ALIBABA CLOUD

Alibaba Cloud

FPGA as a Service
User Guide

Document Version: 20220620

Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloud-authorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.

2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.

3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.

4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).

5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.

6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example
Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.
Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.
Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.
Note	A note indicates supplemental instructions, best practices, tips, and other content.	Note: You can use Ctrl + A to select all files.
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click **Settings**> **Network**> **Set network type**.
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click **OK**.
`Courier font`	Courier font is used for commands	Run the `cd /d C:/window` command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	This format is used for an optional value, where only one item can be selected.	`ipconfig [-all|-t]`
{} or {a|b}	This format is used for a required value, where only one item can be selected.	`switch {active|stand}`

Table of Contents

1.Quick reference

Common operations on ECS instancesECS learning path

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. When you use ECS, you may need to perform various operations on different resources, such as connecting to instances, replacing operating systems, resizing disks, upgrading or downgrading instance configurations, and using snapshots or images. This topic describes common operations on ECS resources.

Limits

- For information about the usage notes of ECS instances, see Usage notes.
- For information about the limits of ECS resources, see Limits and View and increase instance quotas.
-

Create and manage ECS instances

- You can perform the following steps to manage the lifecycle of an ECS instance:
  i. Create an instance by using the wizard
  ii. Connect to an ECS instance
  iii. Stop an instance
  iv. Release an instance
- If the instance type or network configuration of your instance cannot meet your business requirements, you can change the instance type, IP address, and maximum public bandwidth.
  - Subscription instances:
    - Upgrade the instance types of subscription instances
    - Downgrade the configurations of an instance during renewal
  - Pay-as-you-go instances:
    - Change the instance type of a pay-as-you-go instance
    - Modify the bandwidth configurations of pay-as-you-go instances
  - IP addresses of ECS instances:
    - Change the public IP address of an instance
    - Convert the public IP address of a VPC-type instance to an EIP
- If the operating system of your instance cannot meet your requirements, you can replace the operating system. For more information, see Change the operating system.
- You can use the following features to manage ECS instances in a fine-grained manner:
  - User data
  - Instance metadata
  - Instance identity
  - Instance RAM roles

Manage the billing of instances

- Subscription instances:

  You can use one of the following methods to renew subscription instances:

  - Manually renew an instance
  - Enable auto-renewal for an instance
  - Downgrade the configurations of an instance during renewal

- Pay-as-you-go instances:

  You can enable the economical mode for pay-as-you-go instances. For more information, see Economical mode.

- Change the billing methods of instances:
  - Change the billing method of an ECS instance from pay-as-you-go to subscription
  - Change the billing method of an instance from subscription to pay-as-you-go

Improve cost-effectiveness

- You can purchase preemptible instances to reduce costs and use preemptible instances in conjunction with Auto Provisioning for automated provisioning of instances. For more information, see Create an auto provisioning group and Create a preemptible instance.
- You can purchase reserved instances to improve the flexibility of paying for instances and reduce costs. For more information, see Purchase reserved instances.

Create and manage disks

If you want to use a disk as a data disk, you can perform the following steps:

1. Create a disk.
2. Attach a data disk.
3. Partition and format a data disk on a Linux instance or Partition and format a data disk on a Windows instance.
4. Create a snapshot to back up data. For more information, see Create a snapshot of a disk.
5. If the storage capacity of an existing disk cannot meet your requirements, resize the disk. For more information, see the following topics:
   - Resize disks online for Linux instances
   - Resize disks offline for Linux instances
   - Resize disks online for Windows instances
   - Resize disks offline for Windows instances
6. If a data error occurs on a disk, use a snapshot created at a specific point in time to roll back the disk. For more information, see Roll back a disk by using a snapshot.
7. If you want to restore a disk to its initial state, re-initialize the disk. For more information, see Re-initialize a data disk.
8. Detach a data disk.
9. Release a disk.

Create and manage snapshots

You can perform the following steps to use a snapshot:

1. Create a snapshot. You can use one of the following methods to manually or automatically create a snapshot:
   - Create a snapshot of a disk.
   - Use an automatic snapshot policy to automatically create snapshots on a regular basis. For more information, see Apply or disable an automatic snapshot policy.
2. 查看快照容量.
3. Delete snapshots that are no longer needed to save storage space. For more information, see Reduce snapshot fees.

The following section describes the use scenarios of snapshots:

- To copy or restore data, you can use a snapshot to create or roll back a disk. For more information, see Create a disk from a snapshot and Roll back a disk by using a snapshot.
- To deploy an environment, you can use a system disk snapshot to create a custom image and then use the custom image to create instances. For more information, see Create a custom image from a snapshot and Create an ECS instance by using a custom image.

Create and manage custom images

Only custom images can be managed in the ECS console. You can use a custom image to deploy a business environment in a quick manner. You can use one of the following methods to obtain a custom image:

- Create a custom image from a snapshot.
- Create a custom image from an instance.
- Use Packer to create a custom image.
- Copy custom images across regions. For more information, see Copy custom images.
- Share custom images across accounts. For more information, see the "Share a custom image" section in Share or unshare custom images.
- Import custom images.
- Create and import an on-premises image by using Packer.

You can export custom images to back up environments. For more information, see Export a custom image

Create and manage security groups

You can perform the following steps to create and manage a security group.

1. Create a security group.
2. Add a security group rule.
3. Add an ECS instance to a security group.
4. Delete a security group rule.
5. Delete a security group.

You can clone a security group across regions and network types to simplify business deployment. For more information, see Clone a security group.

If new security group rules disrupt your online business, you can restore all or some of the security group rules. For more information, see Restore security group rules.

Create and attach instance RAM roles

You can perform the following steps to create and attach an instance RAM role.

1. (Optional) Authorize a RAM user to manage an instance RAM role. For more information, see Authorize a RAM user to manage an instance RAM role.
2. Create and attach an instance RAM role. For more information, see Attach an instance RAM role to an ECS instance.
3. Replace the instance RAM role based on your requirements. For more information, see Replace an instance RAM role.

Create and manage SSH key pairs

You can perform the following steps to create and manage an SSH key pair:

1. Create an SSH key pair or Import an SSH key pair.
2. Bind an SSH key pair to an instance.
3. Connect to a Linux instance by using an SSH key pair.
4. Unbind an SSH key pair.
5. Delete an SSH key pair.

Create and manage ENIs

You can perform the following steps to create and manage an elastic network interface (ENI).

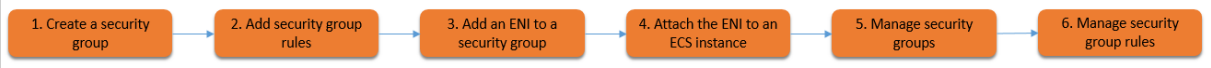

1. Create an ENI.
2. Bind an ENI to an instance or Bind an ENI when you create an instance.
3. (Optional) Configure a secondary ENI.
4. Assign secondary private IP addresses.
5. Unbind an ENI.
6. Delete an ENI.

Use tags

You can use tags to manage resources to enhance efficiency. You can perform the following steps to use a tag:

1. Add a tag.
2. Search for resources by tag.
3. Delete or unbind a tag.

Create and manage launch templates

Launch templates help you create ECS instances that have the identical configurations. You can perform the following steps to use a launch template:

1. Create a launch template.

2. Create a launch template version.
3. Delete a launch template and a template version.

Create and manage deployment sets

Deployment sets help you implement high availability for underlying applications. You can perform the following steps to use a deployment set:

1. Create a deployment set.
2. Create an ECS instance in a deployment set.
3. Change the deployment set of an instance.
4. Delete a deployment set.

Use Cloud Assistant

Cloud Assistant allows you to send remote commands to ECS instances without the need to use jump servers. You can perform the following steps to use Cloud Assistant:

1. (Optional) Manually install and configure the Cloud Assistant client on some ECS instances. For more information, see Install the Cloud Assistant client.
2. Create a command.
3. Run a command.
4. Query execution results and fix common problems.

2.Create an FPGA-accelerated instance

f3 instances use Xilinx FPGAs, whereas f1 instances use Intel FPGAs. When you create an FPGA-accelerated instance, you can use an image that is pre-installed with the FPGA development environment. This topic describes the parameters of which you must take note when you create FPGA-accelerated instances.

Prerequisites

The following preparations are made to create an Elastic Compute Service (ECS) instance:

1. Create an account and complete the account information.
   - Create an Alibaba Cloud account. For more information, see Sign up with Alibaba Cloud.
   - To purchase ECS instances within mainland China regions, complete real-name verification. For more information, see Real-name verification FAQ.
2. Alibaba Cloud provides a default virtual private cloud (VPC) in each region. If you do not want to use the default VPC, you can create a VPC and a vSwitch within the region in which to create the instance. For more information, see Create an IPv4 VPC.
3. Alibaba Cloud provides a default security group in each region. If you do not want to use the default security group, you can create a security group in the region in which to create the instance. For more information, see Create a security group.

If you need other extended features, you must complete the corresponding preparations:

- To bind an SSH key pair when you create a Linux instance, you must create the SSH key pair in the specified region. For more information, see Create an SSH key pair.
- To add user data for the instance, you must first prepare user data. For more information, see Overview of ECS instance user data.
- To associate an ECS instance with a Resource Access Management (RAM) role, you must create the RAM role, attach permission policies to the role, and then bind the role to the instance. For more information, see Attach an instance RAM role to an ECS instance.

f3 instances must use an image that provides the Xilinx development environment. The image can only be shared to users by the FPGA as a service (FaaS) team. To obtain the image, submit a ticket.

f1 instances must use an image that provides the Intel development environment. The image can only be shared to users by the FaaS team. To obtain the image, submit a ticket.

Procedure

This section describes the parameters of which you must take note when you create FPGA-accelerated instances by using the ECS console. For information about other parameters, see Create an instance by using the wizard.

In addition to using the ECS console, you can use a variety of other methods to create FPGA-accelerated instances. For more information, see Provisioning methods of ECS instances.

1. Go to the Custom Launch tab of the instance buy page in the ECS console.
2. Complete the settings in the Basic Configurations step and click **Next**.

   The following table describes the parameters that you must configure in the Basic Configurations

step.

Parameter	Description
Region	FPGA-accelerated instance types are available only within specific regions and zones. For more information, visit the ECS Instance Types Available for Each Region page. Set Billing Method and enter an instance type name to search for the instance type. **Note** The types of resources that you can purchase are determined by your instance quotas. For more information, see View and increase instance quotas.
Instance Type	○ Set **Architecture** to **Heterogeneous Computing**. ○ Set **Category** to **Compute Optimized Type with FPGA**.
Image	○ To create f3 instances, select **Shared Image**. and then select the image that was shared to you and pre-installed with the Xilinx development environment. ○ To create f1 instances, select **Shared Image** and then select the image that was shared to you and pre-installed with the Intel development environment. **Note** The images pre-installed with the Intel development environment are available only as shared images. These images are pre-installed with the development environments of Quartus 17.0, VCS 2017.3, and DCP SDK. You can view files of these development environments in the *opt* directory.
Storage	To create f3 instances, we recommend that you allocate a 200 GiB ultra disk as the system disk.

3. Complete the settings in the Networking step and click **Next**.

   FPGA-accelerated instances can be deployed only in VPCs.
4. Complete the settings in the System Configurations step and click **Next**.
5. Complete the settings in the Grouping step and click **Next**.
6. Check your configurations, read and select *ECS Terms of Service and Product Terms of Service*, and then click Create Order or Create Instance.

What's next

After an f1 instance is created, you can perform the following steps to check whether the license is configured:

1. Connect to the instance.
2. Run the following command to check whether the license is configured:

```
echo $LM_LICENSE_FILE
```

If the license is configured, the value of the LM_LICENSE_FILE variable is displayed. If the license is not configured, no value is displayed.

faascmd is required to use resources on the FaaS platform. For more information about faascmd, see Overview.

