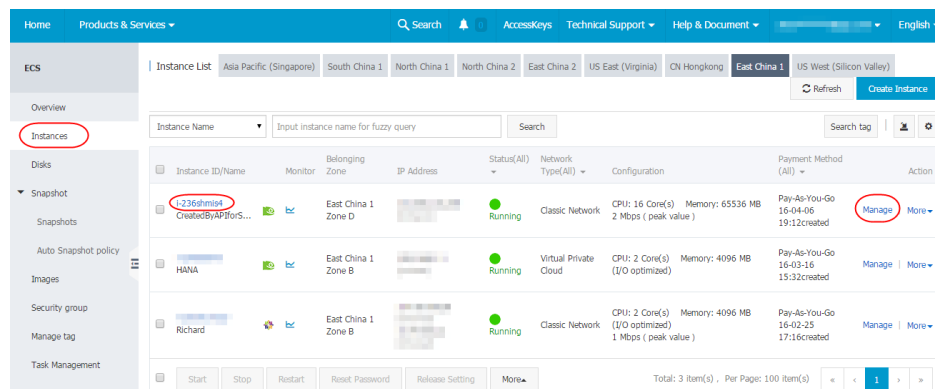
From the Instances menu

Log on to the ECS console.

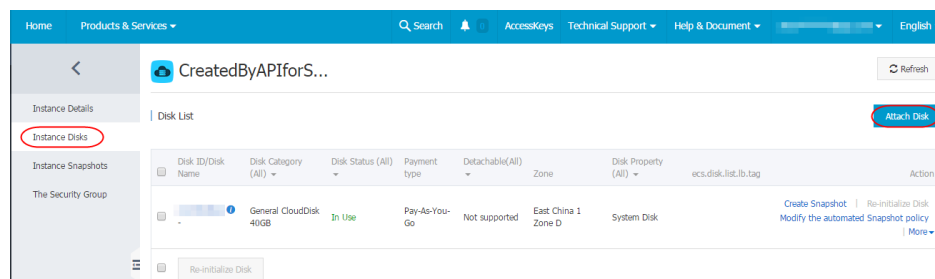
Click **Instances** in the left navigation bar.

Select your desired region.

Click the name of the instance for attachment or click **Manage**.



Click **Instance Disks** in the left navigation bar. The disks already attached to the instance is displayed.



Click **Attach Disk** on the right side of the page. Select **Available Devices** and **Target Disk** to attach the disk.

(Optional) Set whether disks are to be released with instances and whether snapshots are to be deleted with disks.

- **Release disk with instance**

When you release the instance, the disk will be released together.

- **Delete automatic snapshots when releasing disk**

When you release the disk, all auto snapshots will be deleted. However, the

snapshot you manually created will be retained. This option is not recommended.

The screenshot shows the 'Attach Disk' dialog box. At the top, it says 'Your Instance:' followed by a blurred instance ID. Below that, it states 'The instance still has 4 available devices.' The main section is titled 'Available Devices:' and shows a 'Default device' with a blue 'i' icon. Below this is a '*Target Disk:' dropdown menu with a blurred disk ID. There are two checkboxes: 'Release Disk with Instance' and 'Delete automatic snapshot when releasing disk', both of which are currently unchecked. An 'OK' button is located below the checkboxes. At the bottom of the dialog, there is an 'Important Note' box with orange text: 'After the cloud disk is attached, you must log in to the instance, then format and mount the new partition. Operation Guide: Format Partition/Mount Data Disk'. At the very bottom of the dialog, there are 'Attach' and 'Cancel' buttons.

After attaching a disk, you need to log on to the instance to format disk partitions and attach new partitions. For details, refer to **Next step** at the bottom of this section.

From the Disks menu

Log on to the ECS console.

Click **Disks** in the left navigation bar.

Select your desired region.

Find the disk to attach. The disk status must be **Available**.

Click **More > Attach**. Select the target instance and release action:

- **Release disk with instance**

When you release the instance, the disk will be released together.

- **Delete automatic snapshots when releasing disk**

When you release the disk, all auto snapshots will be deleted. However, the snapshot you manually created will be retained. This option is not recommended.

After attaching a disk, you need to log on to the instance to format disk partitions and attach new partitions. For details, refer to **Next step** at the bottom of this section.

Next step for Linux

After attaching a disk, you must log on to the instance to format disk partitions and mount new partitions. For detailed instructions, see [Format and mount a data disk](#).

Next step for Windows

After attaching a disk, you must log on to the instance to format disk partitions. For detailed instructions, see [Format a data disk](#).

ECS supports the detachment of Basic Cloud Disks, Ultra Cloud Disks, and SSD Cloud Disks serving as data disks. System disks cannot be detached. Detach disks on the **Instances** page or the **Disks** page.

Notes

Only the data disks in the **In Use** status can be detached.

Ensure the following based on the operating system of the instance:

- For a Linux instance, log on to the instance and run the `umount` command to unmount the data disk. After the command runs, log on to the ECS console and detach the disk.
- For a Windows instance, stop read and write operations on all file systems of the disk to ensure data integrity. Otherwise, data being read and written will be lost.

Procedures

On the Instances page

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Click the instance ID go to the **Instance Details** page.

In the left-side navigation pane, click **Instance Disks**, and you can see all the disks, both system disks and data disks, associated with the instance.

Find the disk to be detached and select **More > Detach**.

In the **Detach Disk** dialog box, read the note and click **Confirm Detaching**.

You have successfully detached the data disk from the instance.

On the Disks page

Log on to the ECS console.

In the left-side navigation pane, click **Disks**.

Select a region.

Find the disk to be detached and select **More > Detach**.

In the **Detach Disk** dialog box, read the note and click **Confirm Detaching**.

You have successfully detached the data disk from the instance.

After the procedures

You have to pay for the data disk even though it is not attached to an instance. So, if you do not need the data disk any more, **release** it.

By changing the system disk, you can change the operating system to your custom image, which may be:

- A custom image created by using an instance
- A custom image created by using a snapshot
- A custom image copied from another region
- An imported image
- A custom image shared by another account, which is a shared image

When changing a system disk, if you want to keep the operating system, environment configurations, and/or data, you can use the instance or a system disk to create a custom image, and then use the image to change the system disk. This document uses this scenario to describe how to change the image on a system disk to a custom image.

For instructions about how to change the OS to a public image, see [Change the system disk](#)

(public image).

Changing the system disk will not change your instance IP address.

Note: Regions that are not in mainland China do not support replacement between Linux and Windows. A Linux or Windows can be only replaced by a different version of the same operating system type.

Warning:

- Stopping an instance may disrupt traffic.
- Redeploying the runtime environment on the new system disk is required once it is stopped and this may disrupt traffic.
- Automatic snapshots and data from your original system disk will be lost once the system disk is replaced. Ensure that all necessary data has been backed up in advance. If you want to retain auto snapshots, see [disable releasing auto snapshots with disk](#).
- Manually created snapshots from the original system disk are retained but cannot be used to roll back the new system disk because the disk ID is changed. You can use the retained snapshots to create custom images.
- The original system disk will be deleted once the system disk is replaced.
- Ensure that the system disk has at least 1 GB of free space. Otherwise, the instance may fail to start after you change the system disk.

Procedure of changing the system disk

A complete procedure of changing the system disk includes the following steps:

- Step 1. Back up the current system disk by creating a snapshot.
- Step 2. Create an image from the snapshot.
- Step 3. Change the system disk and selecting a new OS.
- Step 4. Set auto snapshot policies for the new system disk.

To retain enough snapshot quota for the auto snapshot policy of the new disk, please delete unwanted snapshots before proceeding. If you wish to change the OS and do not need to retain the data from the current system disk, you can proceed directly to Step 3 of this section.

Step 1: Back up the current system disk by creating a snapshot

Note:

- Skip this step if you do not want to retain the data in the system disk.
- Do not create the snapshot during busy hours.
- It takes about 40 minutes to create a snapshot of 40 GB. Ensure you reserve sufficient time for it.
- Make sure the system disk has at least 1 GB free space. Otherwise, the instance may fail to

start after you change the system disk.

Log on to the ECS console.

Click **Instances** in the left navigation bar.

Select your desired region.

Select the instance containing the system disk with the OS to be changed.

Click **Instance Disks** on the left navigation bar.

Select the system disk with the OS to be changed, and then click **Create Snapshot**.

Enter a name for the snapshot.

Click **Instance Snapshots** on the left navigation bar to check the progress and status of the snapshot.

Step 2: Create an image from the snapshot

Note:

- Skip this step if you do not want to retain the data in the system disk.
- If you do want to retain the data in the current system disk, you need to create an image to replicate the data in the system disk.
- Make sure the system disk has at least 1 GB free space. Otherwise, the instance may fail to start after you change the system disk.

Select the snapshot created in Step 1 and click **Create Custom Image**.

Enter a name and description for the image.

Return to the navigation bar, and then click **Images**.

You can now check the process and status of the new image.

Step 3: Change the system disk

To change a system disk:

Log on to the ECS console.

Click **Instances** in the left navigation bar.

Select your desired region.

Stop the instance before changing the system disk. To do this, select the instance for which you wish to replace the system disk in the instance list and click **Stop**.

Once the instance is stopped, click **More > Change System Disk**.

Click **Custom Image** and select the image created in Step 2.

Click **Confirm Change**. Any expenditure that may have occurred will need to be paid at this time.

Step 4: Set auto snapshot policies for the new system disk

After changing the system disk, any auto snapshot policies you have set will no longer work for the new system disk, because the disk ID has changed. In this scenario, you must configure auto snapshot policies for the new system disk. For more information, see [Set auto snapshot policies for disks](#).

Step 5: Attach data disks (for Linux instances only)

For Linux instances, you must re-attach the data disks after changing the system disk, but you do not need to partition them. For more information, see [Attach a data disk](#).

By changing the system disk, you can change the operating system to a public image, for example, from Windows Server 2003 to Windows 2012.

Note: Regions that are not in mainland China do not support replacement between Linux and Windows. A Linux or Windows can be only replaced by a different version of the same operating system type.

Considerations for changing the system disk

Risks

- This operation requires you to stop your instance, which means interruption of your business. Therefore, perform this operation with caution.

- After replacement, you must redeploy the runtime environment on the new system disk. There is a possibility of a long interruption of your business. Therefore, perform this operation with caution.
- Replacing the system disk will result in the loss of the automatic snapshots and data on your original system disk. Make necessary backup in advance.

Note:

- To retain enough snapshot quota for the auto snapshot policy of the new disk, you can delete unwanted snapshots.
- Changing the system disk will not change your instance IP address.
- Manually created snapshots are retained after the replacement. However, since the disk ID is changed, you can no longer use the manually created snapshots on the original system disk to roll back the new system disk. The retained snapshots can be used to create custom images.
- After the system disk is replaced, the original system disk will be deleted.

Retain automatic snapshots

By default, the automatic snapshots will be released along with the disk. If you want to keep the automatic snapshots, see [Configure releasing auto snapshots together with disk](#).

Procedure of changing the system disk

A complete procedure of changing the system disk includes the following steps:

1. Create a snapshot for the current system disk.
2. Change the system disk.
3. Set automatic snapshot policies for the new system disk.
4. Attach data disks (for Linux instances only).

Step 1: Create a snapshot based on the current system disk

Skip this step if you do not want to retain the data on the system disk.

Do not create the snapshot during busy hours. It takes about 40 minutes to create a snapshot of 40GB. Ensure you reserve sufficient time for it.

Note: Make sure the system disk has at least 1GB free space; otherwise, the instance may fail to start after you change the system disk.

Log on to the ECS console.

Click **Instances** on the left navigation bar. Then click the region.

Click the instance for which you want to change the system disk.

Click **Instance Disks** on the left navigation bar.

Find the system disk you want to change, and then click **Create Snapshot**.

Enter a name for the snapshot.

Click **Snapshots** on the left navigation bar, you can check the progress and status of the snapshot.

Step 2: Change the system disk

To change a system disk:

Log on to the ECS console.

Click **Instances** on the left navigation bar. Then, select a region at the top of the page.

Stop the instance before changing the system disk. In the instance list, select the instance for system disk replacement and click **Stop** at the bottom.

After the instance is stopped, on the right end of the instance table, click **More > Change System Disk**.

A dialog box showing the considerations is displayed. Read the considerations carefully, and then confirm the operation.

Select a public image.

Set the password for Administrator or root.

Click **Pay Now**. Pay for the expenditure that incurred, if any.

An important message is prompted. Read it carefully. After confirming everything is correct, click **OK**.

Step 3: Configure the automatic snapshot policies

After changing the system disk, the automatic snapshot policies you have set will no longer work for the new system disk, because the disk ID has changed. In this case, you need to configure automatic snapshot policies for the new system disk. For more information, see [Set automatic snapshot policies for disks](#).

Step 4: Attach data disks (for Linux instances only)

For Linux instances, after you have changed the system disk, you need to attach the data disks again, but you don't need to partition them. For more information, see [Attach a data disk](#).

Resize cloud disks

You are provided with the **Resize Disk** feature to extend the capacity of data disks. The capacity of the disk after resizing varies by the disk types.

Data disk type	Current capacity	Capacity after resizing
Basic cloud disk	Any	2000 GB
SSD cloud disk or ultra cloud disk	= < 2048 GB	2048 GB
SSD cloud disk or ultra cloud disk	>2048 GB	Cannot be resized

According to the operating system of the instance, perform different steps:

- To resize a data disk attached to a Windows instance, see [Windows _ Resize a data disk](#).
- To resize a data disk attached to a Linux instance, see [Linux _ Resize a data disk](#).

If you want to resize a data disk attached to a Windows instance, see [Windows _ Resize a data disk](#).

As your business grows, the current capacity of your data disks may not be able to meet your data storage needs. You can use the **Resize Disk** feature to resize your data disks as necessary.

Note:

- We recommend that you manually **create a snapshot** to back up your data before

resizing your data disk.

- You can resize a data disk when the data disk is either in the **Available** status or in the **In Use** status.
- If a snapshot is being created for a data disk, you cannot resize the data disk.
- You can resize data disks, but not system disks or local disks.
- Resize the data disks that are attached to an instance only when the instance is in the **Running** or **Stopped** status. You must restart the instance in the ECS console to apply the changes. This action causes your instance to stop working and may cause your business to be interrupted, so please proceed with caution.

This example uses a data disk of the ultra cloud disk type and an ECS instance running 64-bit CentOS 7.3 to describe how to resize a data disk and extend the available capacity.

To resize a data disk, follow these steps:

Step 1. Resize a data disk in the ECS console.

Step 2. Log on to the instance to enable the available capacity.

Step 1. Resize a data disk in the ECS console

To resize a data disk in the ECS console, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Block Storage** > **Cloud Disks**, and select a region.

If the data disk you want to resize has been attached to an instance, in the left-side navigation pane, click the **Instances**, find the corresponding instance, go to the instance details page, and click **Instance Disks**.

Find the disk to be resized, and in the **Action** column, choose **More** > **Resize Disk**.

On the **Resize Disk** page, set **Capacity after resizing** (In this example, 30 GB). The capacity after resizing must be larger than the current capacity.

When the cost is calculated, click **Confirm to resize**.

(Optional.) If your data disk is attached to an instance, restart the instance in the ECS console to make the disk resize take effect.

After the disk resizing is completed in the console,

If the data disk is attached to an instance, log on to the instance to enable the extended storage space.

If the data disk is not attached to an instance, attach the disk to an instance in the console first, and then proceed depending on the data disk:

- If it is a brand new data disk, format and mount the data disk.
- If it has been formatted and partitioned, log on to the instance to enable the extended storage space.

Step 2. Log on to the instance to enable the available capacity

In this example, the data disk is attached to a Linux instance running the 64-bit CentOS 7.3. The data disk before resizing has only one primary partition (/dev/vdb1, ext4 file system), the mount point of the file system is /resizetest, and after resizing is completed, the data disk still has only one primary partition.

Log on to the instance.

Run the `umount [file system name]` command to unmount the primary partition.

```
umount /dev/vdb1
```

Run the `df -h` command to check whether the unmounting is successful. If you do not see the /dev/vdb1 information, unmounting is successful. The following is the sample output.

```
[root@iXXXXXX ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.5G 36G 4% /
devtmpfs 487M 0 487M 0% /dev
tmpfs 497M 0 497M 0% /dev/shm
tmpfs 497M 312K 496M 1% /run
tmpfs 497M 0 497M 0% /sys/fs/cgroup
tmpfs 100M 0 100M 0% /run/user/0
```

Run the `fdisk` command to delete the original partition and create a new partition:

If you use the parted tool to manipulate partitions, you cannot use it in conjunction with fdisk. Otherwise, this results in an inconsistent first sector of the partition.

Instructions on how to use the parted tool can be found [here](#).

- i. Run the fdisk -l command to list the partition details and record the final size of the data disk and its first sector before resizing.
- ii. Run the fdisk [device name of data disk] command. In this example, the device name is /dev/vdb.
- iii. Type d and press the Enter key to delete the original partition.

Deleting a partition does not cause loss of data in the data disk.

- iv. Type n and press the Enter key to start creating a new partition.
- v. Type p and press the Enter key to create a primary partition. In this example, you are creating a single-partition data disk, so it is sufficient to create one primary partition.

If you want to create more than four partitions, create at least one extended partition, that is, type e.

- vi. Type the partition number and press the Enter key. In this example, only one partition is to be created, so type 1.
- vii. Type a number for the First sector: For data consistency, the number for the First sector must be identical with that of the original partition. In this example, press the Enter key to use the default value of 1.

If you find that the First sector is not identical with the recorded one, you may have used the parted tool for partitioning. In that case, stop the current fdisk operation and use **parted** to start over again.

- viii. Type a number for the last sector: Because only one partition is to be created in this example, press the Enter key to use the default value.
- ix. Type wq and press the Enter key to start partitioning.

```
[root@iXXXXXX ~]# fdisk /dev/vdb
Welcome to fdisk (util-linux 2.23.2).
Changes will remain in memory only, until you decide to write them.
Be careful before using the write command.
Command (m for help): d
Selected partition 1
Partition 1 is deleted
Command (m for help): n
Partition type:
 p primary (0 primary, 0 extended, 4 free)
 e extended
Select (default p):
```

```

Using default response p
Partition number (1-4, default 1):
First sector (2048-62914559, default 2048):
Using default value 2048
Last sector, +sectors or +size{K,M,G} (2048-62914559, default 62914559):
Using default value 62914559
Partition 1 of type Linux and of size 30 GiB is set
Command (m for help): wq
The partition table has been altered!
Calling ioctl() to re-read partition table.
Syncing disks.

```

If you are using the parted tool, type `p` to list the current partition details. If any partition is displayed, use `rm + serial number` to delete the original partition table, then run the `unit s` command to specify the start unit, calculated by the number of sectors, and finally run the `mkpart` command to create it, as shown in the following figure.

```

[root@:~]# parted /dev/xvdb
GNU Parted 3.1
Using /dev/xvdb
Welcome to GNU Parted! Type 'help' to view a list of commands.
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 5369MB
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
(parted) unit s
(parted) mkpart primary ext3 56 5369MB
Warning: The resulting partition is not properly aligned for best performance.
Ignore/Cancel? i
(parted) p
Model: Xen Virtual Block Device (xvd)
Disk /dev/xvdb: 10485760s
Sector size (logical/physical): 512B/512B
Partition Table: gpt
Disk Flags:

Number  Start  End  Size  File system  Name  Flags
1       56s    10485726s  10485671s  ext3        primary

```

(Optional) For some operating systems, the file system may be automatically mounted to the mount point after partitioning. We recommend that you run the `df -h` command to check the disk space and usage. Run the `umount [file system name]` to unmount the file system again.