Related information

References

- RunInstances
- Use RTL Compiler on an f1 instance
- Use OpenCL on an f1 instance
- Use the RTL design on an f3 instance
- Use OpenCL on an f3 instance

3.Connect to an FPGA-based ECS instance

3.1. Overview

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are connected to in the same way you connect to other ECS instances. Alibaba Cloud provides a variety of methods to connect to an ECS instance, including VNC, Workbench and third-party client tools. You can choose a method to connect to your instance based on the operating system of your instance, the operating system of your device, and the operations that you want to perform.

Connection methods

Operating system of your instance	Operating system of your device	Connection method
Linux	Windows	• Workbench For information about how to connect to an instance by using a password or a key as the credential, see Connect to a Linux instance by using a password or key. • VNC For more information, see Connect to a Linux instance by using a password. • Client tools such as PuTTY ◦ For information about how to connect to an instance by using an SSH key pair as the credential, see Use an SSH key pair to connect to a Linux instance from a Windows device. ◦ For information about how to connect to an instance by using a username and password as the credential, see Use a username and password to connect to a Linux instance from a Windows device.

Operating system of your instance	Operating system of your device	Connection method
	UNIX-like operating systems such as Linux and macOS	• Workbench For information about how to connect to an instance by using a password or a key as the credential, see Connect to a Linux instance by using a password or key. • VNC For more information, see Connect to a Linux instance by using a password. • SSH commands ◦ For information about how to connect to an instance by using an SSH key pair as the credential, see Use an SSH key pair to connect to a Linux instance from a device that supports SSH commands (configure information by using commands). ◦ For information about how to connect to an instance by using a username and password as the credential, see Use a username and password to connect to a Linux instance from a Linux or Mac OS X device.
	Operating systems of mobile devices, such as iOS and Android	Apps such as SSH Control Lite and JuiceSSH For more information, see Connect to a Linux instance from a mobile device.
	Windows	• Workbench For information about how to connect to an instance by using a password or a key as the credential, see Connect to a Windows instance by using a password or key. • VNC For more information, see Connect to a Windows instance by using a password. • Client tools such as Remote Desktop Connection (formerly called MSTSC) For more information, see Connect from a local client that runs a Windows operating system.

Operating system of your instance	Operating system of your device	Connection method
Windows	Linux	• Workbench For information about how to connect to an instance by using a password or a key as the credential, see Connect to a Windows instance by using a password or key. • VNC For more information, see Connect to a Windows instance by using a password. • Client tools such as rdesktop For more information, see Connect from a local client that runs a Linux operating system.
	macOS	• Workbench For information about how to connect to an instance by using a password or a key as the credential, see Connect to a Windows instance by using a password or key. • VNC For more information, see Connect to a Windows instance by using a password. • Client tools such as Microsoft Remote Desktop Connection for Mac For more information, see Get started with the macOS client.
	Operating systems of mobile devices, such as iOS and Android	Apps such as Microsoft Remote Desktop For more information, see Connect to a Windows instance from a mobile device.

Note

- Except for Workbench and VNC, all connection tools require that instances that you want to connect have public IP addresses or elastic IP addresses (EIPs).
- After a Windows instance is created, it takes 2 to 3 minutes to initialize the operating system. Do not restart the instance while it is being initialized. After a non-I/O optimized Windows instance is created, it takes 10 minutes to initialize the operating system. Do not connect to the instance while it is being initialized.

Comparison of connection tools

The following table compares the advantages of VNC, Workbench, and other third-party client tools.

Item	Workbench	VNC	Third-party client tool
Assignment of a public IP address or an EIP to the instance	Optional. **Note** Workbench cannot be used to troubleshoot network configuration exceptions, such as firewalls being enabled by mistake.	Optional. VNC can be used to troubleshoot network configuration exceptions, such as firewalls being enabled by mistake.	Required.
Enabling services such as SSH on the instance	Required.	Optional. VNC can be used to troubleshoot SSH service exceptions, such as SSHD being disabled.	Required.
Logons by using the ECS console	Supported.	Supported.	Not supported. The local client must be installed.
Independence of the instance operating system	Workbench can be used to connect to both Linux and Windows instances.	VNC can be used to connect to both Linux and Windows instances.	Depends on the client tool. The third-party client tools can be used to connect to Linux or Windows instances.
Simultaneous logons by multiple operating system users to a single instance	Supported.	Not supported.	Depends on the client tool.
Ease of interaction	Workbench supports copying and pasting text.	VNC does support copying and pasting text. To copy or paste text, use the feature for copying long commands.	Depends on the client tool.
Visibility into Linux system file resources	Supported.	Not supported.	Depends on the client tool.
Permissions to control and modify hardware	Not supported.	Supported. VNC can be used to manage resources such as BIOS and troubleshoot exceptions such as system startup failures.	Not supported.

Item	Workbench	VNC	Third-party client tool
Terminal configurability	Supported, but depends on the capabilities that Workbench provides.	Not supported.	Supported, but depends on the capabilities that the client tool provides.

3.2. Use Workbench to connect to an instance

3.2.1. Connect to a Linux instance by using a password or key

Workbench allows multiple users to connect to a single Elastic Compute Service (ECS) instance at the same time and provides a GUI for users to manage files in Linux instances. Workbench is more efficient and convenient than Virtual Network Console (VNC).

Prerequisites

- A logon password is set for or a key pair is bound to the Linux instance to which you want to connect.
- The instance is in the **Running** state.
- Security group rules are added to allow the IP addresses related to the Workbench service to access the instance. For more information about the security group rules, see the Add security group rules to allow Workbench access to a Linux instance section.

Context

By default, a Workbench remote session persists for 6 hours. If you do not perform operations for 6 hours, the remote connection is closed. You must reconnect to the instance.

Workbench can be used to connect to ECS instances over one of the following protocols:

- SSH: By default, Linux instances are connected by using SSH. SSH can also be used to connect to Windows instances on which a GNU-like system such as Cygwin is installed. For information about how to connect to a Linux instance over SSH, see the Connect to a Linux instance over SSH section.
- Remote Desktop Protocol (RDP): By default, Windows instances are connected by using RDP. RDP can also be used to connect to Linux instances on which remote desktop services are enabled. For information about how to connect to a Linux instance over RDP, see the Connect to a Linux instance over RDP section.

  > **Note** If you want to connect to an instance over RDP, make sure that the public bandwidth is at least 5 Mbit/s. If the public bandwidth is less than 5 Mbit/s, the remote desktop freezes.

You can use the GUI provided by Workbench to manage files in your Linux instances in a visual manner. For more information, see Use Workbench to manage files in a Linux instance.

Connect to a Linux instance over SSH

1.
2.
3.
4. On the **Instances** page, find the instance to which you want to connect, and click **Connect** in the **Actions** column.
5. In the **Connection and Command** dialog box, click **Connect** in the **Workbench Connection** section.
6. In the **Instance Login** dialog box, specify parameters.

   The following table describes the required parameters in the dialog box.

Parameter	Description
Instance	The information of the current instance is automatically populated. You can also manually enter the IP address or name of another instance.
Connection	◦ To connect to instances that are located in VPCs, you can use their public or private IP addresses. ◦ To connect to instances that are located in the classic network, you can use their public or internal IP addresses.
Username, **Password**, and **Private Key**	Enter a username such as root and select an authentication method. The following authentication methods are supported: ◦ **Password-based**: Enter the password of your specified username. ◦ **Certificate-based**: Enter or upload a certificate. If the certificate is encrypted, enter its key passphrase.

   In the lower part of the dialog box, click **More Options** to show the optional parameters described in the following table.

Parameter	Description
Resource Group	By default, **All** is selected. You can manually select a resource group from the drop-down list.
Region	By default, **All** is selected. You can manually select a region from the drop-down list.
Protocol	By default, **Terminal Connection (SSH)** is selected.
Port	When Protocol is set to **Terminal Connection (SSH)**, this parameter is automatically set to 22.
Language	Select your preferred language. The selected language affects the outputs of the instance. We recommend that you select **Default** for Workbench to detect the language settings of the instance and to make configurations accordingly.

Parameter	Description
Character Set	Select your preferred character set. The selected character set affects the outputs of the instance. We recommend that you select **Default** for Workbench to detect the character set settings of the instance and to make configurations accordingly.

7. Click **OK**.

If all of the requirements specified in the prerequisites are met but the instance cannot be connected, perform the following checks on the instance:

- Check whether the sshd service (such as sshd in Linux) is enabled. If not, enable the sshd service.
- Check whether the required terminal connection port (typically port 22) is enabled. If not, enable the port.
- If you log on to the Linux instance as the root user, make sure that `PermitRootLogin yes` is configured in the */etc/ssh/sshd_config* file. For more information, see the Enable root logon over SSH on a Linux instance section.

Connect to a Linux instance over RDP

1.
2.
3.
4. On the **Instances** page, find the instance to which you want to connect, and click **Connect** in the **Actions** column.
5. In the **Connection and Command** dialog box, click **Connect** in the **Workbench Connection** section.
6. In the **Instance Login** dialog box, specify parameters.
   i. In the lower part of the dialog box, click **More Options**.
   ii. Set **Protocol** to **Remote Desktop (RDP)**.
   iii. In the message that appears, click **OK**.

iv. Specify the parameters described in the following table.

Parameter	Description
Resource Group	By default, **All** is selected. You can manually select a resource group from the drop-down list.
Region	By default, **All** is selected. You can manually select a region from the drop-down list.
Instance	The information of the current instance is automatically populated. You can also manually enter the IP address or name of another instance.
Connection	■ To connect to instances that are located in VPCs, you can use their public or private IP addresses. ■ To connect to instances that are located in the classic network, you can use their public or internal IP addresses.
Port	When Protocol is set to **Remote Desktop (RDP)**, this parameter is automatically set to 3389.
Username and **Password**	Enter a username, such as Administrator, and its password.

7. Click **OK**.

If all of the requirements specified in the prerequisites are met but the instance cannot be connected, perform the following checks on the instance:

- Check whether a remote desktop service (such as xfreerdp installed on Linux) is enabled. If not, enable a remote desktop service.
- Check whether the required remote desktop port (typically port 3389) is enabled. If not, enable the port.
- If you log on to the Linux instance as the root user, make sure that `PermitRootLogin yes` is configured in the */etc/ssh/sshd_config* file. For more information, see the Enable root logon over SSH on a Linux instance section.

Enable root logon over SSH on a Linux instance

In some Linux systems, sshd disables root logon by default. If this occurs, when you attempt to connect to an instance as the root user over SSH, you are prompted that your username or password is incorrect. To enable root logon over SSH, perform the following operations.

1. Connect to a Linux instance by using a password with VNC
2. Open the SSH configuration file.

```
vi /etc/ssh/sshd_config
```

3. Change `PermitRootLogin no` to `PermitRootLogin yes`.
4. Press the Esc key and enter **:wq** to save the change.
5. Restart sshd.

```
service sshd restart
```

Add security group rules to allow Workbench access to a Linux instance

This section describes how to add rules to security groups of different network types in the ECS console to allow Workbench access to a Linux instance.

- If you want to connect to a Linux instance in a VPC, find a security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Inbound** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
N/A	**Inbound**	**Allow**	◦ If port 22 is enabled by default on the Linux instance, select **SSH (22)**. ◦ If you have manually enabled other ports on the Linux instance, select **Custom TCP**.	◦ If port 22 is enabled by default on the Linux instance, **22/22** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Linux instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	◦ If you want to connect to the instance by using its public IP address, specify 161.117.90.22/24. The public IP address can be the public IP address that is automatically assigned to the instance or an elastic IP address (EIP) that is associated with the instance. ◦ If you want to connect to the instance by using its private IP address, specify 100.104.0.0/16. **Note** You can also specify 0.0.0.0/0 as the authorization object to allow inbound access from all IP addresses. However, this imposes security risks. Proceed with caution.

- If you want to connect to a Linux instance in the classic network over the Internet, find a security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Internet Ingress** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
Public	**Inbound**	**Allow**	◦ If port 22 is enabled by default on the Linux instance, select **SSH (22)**. ◦ If you have manually enabled other ports on the Linux instance, select **Custom TCP**.	◦ If port 22 is enabled by default on the Linux instance, **22/22** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Linux instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	If you want to connect to the instance by using its public IP address, specify 161.117.90.22/24. The public IP address can be the public IP address that is automatically assigned to the instance or an EIP that is associated with the instance. **Note** You can also specify 0.0.0.0/0 as the authorization object to allow inbound access from all IP addresses. However, this imposes security risks. Proceed with caution.

- If you want to connect to a Linux instance in the classic network over the internal network, security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Internal Network Ingress** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
N/A	**Inbound**	**Allow**	◦ If port 22 is enabled by default on the Linux instance, select **SSH (22)**. ◦ If you have manually enabled other ports on the Linux instance, select **Custom TCP**.	◦ If port 22 is enabled by default on the Linux instance, **22/22** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Linux instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	If you want to connect to the instance by using its internal IP address, specify 11.195.184.0/24 and 11.246.55.0/24. **Notice** High security risks may arise if you specify 0.0.0.0/0 as the authorization object. We recommend that you do not specify 0.0.0.0/0.

3.2.2. Connect to a Windows instance by using a password or key

Workbench allows multiple users to connect to a single Elastic Compute Service (ECS) instance at the same time. Workbench is more efficient and convenient than Virtual Network Console (VNC).

Prerequisites

- A logon password or a key is configured for the Windows instance to which you want to connect.

  Note The ECS console cannot be used to bind key pairs to Windows instances. If you want to use a key to log on to a Windows instance, you can enable the sshd service (such as Cygwin SSHD or WinSSHD in Windows) and configure a key on the instance. For more information about how to enable the sshd service in Windows, see Get started with OpenSSH.

- The instance is in the **Running** state.
- Security group rules are added to allow the IP addresses related to the Workbench service to access the instance. For more information, see Add security group rules to allow Workbench access to a Windows instance.

Context

By default, a Workbench remote session persists for 6 hours. If you do not perform operations for 6 hours, the remote connection is closed. You must reconnect to the instance.

Workbench can be used to connect to ECS instances over one of the following protocols:

- Remote Desktop Protocol (RDP): By default, Windows instances are connected by using RDP. RDP can also be used to connect to Linux instances on which remote desktop services are enabled. For information about how to connect to a Windows instance over RDP, see the Connect to a Windows instance over RDP section.

  > **Note** If you want to connect to an instance over RDP, make sure that the public bandwidth is at least 5 Mbit/s. If the public bandwidth is less than 5 Mbit/s, the remote desktop freezes.

- SSH: By default, Linux instances are connected by using SSH. SSH can also be used to connect to Windows instances on which a GNU-like system such as Cygwin is installed. For information about how to connect to a Windows instance over RDP, see the Connect to a Windows instance over SSH section.

Connect to a Windows instance over RDP

1.
2.
3.
4. On the **Instances** page, find the instance to which you want to connect, and click **Connect** in the **Actions** column.
5. In the **Connection and Command** dialog box, click **Connect** in the **Workbench Connection** section.
6. In the **Instance Login** dialog box, specify parameters.

   The following table describes the required parameters in the dialog box.

Parameter	Description
Instance	The information of the current instance is automatically populated. You can also manually enter the IP address or name of another instance.
Connection	○ To connect to instances in virtual private clouds (VPCs), you can use the public or private IP addresses of the instances. ○ To connect to instances in the classic network, you can use their public or internal IP addresses.
Username and **Password**	Enter a username, such as Administrator, and its password.

   In the lower part of the dialog box, click **More Options** to show the optional parameters described in the following table.

Parameter	Description
Resource Group	By default, **All** is selected. You can manually select a resource group from the drop-down list.
Region	By default, **All** is selected. You can manually select a region from the drop-down list.
Protocol	By default, **Remote Desktop (RDP)** is selected.
Port	When Protocol is set to **Remote Desktop (RDP)**, this parameter is automatically set to 3389.

7. Click **OK**.

If all of the requirements specified in the prerequisites are met but the instance cannot be connected, perform the following checks on the instance:

- Check whether a remote desktop service (such as Remote Desktop Services in Windows) is enabled. If not, enable a remote desktop service.
- Check whether the required remote desktop port (typically port 3389) is enabled. If not, enable the port.
- If you log on to the Windows instance as a non-administrator user, the user must belong to the Remote Desktop Users group.

Connect to a Windows instance over SSH

1.
2.
3.
4. On the **Instances** page, find the instance to which you want to connect, and click **Connect** in the **Actions** column.
5. In the **Connection and Command** dialog box, click **Connect** in the **Workbench Connection** section.
6. In the **Instance Login** dialog box, specify parameters.
   i. In the lower part of the dialog box, click **More Options**
   ii. Set **Protocol** to **Terminal Connection (SSH)**.
   iii. In the Confirm message, click **OK**.

iv. Specify the parameters described in the following table.

Parameter	Description
Resource Group	By default, **All** is selected. You can manually select a resource group from the drop-down list.
Region	By default, **All** is selected. You can manually select a region from the drop-down list.
Instance	The information of the current instance is automatically populated. You can also manually enter the IP address or name of another instance.
Connection	■ To connect to instances in VPCs, you can use the public or private IP addresses of the instances. ■ To connect to instances in the classic network, you can use their public or internal IP addresses.
Port	When Protocol is set to **Terminal Connection (SSH)**, this parameter is automatically set to 22.
Username, **Password**, and **Private Key**	Enter a username such as root and select an authentication method. The following authentication methods are supported: ■ **Password-based**: Enter the password of your specified username. ■ **Certificate-based**: Enter or upload a certificate. If the certificate is encrypted, enter its key passphrase.
Language	Select your preferred language. The selected language affects the outputs of the instance. We recommend that you select **Default** for Workbench to detect the language settings of the instance and make configurations accordingly.
Character Set	Select your preferred character set. The selected character set affects the outputs of the instance. We recommend that you select **Default** for Workbench to detect the character set settings of the instance and make configurations accordingly.

7. Click **OK**.

If all of the requirements specified in the prerequisites are met but the instance cannot be connected, perform the following checks on the instance:

- Check whether the sshd service (such as Cygwin SSHD or WinSSHD in Windows) is enabled. If not, enable the sshd service.
- Check whether the required terminal connection port (typically port 22) is enabled. If not, enable the port.
- If you log on to the Windows instance as a non-administrator user, the user must belong to the Remote Desktop Users group.

Add security group rules to allow Workbench access to a Windows instance

This section describes how to add rules to security groups of different network types in the ECS console to allow Workbench access to a Windows instance.

- If you want to connect to a Windows instance in a VPC, find a security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Inbound** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
N/A	**Inbound**	**Allow**	◦ If port 3389 is enabled by default on the Windows instance, select **RDP (3389)**. ◦ If you have manually enabled other ports on the Windows instance, select **Custom TCP**.	◦ If port 3389 is enabled by default on the Windows instance, **3389/3389** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Windows instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	◦ If you want to connect to the instance by using its public IP address, specify 161.117.90.22. The public IP address can be the public IP address that is automatically assigned to the instance or an elastic IP address (EIP) that is associated with the instance. ◦ If you want to connect to the instance by using its private IP address, specify 100.104.0.0/16. **Note** You can also specify 0.0.0.0/0 as the authorization object to allow inbound access from all IP addresses. However, this imposes security risks. Proceed with caution.

- If you want to connect to a Windows instance in the classic network over the Internet, find a security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Internet Ingress** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
Public	**Inbound**	**Allow**	◦ If port 3389 is enabled by default on the Windows instance, select **RDP (3389)**. ◦ If you have manually enabled other ports on the Windows instance, select **Custom TCP**.	◦ If port 3389 is enabled by default on the Windows instance, **3389/3389** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Windows instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	If you want to connect to the instance by using its public IP address, specify 161.117.90.22. The public IP address can be the public IP address that is automatically assigned to the instance or an EIP that is associated with the instance. **Note** You can also specify 0.0.0.0/0 as the authorization object to allow inbound access from all IP addresses. However, this imposes security risks. Proceed with caution.

- If you want to connect to a Windows instance in the classic network over the internal network, find a security group of the instance, go to the **Security Group Rules** page, and then add a rule on the **Internal Network Ingress** tab. The following table describes the parameters to be configured for the rule.

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object

NIC Type	Rule Direction	Action	Protocol Type	Port Range	Priority	Authorization Type	Authorization Object
N/A	**Inbound**	**Allow**	◦ If port 3389 is enabled by default on the Windows instance, select **RDP (3389)**. ◦ If you have manually enabled other ports on the Windows instance, select **Custom TCP**.	◦ If port 3389 is enabled by default on the Windows instance, **3389/3389** is automatically entered after you select the protocol type. ◦ If you have manually enabled other ports on the Windows instance, enter a corresponding port range.	1	**IPv4 CIDR Block**	If you want to connect to the instance by using its internal IP address, specify 161.117.90.22. **Notice** High security risks may arise if you specify 0.0.0.0/0 as the authorization object. We recommend that you do not specify 0.0.0.0/0.

3.3. Use VNC to connect to an instance

3.3.1. Connect to a Linux instance by using a password

If you cannot use Workbench or connection software such as PuTTY, Xshell, and SecureCRT to connect to an Elastic Compute Service (ECS) Linux instance, you can use the **VNC Connection** feature in the ECS console to connect to the Linux instance and view the real-time status of the instance operation interface.

Prerequisites

A logon password is set for the instance.

Note If you have not set a password or forget the password, you can reset the password for the instance. For more information, see Reset the logon password of an instance.

Context

The following passwords are involved when you use VNC to connect to an instance:

- VNC password: the password of management terminals used to connect to the ECS console.
- Instance logon password: the password used to log on to the instance operating system.

By default, a VNC connection session lasts for about 300 seconds. If you do not perform operations within these 300 seconds, the connection to the instance is automatically closed. You must connect to the instance again.

If you cannot use Workbench or connection software to connect to your instance, you can use the **VNC Connection** feature in the ECS console to connect to the instance. After the instance is connected, you can view the status of the instance and perform operations to resolve issues described in the following table.

Scenario	Solution
The instance starts slowly due to self-check on startup.	Check the self-check progress.
The firewall of the instance operating system is enabled by mistake.	Disable the firewall.
The ECS instance is compromised, which causes a high CPU utilization and high bandwidth usage.	Troubleshoot and terminate abnormal processes.

Procedure

The following figure shows how to use VNC to connect to an instance.

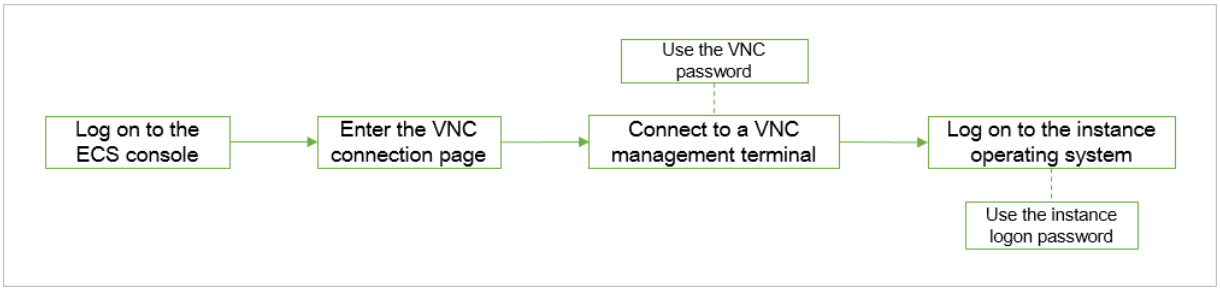

1.
2.
3.
4. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
5. In the **Connection and Command** dialog box, click **Connect** in the **VNC Connection** section.
6. Connect to a VNC management terminal.

   > **Note** In this step, use the VNC password.

   - The first time you connect to a VNC management terminal, perform the following operations:
     a. Change the VNC password. For more information, see the Change the VNC password section in this topic.
     b. In the **Enter VNC Password** dialog box, enter the new password.
     c. Click **OK**.