Check the file system and resize the file system.

```

e2fsck -f /dev/vdb1 # check the file system
resize2fs /dev/vdb1 # resize the file system

```

Note:

- Running the `e2fsck` command is time-consuming because the system needs to check and revise the file system metadata during that process, so be patient.
- Properly running the `e2fsck` command and the `resize2fs` command does not cause data loss.

The following is the sample output.

```
[root@iXXXXXX ~]# e2fsck -f /dev/vdb1
e2fsck 1.42.9 (28-Dec-2013)
Pass 1: Checking inodes, blocks, and sizes
Pass 2: Checking directory structure
Pass 3: Checking directory connectivity
Pass 4: Checking reference counts
Pass 5: Checking group summary information
/dev/vdb1: 11/1835008 files (0.0% non-contiguous), 159218/7339776 blocks
[root@iXXXXXX ~]# resize2fs /dev/vdb1
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/vdb1 to 7864064 (4k) blocks.
The filesystem on /dev/vdb1 is now 7864064 blocks long.
```

Mount the resized file system to the original mount point (in this example, `/resizetest`).

```
mount /dev/vdb1 /resizetest
```

Run the `df -h` command to check disk storage space and usage. If the correct information about the resized file system is displayed, the mounting is successful and the resized file system is ready for use.

After the mounting is completed, you can use the resized file system without restarting the instance.

The following is the sample output.

```
[root@iXXXXXX ~]# df -h
Filesystem Size Used Avail Use% Mounted on
/dev/vda1 40G 1.5G 36G 4% /
devtmpfs 487M 0 487M 0% /dev
tmpfs 497M 0 497M 0% /dev/shm
tmpfs 497M 312K 496M 1% /run
tmpfs 497M 0 497M 0% /sys/fs/cgroup
tmpfs 100M 0 100M 0% /run/user/0
/dev/vdb1 30G 44M 28G 1% /resizetest
```

If you want to resize a data disk attached to a Linux instance, see [Linux _ Resize a data disk](#).

As your business grows, the current capacity of your data disks may not be able to meet your data storage needs. You can use the **Resize Disk** function to resize your data disks as necessary.

Note:

- Resize the data disks that are attached to the instance only when the instance is in the **Running** or **Stopped** status. The changes are applied when you restart the instance in the ECS console. This action stops your instance from working and interrupts your business. Hence, proceed with caution.
- We recommend that you manually **create a snapshot** to back up your data before resizing a data disk.
- You can resize a data disk when the data disk is either in the **Available** status or in the **In Use** status.
- If a snapshot is being created for a data disk, you cannot resize the data disk.
- You can resize data disks, but not system disks or local disks.

This example uses a data disk of the ultra cloud disk type and an ECS instance running 64-bit Windows Server 2008 R2 Enterprise Edition to show how to resize a data disk and extend the available capacity. In this example, the current disk capacity is 20 GB, and we resize it to 24 GB.

To resize a data disk, follow these steps:

Step 1. Resize a data disk in the ECS console

Step 2. Log on to the instance to enable the extended storage space

Step 1. Resize a data disk in the ECS console

To resize a data disk in the ECS console, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Block Storage** > **Cloud Disks**, and select a region.

Note:

If the data disk you want to resize is attached to an instance, click **Instances** in the left-side navigation pane, find the instance, go to the **Instance Details** page, and then click **Instance Disks**.

Find the disk to be resized, and in the **Action** column, choose **More** > **Resize Disk**.

On the **Resize Disk** page, set **Capacity after resizing** (In this example, 24 GB). The capacity after resizing must be larger than the current capacity.

When the cost is calculated, click **Confirm to resize**.

(Optional). If your data disk is attached to an instance, restart the instance in the ECS console to make the disk resize take effect.

Once the data disk resizing completes, you can do the following:

If the data disk is attached to an instance, log on to the instance to enable the extended storage space.

If the data disk is not attached to an instance, attach the disk to an instance in the console first, and then proceed depending on the data disk:


- If it is not formatted or partitioned, format and mount the data disk.
- If it is formatted and partitioned, log on to the instance to enable the extended storage space.

Step 2. Log on to the instance to enable the extended storage space

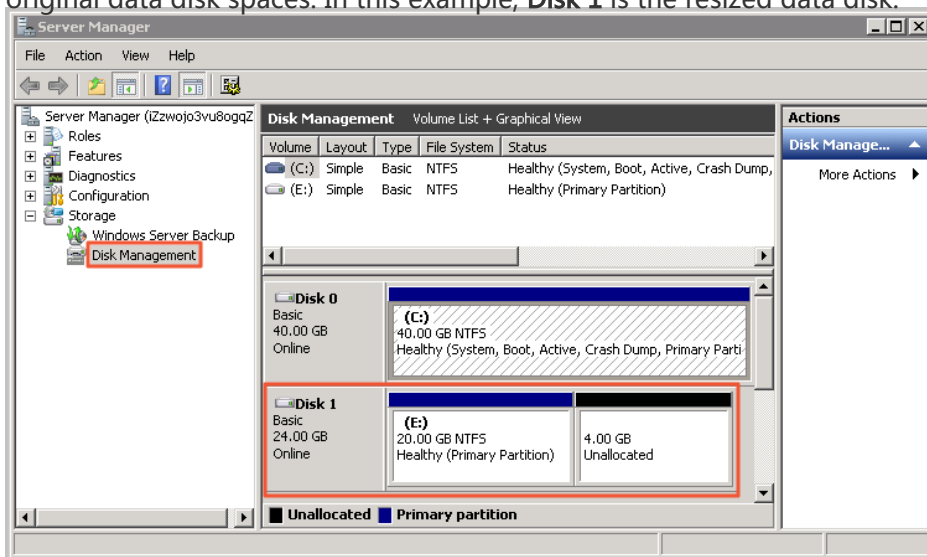
To resize a data disk within the instance, follow these steps:

Log on to the instance.


On the Windows Server desktop, click the **Server Manager** icon



In the left-side navigation pane of **Server Manager**, choose **Storage > Disk Management**. In the disk management area, you can see the relationship between the new and the original data disk spaces. In this example, **Disk 1** is the resized data disk.

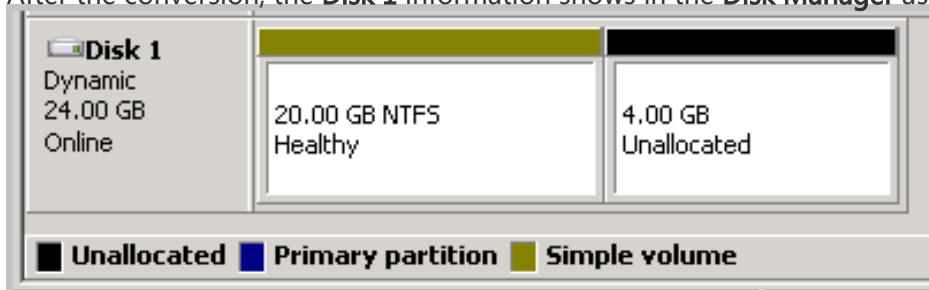


Right click **Disk 1**, select **Convert to Dynamic Disk**, and follow the wizard to convert a basic disk to a dynamic disk.

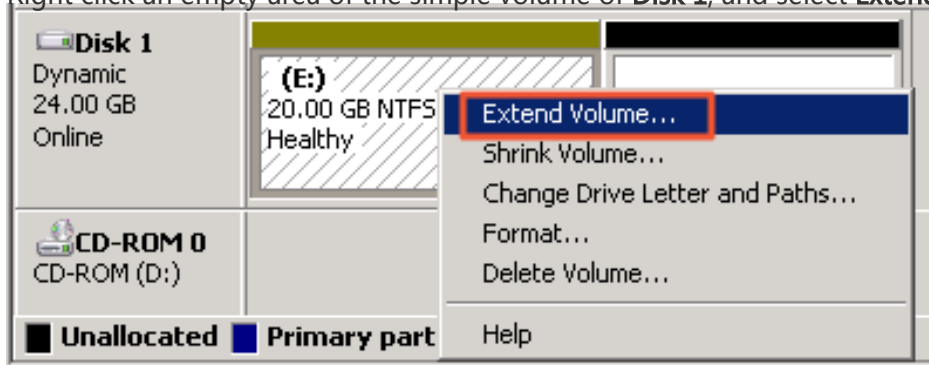
**Note:**

Converting a basic disk into a dynamic disk unmounts the disk from the system. Applications installed on the data disk, if any, are temporarily unavailable during the conversion process. The conversion process does not cause any data loss.


After the conversion, the **Disk 1** information shows in the **Disk Manager** as follows.



Right click an empty area of the simple volume of **Disk 1**, and select **Extend Volume**.



Follow the **Extend Volume Wizard** to extend the volume. When the wizard is complete, the new data disk space is automatically merged into the original volume and the **Disk 1** information showed in the **Disk Manager** as follows.



Note:

On Windows Server 2003, the extended storage space is added to the data disk but it is displayed as a separate volume in **Disk Manager**. On Windows Server 2003, one separate volume is created for each expansion and is not merged into the original volume, which does not affect the availability of the extended storage space.

You have resized a data disk successfully and the extended storage space is ready for use.

Re-initialize a disk

Disk re-initialization restores a disk to its initial state and settings.

Warning: Re-initializing a disk will erase all data on that disk. Ensure that you have backed up all necessary data before proceeding.

Note:

- The operating system and version of the instance is retained, and will be restored to its initial state and settings.
- The IP address of the instance will not change. The data on the original system disk will be cleared, but the automatic backup of snapshots on the instance will be retained, and can be used to roll back the applications on the instance.
- If you are re-initializing a data disk, you do not need to attach it after re-initialization.

Perform the following steps to re-initialize a data disk:

Log on to the ECS console.

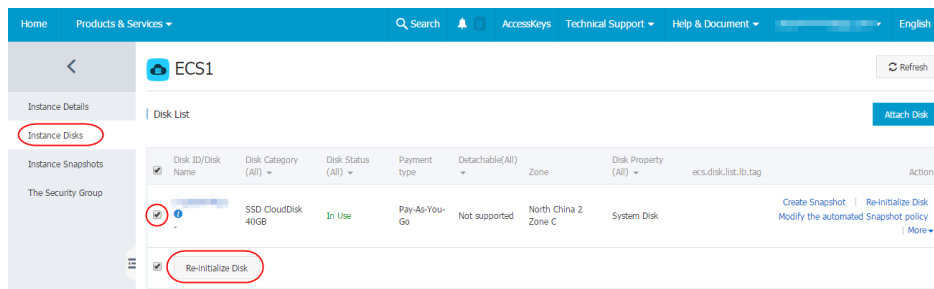
Click **Instances** in the left navigation bar.

Select your desired region.

Select the instance for disk re-initialization, and click **Stop**.

Select the instance name and then click **Instance Disks** in the left navigation bar.

Select one or more disks to re-initialize and click **Re-initialize Disk**.



Enter a new login password once the re-initialization is finished, and then click **Confirm Re-initializing Disk**.

Roll back a disk

Disk rollback restores a disk to a state and setting from previous point in time.

Note:

- Snapshot rollback is a permanent action and cannot be reversed. Once rollback is completed, the original data cannot be restored. You are recommended to proceed with caution.
- Disk rollback can only be performed when the instance is completely stopped.

To rollback a disk, perform the following:

Log on to the ECS console.

Click **Instances** in the left navigation bar.

Select your desired region.

Select the instance for disk rollback, and click **Stop**.

Click the instance name. Then, click **Snapshots** > **Snapshots** in the left navigation bar.

Select the snapshot for rollback. You can select only one snapshot at a time.

Click **Disk Rollback**.

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

View monitoring information of a disk

To view the monitoring information of a disk, such as IOPS and BPS, perform the following:

Log on to the ECS console.

Select a disk to view monitoring information by using one of the following methods:

- Click the instance that the disk is attached to from the instance list page and click **Instance Disks**.
- Locate the disk from the **Disks** list.

Click **Disk Monitoring** to view the monitoring information of the selected disk.

Note: You can select pre-set time segments to initiate regular monitoring periods, such as 1 hour, 6 hours, 1 day, and 7 days. You can also set custom monitoring start and end times.

If you no longer need a data disk, you can release it. Otherwise, we continue charging you for it.

Warning: Releasing a data disk is a permanent action and cannot be reversed. Once deleted, the original data on the data disk cannot be restored. We recommend that you proceed with caution.

Prerequisites

- Make sure that the data disk is in the **Available** status. If the data disk is in the **In Use** status, detach it from the instance in the ECS console.
- Make sure that you have backed up all the data you need in the data disk. You can create snapshots for backup.

Procedure

To release a data disk, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Disks**.

Select a region.

Find the disk that you want to release, and in the **Action** column, select **More > Release**.

In the **Release** dialog box, read the note and click **Confirm Release**.

Now you have successfully released the data disk.

The billing method of a cloud disk depends on how it is created:

For cloud disks created with Subscription (monthly or yearly subscription) instances, upfront payment is required for the service to be ready for use. For more information, see [Subscription](#).

For cloud disks created jointly with Pay-As-You-Go instances or separately created are billed on a Pay-As-You-Go basis. For more information, see [Pay-As-You-Go](#).

You can change the billing method of a cloud disk, as shown in the following table.

Conversion of billing methods	Features	Effective date	Suitable for
Subscription —> Pay-As-You-Go	Renew and downgrade configurations	Effective from the next billing cycle	Subscription cloud disks attached to Subscription instances. The billing method of the system disk cannot be changed.
Pay-As-You-Go —> Subscription	Upgrade configurations	Effective immediately	Pay-As-You-Go data disks attached to Subscription instances. The billing method of the system disk cannot be changed.
	Switch from Pay-As-You-Go to Subscription		The system disks and data disks attached to the Pay-As-You-Go instances.

Snapshots

You can create instance snapshots to save the system state from a certain point in time for data backup or to create images.

Note:

- Creation of the first snapshot will take relatively longer than subsequent snapshots due to the first snapshot being a full snapshot. However, depending on the amount of changed data since previous snapshots, the length of time for each snapshot creation may vary.
- Creating snapshots of a disk may reduce disk performance.
- We recommend that you not create snapshots during peak traffic hours.
- Manually created snapshots, unlike automatic snapshots, will be retained until they are manually deleted.

To create a snapshot, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Cloud Disks**.

Select a region.

Select a system or data disk for which you want to create a snapshot. You can only select one disk at a time.

Click **Create Snapshot**.

Enter a name for the snapshot and click **OK**.

To view all snapshots, go to left-side navigation pane and click **Snapshots**.

To view all snapshots, click **Snapshots** under **Snapshots & Images** in the left-side navigation pane.

An automatic snapshot policy is a set of defined parameters for automatically creating snapshots.

You can create a maximum of 100 automatic snapshot policies in each region.

Create an automatic snapshot policy

To create an automatic snapshot policy, perform the following:

Log on to the ECS console.

In the left-side navigation pane, select **Snapshots & Images > Automatic Snapshot Policies**.

Click **Create Automatic Snapshot Policy**.

Define automatic snapshot policy parameters:

- **Automatic Snapshot Policy**

This parameter is the name of the automatic snapshot policy. It must contain 2 ~ 128 characters and begin with English letters or Chinese characters. It can include digits and periods (.), underscores (_), and hyphens (-).

- **Time**

Defines the time of day for automatically creating snapshots. There are 24 snapshot creation points available between 00:00 and 23:00.

- **Repeated day**

There are seven available repetition day configurations.

- **Retention period**

Defines the number of days a snapshot can be retained. This parameter can be set between 1–65536 days, or permanently. By default, it is set to 30 days.

Create Automatic Snapshot Policy

×

- ECS Snapshot 2.0 data service provides a quota of 64 snapshots for each disk. When the maximum number of snapshots for a disk has been reached, the oldest automatic snapshot generated by the automatic snapshot policy will be deleted when you create a new snapshot.
- If the time for creating a new snapshot exceeds the interval between two automatic snapshot points in time (due to reasons such as huge disk data volume) the next point in time will be skipped automatically without a snapshot being created. For example, if a user sets 9:00, 10:00, and 11:00 as the automatic snapshot points in time, and it takes 70 minutes to create the snapshot scheduled for 9:00, or it is to be completed at 10:10, the scheduled snapshot for 10:00 will be skipped and the next snapshot will be created at 11:00.
- The current snapshot policy execution time defaults to the Eastern Time Zone (UTC + 8). You should adjust this time policy according to the actual business requirements.

*Automatic Snapshot Schedule:

It must contain 2-128 characters and begin with English letters or Chinese characters. It can include numbers and the characters ".", "_", and "-".

*Time:

☐ 00:00 ☐ 01:00 ☐ 02:00 ☐ 03:00 ☐ 04:00 ☐ 05:00
☐ 06:00 ☐ 07:00 ☐ 08:00 ☐ 09:00 ☐ 10:00 ☐ 11:00
☐ 12:00 ☐ 13:00 ☐ 14:00 ☐ 15:00 ☐ 16:00 ☐ 17:00
☐ 18:00 ☐ 19:00 ☐ 20:00 ☐ 21:00 ☐ 22:00 ☐ 23:00

*Repeated Day:

☐ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday
☐ Saturday ☐ Sunday

Retention Period:

☒ Custom Duration

30

 day(s) The retention period
☐ Permanent

OK

Cancel

Click OK.

Delete an automatic snapshot policy

If you no longer need an automatic snapshot policy, navigate to the policy you want to delete and then click **Delete Automatic Snapshot Policy**.

You can apply an automatic snapshot policy to disks according to your business needs.

Note:

- Creating snapshots may disturb read and write operations on your disk. We recommend that

you set the creation time of automatic snapshots to periods when service load is low to reduce effects on your service.

- Automatic snapshot policies cannot be applied to basic cloud disks when they are not in use.
- Snapshots that are manually created do not conflict with automatic snapshots. However, if an automatic snapshot is being created on a disk, you must wait for it to finish before manually creating a snapshot.

You can apply an automatic snapshot policy to a disk through either of the following:

Cloud Disks menu

For applying an automatic snapshot policy to a specific disk.

Snapshots & Images menu

For applying a unified automatic snapshot policy to several or all disks.

From the Cloud Disks menu

To apply an automatic snapshot policy through the **Cloud Disks** menu, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Cloud Disks**.

Select a region.

Select the disk for which you want to execute the policy and click **Automatic Snapshot Policy**.

Enable the automatic snapshot function and select the desired snapshot policy.

Click **OK**.

From the Snapshots & Images menu

To apply or disable an automatic snapshot policy, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Snapshots & Images** > **Automatic Snapshot Policy**.

Select a region.

Select the automatic snapshot policy you want to apply and click **Set Disk**.

- To enable the automatic snapshot policy, select the **Disk without Preset Policy** tab to view the disks. Select the disks in which you want to enable the policy, and then click **Enable the Automatic Snapshot**.

Modify the automated Snapshot policy ✕

After you enable the automated snapshot policy, your Snapshot will be managed according to the automated Snapshot policy.

☒ Disk without preset policy ☐ Disk with preset policy

Disk Name Please enter disk name for fuzzy query

<input checked="" type="checkbox"/>	Disk ID/Disk Name	Disk Category (All) ▾	Disk Property (All) ▾	Action
<input checked="" type="checkbox"/>	d-28eyf2ur4 ⓘ	General CloudDisk 40GB	System Disk	<input type="button" value="Enable autosnapshot"/>
<input checked="" type="checkbox"/>	Enable autosnapshot			Total: 1 item(s) , Per Page: 20 item(s) « < 1 > »

- To disable the automatic snapshot policy, select the **Disk with Preset Policy** tab to view the disks. Select the disks in which you want to disable the policy, and then click **Disable the Automatic Snapshot**.

Modify the automated Snapshot policy ✕

After you enable the automated snapshot policy, your Snapshot will be managed according to the automated Snapshot policy.