- If you are not connecting to a VNC management terminal for the first time, perform the following operations:
  a. In the **Enter VNC Password** dialog box, enter the password.
  b. Click **OK**.

7. Log on to the instance operating system.

   > **Note** In this step, use the instance logon password.

   i. Enter the username *root* and press the Enter key.
   ii. Enter the logon password of the instance and press the Enter key.

   > **Note** The characters of the password are hidden when you enter the password. After you enter the password, press the Enter key.

   You can switch between up to 10 different VNC management terminals when you connect to the Linux instance. The default terminal is **CTRL+ALT+F1**. For example, you can choose **Send Remote Call > CTRL+ALT+F2** to switch to **CTRL+ALT+F2**. A persistent black screen indicates that the instance is in sleep mode. Press a key to wake up the instance.

Change the VNC password

The first time you connect to the VNC management terminal, you must change the VNC password. You can also change the VNC password when you forget the password or when you want to update the password.

> **Notice** After you change the VNC password for a non-I/O optimized instance, you must restart the instance in the ECS console for the new password to take effect. Before you restart the instance, you must stop it. This can lead to service interruption. Proceed with caution.

1. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
2. In the **Connection and Command** dialog box, click **Connect** in the **VNC Connection** section.
3. In the **Enter VNC Password** dialog box, click **Reset VNC Password**.
4. In the **Reset VNC Password** dialog box, enter and confirm the new password, and then click **OK**.
5. (Optional) If the instance is a non-I/O optimized instance, restart the instance.

   For more information, see Restart an instance.

Copy long commands

If you want to copy a long-text item such as a download URL from your computer to the instance, you can use the command copy feature.

1. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
2. Connect to a VNC management terminal.
3. In the upper-left corner of the interface, click **Enter Copy Commands**.
4. In the **Copy and Paste Commands** dialog box, enter the content to be copied and click **OK**.

3.3.2. Connect to a Windows instance by using a password

If you cannot use Workbench or connection software such as Remote Desktop Connection (RDC) and rdesktop to connect to an Elastic Compute Service (ECS) Windows instance, you can use the **VNC Connection** feature in the ECS console to connect to the instance and view the real-time status of the instance operating interface.

Prerequisites

A logon password is set for the instance.

> **Note** If you have not set a password or forget the password, you can reset the password for the instance. For more information, see Reset the logon password of an instance.

Context

The following passwords are involved when you use VNC to connect to an instance:

- VNC password: the password of management terminals used to connect to the ECS console.
- Instance logon password: the password used to log on to the instance operating system.

By default, a VNC connection session lasts for about 300 seconds. If you do not perform operations within these 300 seconds, the connection to the instance is automatically closed. You must connect to the instance again.

If you cannot use Workbench or connection software to connect to your instance, you can use the **VNC Connection** feature in the ECS console to connect to the instance. After the instance is connected, you can view the status of the instance and perform operations to resolve issues described in the following table.

Scenario	Solution
The instance starts slowly due to self-check on startup.	Check the self-check progress.
The firewall of the instance operating system is enabled by mistake.	Disable the firewall.
The ECS instance is compromised, which causes a high CPU utilization and high bandwidth usage.	Troubleshoot and terminate abnormal processes.

Procedure

The following figure shows how to use VNC to connect to an instance.

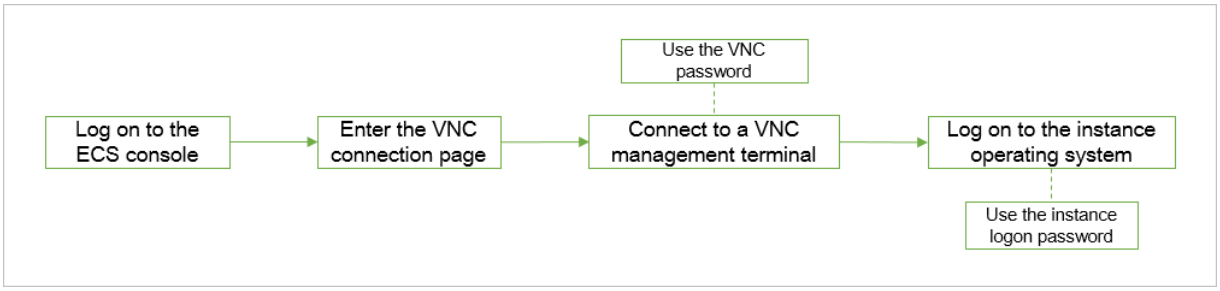

1.
2.
3.
4. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
5. Connect to a VNC management terminal.

   > **Note** In this step, use the VNC password.

   - The first time you connect to a VNC management terminal, perform the following operations:
     a. Change the VNC password. For more information, see the Change the VNC password section in this topic.
     b. In the **Enter VNC Password** dialog box, enter the new password.
     c. Click **OK**.
   - If you are not connecting to a VNC management terminal for the first time, perform the following operations:
     a. In the **Enter VNC Password** dialog box, enter the password.
     b. Click **OK**.
6. In the upper-left corner of the **VNC** page, choose **Send Remote Call > CTRL+ALT+DELETE**.

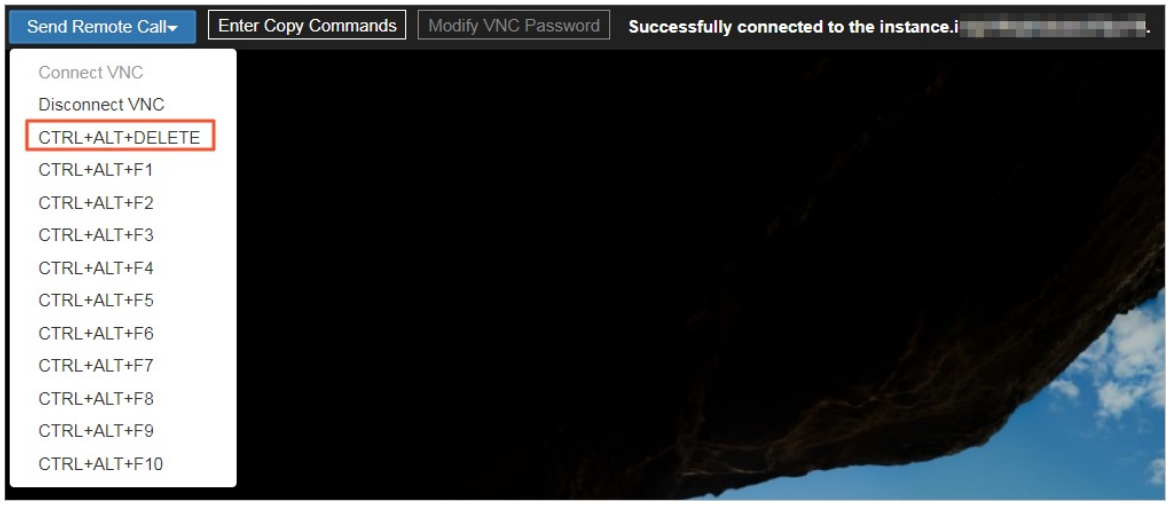

7. Select an account, enter the instance password, and then press the Enter key.

   By default, the Administrator account is available.

Change the VNC password

The first time you connect to the VNC management terminal, you must change the VNC password. You can also change the VNC password when you forget the password or when you want to update the password.

> **Notice** After you change the VNC password for a non-I/O optimized instance, you must restart the instance in the ECS console for the new password to take effect. Before you restart the instance, you must stop it. This can lead to service interruption. Proceed with caution.

1. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
2. In the **Connection and Command** dialog box, click **Connect** in the **VNC Connection** section.
3. In the **Enter VNC Password** dialog box, click **Reset VNC Password**.
4. In the **Reset VNC Password** dialog box, enter and confirm the new password, and then click **OK**.
5. (Optional) If the instance is a non-I/O optimized instance, restart the instance.

   For more information, see Restart an instance.

Copy long commands

If you want to copy a long-text item such as a download URL from your computer to the instance, you can use the command copy feature.

1. On the **Instances** page, find the instance to be connected and click **Connect** in the **Actions** column.
2. Connect to a VNC management terminal.
3. In the upper-left corner of the interface, click **Enter Copy Commands**.
4. In the **Copy and Paste Commands** dialog box, enter the content to be copied and click **OK**.

FAQ

For more information about how to adjust the resolution of the Windows desktop, see How do I adjust the desktop resolution of a Windows instance?.

4.faasutil

4.1. Obtain faasutil

faasutil is a next-generation command line tool provided by Alibaba Cloud FPGA as a Service (FaaS) and can make FPGA-accelerated instances easier to use, and more stable, secure, and scalable. You can use faasutil to run simple commands to configure environments, and generate and load FPGA images without caring about the underlying implementation. This topic describes how to obtain faasutil.

Prerequisites

An FPGA-accelerated instance is created and you are connected to the instance.

Obtain faasutil

1. Make sure that the */dev/virtio-ports/FaaS.agent* file exists in the FPGA-accelerated instance.

   If the file does not exist, the FPGA-accelerated instance does not support faasutil. Use faascmd instead. For more information, see Overview.

2. Download faasutil.

   ```
   wget https://fpga-tools.oss-cn-shanghai.aliyuncs.com/faascmd_bin/v2_0/faasutil
   ```

3. Grant the execute permissions on faasutil.

   ```
   chmod +x faasutil
   ```

4.2. Use faasutil

This topic describes how to use faasutil commands. This topic also provides examples for these commands.

Prerequisites

- faasutil is obtained. For more information, see Obtain faasutil.
- A raw file is prepared. For more information, see the term description in the Context section in this topic.
- An Object Storage Service (OSS) bucket is created. The bucket is used to upload raw files. For more information, see Create buckets.
- An AccessKey pair is created for a Resource Access Management (RAM) user. For more information, see Obtain an AccessKey pair.

  > **Note** FaaS supports only the AccessKey pair of RAM users. You can only access Functions as a Service (FaaS) using the Accesskey pair as an RAM user. This minimizes the risk that the AccessKey pair is exposed.

Context

The following section describes the key terms used in this topic:

- FPGA-accelerated instance: the FPGA-accelerated instance that you create by using the Alibaba

Cloud console or calling API operations. FPGA-accelerated instances include computing resources, Elastic Compute Service (ECS) images, and disks.

- ECS image: the image used by an FPGA-accelerated instance. FPGA-accelerated instances provide FPGA acceleration capabilities and the same user experience as regular ECS instances.
- Raw file: the acceleration program that you developed for FPGA devices and is used to compile and generate FPGA images. For example, the raw file of an Intel FPGA is in the *.gbs* format. The raw file of a Xilinx FPGA is a tar package that is generated from processed scripts.
- FPGA image: the image used by an FPGA device. A raw file is compiled to generate an FPGA image. After you load the FPGA image to an FPGA device, the FPGA device can provide acceleration services based on your design. Each FPGA image has a UUID.

When you use faasutil, take note of the following items:

- Use the following method to run faasutil commands: `./faasutil [command]`.
- faasutil commands and their parameters are case-sensitive. For example, errors occur when you run the `./faasutil Create_Image --Object=faasutiltest-forcompiling.tar.gz --Shell=f30010 --FPGAType=xilinx --Name=faasutiltest-image` command. The command and its parameters must be spelled in the same way as they are defined.
- Make sure that no spaces exist between parameters, equal signs (=), and values. For example, error occurs when you run the `./faasutil create_image --object=faasutiltest-forcompiling.tar.gz --shell = f30010 --fpgatype=xilinx --name=faasutiltest-image` command. You must remove the spaces before and after the equal sign (`=`) of the shell parameter.

faasutil provides commands that can be used to perform the following operations:

- Obtain help information
  - View all commands
  - View specified commands
- Configure the environment
  - Configure user information
  - Grant the read permissions on OSS buckets
  - View the read permissions on OSS buckets
  - Delete the read permissions on OSS buckets
  - Obtain information about FPGA-accelerated instances
  - Obtain the status information of an FPGA device
  - Upload files to an OSS bucket
  - View files in an OSS bucket
  - Download files to an FPGA-accelerated instance
- Manage FPGA images
  - Create tasks to compile FPGA images
  - Query compilation logs
  - View FPGA images
  - Modify FPGA image information
  - Bind FPGA images to ECS images
  - Copy FPGA images across regions

- Load FPGA images
- Query the loading status of FPGA images
- Delete FPGA images

View all commands

View the faasutil version and the list of commands.

Definition:

- `faasutil --help`
- `faasutil -h`

Example:

```
[root@i**** ~]# ./faasutil -h
faasutil [command] [optons]...
  DESCRIPTION
    faasutil is a fpga image management tool.
    version : 2.0.3a
    date    : 2020-07-08-11:04
  command list(use "faasutil [command] --help" for more infomation):
    config                  Config faasutil enviroment.
    list_instances          Print the instances base infomation.
    fpga_status             Check pointed fpga status.
    auth                    Authorize FAAS to read the oss bucket your specified.
    list_objects            List objects of the bucket you specified.
    put_object              Put local file to the bucket object, FAAS need download it from
oss.
    get_object              Get object from oss bucket.
    list_policy             List the detailed information about oss policy authorize to FAA
S.
    delete_policy           Delete the oss policy authorize to FAAS.
    list_images             List the fpga images
    create_image            Create fpga image.
    modify_image            Modify fpga image.
    delete_image            Delete fpga image.
    bind_image              Bind the FPGA image to ECS image.
    download_image          Download FPGA image to faas instance.
    download_status         The status of download image operation.
    copy_fpga_image         Copying images across regions.
    query_log               Get compiler log.
```

View specified commands

View the description of specified commands and their parameters.

Definition:

- `faasutil [command name] --help`
- `faasutil [command name] -h`

Example:

```
[root@i**** ~]# ./faasutil config --help
  COMMAND
    config
  SYNOPSIS
    config --id=[accessid] --key=[accesskey] --token=[sts-token] --bucket=[bucket] --region
Id=[regionId] --ossEndpoint=[ossEndpoint] --popEndpoint=[popEndpoint]
  DESCRIPTION
    Config faasutil enviroment.
    /*Using config to save the value of these follow options.*/
    /*You can also set these options' value follow any command which need to specify.*/
  OPTIONS
    required
    --id
        User access id , using config to save this information.
    --key
        User access key, using config to save this information.
    optional
    --token
        User ststoken  , using config to save this information.
    --bucket
        Bucket used for upload image.
    --regionId
        The region of pop endpoint and oss endpoint.
        Default using ecs region of public network.
    --ossEndpoint
        Oss endpoint info.
        Specify the oss endpoint if you want to get/put files in different region or vpc ne
twork.
    --popEndpoint
        Pop endpoint info.
    --instanceId
        ECS instance ID.
        Default using the local ECS ID.
    --ecsImageId
        ECS image ID.
        Default using the image ID of the local ECS ID.
    --shell
        Specify the shell version you used.
        Default using the shell version of the ECS fpga card.
    --fpgatype
        Specify the FPGA device type.
        Default using the FPGA device type of belong to the local ECS.
```

Configure user information

Write user information to the default configuration file. Take note of the following parameters:

- AccessKey ID and AccessKey secret: required. The AccessKey ID and AccessKey secret are used for authentication.
- OSS bucket name: optional. The name of the OSS bucket to which raw files are uploaded. We recommend that you specify this parameter. This way, you do not need to enter the OSS bucket name when you run commands that require this parameter.

Definition:

```
faasutil config --id=[accessid] --key=[accesskey] --token=[sts-token] --bucket=[bucket] --r
egionId=[regionId] --ossEndpoint=[ossEndpoint] --popEndpoint=[popEndpoint]
```

> **Note** You can run the **config** command to view information that can be configured.

Example:

```
[root@i**** ~]# ./faasutil config --id=L**** --key=v**** --bucket=yk****
Your configuration is saved into /root/.faascredentials .
1.057(s) elapsed
```

Grant the read permissions on OSS buckets

Grant the read permissions on OSS buckets for FPGA-accelerated instances to read files from OSS buckets.

Definition:

```
faasutil auth
```

Example:

```
[root@i**** ~]# ./faasutil auth
{
  "Statement": [
    {
      "Action": [
        "oss:GetObject",
        "oss:GetObjectMeta",
        "oss:ListObjects"
      ],
      "Effect": "Allow",
      "Principal": [
        "26143530338319****"
      ],
      "Resource": [
        "acs:oss:*:*:yk****/*"
      ]
    }
  ],
  "Version": "1"
}
0.223(s) elapsed
```

View the read permissions on OSS buckets

View the granted read permissions on OSS buckets.

Definition:

```
faasutil list_policy
```

Example:

```
[root@i**** ~]# ./faasutil list_policy
Get policy text:
{
  "Statement": [
    {
      "Action": [
        "oss:GetObject",
        "oss:GetObjectMeta",
        "oss:ListObjects"
      ],
      "Effect": "Allow",
      "Principal": [
        "26143530338319****"
      ],
      "Resource": [
        "acs:oss:*:*:yk****/*"
      ]
    }
  ],
  "Version": "1"
}
0.203(s) elapsed
```

Delete the read permissions on OSS buckets

Delete the read permissions on OSS buckets to prohibit FPGA-accelerated instances from reading files from OSS buckets.

Definition:

```
faasutil delete_policy
```

Example:

```
[root@i**** ~]# ./faasutil delete_policy
0.210(s) elapsed
```

Obtain information about FPGA-accelerated instances

Obtain information about FPGA-accelerated instances, such as instance IDs, ECS images, and attached FPGA devices.

Definition:

```
faasutil list_instances
```

Example:

```
[root@i**** ~]# ./faasutil list_instances
instance_id          : i-uf616ov8zzwmiijb****
image_id             : centos_7_8_x64_20G_alibase_20200914.vhd
regionId             : cn-shanghai
fpga bdf             : ['00:09.0', '00:08.0']
shell version        : f30010
requestId            : x-7762368991973****
fpga type            : xilinx
1.054(s) elapsed
```

In this example, an FPGA-accelerated instance is attached with two FPGA devices whose bdf is set to 00:09.0 and 00:08.0.

Obtain the status information of an FPGA device

Obtain the status information based on bdf of an FPGA device. You can run the **list_instances** command to view bdf of an FPGA device.

Definition:

```
faasutil fpga_status --bdf=[bdf]
```

Example:

```
[root@i**** ~]# ./faasutil fpga_status --bdf=00:09.0
clock0               : 240
clock1               : 189
fpgaStatus           : valid
requestId            : x-1093970844245****
shellUUID            : f30010
imageuuid            : NULL
firewall             : ['mgmt : GOOD', 'user : GOOD', 'xdma : GOOD', 'dma : GOOD']
ddr                  : ['ddr0 : online', 'ddr1 : online', 'ddr2 : online', 'ddr3 : online']
message              : FPGA is working.
1.025(s) elapsed
```

This example shows the status information of an FPGA device whose bdf is set to *00:09.0*. imageuuid specifies the UUID of the specified FPGA image. For a newly purchased FPGA-accelerated instance, the value of imageuuid is *null*.

Upload files to an OSS bucket

Upload files such as raw files on an FPGA-accelerated instance to an OSS bucket.

Definition:

```
faasutil put_object --bucket=[bucket_name] --object=[object_name] --file=[file_name]
```

- *[bucket_name]*: the name of the OSS bucket.

  > **Note** If you have run the **config** command to specify an OSS bucket in the default configuration file, you do not need to specify the bucket parameter.

- *[file_name]*: the name of the file to be uploaded.
- *[object_name]*: the displayed name of the uploaded file in the OSS bucket.

Example:

```
[root@i**** ~]# ./faasutil put_object --object=faasutiltest-forcompiling.tar.gz --file=faas
utiltest.tar.gz
100%
539.171(s) elapsed
```

View files in an OSS bucket

View files in an OSS bucket.

Definition:

```
faasutil list_objects
```

Example:

```
[root@i**** ~]# ./faasutil list_objects
faasutiltest-forcompiling.tar.gz
0.288(s) elapsed
```

Download files to an FPGA-accelerated instance

Download files from an OSS bucket to an FPGA-accelerated instance.

Definition:

```
faasutil get_object --bucket=[bucket_name] --object=[object_name] --file=[file_name]
```

- *[bucket_name]*: the name of the OSS bucket.

> **Note** If you have run the **config** command to specify an OSS bucket in the default configuration file, you do not need to specify the bucket parameter.

- *[file_name]*: the name of the file to be downloaded.
- *[object_name]*: the displayed name of the downloaded file in the FPGA-accelerated instance.

Example:

```
[root@i**** ~]# ./faasutil get_object --object=faasutiltest-forcompiling.tar.gz --file=faas
utiltest-fordeveloping.tar.gz
100%
179.909(s) elapsed
```

Create tasks to compile FPGA images

Use raw files to compile and generate FPGA images.

> **Note** Compilation takes some time. Wait until the compilation is complete. If you have already created a compilation task, you must wait at least 30 minutes before you can create another compilation task, regardless of whether the compilation is complete.

Definition:

```
faasutil create_image --object=[object_name] --shell=[shell_ver] --fpgatype=[fpga_type]
```

- *[object_name]*: the name of the raw file used for compilation. Make sure that the raw file has been uploaded to the OSS bucket.
- *[shell_ver]*: the shell version of the FPGA device. You can run the **fpga_status** command to view the shell version of the FPGA device.
- *[fpga_type]*: the type of the FPGA device. Xilinx FPGAs (*xilinx*) and Intel FPGAs (*intel*) are supported. Specify this parameter based on the type of your FPGA device.

Example:

```
[root@i**** ~]# ./faasutil create_image --object=faasutiltest-forcompiling.tar.gz --shell=f
30010 --fpgatype=xilinx --name=faasutiltest-image
{
  "CreateTime": "2020-12-04 14:48:13",
  "Description": "",
  "FpgaImageUniqueId": "xilinxb80832c9-1cd1-49a2-a7b4-7d3a5a4b****",
  "Name": "faasutiltest-image",
  "RequestId": "2E996CE5-F550-43E7-8016-DAF0D32A96BF",
  "ShellUniqueId": "f30010",
  "State": "queueing"
}
0.555(s) elapsed
```

This example shows that the FpgaImageUniqueId parameter is returned for an FPGA image after an FPGA image compilation task is created. In this example, the value of State is *queueing*, which indicates that the task is waiting for compilation. You can run the **list_images** command and view the status of the task to which FpgaImageUniqueId corresponds. When State becomes *success*, the compilation is complete.

Query compilation logs

Query the logs generated when the specified FPGA image compilation task is executed. *[fpgaImageUniqueId]* specifies the UUID of an FPGA image returned after an FPGA image compilation task is created.

Definition:

```
faasutil query_log --fpgaImageUniqueId=[fpgaImageUniqueId]
```

Example:

```
[root@i**** ~]# ./faasutil query_log --fpgaImageUniqueId=xilinxb80832c9-1cd1-49a2-a7b4-7d3a
5a4b****
{
  "Log": "http://aliyun-faas-images-cn-shanghai.oss-cn-shanghai-internal.aliyuncs.com/27095
692903785****%2Fxilinxb80832c9-1cd1-49a2-a7b4-7d3a5a4b****_log?Expires=160706****&OSSAccess
KeyId=L****&Signature=1****",
  "RequestId": "07A744F5-422D-49AA-B69B-1AD898BBD1D8"
}
0.194(s) elapsed
```

In this example, a URL is displayed in the command output. You can access the URL to obtain logs.

> **Note** If you do not want to view logs in an FPGA-accelerated instance, remove `-internal` from the URL.

View FPGA images

View existing FPGA images. You can filter images by image owner. *[owner]* specifies the image owner. Alibaba Cloud Marketplace images (*market*) and local images (*self*) are supported.