☐ Disk without preset policy ☒ Disk with preset policy

Disk Name Please enter disk name for fuzzy query

<input checked="" type="checkbox"/>	Disk ID/Disk Name	Disk Category (All) ▾	Disk Property (All) ▾	Action
<input checked="" type="checkbox"/>	d-28eyf2ur4 ⓘ	General CloudDisk 40GB	System Disk	<input type="button" value="Disable autosnapshot"/>
<input checked="" type="checkbox"/>	Disable autosnapshot			Total: 1 item(s) , Per Page: 20 item(s) « < 1 > »

Retain automatic snapshots when releasing disk

To retain auto snapshots when releasing the disk, perform the following:

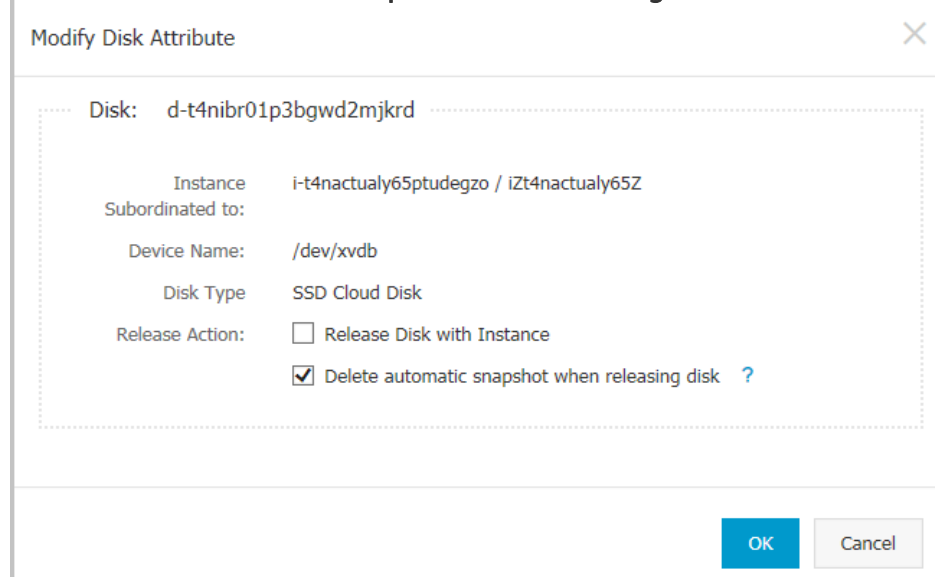
Log on to the ECS console.

In the left-side navigation pane, click **Cloud Disks**.

Select a region.

Select the disk that you want to configure and click **More > Modify Attributes**.

Deselect **Delete automatic snapshots when releasing disk**, and then click **OK**.



Modify Disk Attribute

Disk: d-t4nibr01p3bgwd2mjkrd

Instance Subordinated to: i-t4nactualy65ptudegzo / iZt4nactualy65Z

Device Name: /dev/xvdb

Disk Type: SSD Cloud Disk

Release Action: ☐ Release Disk with Instance ☒ Delete automatic snapshot when releasing disk ?

OK Cancel

You may want to delete a snapshot or an automatic snapshot policy.

Delete a snapshot

When you no longer need a snapshot, or you have reached your snapshot quota, you can delete snapshots to free up space.

Note:

- Deleting a snapshot is a permanent action and cannot be reversed. Once deletion is completed, the original snapshot cannot be restored. Proceed with caution.
- If a snapshot has been used to create a custom image, you must delete the associated image before you can delete the snapshot.

To delete a snapshot, perform the following:

Log on to the ECS console.

In the left-side navigation pane, select **Snapshots & Images > Snapshots**.

Select a region.

Select the snapshots you want to delete.

Click **Delete** at the bottom of the window.

Click **OK**.

Delete an automatic snapshot policy

Log on to the ECS console.

In the left-side navigation pane, select **Snapshots & Images > Automatic Snapshot Policy**.

Select a region.

Find the target automatic snapshot policy, and in the **Action** column, click **Delete Automatic Snapshot Policy**.


In the dialog box, confirm information and click **OK**.

Images

Create custom image

Custom images help you run ECS instances effectively by allowing you to create multiple ECS instances with identical OS and environment data to meet scaling requirements.

Custom images are based on ECS disk snapshots. You can set up identical or different configurations for ECS instances that are created from images.



Considerations

Custom images are subject to the following restrictions:

- If the ECS used for creating a custom image expires, or the data is erased (that is, the system disk used for the snapshot expires or is released), the custom image and the ECS instances created from the custom image will not be affected. However, auto snapshots will be cleared when an ECS instance is released.
- You can upgrade the CPU, memory, bandwidth, hard drive, and, other configurations of ECS instances activated using a custom image.
- Custom images cannot be used across regions.
- A custom image applies to any ECS payment method, either yearly/monthly subscription or Pay-As-You-Go. Custom images for ECS instances under yearly/monthly subscription plans can be used to create Pay-As-You-Go instances, and vice versa.

When creating custom images on a Linux instance, adhere to the following:

- Do not load data disk information in the `/etc/fstab` file. Otherwise, instances created using this image will not start.
- It is recommended that you unmount all data disks before taking a snapshot and creating an image. Otherwise, ECS instances that are created based on this custom image may not start.
- Modifying the default login user name **root** is not allowed (Linux OS only).

Operating procedure

To create a custom image from a snapshot, perform the following:






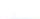




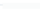
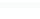
Log on to the ECS console.

Click **Snapshots** > **Snapshots** in the left navigation bar.

Select your desired region.

Select a snapshot with the disk attribute of **System Disk** and click **Create Custom Image**.

Note: Data disks cannot be used to create custom images.

<div>Elastic Computing Se...</div> <div>Overview</div> <div>Instances</div> <div>Disks</div> <div>Snapshots</div> <div>Snapshots</div> <div>Automatic Snapshot P...</div> <div>Images</div> <div>Security Groups</div> <div>Manage Tags</div> <div>Operation Logs</div>	test2		Disk		16:29:48		Create Custom Image	
	<input type="checkbox"/>	  jgf	20G	Data Disk	2016-12-27 16:29:34	100%	Success	Disk Rollback Create Custom Image
	<input type="checkbox"/>	  s0r5oi	40G	System Disk	2016-12-21 11:12:08	100%	Success	Disk Rollback Create Custom Image
	<input type="checkbox"/>	  az680	40G	System Disk	2016-12-13 11:07:47	100%	Success	Disk Rollback Create Custom Image
	<input type="checkbox"/>	  3gi3	40G	System Disk	2016-11-25 08:57:49	100%	Success	Disk Rollback Create Custom Image
	<input type="checkbox"/>	 	37G	Data Disk	2016-08-05 13:38:07	100%	Success	Disk Rollback Create Custom Image
	<input type="checkbox"/>	 	40G	System Disk	2016-03-14 16:00:02	100%	Success	Disk Rollback Create Custom Image

In the displayed dialog box, you can view the snapshot ID. Enter a name and description for the custom image.

(Optional) Click **Add Data Disk Snapshot** to select multiple snapshots of data disks for the image.

Note: If the snapshot disk capacity is left blank, an empty disk will be created with the default capacity of 5 GB. If you select available snapshots, the disk size is the same as the size of these snapshots.

Create Custom Image

When creating a custom image with Linux system, please do not load data disk information in the /etc/fstab file. Otherwise, you cannot launch the instance created through the image.

System Snapshot ID:

* Image Name:

It must contain 2-128 characters and begin with English letters or Chinese characters. It can include numbers and the characters ".", "_", and "-".

* Image Description:

It must contain 2-256 characters and it cannot begin with http:// or https://

Add Data Disk Snapshot

Snapshot Details:

Snapshot ID	Device Name:	Disk Capacity:	Action
Snapshot ID	/dev/xvda	GB	Delete

Add

1. Leaving the snapshot ID blank will create an empty disk. Default capacity: 5 GB, with up to 2,000 GB supported.

2. If a snapshot ID is selected, the default disk capacity will be the snapshot capacity.

3. If the device name is blank, it will be randomly allocated.

Create

Cancel

Click **Create**. The custom image is successfully created.

(Optional) To view images you have created, select **Images** in the left navigation bar.

FAQ for images of Linux instances

How to unmount a disk and delete disk table data?

If `/dev/hda5` is attached to `/mnt/hda5`, run any of the following three commands to detach the file system:

```
umount /dev/hda5
umount /mnt/hda5
umount /dev/hda5 /mnt/hda5
```

`/etc/fstab` is an important configuration file in Linux. It contains file system details and storage devices attached at startup.

If you do not want to mount a specified partition when starting the VM, delete the corresponding lines from `/etc/fstab`. For example, you can delete the following statement to disconnect `xvdb1` at startup: `/dev/xvdb1 /leejd ext4 defaults 0 0`.

How to determine whether a data disk is detached and a custom image can be created?

You must ensure that the auto attach data disk statement line has been deleted from the `fstab` file.

Use the `mount` command to view information on all mounted devices. Ensure that the execution results do not contain the information of the data disk partition.

Relevant configuration files

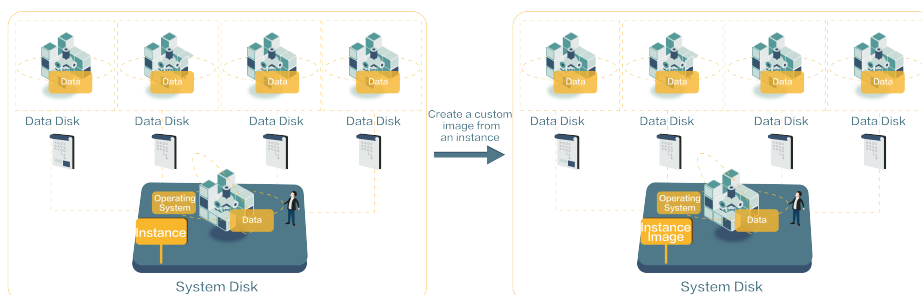
Before creating an image, ensure the key configuration files from the following table have not been modified; otherwise, the new instance will not be able to start.

Configuration File	Purpose	Risks if changed
<code>/etc/issue*</code> , <code>/etc/*-release</code> , <code>/etc/*_version</code>	For system release and version	Modifying <code>/etc/issue*</code> will make the system release version unidentifiable, and cause instance creation failure.
<code>/boot/grub/menu.lst</code> , <code>/boot/grub/grub.conf</code>	For system boot	Modifying <code>/boot/grub/menu.lst</code> will result in kernel loading failure, and the system will not be able to start.
<code>/etc/fstab</code>	For mounting partitions during boot.	Modifying it will cause partition mounting failure, and the system will not be able to start.
<code>/etc/shadow</code>	For storing system passwords.	If this file is set to read-only, the password file cannot be edited, and instance creation

		will fail.
/etc/selinux/config	For system security policies	Modifying /etc/selinux/config and enabling SELinux will result in start failure.

You can create a custom image using an ECS instance, namely, you fully copy all its disks and pack them into an image.

During this process, snapshots are automatically created for all disks of the instance, including the system disk and data disks. All the created snapshots compose a new custom image. See the following picture.



In addition, you can create a custom image based on the snapshot, see [Create a custom image using a snapshot](#) for instruction.

Prerequisites

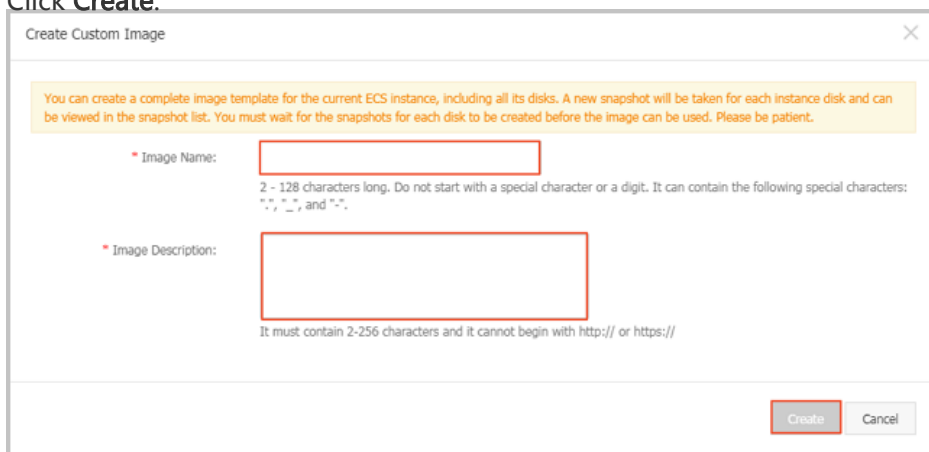
- To prevent the data privacy breach, make sure you delete all the confidential data in the ECS instance before creating a custom image.
- During creation, do not change the status of the instance. Do not stop, start, or restart the instance.
- If your custom image contains the data on the data disk, new data disk along with the ECS instance are created together. The data on the data disk duplicates the data disk snapshot in your custom image according to the mount device.
- You can export custom images that contain data of data disks.
- You cannot use a custom image which contains the data on the data disk to replace the system disk.

Procedure

See the following steps to create a custom image.

1. Log on to the ECS console.
2. Click **Instances** on the left-side navigation pane.
3. Select a region.

4. Select an instance, and then choose **More > Create Custom Image**.
5. Enter the name and description.
6. Click **Create**.



Follow-up operation

After creating the custom image, you may want to **Create an instance using a custom image**.

Packer is a convenient open-source tool to create custom images. It runs on major operating systems. This document provides information about how to install and use Packer. With Packer, you can easily create a custom image by using only one or two lines of commands.

Prerequisites

You must have the AccessKey ready. For more information, see **Create AccessKey**.

Note:

The AccessKey has a high level of account privileges. To avoid improper operations and data breach, we recommend that you **create a RAM user**, and act as a RAM user to **create your AccessKey**.

Create a custom image

Step 1. Install Packer

Go to the official download page of Packer where you can choose and download the version of Packer for your operating system. Follow these steps or visit the official installation page of Packer for how to install Packer.

To install Packer on a Linux server

Take CentOS 6.8 64-bit as an example:

1. Connect and log on to the Linux server. If the server you want to connect to is an ECS Linux instance, see [Connect to a Linux instance](#).
2. Run `cd /usr/local/bin` to go to the `/usr/local/bin` directory.

Note:

The `/usr/local/bin` directory is an environment variable directory. You can install Packer to this directory or another directory that has been added to the environment variable.

3. Run `wget https://releases.hashicorp.com/packer/1.1.1/packer_1.1.1_linux_amd64.zip` to download the Packer installer. You can visit the [official download page of Packer](#) to download installers for other versions of Packer.
4. Run `unzip packer_1.1.1_linux_amd64.zip` to unzip the package.
5. Run `packer -v` to verify Packer's installation status.
If the Packer version number is returned, you have successfully installed Packer.
If error command not found prompt is returned, Packer has not been correctly installed.

To install Packer on a Windows server

Take Windows Server 2012 64-bit as an example:

1. Connect and log on to the Windows server. If the server you want to connect to is an ECS Windows instance, see [Connect to a Windows instance](#).
2. Open the [official download page of Packer](#) and select an appropriate Packer installer for 64-bit Windows.
3. Unzip the package to a specified directory and install Packer.
4. Define the directory for Packer in the PATH environment variable.
 - i. Open the **Control Panel**.
 - ii. Select **All Control Panel Items > System > Advanced System Settings**.
 - iii. Click **Environment Variable**.
 - iv. Find **Path** in the system variable list.
 - v. Add the Packer installation directory to the **Variable Value**, such as `C:\Packer` as seen in this example. Separate multiple directories with half-width semicolons (;). Click **OK**.
5. Run `packer.exe -v` in CMD to verify Packer's installation status.
If the Packer version number is returned, you have successfully installed Packer.
If error command not found prompt is returned, Packer has not been correctly installed.

Step 2. Define a Packer template

To create a custom image by using Packer, firstly, create a JSON format template file. In the template, specify the **Alibaba Cloud Image Builder** and **Provisioner** for the custom image to be created. Packer has diverse provisioners for you to choose from when configuring the content

generation mode of the custom image.

In the following alicloud JSON file, we have used the **Shell** provisioner as an example to illustrate how to define a Packer template.

Create a JSON file named alicloud and paste the following content:

```
{
  "variables": {
    "access_key": "{{env `ALICLOUD_ACCESS_KEY`}}",
    "secret_key": "{{env `ALICLOUD_SECRET_KEY`}}",
  },
  "builders": [{
    "type": "alicloud-ecs",
    "access_key": "{{user `access_key`}}",
    "secret_key": "{{user `secret_key`}}",
    "region": "cn-beijing",
    "image_name": "packer_basic",
    "source_image": "centos_7_02_64_20G_alibase_20170818.vhd",
    "ssh_username": "root",
    "instance_type": "ecs.n1.tiny",
    "internet_charge_type": "PayByTraffic",
    "io_optimized": "true"
  ]},
  "provisioners": [{
    "type": "shell",
    "inline": [
      "sleep 30",
      "yum install redis.x86_64 -y"
    ]
  }]
}
```

Note:

You must customize the values of the following parameters.

Parameter	Description
access_key	Your AccessKey ID
secret_key	Your AccessKey Secret
region	The region of the temporary instance used to create the custom image. For more information, see Regions and zones .
image_name	The custom image's name
source_image	You can retrieve the basic image name from Alibaba Cloud public image list.
instance_type	Type of the temporary instance generated to create the custom image. For more information, see Instance generations and type families .

internet_charge_type	Internet bandwidth billing method for the temporary instance generated for creating the custom image. Optional values: - PayByTraffic - PayByBandwidth
provisioners	Type of Packer Provisioner used for creating the custom image

Step 3. Create a custom image by using Packer

Follow these step to specify the Packer template file and create a custom image:

1. Run `export ALICLOUD_ACCESS_KEY=your AccessKeyID` to import your AccessKey ID.
2. Run `export ALICLOUD_SECRET_KEY=your AccessKeySecret` to import your AccessKey Secret.
3. Run `packer build alicloud.json` to create the custom image.

The sample runs like follows. The sample creates a custom image containing ApsaraDB for Redis and runs as follows:

alicloud-ecs output will be in this color.

```
==> alicloud-ecs: Prevalidating alicloud image name...
alicloud-ecs: Found image ID: centos_7_02_64_20G_alibase_20170818.vhd
==> alicloud-ecs: Start creating temporary keypair: packer_59e44f40-c8d6-0ee3-7fd8-b1ba08ea94b8
==> alicloud-ecs: Start creating alicloud vpc
-----
==> alicloud-ecs: Provisioning with shell script: /var/folders/3q/w38xx_js6cl6k5mwkrqsnw7w0000gn/T/packer-shell257466182
alicloud-ecs: Loaded plugins: fastestmirror
-----
alicloud-ecs: Total 1.3 MB/s | 650 kB 00:00
alicloud-ecs: Running transaction check
-----
==> alicloud-ecs: Deleting temporary keypair...
Build 'alicloud-ecs' finished.

==> Builds finished. The artifacts of successful builds are:
--> alicloud-ecs: Alicloud images were created:

cn-beijing: m-2ze12578be10a4ovs6r9
```

Next steps

You can use this custom image to create an ECS instance. For more information, see [Create an instance by using a custom image](#).

References

- For more information, visit [Packer-provider](#), the Packer repository of Alibaba Cloud Github.
- See the [Packer Official Documents](#) to learn more about how to use Packer.

Copying an image is the process in which a custom image is copied from one region to another region. When a request of copying a custom image is initiated, Alibaba Cloud copies the snapshot that the custom image is created from the source region to the target region, and then creates a custom image from the copied snapshot in the target region.

The speed of the process of copying the snapshot between regions depends on the network status, and Alibaba Cloud supports processing concurrent requests of copying images and your request maybe in a long queue. Therefore, it may take long time to complete copying.

Copying images across regions allows you to deploy a backup image system, or an identical application environment, in different regions.