Definition:

```
faasutil list_images --owneralias=[owner]
```

Example:

```
[root@i**** ~]# ./faasutil list_images
{
  "FpgaImages": [
    {
      "CreateTime": "2020-12-04 14:48:14",
      "Description": "faasutil test round 1",
      "Encryption": false,
      "FpgaImageUniqueId": "xilinxb80832c9-1cd1-49a2-a7b4-7d3a5a4b****",
      "Name": "faasutiltest-image",
      "OwnerId": "27095692903785****",
      "ShellUniqueId": "f30010",
      "State": "success",
      "Tags": "",
      "UpdateTime": "2020-12-04 15:02:05"
    }
  ],
  "RequestId": "C85DD548-9B3D-463E-BE25-316FF7288CEB"
}
0.087(s) elapsed
```

Modify FPGA image information

Modify the information of an existing FPGA image, such as its name and description.

Definition:

```
faasutil modify_image --fpgaImageUniqueId=[fpgaImageUniqueId] --name=[name] --description=[
description] --tags=[tags]
```

Example:

```
[root@i**** ~]# ./faasutil modify_image --fpgaImageUniqueId=xilinxb80832c9-1cd1-49a2-a7b4-7
d3a5a4b*** --name="faasutiltest-image" --description="faasutil test round 2"
{
  "Description": "faasutil test round 2",
  "FpgaImageUniqueId": "xilinxb80832c9-1cd1-49a2-a7b4-7d3a5a4b****",
  "Name": "faasutiltest-image",
  "RequestId": "AED50829-B84A-42F9-9891-E1EE9821DA33",
  "Tags": ""
}
```

View the modified information:

```
[root@i**** ~]# ./faasutil list_images
{
  "FpgaImages": [
    {
      "CreateTime": "2020-12-04 14:48:14",
      "Description": "faasutil test round 2",
      "Encryption": false,
      "FpgaImageUniqueId": "xilinxb80832c9-1cd1-49a2-a7b4-7d3a5a4b****",
      "Name": "faasutiltest-image",
      "OwnerId": "27095692903785****",
      "ShellUniqueId": "f30010",
      "State": "success",
      "Tags": "",
      "UpdateTime": "2020-12-04 15:02:05"
    }
  ],
  "RequestId": "C85DD548-9B3D-463E-BE25-316FF7288CEB"
}
0.087(s) elapsed
```

Bind FPGA images to ECS images

After you bind FPGA images to ECS images, you can publish the ECS images to Alibaba Cloud Marketplace. When your ECS images are used by another user, the user can also use the FPFA images bound to the ECS images.

Definition:

```
faasutil bind_image --fpgaImageUniqueId=[fpgaImageUniqueId] --ecsImageId=[ecsImageId]
```

- *[fpgaImageUniqueId]*: the UUID of the FPGA image.
- *[ecsImageId]*: the ID of the ECS image. By default, the ID of the ECS image on the current FPGA-accelerated instance is used.

Example:

```
[root@i**** ~]# ./faasutil bind_image --fpgaImageUniqueId=xilinx5395ad11-edb9-4d93-964d-12c
59a07****
{
  "FpgaImageUniqueId": "xilinx5395ad11-edb9-4d93-964d-12c59a07****",
  "ImageId": "",
  "Message": "publish succeed!",
  "RequestId": "966883DD-6A6F-456A-8894-8FB7FA8A321C"
}
0.119(s) elapsed
```

Copy FPGA images across regions

If you want to use an existing FPGA image in a new region, you do not need to create a new image compilation task in this region. You can copy the existing FPGA image to this region.

Definition:

```
faasutil copy_fpga_image --fpgaImageUniqueId=[fpgaImageUniqueId] --targetRegion=[target reg
ion id]
```

Example:

```
[root@i**** ~]# ./faasutil copy_fpga_image --fpgaImageUniqueId=xilinx5395ad11-edb9-4d93-964
d-12c59a07**** --targetRegion=cn-hangzhou
{
  "CreateTime": "2020-12-04 17:46:04",
  "Description": "",
  "FpgaImageUniqueId": "xilinx5395ad11-edb9-4d93-964d-12c59a07****",
  "Name": "faasutiltest-image",
  "RequestId": "3AAEC35C-2FB2-4DBA-A982-7A8219055D60",
  "ShellUniqueId": "f30010",
  "State": "committed"
}
0.309(s) elapsed
```

Load FPGA images

Load FPGA images to FPGA devices. You can run the **list_instances** command to view the ID of an FPGA-accelerated instance and bdf of an FPGA device. You can run the **list_images** command to view the UUID of an FPGA image.

Definition:

```
faasutil download_image --bdf=[bdf] --fpgaImageUniqueId=[fpgaImageUniqueId] --instanceId=[i
nstanceId] --owneralias=[owner]
```

Example:

```
[root@i**** ~]# ./faasutil download_image --bdf=00:09.0 --fpgaImageUniqueId=xilinx5395ad11-
edb9-4d93-964d-12c59a07****
null
status               : operating
info                 : download task is operating
requestId            : x-84128618168241720
1.165(s) elapsed
```

Query the loading status of FPGA images

Query the loading status of FPGA images on FPGA devices. You can run the **list_instances** command to view bdf of an FPGA device or run the **list_images** command to view the UUID of an FPGA image.

Definition:

```
faasutil download_status --bdf=[bdf] --fpgaImageUniqueId=[fpgaImageUniqueId] --owneralias=[
owner]
```

Example:

```
[root@iZ**** ~]# ./faasutil download_status --bdf=00:09.0 --fpgaImageUniqueId=xilinx5395ad1
1-edb9-4d93-964d-12c59a07****
status               : done
info                 : download task finish
requestId            : x-42850294908763910
time                 : 2020-12-04 17:50:48
1.026(s) elapsed
```

Delete FPGA images

Delete FPGA images that are no longer needed.

> **Note** If the FPGA images have been published to Alibaba Cloud Marketplace, you must specify `--owneralias=market`.

Definition:

```
faasutil delete_image --fpgaImageUniqueId=[fpgaImageUniqueId] --owneralias=[owner]
```

Example:

```
[root@i**** ~]# ./faasutil delete_image --fpgaImageUniqueId=xilinxe27832a0-7c1c-4d6d-80b8-4
35253e9****
{
  "FpgaImageUniqueId": "xilinxe27832a0-7c1c-4d6d-80b8-435253e9****",
  "RequestId": "F2F481F5-B0E3-4009-9AF4-CF4BFC00458A"
}
0.157(s) elapsed
```

View the deletion result:

```
[root@i**** ~]# ./faasutil list_images
{
  "FpgaImages": [],
  "RequestId": "30B73A5B-99FC-4920-8B90-4BA08A075040"
}
0.049(s) elapsed
```

5.faascmd tool

5.1. Overview

faascmd is a command-line tool provided by Alibaba Cloud FPGA as a Service (FaaS). The faascmd script is developed based on SDK for Python.

You can use faascmd to perform the following operations:

- Perform authorization and related operations.
- Manage FPGA images.
- View and upload objects.
- Obtain information about FPGA-accelerated instances.

> **Note** faasutil is a next-generation command-line tool provided by Alibaba Cloud FaaS. faasutil improves the ease-of-use, stability, security, and scalability of FPGA-accelerated instances. For more information, see Obtain faasutil.

5.2. Install faascmd

This topic describes how to download and install faascmd.

Prerequisites

Before you install faascmd, make sure that the following operations are complete:

1. The AccessKey ID and AccessKey secret of the RAM user are obtained. For more information, see Obtain an AccessKey pair.
2. Python 2.7.x is installed and the version of SDK for Python is 2.11.x or later. The following procedure describes how to check the versions of Python and SDK for Python:
   i. Run the `python -V` command to check whether the Python version is 2.7.x.

   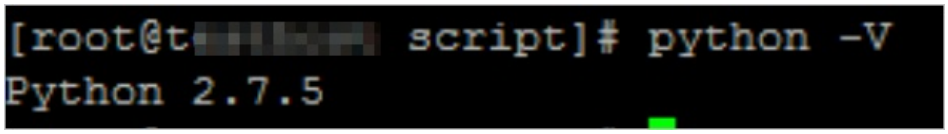

   ii. Run the following commands to install Python modules:

   ```
   pip -q install oss2
   pip -q install aliyun-python-sdk-core
   pip -q install aliyun-python-sdk-faas
   pip -q install aliyun-python-sdk-ram
   ```

iii. Run the following command to check whether the version of aliyun-python-sdk-core is 2.11.0 or later:

```
cat /usr/lib/python2.7/site-packages/aliyunsdkcore/__init__.py
```

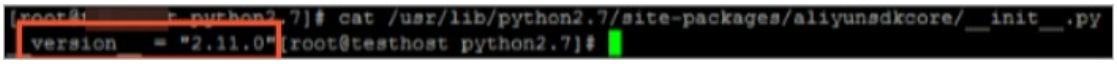

> **Note** If the version is earlier than 2.11.0, run the `pip install --upgrade aliyun-python-sdk-core` command to upgrade aliyun-python-sdk-core to the latest version.

Procedure

1. Log on to the instance. Run the `wget http://fpga-tools.oss-cn-shanghai.aliyuncs.com/faascmd` command in a directory to download faascmd.

   > **Note** Record the directory. When you configure the faascmd tool, you must add the absolute path of the directory where faascmd is installed to the PATH variable.

2. Run the following command to execute permissions to fasscmd.

   ```
   chmod +x faascmd
   ```

5.3. Configure faascmd

Before you use faascmd, you must configure the environment variables and the AccessKey pair of the RAM user.

Procedure

1. After you log on to the instance, run the following command to configure the environment variable PATH:

   ```
   export PATH=$PATH:/<The path of faascmd>/bin
   ```

2. Run the following command to configure your AccessKey ID and AccessKey secret:

   ```
   faascmd config --id=<yourAccessKeyID> --key=<yourAccessKeySecret>
   ```

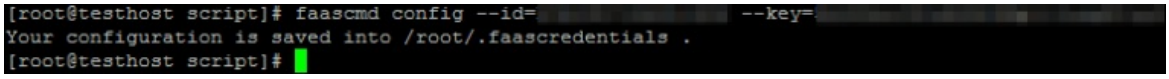

5.4. Use faascmd

This topic describes how to use faascmd commands.

Prerequisites

faascmd is configured. For more information, see Configure faascmd.

Before you use the authorization command, make sure that the following requirements are met:

1. An Object Storage Service (OSS) bucket is created for FPGA as a Service (FaaS) to upload the compiled DCP file.

2. A folder named *compiling_logs* is created in the bucket.

Context

Description of the faascmd command syntax:

- All commands and parameters provided by faascmd are case-sensitive.
- In faascmd commands, extra spaces are not allowed between parameters, `equal signs (=)`, or values.

This topic describes the following faascmd commands:

- Grant permissions to a RAM user
- View an authorization policy
- Delete an authorization policy
- View all objects in an OSS bucket
- Upload an compiled file
- Download an object from an OSS bucket
- Create an FPGA image
- View FPGA images
- Delete FPGA images
- Download an FPGA image
- View the download progress of an FPGA image
- Publish an FPGA image
- Query information of FPGA-accelerated instances

Grant permissions to a RAM user

You can run the `faascmd auth` command to authorize a RAM user to access your OSS buckets as a FaaS administrator.

Command syntax:

```
faascmd auth --bucket=<YourFaasOSSBucketName>
```

Sample code:

```
[root@testhost script]# faascmd auth --bucket=juliabucket
faasRole has existed!
RAMSECTION has existed!
OSSSECTION has existed!
RoleArn: acs:ram::[redacted]:role/faasrole
Create role success
faasPolicy has not existed! Create it Now!
Create policy success
Attach policy to role success
0.459(s) elapsed
```

Note If your Alibaba Cloud account has multiple RAM users, we recommend that you share the OSS bucket with the RAM users. This way, you no longer need to repeatedly modify or overwrite authorization policies.

View an authorization policy

You can run the `faascmd list_policy` command to check whether the specified OSS bucket is added in the authorization policy (faasPolicy).

Command syntax:

```
faascmd list_policy
```

Sample code:

```
[root@testhost script]# faascmd list_policy
VersionId : v1   CreateTime : 2018-11-09T03:22:01Z   IsDefaultVersion : True
{
  "Statement": [
    {
      "Action": "ecs:DescribeInstances",
      "Effect": "Allow",
      "Resource": "acs:ecs:*:*:*"
    },
```

Note Check whether your OSS bucket and the compiling_logs folder in the bucket are displayed in the policy information.

Delete an authorization policy

You can run the `faascmd delete_policy` command to delete an authorization policy (faasPolicy).

Command syntax:

```
faascmd delete_policy
```

Sample code:

```
[root@testhost script]# faascmd delete_policy
Detach faasPolicy from faasRole successfully!!!
Delete the faasPolicy successfully!!!
0.306(s) elapsed
```

Note If your Alibaba Cloud account has multiple RAM users, we recommend that you perform this operation in the RAM console. This helps prevent policies from being accidentally deleted.

View all objects in an OSS bucket

You can run the `faascmd list_objects` command to view all objects in an OSS bucket.

Command syntax:

```
faascmd list_objects
```

Sample code:

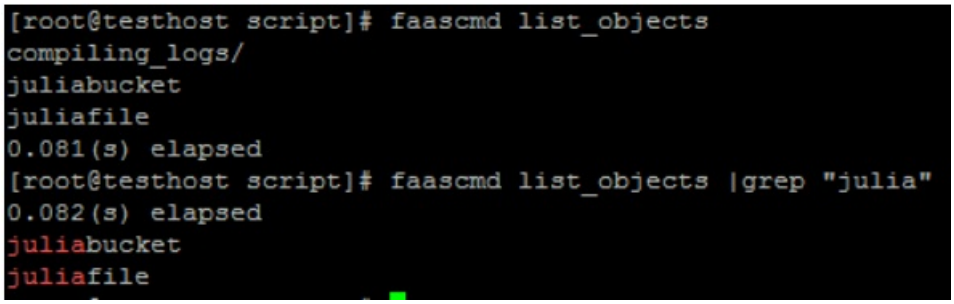

Note You can use this command and the `grep` command to filter objects. Example: `faascmd list_objects | grep "xxx"` .

Upload an compiled file

You can run the `faascmd upload_object` command to upload original copies of files compiled on your local PC to a specified OSS bucket.

Command syntax:

```
faascmd upload_object --object=<NewFileNameInOSSBucket> --file=<YourFilePath>/<FileNameYouWantToUpload>
```

Sample code:

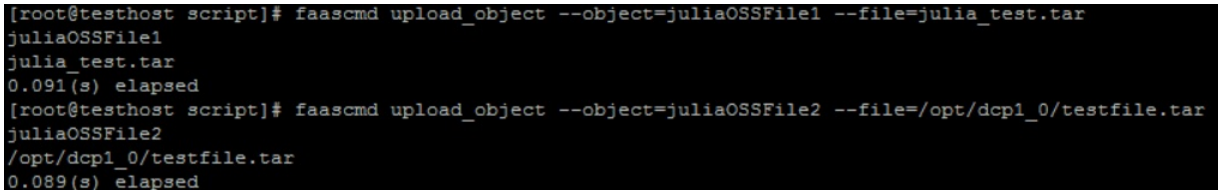

Note

- If the file that you want to upload is stored in the current directory, you do not need to specify a path.
- The locally compiled original files provided by Intel FPGA are in the .gbs format. The locally compiled original files provided by Xilinx FPGA are compressed as packages in the .tar format after script processing.

Download an object from an OSS bucket

You can run the `faascmd get_object` command to download the specified object from an OSS bucket.

Command syntax:

```
faascmd get_object --object=<YourObjectName> --file=<YourLocalPath>/<YourFileName>
```

Sample code:

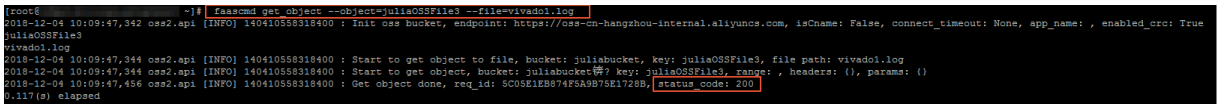

> **Note** If you do not specify a path, the object is downloaded to the current folder.

Create an FPGA image

You can run the `faascmd create_image` command to submit a request to create an FPGA image. If the request succeeds, FpgaImageUUID is returned.

Command syntax:

```
faascmd create_image --object=<YourObjectName>
--fpgatype=<intel/xilinx>  --encrypted=<true/false>
--kmskey=<key/If encrypted is set to true, this parameter is required. Otherwise, this para
meter is optional.>
--shell=<Shell Version/Required> --name=<name/Optional>
--description=<description/Optional> --tags=<tags/Optional>
```

Sample code:

View FPGA images

You can run the `faascmd list_images` command to view the information of all FPGA images that you create.

Command syntax:

```
faascmd list_images
```

Sample code:

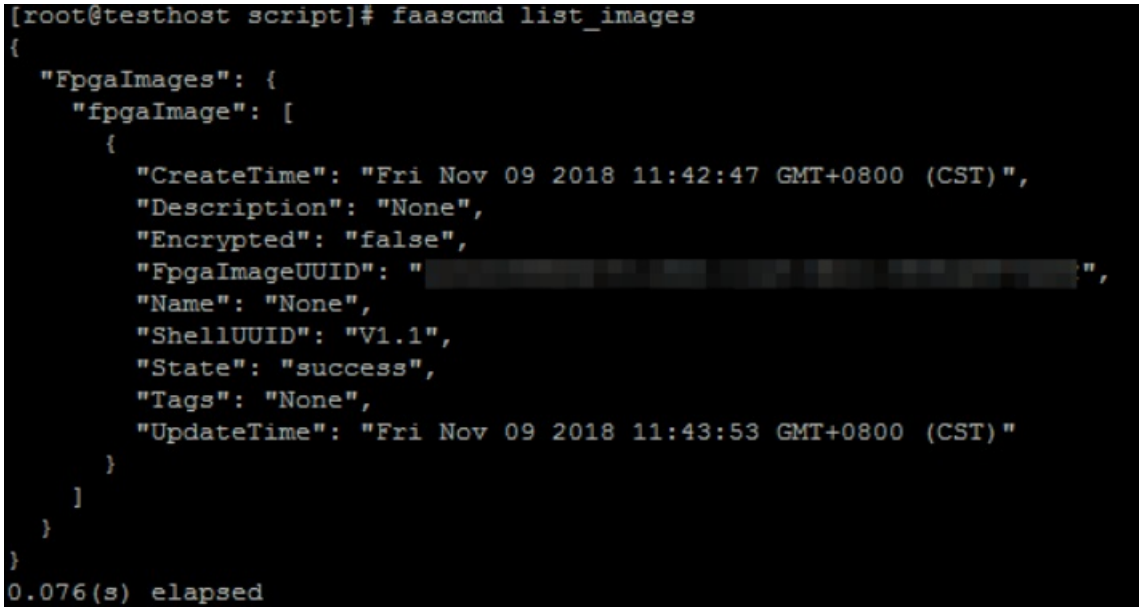

> **Note** Each RAM user can have up to 10 FPGA images.

Delete FPGA images

You can run the `faascmd delete_image` command to delete an FPGA image.

Command syntax:

```
faascmd delete_image --imageuuid=<yourImageuuid>
```

Sample code:

```
[root@testhost script]# faascmd delete_image --imageuuid=
{"Status":200,"FpgaImageUUID":"                                    ","Message":"delete succeed!"}
0.143(s) elapsed
```

Download an FPGA image

You can run the `faascmd download_image` command to submit a request to download an FPGA image.

Command syntax:

```
faascmd download_image  --instanceId=<YourInstanceId>
--fpgauuid=<Yourfpgauuid> --fpgatype=<intel/xilinx>
--imageuuid=<YourImageuuid> --imagetype=<afu>
--shell=<YourImageShellVersion>
```

Sample code:

```
faascmd download_image --instanceId=XXXXX --fpgauuid=XXXX --fpgatype=intel --imageuuid=XXXX
```

View the download progress of an FPGA image

You can run the `faascmd fpga_status` command to view the status of the current FPGA board card or the download progress of the FPGA image.

Command syntax:

```
faascmd fpga_status --fpgauuid=<Yourfpgauuid> --instanceId=<YourInstanceId>
```

Sample code:

```
[root@testhost script]# faascmd fpga_status --fpgauuid=                    --instanceId=
{"shellUUID":"V1.0","FpgaImageUUID":"                                    ","FpgaUUID":"                    ",
askStatus":"invalid","Encrypted":"false"}
0.310(s) elapsed
```

Publish an FPGA image

You can run the `faascmd publish_image` command to submit a request to publish an FPGA image.

Command syntax:

```
faascmd publish_image --imageuuid=<YourImageuuid> --imageid=<YourInstanceImageid>
```

> **Note**
>
> - imageuuid specifies the ID of the FPGA image that you want to publish to Alibaba Cloud Marketplace. You can run the `faascmd list_images` command to view this image ID.
> - imageid specifies the image ID of the current FPGA-accelerated instance. You can go to the product details page in the ECS console to view the image ID.

Query information of FPGA-accelerated instances

You can run the `faascmd list_instances` command to query the basic information of an FPGA-accelerated instance. This information includes the instance ID, FPGA board card information, and shell version.

Command syntax:

```
faascmd list_instances --instanceId=<YourInstanceId>
```

Sample code:

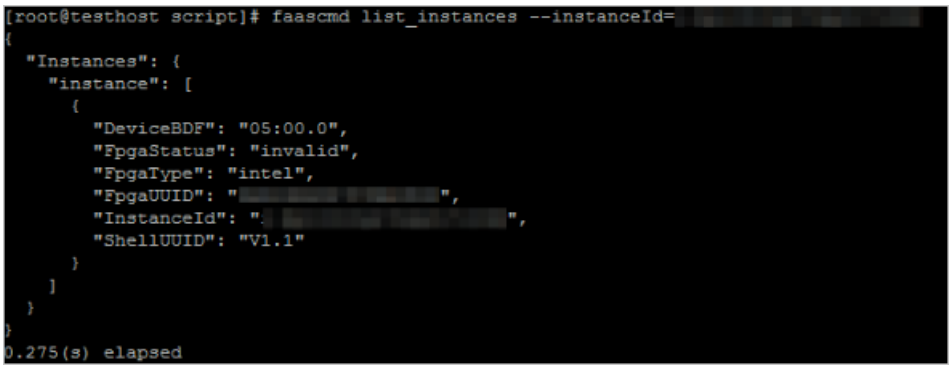

5.5. FAQ

This topic provides answers to some frequently asked questions about the faascmd tool.

- FAQ
  - What do I do if the "Name Error:global name'ID' is not defined." error message is returned?
  - What do I do if the "SDK.InvalidRegionId. cannot find endpoint to access."error message is returned?
  - What do I do if the "HTTP Status:" 404 Error: EntityNotExist. Role Error. The specified Role not exists . error message is returned?
  - When I attempt to download an FPGA image, the "HTTP Status:404 Error:SHELL NOT MATCH. The image Shell is not match with fpga Shell! Request ID:D7D1AB1E-8682-4091-8129-C17D54FD10D4" error message is returned. What do I do?
  - When I attempt to download an FPGA image, the "HTTP Status:503 Error:ANOTHER TASK RUNNING . Another task has not finished yet, please retry later! Request ID: 5FCB6F75-8572-4840-9BDC-87C57174F26D" error message is returned. What do I do?
  - When I run the faascmd list_images command, an error message is returned which indicates that the image is in the failed state. What do I do?
- Error codes

What do I do if the "Name Error:global name'ID' is not defined." error message is returned?

Cause: faascmd cannot obtain your AccessKey ID or AccessKey secret.

Solution: Run the `faascmd config` command to save the AccessKey ID and AccessKey secret that you entered to the */root/.faascredentials* file.