To copy a custom image, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Snapshots & Images > Images**.

Select a region.

Select the custom image you want to copy, and in the **Action** column, click **Copy Image**.


If your custom image is larger than 100 GB, when you click **Copy Image**, you are directed to open a ticket to complete copying the image.

In the **Copy Image** dialog box, the ID of the selected image is displayed, and you have to complete the configurations:

- Select the target region. Currently, copying images is only allowed between regions in mainland China.
- Specify a name for the image to be displayed in the target region, and give a short description of the image to ease future management.
- Click **OK**.

Click the target region and check the progress. When 100% progress is displayed, the image is copied successfully.

When the progress is not 100% and the status of the image is **Creating**, you can click **Cancel Copy** to cancel the copying process. After the process is canceled, the image information is removed from the target region.



When the status of the image is **Available**, you can use the custom image to create an ECS instance or change a system disk. You can check the snapshot for creating the custom image in the **Snapshot List** in the target region.

For more information about the process of copying images, refer to [Image copy FAQ](#).

You can share your custom images with other users. Through the ECS console or ECS API, you can query images shared by other accounts with your own account, and select images shared by other accounts to create ECS instances.

Note:

- The integrity or security of images is not guaranteed. Make sure that you use only images shared by trusted accounts.
- Before using shared images to create ECS instances, log on to the ECS instances to which the shared images belong and verify that the images are secure and complete.
- Before sharing an image, make sure that no confidential data is accessible on the disks to be shared.

Considerations

Restrictions

- One image can be shared with a maximum of **50** accounts.
- Shared images do not count towards your image quota.
- Shared images can only be used to create instances in the same region as the source image.
- Only image owners can share images with other accounts.

Impact of deleting shared images

- You can delete a custom image even you have shared it with other accounts. Before deleting the shared image, however, you must unassociate it from other accounts.
- If you delete an account that has shared a custom image, the users who are using the shared image can no longer find the image through the ECS console or ECS API, or use the image to create ECS instances.

- Deleting shared custom images may cause system disk re-initialization to fail for ECS instances created from these images.

Procedure

Log on to the ECS console.

In the left-side navigation pane, click **Images**.

Select a region.

Select the image you want to share.

Note: The image type must be **Custom Image**.

Click **Share Image**.

In the displayed dialog box, select the **Account Type** and enter the account ID you want to share the image with.

Note: The **Account ID** can be obtained from the **Account Management > Security Settings** on the Alibaba Cloud website by logging on to Security Settings.

Click **Share Image**.

View accounts using your shared images

You can view which accounts are using your shared images.

To view accounts using your shared images, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Images**.

Select a region.

Select the image you want to check.

Click **Share Image**. A list of the accounts using the selected image is displayed.

Cancel the sharing of an image

You can cancel the sharing of an image to specific accounts at any time.

Note: When the sharing of an image is cancelled:

- Any accounts currently using the image will no longer be able to use the image. Therefore, you must disassociate the image from other accounts before cancelling it being shared.
- Any instances using the image, including instances of other accounts using the shared image, will not be able to reinitialize the system disk.

Procedure

To cancel the sharing of an image, perform the following:

Log on to the ECS console.

In the left-navigation pane, click **Images**.

Select a region.

Select the image you want to cancel sharing.

Click **Share Image**.

Click **Unshare** next to the account with which you want to stop sharing the image.

View the shared images you are using

You can view a list of the shared images from other accounts that you are using.


To view a list of the shared images you are using, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Images**.

Select a region.

In the image type dropdown, select **Shared Image** as the **Image Type**.



A list of the shared images you are using will be displayed.

Import images

To guarantee the usability of an imported image and to improve the efficiency of importing an image, pay attention to the followings before importing an image.

The notes vary by the operating system of your instance:

- For a Linux image
- For a Windows image

Notes for importing a Linux image

When importing a Linux image, pay attention to the following notes.

Limits

Image import does not support the use of multiple network interfaces or IPv6 addresses.

Passwords must be 8–30 characters in length and contain three types of characters (uppercase or lowercase letters, digits, and special characters).

You must install the XEN and KVM virtualization platform drivers.

The firewall is disabled, and port 22 is enabled by default.

DHCP is enabled in the image.

We recommend that you install cloud-init to guarantee the successful configuration of hostname, NTP source, and yum source.

Notes

If you want to import a Linux image, you must pay attention to the notes listed in the table.

Item	Images of standard operating systems	Images of non-standard operating system
Definition	<p>The official distribution editions of operating systems supported by Alibaba Cloud, including:</p> <ul style="list-style-type: none">- CentOS 5,6,7- Ubuntu 10,12,13,14- Debian 6,7- OpenSUSE 13.1- SUSE Linux 10,11,12- CoreOS 681.2.0+	<p>The non-standard operating system refers to either of the followings:</p> <ul style="list-style-type: none">- The operating system that are not included in the list of operating systems currently supported by Alibaba Cloud- A standard operating system that fails to comply with the requirements for a standard operating system in terms of critical system configuration files, system basic environment, and applications. <p>If you want to use an image of a non-standard operating system, you are only allowed</p>

		<p>to choose:</p> <ul style="list-style-type: none">- Customized Linux: Customized image. If you import an image of this type of operating system, Alibaba Cloud conducts necessary network or password configuration according to the pre-defined configuration norms. For detailed configuration information, see Configuration of Customized Linux.- Others Linux: Alibaba Cloud identifies all of these images as other system types. If you import an image of such operating system, Alibaba Cloud does not perform any processing on the created instance. After completing the instance creation, you must connect to the instance by using the Connect function in the console and then manually configure the IP address, the router, and the
--	--	--

		password.
Critical system configuration files	<ul style="list-style-type: none"> - Do not modify /etc/issue*. If it is modified, the distribution of the system cannot be properly recognized and the system creation fails. - Do not modify /boot/grub/menu.lst . If it is modified, the system may fail to start up. - Do not modify /etc/fstab. If it is modified, an exception may occur preventing partitions from being loaded, leading to system startup failure. - Do not modify /etc/shadow to read-only. If it is modified, the password file cannot be modified and the system startup fails. - Do not enable SELinux by modifying /etc/selinux/config. If it is modified, the system may fail to start up. 	Fails to comply with the requirements for a standard operating system.
Requirements for system basic environments	<ul style="list-style-type: none"> - Do not adjust the partition of the system disk. 	Fails to comply with the requirements for a standard operating system.

	<p>Currently only a single root partition is supported.</p> <ul style="list-style-type: none"> - Make sure that the system disk has sufficient free space. - Do not modify critical system files, such as /sbin, /bin, or /lib*. - Before importing an image, confirm the integrity of the file system. - File system: File systems of xfs, ext3, and ext4 for Linux images are supported. MBR is used. 	
Applications	Do not install qemu-ga in an imported image. If it is installed, some of the services that Alibaba Cloud needs may become unavailable.	Fails to comply with the requirements for a standard operating system.
File format	Currently, images in only RAW and VHD formats are supported. If you want to import images in other formats, use a tool to convert the format before importing the image. We recommend that you import images in a VHD format, which has a smaller transmission capacity.	
File size	Setting the system disk size when importing an image: We recommend that you configure the system disk size for importing based on the virtual file size (not the usage) of the image. The size of the disk for importing must be from 40 GB to 500 GB.	

Notes for importing a Windows image

When importing a Windows image, pay attention to the following notes.

Limits

Passwords must be 8–30 characters in length and contain three types of characters (uppercase or lowercase letters, digits, and special characters).

The firewall is disabled, and port 3389 is enabled by default.

Distribution editions of Windows operating system

You are allowed to import the following distribution editions of Windows operating system:

Microsoft Windows Server 2016

Microsoft Windows Server 2012 R2 (Standard Edition)

Microsoft Windows Server 2012 (Standard Edition, Data Center Edition)

Microsoft Windows Server 2008 R2 (Standard Edition, Data Center Edition, Enterprise Edition)

Microsoft Windows Server 2008 (Standard Edition, Data Center Edition, Enterprise Edition)

Microsoft Windows Server 2003 (Standard Edition, Data Center Edition, Enterprise Edition), including R2 and with Service Pack 1 (SP1)

Note: Windows XP, Windows 7 (both Professional Edition and Enterprise Edition), Windows 8, and Windows 10 are not supported.

Requirements for system basic environments

System disks with multiple partitions are supported.

Make sure that the system disk has sufficient free space.

Do not modify critical system files.

Before importing an image, confirm the integrity of the file system.

File system: Only NTFS file system and MBR is supported.

Applications

Do not install qemu-ga in an imported image. If it is installed, some of the services that Alibaba Cloud needs may become unavailable.

Size and format

Currently, the images in only RAW and VHD formats are supported. If you want to import images in other formats, use a tool to convert the format before importing the image. We recommend that you import images in a VHD format, which has a smaller transmission capacity.

Setting the system disk size when importing an image: We recommend that you configure the system disk size for importing based on the virtual file size (not the usage) of the image. The size of the disk for importing must be from 40 GB to 500 GB.

To guarantee the successful configuration of the hostname, NTP source, and yum source of the imported image, we recommend that you install cloud-init in your on-premise server, virtual machine, or cloud host before importing an image.

Currently, cloud-init supports the following Linux distributions:

- CentOS
- Debian
- Fedora
- FreeBSD
- Gentoo
- RHEL (Red Hat Enterprise Linux)
- SLES (SUSE Linux Enterprise Server)
- Ubuntu

Prerequisites

Make sure that you have installed the following programs:

git: For downloading the source code package of cloud-init

Command: `yum install git`

python2.7: The basis of running and installing cloud-init

Command: `yum install python`

pip: For installing the libraries that are missing from python2.7 but cloud-init depends on

Command: `yum install python-pip`

In this document, we use yum as an example to describe the installation. If you are using zypper or apt-get to manage packages, the installation methods are similar.

Install cloud-init

Follow these steps to install cloud-init:

Connect to your on-premise server, virtual machine or cloud host. If your cloud host is an ECS instance, see [Connect to a Linux instance](#).

Run `git clone https://git.launchpad.net/cloud-init` to download the source code package of cloud-init from its official website.

Run `cd cloud-init` to change the directory to cloud-init.

Run `python setup.py install` to install setup.py, which is the installation file of cloud-init.

Run `vi /etc/cloud/cloud.cfg` to modify configuration file cloud.cfg.

```
# The top level settings are used as module
# and system configuration.

# A set of users which may be applied and/or used by various modules
# when a 'default' entry is found it will reference the 'default_user'
# from the distro configuration specified below
users:
  - default

# If this is set, 'root' will not be able to ssh in and they
# will get a message to login instead as the default $user
disable_root: true

# This will cause the set-update hostname module to not operate (if true)
preserve_hostname: false

# Example datasource config
# datasource:
#   Ec2:
#     metadata_urls: [ 'blah.com' ]
#     timeout: 5 # (defaults to 50 seconds)
#     max_wait: 10 # (defaults to 120 seconds)

# The modules that run in the 'init' stage
cloud_init_modules:
```

Change the preceding content of cloud_init_modules to the following:

```
# Example datasource config
# The top level settings are used as module
# and system configuration.

# A set of users which may be applied and/or used by various modules
# when a 'default' entry is found it will reference the 'default_user'
# from the distro configuration specified below
users:
- default

user:
name: root
lock_passwd: False

# If this is set, 'root' will not be able to ssh in and they
# will get a message to login instead as the above $user
disable_root: false

# This will cause the set+update hostname module to not operate (if true)
preserve_hostname: false

syslog_fix_perms: root:root

datasource_list: [ AliYun ]

# Example datasource config
datasource:
AliYun:
support_xen: false
timeout: 5 # (defaults to 50 seconds)
max_wait: 60 # (defaults to 120 seconds)
# metadata_urls: [ 'blah.com' ]

# The modules that run in the 'init' stage
cloud_init_modules:
```

Troubleshooting

Note: The missing libraries may vary depending on the operating system. You can use pip to install the missing libraries. After you install the missing libraries, run `python setup.py install` again to install `setup.py`.

1. Library six or library oauthlib is missing.

During installation, the following message may appear. It means the six library is missing from Python. Run `pip install six` to install the six library.

```
File "/root/cloud-init/cloudinit/log.py", line 19, in <module>
```



```
import six
ImportError: No module named six
)
```

During installation, the following message may appear, which means the oauthlib library is missing from Python. Run `pip install oauthlib` to install the oauthlib library.

```
File "/root/cloud-init/cloudinit/url_helper.py", line 20, in <module>
import oauthlib.oauth1 as oauth1
ImportError: No module named oauthlib.oauth1
)
```

2. No library is specified when an error occurs during installation.

If no dependency library is specified according to the error output, you may run `pip install -r requirements.txt` to install all the dependency libraries listed in file `requirements.txt` of `cloud-init`.

Next step

You can Import an image in the ECS console.

Only image files in RAW or VHD format can be imported. If you want to import images in other formats, convert the format before importing the image.

This document introduces how to use the `qemu-img` tool to convert image files into VHD or RAW from other formats, such as RAW, Qcow2, VMDK, VDI, VHD (vpc), VHDX, qcow1, or QED.

You can use different methods to install `qemu-img` and convert the image file format based on operating system of your local computer:

- Windows
- Linux

Windows

To install `qemu-img` on Windows system and convert image file formats, follow these steps:

Download and install `qemu`. Download address: <https://qemu.weilnetz.de/w64/>. Installation path: `C:\Program Files\qemu`.

Perform the following to create an environment variable (For Windows 7):

- Select **Start > Computer**, and right click **Properties**.

- ii. In the left-side navigation pane, click **Advanced system settings**.
- iii. In the **System Properties** dialog box, click the **Advanced** tab and click **Environment Variables**.
 - a. In the **Environment Variables** dialog box, in the **System variables**, find the **Path** variable, and click **Edit**. If the **Path** variable does not exist, click **New**.
 - b. Add a variable value:
 - a. In the **Edit System Variable**: Add `C:\Program Files\qemu` to the **Variable value**. Different variable values are separated with semicolon (;).
 - b. In the **New System Variable**: Enter `Path` as the **Variable name**, and enter `C:\Program Files\qemu` as the **Variable value**.

Open **Command Prompt** in Windows and run the `qemu-img --help` command. If it is displayed successfully, the installation was successful.

In the **Command** prompt, run the `cd [directory of the source image file]` command to change the directory. For example, `cd D:\ConvertImage`.

Run the following command in **Command** prompt to convert the image file format:

```
qemu-img convert -f raw -O qcow2 centos.img centos.qcow2
```

The command parameters are described as follows:

- `-f` is followed by the source image format.
- `-O` (uppercase is required) is followed by the converted image format, the source file name, and the target file name.

Linux

To install `qemu-img` and convert the image file format, follow these steps:

Install `qemu-img`, for example:

- For Ubuntu, run the command: `apt install qemu-img`.
- For CentOS, run the command: `yum install qemu-img`.

Run the following command to convert the image file format.

```
qemu-img convert -f raw -O qcow2 centos.img centos.qcow2
```

The command parameters are described as follows:

- `-f` is followed by the source image format.
- `-O` (uppercase is required) is followed by the converted image format, the source file name, and the target file name.

You can import image files to the ECS environment to create custom images. You can then use these images to create ECS instances.

Prerequisites

- See [Notes for importing custom images](#), [Configuration of Customized Linux](#) and [Convert image file format](#) for the restrictions and requirements when you import the custom image.
- You must enable the OSS service, and make sure that you must grant the official ECS service account access to your OSS.
- You can only import an image file to a region from OSS in the same region. The image and the OSS must belong to one account.
- You can use an OSS third-party tool client, OSS API or OSS SDK, to upload the file to a bucket in the same region as the ECS custom image to import. See [Multipart upload](#) to upload an image file that is larger than 5 GB.

Procedure

Log on to the ECS console.

In the left-side navigation pane, choose **Snapshots and Images > Images**.

Click **Import Image**.

Click **Confirm Address** on the third items of **How to import an image**.

Click **Confirm Authorization Policy**.

In the left-side navigation pane, choose **Snapshots and Images > Images**.

Choose a region.

Click **Import Image**, and enter the following information in the pop-up window.

Region of image

Select the region where you want to deploy the application.

OSS Object Address

Copy the object address taken from the OSS console.

Image Name

The length must be 2 to 128 characters. It can contain uppercase letters, lowercase letters or Chinese characters. It cannot contain numbers, underscores (_), or hyphens (-).

Operating System

Supported OS releases are:

- Windows
- Linux

System Disk Size

- For Windows system: 40 - 500 GB.
- For Linux system: 20- 500 GB.

System Architecture

64-bit OS: x86_64. 32-bit OS: i386.

System Platform

Supported operating system releases are:

- Windows: Windows Server 2003, Windows Server 2008, and Windows Server 2012.
- Linux: CentOS, SUSE, Ubuntu, Debian, FreeBSD, and CoreOS.

Note:

- **(Linux only)** Open a ticket to Alibaba Cloud to confirm the selected edition is supported.
- If your image OS is a custom edition developed on a Linux kernel, open a ticket to Alibaba Cloud.

Image Format

Supports RAW and VHD format. RAW format is recommended.

Note: You cannot use qemu-image to create VHD images.

Image Description

The description of the image.

Click **OK**.

Note: It usually takes 1 to 4 hours to import an image, and the duration of the task depends on the size of your image file and the amount of concurrent tasks. You can view the task progress in the image list of the import region.

You can find and cancel the image import task in the task manager.

Follow-up operation

After you import the custom image, you may want to create an instance by using a custom image.

To avoid failing to start the Linux instances created by using the imported images of your servers, virtual machines, or cloud hosts, an Xen (pv) or virtio driver must be installed on your on-premises image and configured before importing. Follow these steps to check whether you must install the driver manually, and then install and configure the virtio driver for a Linux server if needed.

Images that require no manual installation

After you import images in the ECS console, if the operating systems of your image is listed in the following, Alibaba Cloud automatically processes the virtio driver for you. You can skip to recover the temporary root file system of initramfs or initrd.

- Windows Server 2008
- Windows Server 2012
- Windows Server 2016
- CentOS 6/7
- Ubuntu 12/14/16
- Debian 7/8/9
- SUSE 11/12

Images that require manual installation

For Linux images that are not included in the preceding list, you must install the virtio driver on-premises before importing the images.

To check the availability of virtio driver on a server

Run `grep -i virtio /boot/config-$(uname -r)` to inspect whether the virtio driver is already built in the kernel of your server.

```
[root@1zbpllcneefoj@ec2-vadttlz ~]# grep -i virtio /boot/config-$(uname -r)
CONFIG_VIRTIO_VSOCKETS=m
CONFIG_VIRTIO_VSOCKETS_COMMON=m
CONFIG_VIRTIO_BLK=m
CONFIG_SCSI_VIRTIO=m
CONFIG_VIRTIO_NET=m
CONFIG_VIRTIO_CONSOLE=m
CONFIG_HW_RANDOM_VIRTIO=m
CONFIG_DRM_VIRTIO_GPU=m
CONFIG_VIRTIO=m
# Virtio drivers
CONFIG_VIRTIO_PCI=m
CONFIG_VIRTIO_PCI_LEGACY=y
CONFIG_VIRTIO_BALLOON=m
CONFIG_VIRTIO_INPUT=m
# CONFIG_VIRTIO_MMIO is not set
```

Options:

- **Option 1.** If VIRTIO_BLK and VIRTIO_NET do not exist in the output, the virtio driver is not built in the kernel, and you must install and configure the virtio driver on your server.
- **Option 2.** If the values of parameter CONFIG_VIRTIO_BLK and parameter CONFIG_VIRTIO_NET are y, the virtio driver is already built in the kernel. You can import the image after reading Notes for importing custom images.
- **Option 3.** If the values of parameter CONFIG_VIRTIO_BLK and parameter CONFIG_VIRTIO_NET are m, continue to step 2.

Run `lsinitrd /boot/initramfs-$(uname -r).img | grep virtio` to make sure virtio driver has been compiled in the temporary root file system of initramfs or initrd.

```
[root@1zbpllcneefoj@ec2-vadttlz ~]# lsinitrd /boot/initramfs-$(uname -r).img | grep virtio
Arguments: -f --add-drivers ' xen-blkfront xen-blkfront virtio_blk virtio_pci virtio_console virtio_console'
-rw-r--r-- 1 root root 7628 Sep 13 07:14 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/block/virtio_blk.ko.xz
-rw-r--r-- 1 root root 12828 Sep 13 07:15 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/char/virtio_console.ko.xz
-rw-r--r-- 1 root root 7988 Sep 13 07:16 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/scsi/virtio_scsi.ko.xz
dlnx-r-xr-x 2 root root 0 Oct 24 14:09 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/virtio
-rw-r--r-- 1 root root 4348 Sep 13 07:16 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/virtio/virtio.ko.xz
-rw-r--r-- 1 root root 9488 Sep 13 07:16 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/virtio/virtio_pci.ko.xz
-rw-r--r-- 1 root root 8136 Sep 13 07:16 usr/lib/modules/3.10.0-693.2.2.el7.x86_64/kernel/drivers/virtio/virtio_ring.ko.xz
```

Options:

- **Option 1.** According to the preceding figure, the virtio_blk driver, including its dependency virtio.ko, virtio_pci.ko and virtio_ring.ko, has been compiled in the temporary root file system initramfs. You can directly import the image after reading Notes for importing custom images.
- **Option 2.** If virtio driver is unavailable in the initramfs, you must recover the temporary root file system of initramfs or initrd before importing images or migration.

To recover the temporary root file system

If the virtio driver is supported by the kernel but not compiled in the temporary root file system, you must recover the temporary root file system. Take CentOS as an example:

- CentOS/RedHat 5

```
mkinitrd -f --allow-missing \
--with=xen-vbd --preload=xen-vbd \
--with=xen-platform-pci --preload=xen-platform-pci \
--with=virtio_blk --preload=virtio_blk \
--with=virtio_pci --preload=virtio_pci \
--with=virtio_console --preload=virtio_console \
```

- CentOS/RedHat 6/7

```
mkinitrd -f --allow-missing \
--with=xen-blkfront --preload=xen-blkfront \
--with=virtio_blk --preload=virtio_blk \
--with=virtio_pci --preload=virtio_pci \
--with=virtio_console --preload=virtio_console \
/boot/initramfs-$(uname -r).img $(uname -r)
fi
```

- Debian/Ubuntu

```
echo -e 'xen-blkfront\nvirtio_blk\nvirtio_pci\nvirtio_console' >> \
/etc/initramfs-tools/modules
mkinitramfs -o /boot/initrd.img-$(uname -r)"
```

To compile and install virtio driver

Take Redhat server as an example:

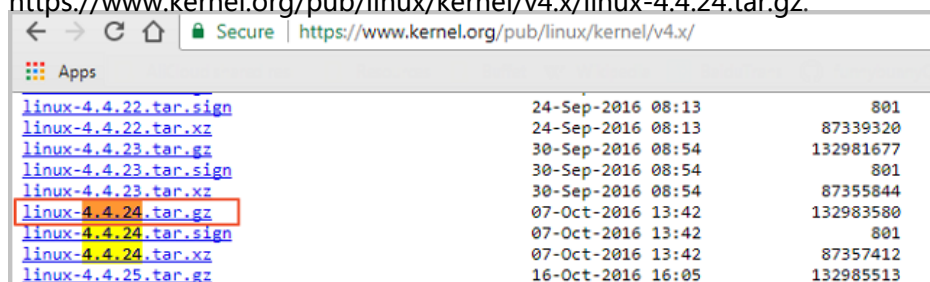
To download the kernel package

1. Run `yum install -y ncurses-devel gcc make wget` to install necessary components to compile the kernel.
2. Run `uname -r` to query the kernel version of your server, such as `4.4.24-2.el7.x86_64`.

```
[root@iZbp1127hr3wi6p2cq9lnbZ ~]# uname -r
4.4.24-2.el7.x86_64
```

3. Visit published Linux Kernel Archives to download the source codes of kernel, for example, the download link of kernel version starting with 4.4.24 is

<https://www.kernel.org/pub/linux/kernel/v4.x/linux-4.4.24.tar.gz>.



File Name	Date	Size
linux-4.4.22.tar.sign	24-Sep-2016 08:13	801
linux-4.4.22.tar.xz	24-Sep-2016 08:13	87339320
linux-4.4.23.tar.gz	30-Sep-2016 08:54	132981677
linux-4.4.23.tar.sign	30-Sep-2016 08:54	801
linux-4.4.23.tar.xz	30-Sep-2016 08:54	87355844
linux-4.4.24.tar.gz	07-Oct-2016 13:42	132983580
linux-4.4.24.tar.sign	07-Oct-2016 13:42	801
linux-4.4.24.tar.xz	07-Oct-2016 13:42	87357412
linux-4.4.25.tar.gz	16-Oct-2016 16:05	132985513

4. Run `cd /usr/src/` to change the directory.

5. Run `wget https://www.kernel.org/pub/linux/kernel/v4.x/linux-4.4.24.tar.gz` to download the installation package.
6. Run `tar -xzf linux-4.4.24.tar.gz` to decompress the package.
7. Run `ln -s linux-4.4.24 linux` to establish a link.
8. Run `cd /usr/src/linux` to change the directory.

To compile the kernel

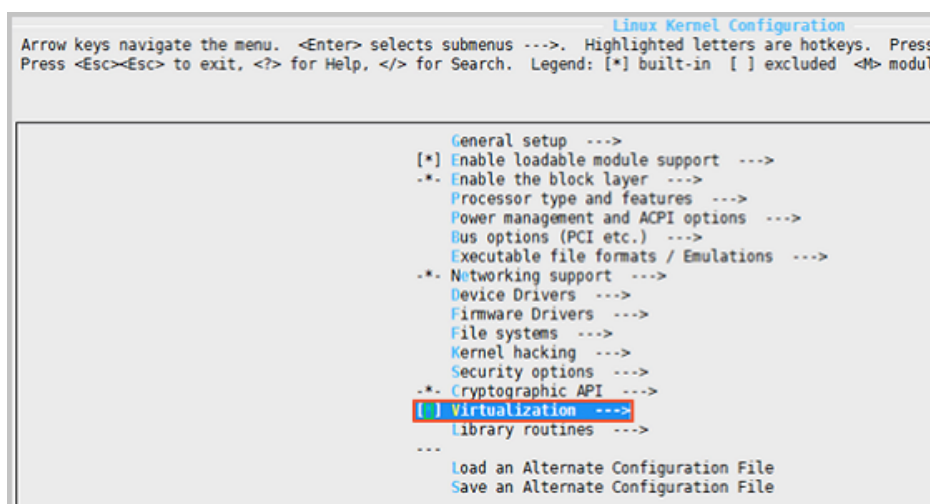
Run the following commands to compile the driver into the kernel.

```
make mrproper
symvers_path=$(find /usr/src/ -name "Module.symvers")
test -f $symvers_path && cp $symvers_path .
cp /boot/config-$(uname -r) ./config
make menuconfig
```

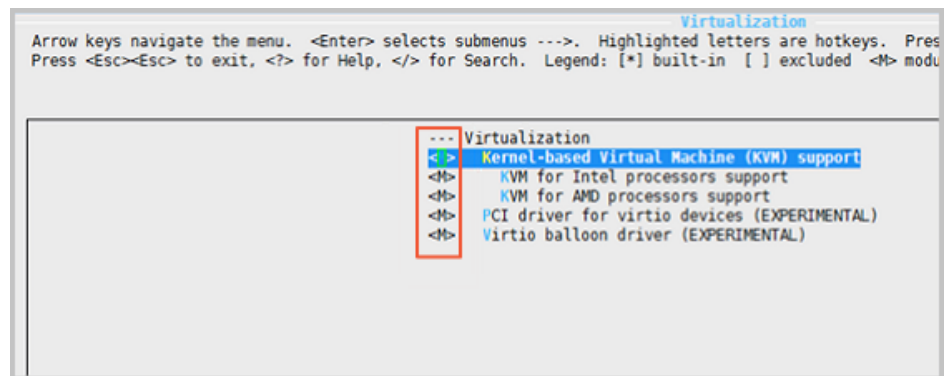
Configure the corresponding settings of virtio driver in the following windows:

Note: Select `*` to build the driver in the kernel, select `m` to compile it as a module.

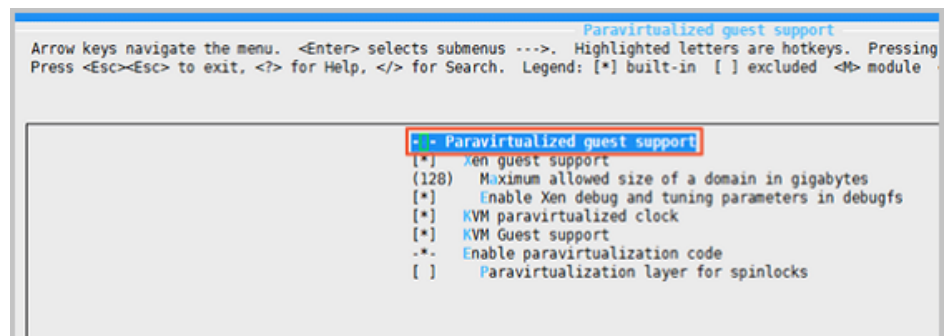
Press the **space bar** to select **Virtualization**.



Make sure that you have selected the option of KVM (Kernel-based Virtual Machine).



Processor type and features --->
 [*] Paravirtualized guest support --->
 --- Paravirtualized guest support
 (128) Maximum allowed size of a domain in gigabytes
 [*] KVM paravirtualized clock
 [*] KVM Guest support



Device Drivers --->
 [*] Block devices --->
 <M> Virtio block driver (EXPERIMENTAL)
 *- Network device support --->
 <M> Virtio network driver (EXPERIMENTAL)

Press the **Esc** key to exit the kernel configuration windows, and save changes to file `.config` according to the dialog box.

Inspect whether all the corresponding settings of virtio driver has been correctly configured or not.

(Optional) If no configuration of virtio driver is settled, run the following commands to edit the file `.config` manually.

```
make oldconfig
make prepare
make scripts
```

```
make
make install
```

Run the following commands to check whether the virtio driver is installed.

```
find /lib/modules/"$(uname -r)" -name "virtio.*" | grep -E "virtio.*"
grep -E "virtio.*" < /lib/modules/"$(uname -r)"/modules.builtin
```

If any of the output includes virtio_blk and virtio_pci.virtio_console, your server has correctly installed the virtio driver.

Follow-up operations

After compiling the virtio driver:

- You can **Import images** to Alibaba Cloud console.
- You can **Migrate your server** to Alibaba Cloud by using Cloud Migration Tool.

Packer is a convenient open-source tool to create on-premises image files. It runs on the most major operating systems.

To create an on-premises image by yourself and then upload it on a cloud platform is a complex process. However, by using Packer, you can create identical on-premises images for multiple platforms from a single source configuration.

Follow these steps to create an on-premises image for CentOS 6.9 on an Ubuntu 16.04 server and to upload it to Alibaba Cloud. To create on-premises images for other operating systems, you can **customize your Packer templates as necessary**.

Prerequisites

You must have the AccessKey ready to fill out the configuration file.

Note:

The AccessKey has a high level of account privileges. We recommend that you first **create a RAM user** and use the RAM account to **create an AccessKey** to prevent data breach.

Before uploading your on-premises images to Alibaba Cloud, you must sign up for Alibaba Cloud OSS.

A sample of creating and importing an on-premises image

Run `egrep "(svm|vmx)" /proc/cpuinfo` to check whether your on-premises server or virtual machine supports KVM. If the following output returns, KVM is supported.

```
pat pse36 clflush dts acpi mmx fxsr sse sse2 ss ht tm pbe syscall nx pdpe1gb rdtscp lm constant_tsc art
arch_perfmon pebs bts rep_good nopl xtopology nonstop_tsc aperfmperf tsc_known_freq pni pclmulqdq
dtes64 monitor ds_cpl vmx est tm2 ssse3 sdbg fma cx16 xtpr pdcm pcid sse4_1 sse4_2 x2apic movbe
popcnt tsc_deadline_timer aes xsave avx f16c rdrand lahf_lm abm 3dnowprefetch epb intel_pt tpr_shadow
vnmi flexpriority ept vpid fsgsbase tsc_adjust bmi1 avx2 smep bmi2 erms invpcid mpx rdseed adx smap
clflushopt xsaveopt xsavec xgetbv1 xsaves dtherm ida arat pln pts hwp hwp_notify hwp_act_window
hwp_epp
flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov
```

Run the following commands to install the KVM:

```
egrep "(svm|vmx)" /proc/cpuinfo # Check whether the CPU of your server supports KVM (Kernel-based
Virtual Machine) or not.
sudo apt-get install qemu-kvm qemu virt-manager virt-viewer libvirt-bin bridge-utils # Install KVM and
related dependencies.
sudo virt-manager # Enable virt-manager.
```

If a GUI runs in the VM console window, you have successfully installed the KVM.

Install Packer.

To install Packer, see [Use Packer to Create a Custom Image](#).

Run the following commands to define a Packer template.

Note:

The on-premises image created in the following configuration is for the CentOS 6.9 operating system only. To create images for other operating systems, **customize** the configuration file as needed.

```
cd /user/local # Switch the directory.
wget https://raw.githubusercontent.com/alibaba/packer-
provider/master/examples/alicloud/local/centos.json # Download file centos.json that is released by
Alibaba Cloud.
wget https://raw.githubusercontent.com/alibaba/packer-
provider/master/examples/alicloud/local/http/centos-6.9/ks.cfg # Download file ks.cfg that is released by
Alibaba Cloud.
```

```
mkdir -p http/centos-6.9 # Create a directory.
mv ks.cfg http/centos-6.9/ # Move file ks.cfg to the http/centos-6.9 directory.
```

Run the following commands to create an on-premises image.

```
export ALICLOUD_ACCESS_KEY= SpecifyYourAccessKeyIDHere # Import your AccessKeyID,
export ALICLOUD_SECRET_KEY= SpecifyYourAccessKeySecretHere # Import your AccessKeySecret.
packer build centos.json # Create an on-premises image.
```

The running result of the sample is as follows.

```
qemu output will be in this color.

==> qemu: Downloading or copying ISO
qemu: Downloading or copying: http://mirrors.aliyun.com/centos/6.9/isos/x86_64/CentOS-6.9-x86_64-
minimal.iso
.....
==> qemu: Running post-processor: alicloud-import
qemu (alicloud-import): Deleting import source https://oss-cn-beijing.aliyuncs.com/packer/centos_x86_64
Build 'qemu' finished.

==> Builds finished. The artifacts of successful builds are:
--> qemu: Alicloud images were created:

cn-beijing: XXXXXXXXX
```

Wait for a few minutes, log on to the ECS console and check your custom image in the image list that is in the corresponding region. In this sample, the region is China North 2 (cn-beijing).

Customize a Packer template

The image file created in the preceding sample is for the CentOS 6.9 operating system only. To create images for other operating systems, you must customize the Packer template.

For example, the following JSON file is customized based on the template to create an image for the CentOS 6.9.

```
{
  "variables": {
    "box_basename": "centos-6.9",
    "build_timestamp": "{{isotime `20060102150405`}}",
    "cpus": "1",
    "disk_size": "4096",
    "git_revision": "__unknown_git_revision__",
    "headless": "",
    "http_proxy": "{{env `http_proxy`}}",
    "https_proxy": "{{env `https_proxy`}}",
```

```

"iso_checksum_type": "md5",
"iso_checksum": "af4a1640c0c6f348c6c41f1ea9e192a2",
"iso_name": "CentOS-6.9-x86_64-minimal.iso",
"ks_path": "centos-6.9/ks.cfg",
"memory": "512",
"metadata": "floppy/dummy_metadata.json",
"mirror": "http://mirrors.aliyun.com/centos",
"mirror_directory": "6.9/isos/x86_64",
"name": "centos-6.9",
"no_proxy": "{{env `no_proxy`}}",
"template": "centos-6.9-x86_64",
"version": "2.1.TIMESTAMP"
},
"builders": [
{
"boot_command": [
"<tab> text ks=http://{{ .HTTPIP }}:{{ .HTTPPort }}/{{user `ks_path`}}<enter> <wait> "
],
"boot_wait": "10s",
"disk_size": "{{user `disk_size`}}",
"headless": "{{user `headless`}}",
"http_directory": "http",
"iso_checksum": "{{user `iso_checksum`}}",
"iso_checksum_type": "{{user `iso_checksum_type`}}",
"iso_url": "{{user `mirror`}}/{{user `mirror_directory`}}/{{user `iso_name`}}",
"output_directory": "packer-{{user `template`}}-qemu",
"shutdown_command": "echo 'vagrant'|sudo -S /sbin/halt -h -p",
"ssh_password": "vagrant",
"ssh_port": 22,
"ssh_username": "root",
"ssh_wait_timeout": "10000s",
"type": "qemu",
"vm_name": "{{user `template`}}.raw",
"net_device": "virtio-net",
"disk_interface": "virtio",
"format": "raw"
}
],
"provisioners": [
{
"type": "shell",
"inline": [
"sleep 30",
"yum install cloud-util cloud-init -y"
]
}
],
"post-processors": [
{
"type": "alicloud-import",
"oss_bucket_name": "packer",
"image_name": "packer_import",
"image_os_type": "linux",
"image_platform": "CentOS",
"image_architecture": "x86_64",
"image_system_size": "40",
"region": "cn-beijing"
}
]
}

```

```

]
}

```

Parameters in a Packer builder

QEMU builder is used in the preceding sample to create a virtual machine image. Required parameters for the builder are as follows.

Parameters	Type	Description
iso_checksum	string	The checksum for the OS ISO file. Packer verifies this parameter before starting a virtual machine with the ISO attached. Make sure you specify at least one of the iso_checksum or iso_checksum_url parameter. If you have the iso_checksum parameter specified, the iso_checksum_url parameter is ignored automatically.
iso_checksum_type	string	The type of the checksum specified in iso_checksum. Optional values: <ul style="list-style-type: none"> - none: If you specify none for iso_checksum_type, the checksumming is ignored, thus none is not recommended. - md5 - sha1 - sha256 - sha512
iso_checksum_url	string	This is a URL pointing to a GNU or BSD style checksum file that contains the ISO file checksum of an operating system. It may come in either the GNU or BSD pattern. Make sure you specify at least one of the iso_checksum or the iso_checksum_url parameter. If you have the iso_checksum parameter specified, the iso_checksum_url parameter is ignored automatically.
iso_url	string	This is a URL pointing to the ISO file and containing the installation image. This URL may be an HTTP URL or a file path:

		<ul style="list-style-type: none"> - If it is an HTTP URL, Packer downloads the file from the HTTP link and caches the file for running it later. - If it is a file path to the IMG or QCOW2 file, QEMU directly starts the file. If you have the file path specified, set parameter <code>disk_image</code> to true.
headless	boolean	By default, Packer starts the virtual machine GUI to build a QEMU virtual machine. If you set <code>headless</code> to True, a virtual machine without any console is started.

For more information about other optional parameters, see [Packer QEMU Builder](#).

Parameters in a Packer provisioner

The provisioner in the preceding [sample](#) contains a Post-Processor module that enables automated upload of on-premises images to Alibaba Cloud. Required parameters for the provisioner are as follows:

Parameters	Type	Description
access_key	string	Your AccessKeyID. The AccessKey has a high privilege. We recommend that you first create a RAM user and use the RAM account to create an AccessKey to prevent data breach.
secret_key	string	Your AccessKeySecret. The AccessKey has a high privilege. We recommend that you first create a RAM user and use the RAM account to create an AccessKey to prevent data breach.
region	string	Select the region where you want to upload your on-premises image. In the sample , the region is cn-beijing. For more information, see Regions and

		zones.
image_name	string	<p>The name of your on-premises image. The value:</p> <ul style="list-style-type: none">- Can contain [2, 128] characters in length.- Must start with an either upper case or lower case letter.- Can contain digits, underscores (_) or hyphens (-).- Cannot start with http:// or https://.
oss_bucket_name	string	<p>Your OSS bucket name. If you specify a bucket name that does not exist, Packer creates a bucket automatically with the specified oss_bucket_name when uploading the image.</p>
image_os_type	string	<p>Image type. Optional values:</p> <ul style="list-style-type: none">- linux- windows
image_platform	string	<p>Distribution of the image. For example, CentOS.</p>
image_architecture	string	<p>The system architecture of the image. Optional values:</p> <ul style="list-style-type: none">- i386- x86_64
format	string	<p>Image format. Optional values:</p> <ul style="list-style-type: none">- RAW- VHD

For more information about other optional parameters, see [Packer AliCloud Import Post-Processor](#).

Next steps

You can use the created image to create an ECS instance. For more information, see [Create an instance by using a custom image](#).

References

For more information about configuration file `ks.cfg`, see [Anaconda Kickstart](#).

For more information about how to use Packer, see [Packer documents](#).

For more information about release information, visit the Packer repository on GitHub [packer](#).

For more information about Alibaba Cloud open source tools, visit Alibaba repository on GitHub [opstools](#).

For more information about Alibaba Cloud and Packer project, visit the Alibaba & Packer repositories on GitHub [packer-provider](#).

Modify custom image names and description

You can modify the names and descriptions of custom images at any time.

To modify the names and descriptions of custom images, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Images**.











Select a region.

Select the image to edit.

Note: The image type must be **Custom Image**.

Modify the image name by hovering the cursor over the image name, and then clicking the **pen icon** that appears.

Modify the description of an image by clicking **Modify Image Description**, and then entering a description.

Image ID/Name	Image Type(Custom Image)	Platform	System Bit	Creation Time	Status	Progress	Action
 	 Custom Image	CentOS	64Bit	2017-03-07 11:41:14	Available	100%	Modify Image Description Related Instances Copy Image Share Image
 	 Custom Image	CentOS	64Bit	2016-11-25 10:19:00	Available	100%	Modify Image Description Related Instances Copy Image Share Image
  	 Custom Image	CENTOS	64Bit	2016-11-24 12:22:09	Available	100%	Modify Image Description Related Instances Copy Image Share Image

Click **OK**.

Delete a custom image

You can delete custom images that you no longer require. To ensure successful deletion, check that you do not currently have any ECS instances created from this custom image.

To delete a custom image, perform the following:

Log on to the ECS console.

Click **Images** in the left-side navigation bar.

Select your desired region.

Select the image you want to delete.

Note: The image type must be **Custom Image**.

Click **Delete**.

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

You can export custom images to a local device for test purposes or to offline private cloud. This document describes the constraints and restrictions of the image export function, and provides instructions on how to export images in the ECS console.

Note:

Exported image is stored in your OSS bucket that is in the same region as the custom images. Traffic fees for OSS storage and downloading will be generated.

Limits

Currently, the image export function has the following constraints and restrictions:

The image export function is usable after it is whitelisted.

You cannot export the custom images that is created by a system disk snapshot from the marketplace.

You can export the custom images that contains four snapshots of data disks at most, and for a single data disk, the maximum volume must be less than 100 GB.

The default format of exported image files is RAW.

Prerequisites

Open a ticket to activate the image export function.

Activate OSS and make sure that the region where your custom images are located has an available OSS bucket. See [Create a bucket](#) to create an OSS bucket.

Export a custom image

Follow these steps to export custom images:

Log on to the ECS console.

Authorize the ECS service to access your OSS resources:

- i. Choose **Snapshots & Images > Images** in the left-side navigation pane.
- ii. Select a region.
- iii. Find the custom image you want to export. In the **Action** column, click **Export Image**.
- iv. In the **Export Image** dialog box, click **Confirm Address** in Step 3 of the prompt message.
- v. In the **Cloud Resource Access Authorization** window, click **Confirm Authorization Policy**. Return to the ECS console homepage.

Choose **Snapshot & Images > Images** in the left-side navigation pane.

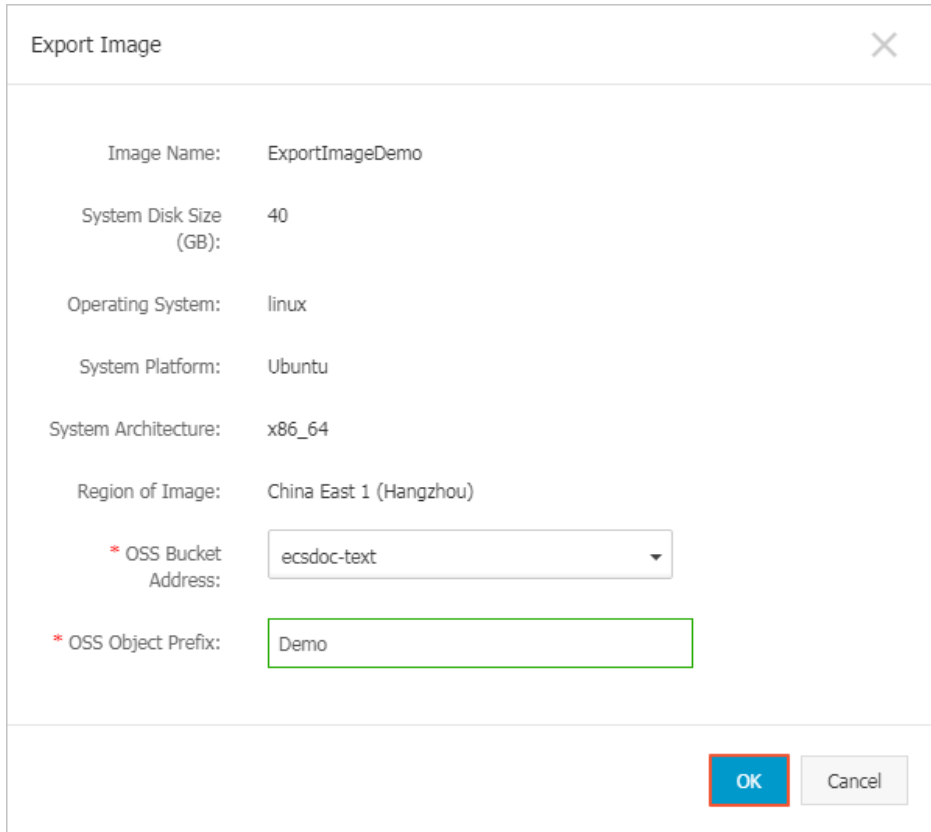
Select a region.

Find the custom image you want to export. In the **Action** column, click **Export Image**.

In the **Export Image** dialog box,

- Select the OSS bucket in the specified region.
- Set the prefix of the object name of the exported image. For example, if you set Demo as the prefix, then the exported image file is named Demo-[automatically generated file name] in the OSS bucket.

Click **OK** to export the image.



The screenshot shows the 'Export Image' dialog box with the following configuration:

Field	Value
Image Name:	ExportImageDemo
System Disk Size (GB):	40
Operating System:	linux
System Platform:	Ubuntu
System Architecture:	x86_64
Region of Image:	China East 1 (Hangzhou)
* OSS Bucket Address:	ecsdcc-text
* OSS Object Prefix:	Demo

At the bottom right, there are two buttons: 'OK' (highlighted with a red border) and 'Cancel'.

The duration of exporting depends on the size of the image file and the number of other export tasks in the queue. Be patient. You can go to the **Manage Tasks** page in the ECS console to query the task progress based on the task ID. When the **Task Status** is **Task Completed**, the image is successfully exported.

To cancel the export task, go to the **Manage Tasks** page and find the task.

Note:

To query the export result, log on to the **OSS console**.

Next step

To download the exported image file, log on to the **OSS console** and **Get object URL**.

Security groups

The following scenarios apply to **Classic Networks**.

Security groups provide security assurance to networks. They can be used to:

- Provide secure intranet communication.
- Block access to instances from specified IP addresses.
- Allow remote logins using only a specified IP address.
- Permit an instance to access only a specified IP address.

Scenario 1: Provide secure intranet communication

In a classic network, you can use security groups for Intranet communication between:

- ECS instances belonging to the same account in the same region.
- ECS instances belonging to different accounts in the same region.

ECS instances belonging to the same account in the same region

By setting security group rules, you can allow classic network instances belonging to the same account in the same region to communicate via intranet.

By default, ECS instances in the same security group can communicate through the intranet, but instances in different security groups cannot communicate through the intranet.

There are two ways to allow communication between ECS instances in different security groups. You can:

- Place the instances in the same security group to allow intranet communication.
- Authorize intranet communication between the two security groups by setting access-type security group rules. In **Authorization Type**, select **Security Group Access**, and then select the security group of each other for the **Authorization Object**.

Instances belonging to different accounts in the same region

By setting security group rules, you can allow classic network ECS instances belonging to different accounts in the same region to communicate via intranet.

To achieve intranet communication between instances belonging to different accounts in the same region, each user must perform the following:

- Add the other user' s security group to its inbound intranet.
- Authorize the ECS instances of the other user' s security group to access all instances in the their account.

Notice: To ensure the security of your instances, when you are configuring an intranet inbound rule for a security group of classic network type, **Security Group Access** is the top priority for **Authorization Type**. If you want to select **Address Field Access**, you must enter an IP address with CIDR prefix, **"/32"** , in the format of a.b.c.d/32. Only IPv4 is supported.

Block access to an ECS instance from an IP address

You can use security groups to block access to an ECS instance or a port of an ECS instance from specified IP addresses.

To block access to an instance from a specific IP address, perform the following:

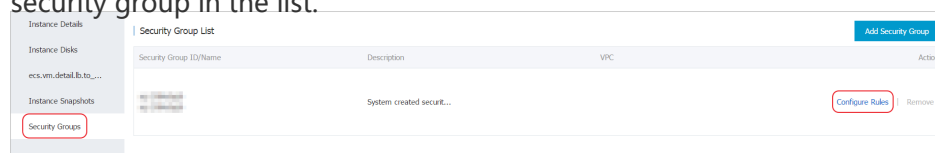
Log on to the ECS console.

Click **Instances** in the left-side navigation pane.


Select your desired region.

Select your desired instance, and click **Manage**.

Select **Security Groups** in the left-side navigation pane, and click **Configure Rules** of one security group in the list.



Click **Internet Inbound** and then click **Add Security Group Rules**.



If you want to drop access from an IP address, on the **Add Security Group Rules** dialog box:

- Select **Drop** for the **Authorization Policy**.
- Select **All** for the **Protocol Type**.
- Select **Address Field Access** for the **Authorization Type** and enter an IP address with or without a CIDR prefix, for example 1.2.3.4 or 1.2.3.4/24, as the **Authorization**

Object.

- Enter **1** for **Priority**.
- Click **OK**.

Add Security Group Rules

NIC: Internet

Rule Direction: Inbound

Authorization Policy: Drop

Protocol Type: All

* Port Range: -1/-1

The value range is 1-65535. For example, "1/200", "80/80". [Tutorial](#)

Authorization Type: Address Field Access

* Authorization Object: e.g. 10.0.0.0/32

Please be cautious when setting authorization objects. Based on different authorization policies, 0.0.0.0/0 indicates that access by all IPs is either allowed or rejected. [Tutorial](#)

Priority: 1

The priority value range is 1 - 100. The default value is 1 which is the highest priority.

OK Cancel

If you want to drop an IP address to access to Port 22 of your ECS instance, on the **Add Security Group Rules** dialog:

- Select **Drop** for the **Authorization Policy**.
- Select **Custom TCP** for the **Protocol Type**.
- Enter **22/22** as the **Port Range**.
- Select **Address Field Access** for the **Authorization Type** and enter an IP address with or without a CIDR prefix, for example 1.2.3.4 or 1.2.3.4/24, as the **Authorization Object**.
- Enter **1** for **Priority**.
- Click **OK**.

Add Security Group Rules

NIC: Internet

Rule Direction: Inbound

Authorization Policy: Drop

Protocol Type: Custom TCP

* Port Range: 22/22 The value range is 1-65535. For example, "1/200", "80/80". [Tutorial](#)

Authorization Type: Address Field Access

* Authorization Object: e.g. 10.0.0.0/32 Please be cautious when setting authorization objects. Based on different authorization policies, 0.0.0.0/0 indicates that access by all IPs is either allowed or rejected. [Tutorial](#)

Priority: 1 The priority value range is 1 - 100. The default value is 1 which is the highest priority.

OK Cancel

Allow remote login from a specific IP address

Take a Linux instance as an example. Configure to allow a specific IP address to access port 22.

A Linux instance is used in the following example. Allow a specific IP address to SSH the instance.

To allow remote login from a specific IP address, perform the following:

Log on to the ECS console.

Click **Instances** in the left-side navigation pane.

Select your desired region.

Select your desired instance, and click **Manage**.

Select **Security Groups** in the left-side navigation pane, and click **Configure Rules** of one

security group in the list.

Click **Internet Inbound** and then click **Add Security Group Rules**.

On the **Add Security Group Rules** dialog box:

- Select **Allow** for the **Authorization Policy**.
- Select **Custom TCP** for the **Protocol Type**.
- Enter **22/22** as the **Port Range**.
- Select **Address Field Access** for the **Authorization Type** and enter an IP address with or without a CIDR prefix, for example 1.2.3.4 or 1.2.3.4/24, as the **Authorization Object**.
- Enter **1** for **Priority**.
- Click **OK**.

The screenshot shows the 'Add Security Group Rules' dialog box with the following configuration:

- NIC:** Internet
- Rule Direction:** Inbound
- Authorization Policy:** Allow
- Protocol Type:** Custom TCP
- * Port Range:** 22/22
- Authorization Type:** Address Field Access
- * Authorization Object:** 1.2.3.4
- Priority:** 1

Informational text on the right side of the dialog box:

- For Port Range: The value range is 1-65535. For example, "1/200", "80/80". [Tutorial](#)
- For Authorization Object: Please be cautious when setting authorization objects. Based on different authorization policies, 0.0.0.0/0 indicates that access by all IPs is either allowed or rejected. [Tutorial](#)
- For Priority: The priority value range is 1 - 100. The default value is 1 which is the highest priority.

Buttons: OK, Cancel

Add another security group rule:

- Select **Drop** for the **Authorization Policy**.
- Select **Custom TCP** for the **Protocol Type**.
- Enter **22/22** as the **Port Range**.

Click **Instances** in the left-side navigation pane.

Select your desired region.

Select your desired instance, and click **Manage**.

Select **Security Groups** in the left-side navigation pane, and click **Configure Rules**.

Click **Internet Outbound** and then click **Add Security Group Rules**.

On the **Add Security Group Rules** dialog box:

- Select **Drop** for the **Authorization Policy**.
- Select **All** for the **Protocol Type**.
- Select **Address Field Access** for the **Authorization Type** and enter **0.0.0.0/0** as the **Authorization Object**.
- Enter **2** for **Priority**.
- Click **OK**.

Add Security Group Rules

NIC: Internet

Rule Direction: Outbound

Authorization Policy: Drop

Protocol Type: All

* Port Range: -1/-1

The value range is 1-65535. For example, "1/200", "80/80". [Tutorial](#)

Authorization Type: Address Field Access

* Authorization Object: 0.0.0.0/0

Please be cautious when setting authorization objects. Based on different authorization policies, 0.0.0.0/0 indicates that access by all IPs is either allowed or rejected. [Tutorial](#)

Priority: 2

The priority value range is 1 - 100. The default value is 1 which is the highest priority.

OK **Cancel**

Add another security group rule:

- Select **Allow** for the **Authorization Policy**.
- Select **All** for the **Protocol Type**.
- Select **Address Field Access** for the **Authorization Type** and enter an IP address with or without a CIDR prefix, for example 1.2.3.4 or 1.2.3.4/24, as the **Authorization Object**.
- Enter **1** for **Priority**.
- Click **OK**.

The screenshot shows the 'Add Security Group Rules' dialog box with the following configuration:

- NIC: Internet
- Rule Direction: Outbound
- Authorization Policy: Allow
- Protocol Type: All
- * Port Range: -1/-1
- Authorization Type: Address Field Access
- * Authorization Object: 1.2.3.4
- Priority: 1

Help text for Port Range: The value range is 1-65535. For example, "1/200", "80/80". [Tutorial](#)

Help text for Authorization Object: Please be cautious when setting authorization objects. Based on different authorization policies, 0.0.0.0/0 indicates that access by all IPs is either allowed or rejected. [Tutorial](#)

Help text for Priority: The priority value range is 1 - 100. The default value is 1 which is the highest priority.

To check that the rules were successfully configured, log on to the instance and runping or telnet. If you do not have access to IP addresses except for the IP address previously specified for access authorization, the configuration is successful.

This document introduces the default rules of security groups that are created by the system and by yourself.

Security groups created by the system

The security group created by the system has only rules for access over all ICMP ports, TCP Port 22,

and TCP Port 3389, of which,

- All ICMP ports are used by network devices, including routers, to send error messages and operational information.
- TCP Port 22 is used to connect to a Linux instance using SSH.
- TCP Port 3389 is used to remotely connect to a Windows instance using Windows Remote Desktop.

Classic network

The default security group created by the system includes rules of:

- **Intranet:** drops inbound traffic on all ports, and accept outbound traffic on all ports.
- **Internet:** accepts outbound traffic on all ports, but only accept inbound traffic on TCP Port 22, TCP Port 3389, and all ICMP ports.

VPC

A security group for VPC only has intranet rules for inbound and outbound traffic.

The default security group created by the system includes rules of:

- **Outbound:** accepts outbound traffic on all ports.
- **Inbound:** accepts inbound traffic on TCP Port 22, TCP Port 3389, and all ICMP ports.

Inbound		Outbound				
Authorization Policy	Protocol Type	Port Range	Authorization Type	Authorization Object	Description	Priority
Allow	Custom TCP	22/22	Address Field Access	0.0.0.0/0	-	110
Allow	Custom TCP	3389/3389	Address Field Access	0.0.0.0/0	-	110
Allow	All ICMP	-1/-1	Address Field Access	0.0.0.0/0	-	110

All the default security group rules have the priority of 110. Priority 110 means that these rules have the lowest priority in the group. When you manually create a security group, only a value from 1 to 100 is valid for Priority.

User-defined security groups

For user-defined security groups, the default rules of the default security group are as follows:

- Allow all for outbound traffic.
- Drop all for inbound traffic, for both intranet and Internet.

Create a security group

A security group functions similarly to virtual firewalls, and is used to set network access controls for one or more ECS instances. When creating instances, you must select a security group. You can also add security group rules to control outbound and inbound network access for all ECS instances in the security group.

To create a security group, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Security Groups**.

Select a region.

Click **Create Security Group**. In the displayed dialog box, enter the following:

Security Group Name

The length must be 2–128 characters. It can contain uppercase letters, lowercase letters, and Chinese characters. It cannot contain numbers, underscores (_), or hyphens (-).

Description

The length must be 2–256 characters. Do not start with http:// or https://.

- Network Type

There are two network types, Classic network and VPC. If you select VPC, you must select a specific VPC. If no VPCs have been created in the current region, you must create one first.

Click **OK**.

You can add security group rules to enable or disable access to and from the Internet, intranet, or private networks for ECS instances in the security group:

- **VPC network**: You only need to set outbound and inbound rules, and do not need different rules for private networks and Internet.
- **Classic network**: It is required to set outbound and inbound rules for Internet and intranet respectively.

Changes to the security group rules are automatically applied to ECS instances in the security group.

Prerequisites

You have created a security group. For more information, see [Create a security group](#).

You know which Internet, intranet, or private network requests need to be allowed or dropped for your instance.

Procedure

To add a security group rule, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security > Security Groups**.

Select a region.

Find the security group to add authorization rules, and in the **Action** column click **Configure Rules**.

On the **Security Group Rules** page, click **Add Security Group Rules**.

(Optional) If you do not need to enable or disable all ports for all protocols, ICMP, or GRE, you can select **Quickly Create Rules**.

In the dialog box, set the following parameters:

- **NIC:**

- If the security group is for VPC, you do not need to select the NIC.
 - If your instances can access the Internet, the rules work for both the Internet and intranet.
 - If your instances cannot access the Internet, the rules only works for the intranet.
- If the security group is for Classic network, you must select **Internet** or **Intranet**.

- **Rule Direction:**

- **Outbound:** ECS instances access other ECS instances over intranet, private networks, or through Internet resources.
- **Inbound:** Other ECS instances in the intranet or private networks and Internet resources access the ECS instance.

- **Authorization Policy:** Select **Allow** or **Drop**.

Note:

Drop policy discards the data packet without returning a response. If two security groups overlap except the authorization policy, the **Drop** rule takes priority over the **Allow** rule.

- **Protocol Type and Port Range:** The port range setting is affected by the selected protocol type. The following table shows the relationship between protocol types and port ranges.

Protocol type	Port range	Scenarios
All	Shown as -1/-1, indicating all ports.	Used in scenarios where both applications are fully mutually trusted.
All ICMP	Shown as -1/-1, indicating all ports.	Used to detect the instance' s network connection status by using the Ping tool.
All GRE	Shown as -1/-1, indicating all ports.	Used for VPN service.
Custom TCP	For custom port ranges, the valid port value is 1–65535, and the valid port range format is Start Port/End Port . A valid port range format must be used for one port. For example, use 80/80 to indicate port 80.	Used to allow or deny one or several successive ports.
Custom UDP		
SSH	Shown as 22/22, the default SSH port 22.	Used for remotely connecting to Linux instances.
TELNET	Shown as 23/23.	Used to remotely log on to instances by using Telnet.
HTTP	Shown as 80/80.	The instance is used as a server for a website or a web application.
HTTPS	Shown as 443/443.	The instance is used as a server for a website or a web application that supports the HTTPS protocol.
MS SQL	Shown as 1433/1433.	The instance is used as a MS SQL server.

Oracle	Shown as 1521/1521.	The instance is used as an Oracle SQL server.
MySQL	Shown as 3306/3306.	The instance is used as a MySQL server.
RDP	Shown as 3389/3389, the default RDP port 3389.	Used for remotely connecting to Windows instances.
PostgreSQL	Shown as 5432/5432.	The instance is used as a PostgreSQL server.
Redis	Shown as 6379/6379.	The instance is used as a Redis server.

Port 25 is disabled by default, and it cannot be enabled by adding security group rules.

- **Authorization Type** and **Authorization Object**: The authorization object affects setting of authorization type. The following table shows the relationship between them.

Authorization type	Authorization object
Address Field Access	Use the IP or CIDR block format such as <i>10.0.0.0</i> or <i>192.168.0.0/24</i> . Only IPv4 addresses are supported. 0.0.0.0/0 indicates all IP addresses.
Security Group Access	<p>Authorize the instances in a security group under your account or another account to access the instances in this security group.</p> <ul style="list-style-type: none"> • Authorize This Account: Select a security group under your account. • Authorize Other Account: Enter the target security group ID and the Account ID. You can view the account ID in Account Management > Security Settings. <p>For VPC network instances, Security Group Access works for private IP addresses only. If you want to authorize Internet IP address access, use Address Field Access.</p>

Note:

To guarantee the security of your instance, when you are configuring an intranet inbound rule for a security group of the Classic network type, **Security Group Access** is the top priority for **Authorization Type**. If you want to select **Address Field Access**, and you want to enter an IP address in the CIDR format, you must enter an IP address in the format of a.b.c.d/32. Only 32 is the valid CIDR prefix.

- **Priority:** 1–100. The smaller the number is, the higher the priority is. For more information on priority, see [Security group rule priority](#).

7. Click **OK** to add the security group rule to the specified security group.

Check whether security group rules takes effect

If you have installed a web service in the instance and added a security group rule in a security group: allow all IP addresses to have inbound access to TCP port 80 of the instance and follow these steps according to your instance OS.

Security group rules usually take effect immediately, though some delays may occur.

Linux instances

If it is a Linux instance in the security group, follow these steps to check whether the security group rule has been activated.

Connect to the ECS instance.

Run the following command to check whether TCP 80 is being listened to.

```
netstat -an | grep 80
```

If the following result is returned, web service for TCP port 80 is enabled.

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

Enter http://public IP address in the address bar of a browser. If the rule has taken effect, you can successfully access the address.

Windows instances

If it is a Windows instance in the security group, follow these steps to check whether the security group rule has been activated.

Connect to the ECS instance.

Run **cmd**, and run the following command to check whether TCP 3389 is being listened to.

```
netstat -aon | findstr :80
```

If the following result is returned, TCP port 3389 is enabled.

```
TCP 0.0.0.0:80 0.0.0.0:0 LISTENING 1172
```

Enter <http://Public IP address> in the address bar of a browser. If the rule has taken effect, you can successfully access the address.

Security group rule priority

The **Priority** of a security group rule can be a number from 1 to 100. The smaller the number is, the higher the priority is.

ECS instances can belong to different security groups. As a result, instances may have multiple security group rules that apply for the same protocol types, port ranges, authorization types, and authorization objects. The rule that takes effect depends on the settings, as shown in the following table. See the **Result** column of the table for more information about why a particular rule is followed.

No.	Security group rules	Priority	Authorization policy	Result
i	A	Equal	Allow	B takes effect. If two security group rules have the equal priority and differ from each other only in the authorization policy, the Drop rule has priority and takes effect.
	B		Drop	
ii	C	1	Allow	C takes effect. The security rule with the higher priority takes effect.
	D	2	Drop	

View the security group list

You can view the security groups on the ECS console at any time.

To view the security groups list, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Security Groups**.

Select a region. A list of all the security groups in the specified region will be displayed.

(Optional) You can select **VPC ID** in the filter input box, and enter a specific VPC ID to search, then all the security groups under this VPC appear.

You can modify the name and description of a security group at any time.

To modify the name and description of a security group, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Security Groups**.

Select a region to display all all the security groups in this region.

Modify attributes of a security group:

- Modify the name: Hover the cursor over the name of a security group, and then click the pen icon that appears.
- Modify the name and description: Click **Modify**, and then enter a new name and description in the dialog box.

Click **OK**.

View the rules of a security group

You can view the rules of a security group at any time.

To view the rules of a security group, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Networks & Security > Security Groups**.

Select a region.

Select a security group.

Click **Configure Rules**. The following security group rule tabs will be displayed for Classic Networks and VPCs:

- For Classic networks
 - Internet Inbound
 - Internet Outbound
 - Intranet Inbound
 - Intranet Outbound
- For VPCs
 - Inbound
 - Outbound

Click a tab to view the security group rules for that type.

You can delete security group rules if you no longer need them.

To delete rules in a security group, perform the following:

Log on to the ECS console.

In the left-side navigation pane, click **Security Groups**.

Select a region.

Find the security group where you want to delete rules, and in the **Action** column, click **Configure Rules**.

On the security group management page, choose the rule direction and find the rule you want to delete.

- If the security group is for Classic network, the rule directions are **Internet Inbound**,

Internet Outbound, Intranet Inbound, and Intranet Outbound.

- If the security group is for VPC network, the rule directions are **Inbound** and **Outbound**.

In the **Action** column, click **Delete**.

On **Delete Security Group Rules** dialog box, read and confirm the notes, and then click **OK**.

You can delete security groups, if you no longer require them.

Note:

- Before deleting a security group, ensure it does not contain instances and is not referenced in the rules of another security group.
- Deleting a security group will delete all its rules.

To delete a security group, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Networks & Security > Security Groups**.

Select a region to display a list of all the security groups in the region.

Select one or more security groups.

Click **Delete**.

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

Alibaba Cloud supports cloning a security group across regions and network types.

Application scenarios

You may need to clone a security group in the following scenarios:

You have created a security group, named SG1, in Region A, and you want to apply the same rules of SG1 to ECS instances in Region B. Then you can clone SG1 to Region B without creating a new security group in Region B.

You have a security group in Classic network, named SG2. You want to apply the rules of SG2 to instances in a VPC. You can clone SG2 and choose VPC as the network type when configuring the cloning. Then in VPC network, you have a new security group that has same rules with SG2.

If you want to apply new security group rules to an ECS instance that are running an online business application, we recommend that you clone the security group as a backup before modifying the rules. If the new security group rules are disadvantageous to the online business application, you can restore the rules completely or partly.

Prerequisites

If you want to change the network type of a security group from Classic to VPC, you have to create a VPC and VSwitch in the target region first.

Procedure

Follow the steps to clone a security group.

Log on to the ECS console.

In the left-side navigation pane, choose **Network & Security > Security Groups**.

On the **Security Group List** page , select the target region.

Find the target security group, and in the **Action** column, click **Clone Security Group**.

In the **Clone Security Group** dialog box, set the new security group information:

- **Destination Region:** Select a region suitable for the new security group. Not all regions are supported now. The supported regions are displayed in the drop-down list.
- **Security Group Name:** Specify a new name for the new security group.
- **Network Type:** Select a network type suitable for the new security group. If VPC is selected, you have to choose one VPC in the drop-down list.

Click **OK**.

The new security group is displayed in the Security Group List.

Restoring security group rules refers to the process of completely or partially restoring the rules in the original security group to those of a target security group. Specifically:

Completely restoring refers to moving the rules that do not exist in the target security group from the original security group and adding the rules that only exist in the target security group to the original security group. After restoration, rules in the original security group are identical with those in the target security group.

Partially restoring refers to adding the rules that only exist in the target security group to the original security group and ignoring the rules that only exist in the original group.

Limits

Restoring security group rules has the following limits:

The original security group and the target security group must be in the same region.

The original security group and the target security group must be of the same network type.

If any system-level security group rules, of which the priority is 110, exist in the target security group, they are not created during restoration. After restoration, the rules in the original security group may be different from what is expected. If you need the system-level security group rules, you have to manually create the rules and set their priority to 100.

Use cases

If you want to apply new security group rules to an ECS instance that is running an online business application, you can clone the former security group as a backup, and then modify the rules inside. If the new security group rules impair the online business application, you can restore the rules fully or partially.

Prerequisites

You must own at least one security group of the same network type in the same region.

Procedure

To restore your security group rules, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Network & Security** > **Security Groups**.

Select a region.

Find the security group you want to restore rules for as the original security group, and in the **Action** column, click **Restore rules**.

In the **Restore rules** dialog box, follow these steps:

- i. Select the **Target security group**: Select a security group as the target security group that must have different rules from the original security group.
- ii. Select a **Restore type**:
 - i. If you want the original security group to have the same rules as the target security group, select **Completely restored**.
 - ii. If you only want to add the rules that only exist in the target security group to the original security group, select **Partially restored**.
- iii. In the **Result preview** area, preview the restoration result:
 - i. Rules highlighted in green only exist in the target security group. No matter whether you choose **Completely restored** or **Partially restored**, these rules are added to the original security group.
 - ii. Rules highlighted in red are the rules that do not exist in the target security group. If **Completely restored** is selected, the system removes these rules from the original security group. If **Partially restored** is selected, the rules are retained in the original security group.

Click **OK**.

The **Restore rules** dialog box is closed automatically after successful creation. In the **Security Group List**, find the original security group you restored the rules for. In the **Action** column, click **Configure Rules** to enter the **Security Group Rules** page and view the updated security group rules.

Key pairs

Limits

The SSH key pair, abbreviated as key pair, applies to Linux instances only.

Alibaba Cloud only supports the creation of 2048-bit RSA key pairs.

- Alibaba Cloud keeps the public key of the key pair.
- After creating the key pair, you must save and keep the private key of the key pair for further use.
- The private key follows the unencrypted PEM-encoded PKCS#8 format.

An Alibaba Cloud account can have a maximum of 500 key pairs per region.

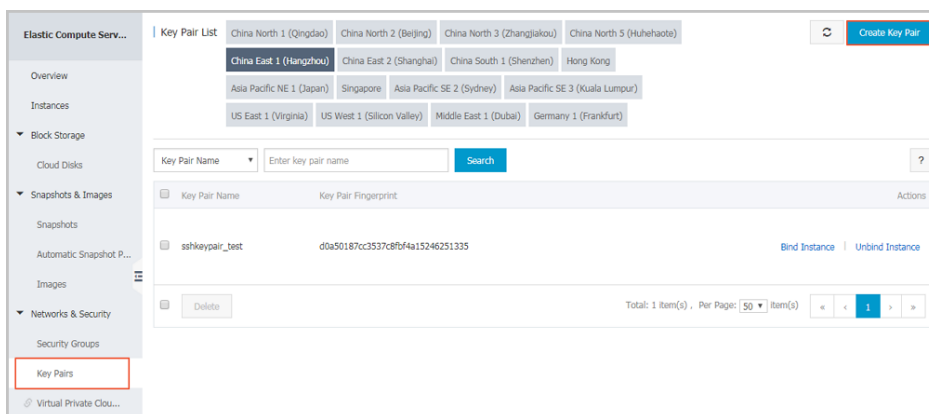
Create an SSH key pair

To create an SSH key pair, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Networks & Security** > **Key Pairs**.

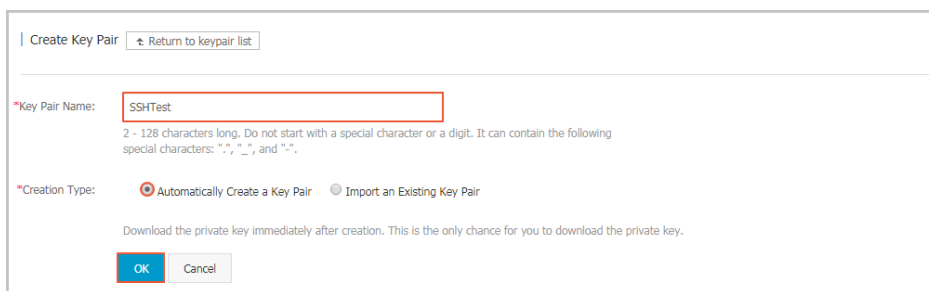
On the **Key Pairs** page, select a region, and click **Create Key Pair**.



On the **Create Key Pair** page, enter a name for the key pair, and select **Automatically Create a Key Pair** for the **Creation Type**.

Note:

The specified key pair name must be unique. It must not match with the existing key pair or a key pair that was deleted when it was still bound to an instance. Otherwise, an error message “The key pair already exists” appears.



Click **OK** to create a key pair.

Note: After a key pair is created, you must download and save the private key for further use. If you do not have the private key, you cannot log on to your ECS instance that is bound to this key pair.

After creating the key pair, you can view the information, including **Key Pair Name** and **Key Pair Fingerprint**, in the key pair list.

Follow-up operation

After creating an SSH key pair, you can bind or unbind it to an ECS instance.

If you prefer to use another key generation tool, you can use it to generate an RSA key pair and then import its public key into Alibaba Cloud. See [Introduction to SSH key pairs](#) for the supported types of imported key pairs.

Note: To make sure your instance's security, keep the private key of the key pair secure and do not import the private key to Alibaba Cloud.

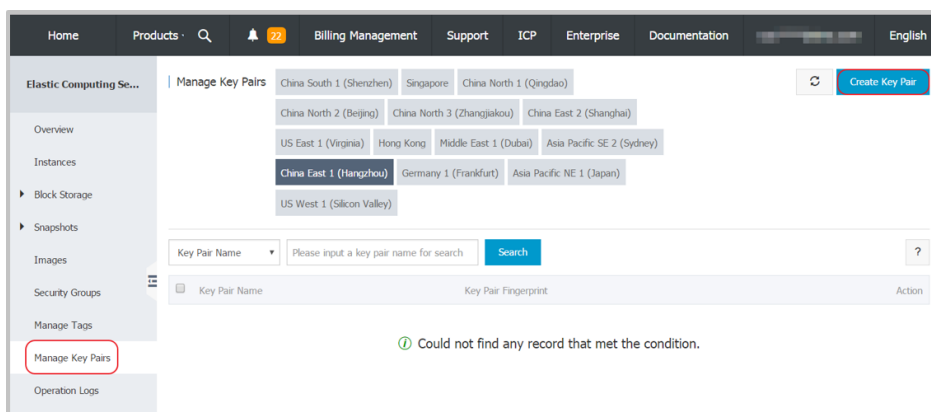
To import an SSH key pair, you must have a key pair that has been generated using another tool, and the public key to be imported into Alibaba Cloud must be Base64-encoded.

To import an SSH key pair, follow these steps.

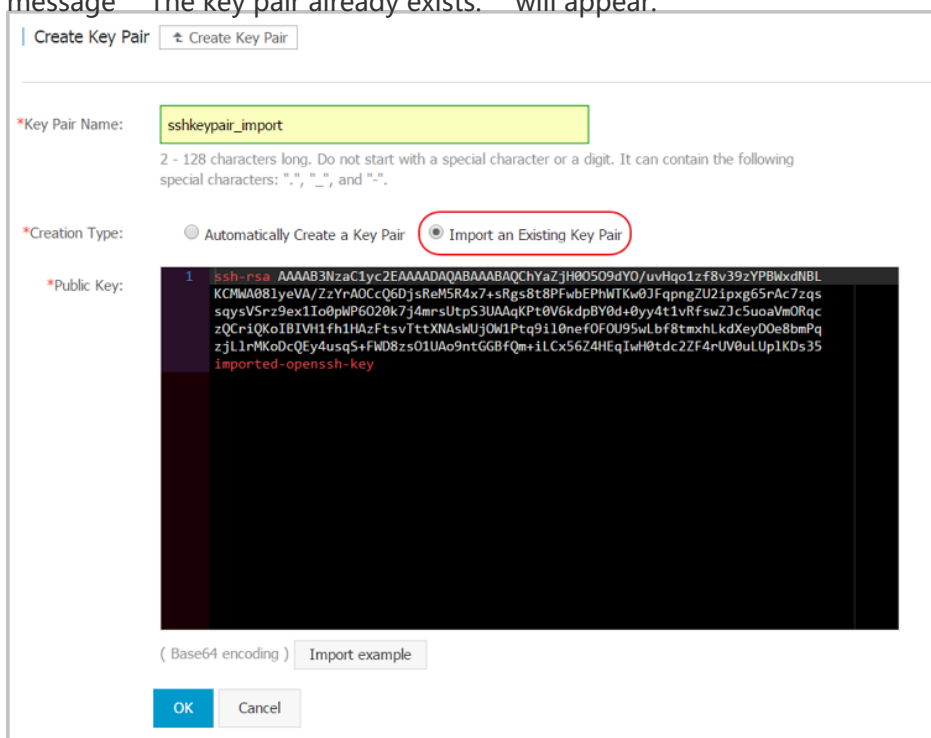
Log on to the ECS console.

In the navigation pane, click **Key Pairs** under **Networks & Security**.

On the **Key Pair** page, select the region in which to import a key pair, and then click **Create Key Pair**.



On the **Create Key Pair** page, enter a name for the key pair, and select **Import an Existing Key Pair**. The specified key pair name must not be the same with that of an existing key pair or a key pair that was deleted when it was still bound to an instance. Otherwise, an error message **"The key pair already exists."** will appear.



Click **OK**.

After creation, you can view the information, including **Key Pair Name** and **Key Pair Fingerprint**, in the key pair list.

Limits

- Only applicable to Linux instance.

- When your ECS instance is in the Running status, **restart** it after you bind an SSH key pair to the instance.
- After unbinding a key pair from an instance, you must **restart** the instance to disable the SSH key pair.
- If you use password-based authentication for Linux logon, the password authentication feature is automatically disabled after the key pair is bound.
- After an SSH key pair is unbound, you must **reset** the instance password for successful connection.
- Except for the non-I/O optimized generation I instance, all the Linux instance in the instance generations and type families support the authentication method of SSH key pair.

Bind an SSH key pair

To bind an SSH key pair to an ECS instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Networks & Security > Key Pairs**.

Select a region.

Choose a key pair, and then click **Bind Instance**. On the **Bind Instance** dialog box, in the **Select Instance** box, select one or more instances, and then click the icon >.

Note:

In the **Select Instance** box, the instance names in gray are either Windows instances or non-I/O-optimized instances of Generation I. Those instances do not support SSH key pairs.

Click **OK**.

Unbind an SSH key pair

To unbind an SSH key pair from an ECS instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, choose **Networks & Security > Key Pairs**.

Select a region.

Choose an SSH key pair, and then click **Unbind Instance**. On the **Unbind Instance** dialog, in the **Select Instance** box, select one or more instances, and then click the icon >.

Click **OK**.

Delete an SSH key pair

If you no longer require a key pair, you can delete it. Note that a deleted key pair is not recoverable. Existing instances that have used the key pair will not be affected, and the deleted key pair name will remain associated to the instance.

See the following steps to delete key pairs.

Log on to the ECS console.

In the left-side navigation pane, click **Key Pairs** under **Networks & Security**.

Select one or more key pairs.

Click **Delete** > **OK**.

Note:

- If you delete a key pair that is still bound to an instance, its name is not available for you to create or import a key pair again. Otherwise, an error message "The key pair already exists." will appear when you are using the same name to create or import a key pair.
- If you delete a key pair that is not bound to an instance, its name is still available for you to create or import a key pair again.

Elastic Network Interfaces

You can create an ENI in the ECS console, and then **attach** it to an instance. This document describes how to create an ENI in the ECS console.

Limits

To create an ENI, you have the following limits:

- Each ENI must be in a VSwitch of a VPC.
- Each ENI must be in one security group.

Prerequisites

Before you create an ENI, finish the following operations:

- Create a VPC and then create a VSwitch in the VPC.
- Create a security group in the same VPC.

Procedure

To create an ENI, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security > Network Interfaces**.

Select a region.

Click **Create**.

In the **Create** dialog box, finish the following configurations:

- Network Interface Name:** Specify a name for the ENI.
- VPC:** Select a VPC.

Note:

When you attach an ENI to an instance, they must be in the same VPC. In addition, after an ENI is created, you cannot change the VPC.

- VSwitch:** Select a VSwitch.

Note:

When you attach an ENI to an instance, they must be in the same zone, but they do not have to be in the same VSwitch. In addition, after an ENI is

created, you cannot change the VSwitch.

- iv. **IP**: Specify an IPv4 address as the private IP address of the ENI. The IPv4 address must be available in the CIDR block of the specified VSwitch. If you do not specify one, a private IP address is automatically assigned to your ENI after the ENI is created.
- v. **SecurityGroup**: Select a security group in the selected VPC.
- vi. **Description**: Give a brief description for the ENI for easing further management.
- vii. Click **OK**.

In the **Network Interfaces** page, refresh the table. When the new ENI is in the **Available** status, it is created successfully.

Follow-up operations

After you create an ENI, perform the following operations:

- Attaching an ENI to an instance
- Modifying attributes of an ENI
- Deleting an ENI

APIs

CreateNetworkInterface

You can attach an ENI to an instance. This document describes how to attach an ENI to an instance in the ECS console.

Limits

To attach an ENI to an instance, you have the following limits:

- You can only attach a secondary ENI to an instance.
- You can only attach an ENI to a VPC instance, and they must be in the same VPC.
- The VSwitches of the ENI and the instance do not have to be the same, but they must be in the same zone.
- The ENI must be in the **Available** status.
- The instance must be in the **Stopped** status.
- An ENI can only be attached to one VPC instance at a time. However, a VPC instance can be associated with multiple ENIs. For more information about the maximum number of ENIs that can be attached to one instance, see [Elastic network interfaces](#).

Prerequisites

Before you attach an ENI to an instance, finish the following operations:

- Create an ENI.
- Make sure the ENI is in the **Available** status.
- Make sure your instance can be associated with more than one ENI, and stop the instance.

Procedure

To attach an ENI to an instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security > Network Interfaces**.

Select a region.

Find an available ENI, and in the **Operations** column, click **Attach**.

In the **Attach** dialog box, select an instance, and then click **OK**.

In the **Network Interfaces** page, refresh the table. When the selected ENI is in the **InUse** status, it is successfully attached to the instance.

Follow-up operations

After an ENI is attached to an instance, you can perform the following operations:

- Detaching the ENI from an instance, and then deleting the ENI
- Modifying attributes of the ENI
- Configuring the ENI, if the ENI cannot be recognized by the operating system of your instance.

APIs

AttachNetworkInterface

You can detach a secondary ENI, but not the primary ENI, from an instance.

Limits

To detach a secondary ENI from an instance, you have the following limits:

- The secondary ENI must be in the **InUse** status.
- The instance must be in the **Stopped** status.

Prerequisites

Your ENI is attached to an instance. Before you detach an ENI from an instance, you must stop the instance.

Procedure

To detach a secondary ENI from an instance, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security > Network Interfaces**.

Select a region.

Note:

Now, the ENI feature is available in the following regions: China North 1, China North 2, China North 3, China East 2, Hong Kong, Asia Pacific SE 1 (Singapore), Asia Pacific SE 3 (Kuala Lumpur), Middle East 1 (Dubai), and Germany 1 (Frankfurt).

Find an ENI in the **InUse** status, and in the **Operations** column, click **Detach**.

In the **Detach** dialog box, confirm the information, and then click **OK**.

In the **Network Interfaces** page, refresh the table. When the selected ENI is in the **Available** status, it is successfully detached from the instance.

Follow-up operations

After an ENI is detached from an instance, you can perform these operations:

- Attaching the ENI to another instance
- Deleting the ENI

- Modifying attributes of the ENI

APIs

DetachNetworkInterface

You can modify attributes of a secondary ENI, but not the primary ENI, of an instance. You can modify the following attributes of an ENI:

- The name of the ENI.
- The security group associated with the ENI. One ENI must be associated with at least one security group. However, it cannot be associated with more than five security groups.
- Description of the ENI.

You can modify attributes of an ENI when it is in the **Available** or the **InUse** status.

This document describes how to modify attributes of an ENI in the ECS console.

Prerequisites

Before you modify attributes of an ENI, create an ENI.

Procedure

To modify attributes of an ENI, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security** > **Network Interfaces**.

Select a region.

Find an ENI, and in the **Operations** column, click **Modify**.

In the **Modify** dialog box, modify the following optional configurations:

- **Network Interface Name**: Specify a new name for the selected ENI.
 - **SecurityGroup**: Select more security groups for the ENI, or remove security groups.
 - **Description**: Give a brief description for the ENI.
- After you finish the modification, click **OK**.

APIs

ModifyNetworkInterfaceAttribute

If you do not require an ENI, you can delete it. But you can only delete a secondary ENI, but not the primary ENI of an instance.

After an ENI is deleted, the primary private IP address of the ENI is released automatically, and the ENI is automatically removed from all associated security groups.

An ENI will be deleted along with an instance if you did not detach it from the instance before you release the instance.

Limits

You can only delete an ENI in the **Available** status.

Prerequisites

If an ENI is attached to an instance, detach it from the instance.

Procedure

To delete an ENI, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security > Network Interfaces**.

Select a region.

Find an available ENI, and in the **Operations** column, click **Delete**.

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

In the **Network Interfaces** page, refresh the table. If the ENI disappears, it is deleted successfully.

APIs

DeleteNetworkInterface

If your instance is running one of the following images, you do not have to configure the Elastic

Network Interfaces (ENI) manually to have them recognized by the OS.

- CentOS 7.3 64-bit
- CentOS 6.8 64-bit
- Windows Server 2016 Data Center Edition 64-bit
- Windows Server 2012 R2 Data Center Edition 64-bit

If your instance is running none of the preceding images, and you want to attach an ENI to your instance, you must manually configure the ENI to be recognizable.

This document uses an instance running CentOS 7.2 64-bit as an example to introduce how to configure an ENI to make the interface recognizable.

Prerequisites

You must attach an ENI to an ECS instance.

Procedure

To configure the ENI, follow these steps:

Use the `DescribeNetworkInterfaces` interface or log on to the ECS console to obtain the following attributes of the ENI: primary private IP address, subnet mask, the default route, and the MAC address. To obtain these attributes in the ECS console, follow these steps:

- Log on to the ECS console.

In the left-side navigation pane, select **Networks & Security** > **Network Interfaces**.

Select a region.

Find a network interface, and obtain its primary private IP address, subnet mask, default route, and MAC address.

Example

```
eth1 10.0.0.20/24 10.0.0.253 00:16:3e:12:e7:27
eth2 10.0.0.21/24 10.0.0.253 00:16:3e:12:16:ec
```

Connect to the ECS instance.

Run the command to generate the configuration file of the network interface.

```
cat /etc/sysconfig/network-scripts/ifcfg-[network interface name in the OS]
```

Note:

- Pay attention to the relation between network interface name in the OS and the MAC address.
- The default route must be set to DEFROUTE=no. Other editions must have the same configuration. Note that running the ifup command may change the active default route configuration after configuring the network interface.

Example:

```
# cat /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1
BOOTPROTO=dhcp
ONBOOT=yes
TYPE=Ethernet
USERCTL=yes
PEERDNS=no
IPV6INIT=no
PERSISTENT_DHCLIENT=yes
HWADDR=00:16:3e:12:e7:27
DEFROUTE=no
```

Follow these steps to start the network interface:

- Run the ifup [network interface name in the OS] command to start the dhclient process, and initiate a DHCP request.

Example:

```
# ifup eth1
# ifup eth2
```

- After a response is received, run the ip a command to check the IP allocation on the network interfaces, which must match with the information displayed on the ECS console.

Example:

```
# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
inet 127.0.0.1/8 scope host lo
valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
```

```

link/ether 00:16:3e:0e:16:21 brd ff:ff:ff:ff:ff:ff
inet 10.0.0.19/24 brd 10.0.0.255 scope global dynamic eth0
valid_lft 31506157sec preferred_lft 31506157sec
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
link/ether 00:16:3e:12:e7:27 brd ff:ff:ff:ff:ff:ff
inet 10.0.0.20/24 brd 10.0.0.255 scope global dynamic eth1
valid_lft 31525994sec preferred_lft 31525994sec
4: eth2: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
qlen 1000
link/ether 00:16:3e:12:16:ec brd ff:ff:ff:ff:ff:ff
inet 10.0.0.21/24 brd 10.0.0.255 scope global dynamic eth2
valid_lft 31526009sec preferred_lft 31526009sec

```

Set the metric for each network interface in the route table. In this example, set the metric parameters of eth1 and eth2 as follows.

Example:

```

eth1: gw: 10.0.0.253 metric: 1001
eth2: gw: 10.0.0.253 metric: 1002

```

- i. Run the following command to set the metric parameters.

Example:

```

# ip -4 route add default via 10.0.0.253 dev eth1 metric 1001
# ip -4 route add default via 10.0.0.253 dev eth2 metric 1002

```

- ii. Run the route -n command to check whether the configuration is successful or not.

Example:

```

# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
0.0.0.0 10.0.0.253 0.0.0.0 UG 0 0 0 eth0
0.0.0.0 10.0.0.253 0.0.0.0 UG 1001 0 0 eth1
0.0.0.0 10.0.0.253 0.0.0.0 UG 1002 0 0 eth2
10.0.0.0 0.0.0.0 255.255.255.0 U 0 0 0 eth0
10.0.0.0 0.0.0.0 255.255.255.0 U 0 0 0 eth1
10.0.0.0 0.0.0.0 255.255.255.0 U 0 0 0 eth2
169.254.0.0 0.0.0.0 255.255.0.0 U 1002 0 0 eth0
169.254.0.0 0.0.0.0 255.255.0.0 U 1003 0 0 eth1
169.254.0.0 0.0.0.0 255.255.0.0 U 1004 0 0 eth2

```

Follow these steps to build a route table:

Note:

We recommend that you use the metric value as the route table name.

Run the command to build a route table.

Example:

```
# ip -4 route add default via 10.0.0.253 dev eth1 table 1001
# ip -4 route add default via 10.0.0.253 dev eth2 table 1002
```

Run the command to check whether the route table is built successfully or not.

Example:

```
# ip route list table 1001
default via 10.0.0.253 dev eth1
# ip route list table 1002
default via 10.0.0.253 dev eth2
```

Follow these steps to configure the policy-based routing:

- i. Run the command to configure policy-based routing.

Example:

```
# ip -4 rule add from 10.0.0.20 lookup 1001
# ip -4 rule add from 10.0.0.21 lookup 1002
```

- ii. Run the `ip rule list` command to view the route rules.

Example:

```
# ip rule list
0: from all lookup local
32764: from 10.0.0.21 lookup 1002
32765: from 10.0.0.20 lookup 1001
32766: from all lookup main
32767: from all lookup default
```


Now, you have finished the ENI configuration.

Tags

You can bind tags to the following resources in the ECS console: ECS instance, storage, snapshot, image, and security group.

Tags have the following limits:

- Each tag has a key-value pair.
- You can bind ten tags to an instance at most. You can bind five tags at most to an instance at a time.
- Every tag key of a resource must be unique. A tag with the same key as an existing one will be overwritten.
- Tag information is not shared across regions. For example, tags created in China East 1 (Hangzhou) are invisible to China East 2 (Shanghai).
- If a tag is unbound and no longer bound to any other resource, the tag will be automatically deleted.

If your account maintains various types of resources that are associated with each other in different ways, you can bind tags to the resources to categorize and manage the resources in a unified manner.

You can bind ten tags to a resource at most. You can bind/unbind five tags at most for the resource each time.

Take the following steps to bind resources with tags:

1. Log on to the ECS console.
2. Select the resource type in the left-side navigation bar for the binding operation, such as **Instance**, **Cloud Disks**, **Snapshot**, **Image**, and **Security Groups**.
3. Select a region.
4. Select the resources in the resource list to bind tags.
5. Click **Edit Tags** at the bottom of the resource list.

Choose **More > Edit Tags** at the bottom of the resource list if the selected resources are **Instance**.

6. Select or cross off tags in the dialog box:
 - Click **Available Tags** and select available tags in the tag list for the selected

resource.

- Click **Create** and set **Key** and **Value** if no tags are available for the selected resource:
 - **Key** is mandatory whereas **Value** is optional.
 - **Key** cannot start with aliyun, http://, or https://. The key is case-insensitive and can contain up to 64 characters.
 - **Value** cannot start with http:// or https://. The value is case-insensitive and can contain up to 128 characters. It can be empty.
 - Any tag **Key** of a resource must be unique. A tag with the same key as an existing one will be overwritten.
- **Available Tags** and **Create** are grayed out if the selected resources are already bound with 10 tags. You need to unbind some tags before binding new tags.

7. Click **Confirm**.

To check if tags are successfully bound, use the **Edit Tags** function of the resource or click **Tags** in the left-side navigation bar of the ECS console. You can click **Tags** with a tag symbol at the top of the resource list to filter resources.

You can unbind a tag from the resource if the tag is no longer applicable to resource management. After a tag is unbound and is no longer bound to any other resource, the tag will be automatically deleted.

- The **Delete Tags** function unbinds one or more tags from an instance at a time.

Currently, this function is only available for instances. It is unavailable for other resource types.

- The **Edit Tags** function unbinds tags one by one. You can unbind five tags from a resource each time.

Unbind tags from instances using the tag deletion function

Currently, the **Delete Tags** function is only available for instances.

See the following steps to delete tags:

1. Log on to the ECS console.
2. Click **Instance** in the left-side navigation pane.
3. Select a region.
4. Select the instance(s) from which you want to unbind tags in the instance list. You can also filter instances by tag and select the expected instance.
5. Choose **More > Delete Tags** at the bottom of the resource list.
6. In the **Delete Tags** dialog box, enter the **Tag Key** of the tags you want to unbind.
7. Click **OK** to complete tag unbinding.



To check whether the tags are successfully unbound, use the **Edit Tags** function of the instance or click **Tags** in the left-side navigation pane of the ECS console.

Unbind tags from resources using the tag edit function

The **Edit Tags** function unbinds one or more tags from a resource.

See the following steps to unbind tags:

1. Log on to the ECS console.
2. In the left-side navigation pane, select the resource type for the unbinding operation, such as **Instance**, **Cloud Disks**, **Snapshots**, **Images**, or **Security Groups**.
3. Select a region.
4. In the resource list, select the resource from which you want to unbind tags. You can also filter resources by tag and select the expected resource.
5. Click **Edit Tags** at the bottom of the resource list.
6. In the **Edit Tags** dialog box, click the deletion icon next to a tag.
7. Click **Confirm** to complete tag unbinding.

To check whether the tags are successfully unbound, use the **Edit Tags** function of the resource or click **Tags** in the left-side navigation pane of the ECS console.

After tags are bound to resources, you can use the following two methods to filter resources by tag.

Filter resources by resource list

See the following steps to filter resources:

1. Log on to the ECS console.
2. In the left-side navigation pane, select the resource type you want to view, such as **Instances, Cloud Disks, Snapshots, Images, or Security Groups**.
3. Select a region.
4. Click **Tag** at the top of the resource list.
 - Click a key to filter out the resources that are bound with this key, which may have multiple values.
 - Click a key and value to filter out the resources that are bound with this key-value pair (tag).

The console returns a list of resources that are bound with the key or key-value pair.

Filter resources by tags

See the following steps to filter resources:

1. Log on to the ECS console.
2. Click **Tags** in the left-side navigation pane.
3. Select a region.
4. Enter a key in the search box and click **Search**.

The console returns a list of resources that are bound with the key.

You can monitor the operating statuses of instances to ensure optimal performance.

You can monitor the status of instances using the following two portals:

- Instance Details page
- CloudMonitor

Monitor the status of an instance by using Instance Details page

To monitor the status of an instance by using Instance Details, follow these steps:

Log on to the ECS console.

In the left-side navigation pane, click **Instances**.

Select a region.

Click an instance to go to the **Instance Details** page.

On the **Instance Details** page, you can view the monitoring information, including CPU utilization and outbound/inbound network traffic information.

- Information about CPU monitoring:
 - For Linux instances, use the `top` command to view CPU usage details. Log on to the instance and run the `top` command in the command line. Then, press Shift+P key to list programs by CPU utilization to view which processes are using the most CPU resources.
 - For Windows instances, use the **Task Manager** on an instance to view the CPU utilization to view which programs are using the CPU resources of the server.
- The displayed monitoring data shows the Internet traffic of the instance in Kbps (1 Mbps = 1,024 Kbps). The monitoring data shows inbound and outbound instance traffic. For 1 Mbps of bandwidth, the bandwidth is working at full capacity if the outbound network traffic reaches 1,024 Kbps.

CloudMonitor

To install a CloudMonitor, follow these steps:

In the Alibaba Cloud Console, choose **Products and Services > CloudMonitor**.

In the left-side navigation pane, click **Host Monitoring**, and then select the name of the instance you want to monitor.

Click **Batch Install** to monitor the instance OS, click **Monitoring Chart** to view basic parameters, and click **Alarm Rules** to set alarm rules.

Note: For more information about CloudMonitor, see [CloudMonitor Product Documentation](#).