What do I do if the "SDK.InvalidRegionId. cannot find endpoint to access."error message is returned?

Cause: faascmd cannot obtain the endpoint of FPGA as a Service (FaaS).

Solution: Perform the following steps to check whether faascmd configurations meet the specified requirements:

- Run the `python -V` command to check whether the version of Python that is installed is 2.7.x.
- Run the `which python` command to check whether the default installation path of Python is `/usr/bin/python`.
- Run the `cat /usr/lib/python2.7/site-packages/aliyunsdkcore/__init__.py` command to check whether the version of aliyunsdkcore that is installed is 2.11.0 or later.

> **Note** If the aliyunsdkcore version is earlier than 2.11.0, run the `pip install --upgrade aliyun-python-sdk-core` command to upgrade aliyunsdkcore to the latest version.

What do I do if the "HTTP Status:" 404 Error: EntityNotExist. Role Error. The specified Role not exists . error message is returned?

Cause: AliyunFAASDefaultRole does not exist in your Alibaba Cloud account.

Solution: Log on to the RAM console to check whether AliyunFAASDefaultRole exists.

- If AliyunFAASDefaultRole does not exist, run the `faascmd config` and `faascmd auth` commands to create the role and grant permissions to the role.
- If AliyunFAASDefaultRole exists, submit a ticket.

When I attempt to download an FPGA image, the "HTTP Status:404 Error:SHELL NOT MATCH. The image Shell is not match with fpga Shell! Request ID:D7D1AB1E-8682-4091-8129-C17D54FD10D4" error message is returned. What do I do?

Cause: The shell versions of the target FPGA image and the specified FPGA do not match.

Solution: Perform the following steps:

- Run the `faascmd list_instances --instance=xxx` command to check the shell version of the current FPGA.
- Run the `faascmd list_images` command to check the shell version of the specified FPGA image.

> Note
> - If the two shell versions are different, you must create a new FPGA image of the same shell version as the FPGA. After the new FPGA image is created, download the image.
> - If the two shell versions are the same, submit a ticket.

When I attempt to download an FPGA image, the "HTTP Status:503 Error:ANOTHER TASK RUNNING . Another task has not finished yet, please retry later! Request ID: 5FCB6F75-8572-4840-9BDC-87C57174F26D" error message is returned. What do I do?

Cause: The FPGA is still in the operating state due to an unexpected failure or interruption of the download request that you submitted.

Solution: We recommend that you wait 10 minutes until the download task ends. Then, you can resubmit an image download request.

> Note If this issue persists, submit a ticket.

When I run the faascmd list_images command, an error message is returned which indicates that the image is in the failed state. What do I do?

Run the following commands to obtain the compilation logs for troubleshooting:

```
faascmd list_objects|grep vivado
faascmd get_object --object=<YourObjectName> --file=<YourLocalPath>/vivado.log  #If no path is specified, the compilation log is downloaded to the current folder.
```

Error codes

Http Code	Error code	Error message	Description	Scope
400	PARAMETER INVALIDATE	Specify parameters are invalid.	The error message returned because input parameters are invalid.	
500	InternalError	The request processing has failed due to some unknown error.	The error message returned because an unknown error has occurred. Submit a ticket.	
404	InvalidProduct .NotFound	Cannot find product according to your specified domain.	The error message returned because the FaaS service does not exist. Check whether the endpoint configurations of the Python Core SDK are correct.	

Http Code	Error code	Error message	Description	Scope
404	InvalidAccessKeyId.NotFound	Specified access key is not found.	The error message returned because the specified AccessKey ID does not exist.	• All faascmd commands • All API operations
400	InvalidAccessKeyId.Inactive	Specified access key is disabled.	The error message returned because the specified AccessKey ID is unavailable.	
400	InvalidSecurityToken.Expired	Specified SecurityToken is expired.	The error message returned because the specified SecurityToken is expired.	
400	InvalidSecurityToken.Malformed	Specified SecurityToken is malformed.	The error message returned because the specified SecurityToken is malformed.	
400	InvalidSecurityToken.MismatchWithAccessKey	Specified SecurityToken mismatch with the AccessKey.	The error message returned because the specified security token and AccessKey pair do not match.	
403	NoPermisson	You are not authorized to do this action.	The error message returned because you are not authorized to perform this operation.	• faascmd command: auth • API operation: auth
401	IMAGE NUMBER EXCEED	The user is allowed to have no more than 30 images.	The error message returned because the number of images has reached the upper limit of 30. Delete the images that you no longer need and try again.	
503	FREQUENCY ERROR	CreateFpgaImage task is allowed to take every half an hour.	If you submit a request to create an image, you must wait at least 30 minutes before you can submit another request to create an image.	
404	SHELL NOT SUPPORT	The shellUUID is not supported, please check your input shellUUID.	The error message returned because the specified shell version is not supported.	

Http Code	Error code	Error message	Description	Scope
404	EntityNotExist.RoleError	The specified Role not exists.	The error message returned because your account does not have the AliyunFAASDefaultRole role.	• faascmd command: create_image • API operation: CreateFpgaImage
403	AccessDeniedError	The bucket you visit does not belong to you.	The error message returned because the FaaS service role does not have permissions to access the current bucket.	
403	CALLER TYPE NOT SUPPORT	The callerType is not supported, please use sub user's AK.	The error message returned because the specified user identity credentials are not supported. Only the identity credentials of RAM users are supported.	
404	NoSuchBucketError	The specified bucket does not exist.	The error message returned because the specified OSS bucket does not exist. Check whether the specified bucket name is correct.	
404	OSS OBJECT NOT FOUND	The specified oss object does not exist.	The error message returned because the specified OSS object does not exist or because you have not authorized the FaaS RAM role to access the object.	

404	IMAGE NOT FOUND	The specify image does not found.	The error message returned because the specified FPGA image does not exist.	• faascmd command: delete_image
				• API operations:
				  ◦ DeleteFpgaImage
				  ◦ DeletePublishFpgaImage

Http Code	Error code	Error message	Description	Scope
401	NOT AUTHORIZED	You are not allowed to access this instance.	The error message returned because you are not authorized to access the specified instance. Check whether the permission policy attached to your account includes the permission to call the DescribeInstances operation.	• faascmd command: list_instances • API operation: DescribeFpgaInstances
403	CALLER TYPE NOT SUPPORT	The callerType is not supported.	The error message returned because the specified user identity credentials are not supported. Only the AccessKey pairs of RAM users and STS tokens are supported.	
404	INSTANCE INVALIDATE	The instance you specify is not FPGA type.	The error message returned because the specified instance is not an FPGA-accelerated instance. If the instance is an FPGA-accelerated instance, submit a ticket.	
401	NOT AUTHORIZED	You are not allowed to access this instance.	The error message returned because the specified instance ID does not exist. Check the input parameters.	• faascmd command: fpga_status • API operation: DescribeLoadTaskStatus
404	FPGA NOT FOUND	The fpga you specify is not found.	The error message returned because the specified fpgauuid value does not exist. Check the input parameters.	
503	ANOTHER TASK RUNNING	Another task is running, user is allowed to take this task half an hour.	The error message returned because the image download task that you submitted is in the operating state.	
401	IMAGE ACCESS ERROR	You are not allowed to access this fpga image.	The error message returned because the specified image does not belong to your account.	

Http Code	Error code	Error message	Description	Scope
401	YOU HAVE NO ACCESS TO THIS INSTANCE	You are not allowed to access this instance.	The error message returned because the specified instance does not belong to your account.	• faascmd command: download_image • API operation: LoadFpgaImage
404	IMAGE NOT FOUND	The fpga image you specify is not found.	The error message returned because the specified FPGA image does not exist.	
404	FPGA NOT FOUND	The fpga you specify is not found.	The error message returned because the specified FPGA-accelerated instance does not exist.	
404	SHELL NOT MATCH	The imageShell is not match with fpgaShell.	The error message returned because the shell version of the specified image does not match the shell version of the specified FPGA-accelerated instance.	
403	ASSUME ROLE USER NOT SUPPORT	AssumeRoleUser only support loading market fpga images.	The error message returned because an STS token is used to download an FPGA image that is not an Alibaba Cloud Marketplace image. STS tokens can be used to download only Alibaba Cloud Marketplace images.	
404	Image not in success state	The fpga image you specify is not in success state.	The error message returned because the specified FPGA image is not in the success state. You can download only the images that are in the success state.	
404	FPGA IMAGE STATE ERROR	The specify fpga image is not in success state.	The error message returned because the specified FPGA image is not in the success state.	• faascmd command: publish_image

Http Code	Error code	Error message	Description	Scope
404	FPGA IMAGE NOT FOUND	The specify fpga image does not found.	The error message returned because the specified image does not exist or does not belong to your account.	• API operation: PublishFpgaImage

6.Stop an instance

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. This topic describes how to stop an ECS instance and how to enable economical mode for instances that are located in virtual private clouds (VPCs).

Prerequisites

The instance that you want to stop is in the **Running** state.

> **Note** If you stop an instance, services that are running on an instance are interrupted. Proceed with caution when you perform this operation.

Context

The billing of a subscription instance is not affected when you stop the instance.

The billing of a pay-as-you-go instance may be affected when you stop the instance. This depends on whether economical mode is enabled for the instance.

- Pay-as-you-go instances in the classic network do not support economical mode and continue to be billed after they are stopped. Billing stops only when the instances are released. For more information, see Release an instance.
- Pay-as-you-go instances in VPCs support economical mode.
  - If economical mode is disabled for a pay-as-you-go instance in a VPC, the instance continues to be billed after it is stopped.
  - If economical mode is enabled for a pay-as-you-go instance in a VPC, the vCPUs, memory, and public IP address of the instance are no longer billed after the instance is stopped. Other resources continue to be billed. For more information, see Economical mode.

Stop a subscription instance

1.
2.
3.
4. Use one of the following methods to stop subscription instances:
   - To stop a single instance at a time, find the instance and choose **More > Instance Status > Stop** in the **Actions** column.
   - To stop multiple instances at a time, select the instances and click **Stop** in the lower part of the Instances page.
5. Configure Stopped By. Valid values:
   - **Stop**: stops the instance by shutting it down properly.
   - **Force Stop**: forcibly stops the instance. Forcible stop is equivalent to a physical shutdown and may cause data loss if instance data has not been written to disks.
6. Click **OK**.

Stop a pay-as-you-go instance

The procedures to stop preemptible instances are the same as those to stop pay-as-you-go instances. However, more factors affect the startup of stopped preemptible instances. For more information, see Stop a preemptible instance.

1.
2.
3.
4. Use one of the following methods to stop pay-as-you-go instances:
   - To stop a single instance at a time, find the instance and choose **More > Instance Status > Stop** in the **Actions** column.
   - To stop multiple instances at a time, select the instances and click **Stop** in the lower part of the Instances page.
5. Configure Stopped By and Stop Mode.
   - For a pay-as-you-go instance in the classic network:
     a. Configure Stopped By. Valid values:
        - **Stop**: stops the instance by shutting it down properly.
        - **Force Stop**: forcibly stops the instance. Forcible stop is equivalent to a physical shutdown, and may cause data loss if instance data has not been written to disks.
     b. Click **OK**.
   - For a pay-as-you-go instance in a VPC:
     a. Configure Stopped By. Valid values:
        - **Stop**: stops the instance by shutting it down properly.
        - **Force Stop**: forcibly stops the instance. Forcible stop is equivalent to a physical shutdown, and may cause data loss if instance data has not been written to disks.
     b. Configure Stop Mode. Valid values:
        - **Standard Mode**: The resources of the instance are retained and continue to be billed after the instance is stopped.
        - **Economical Mode (Formerly Known as No Fees for Stopped Instances Mode)**: After the instance is stopped, its computing resources (vCPUs and memory) are released and no longer billed. The cloud disks (including the system disk and data disks), elastic IP addresses (if any), and bandwidth continue to be billed. The public IP address is recycled and the private IP address is retained.
     c. Click **OK**.

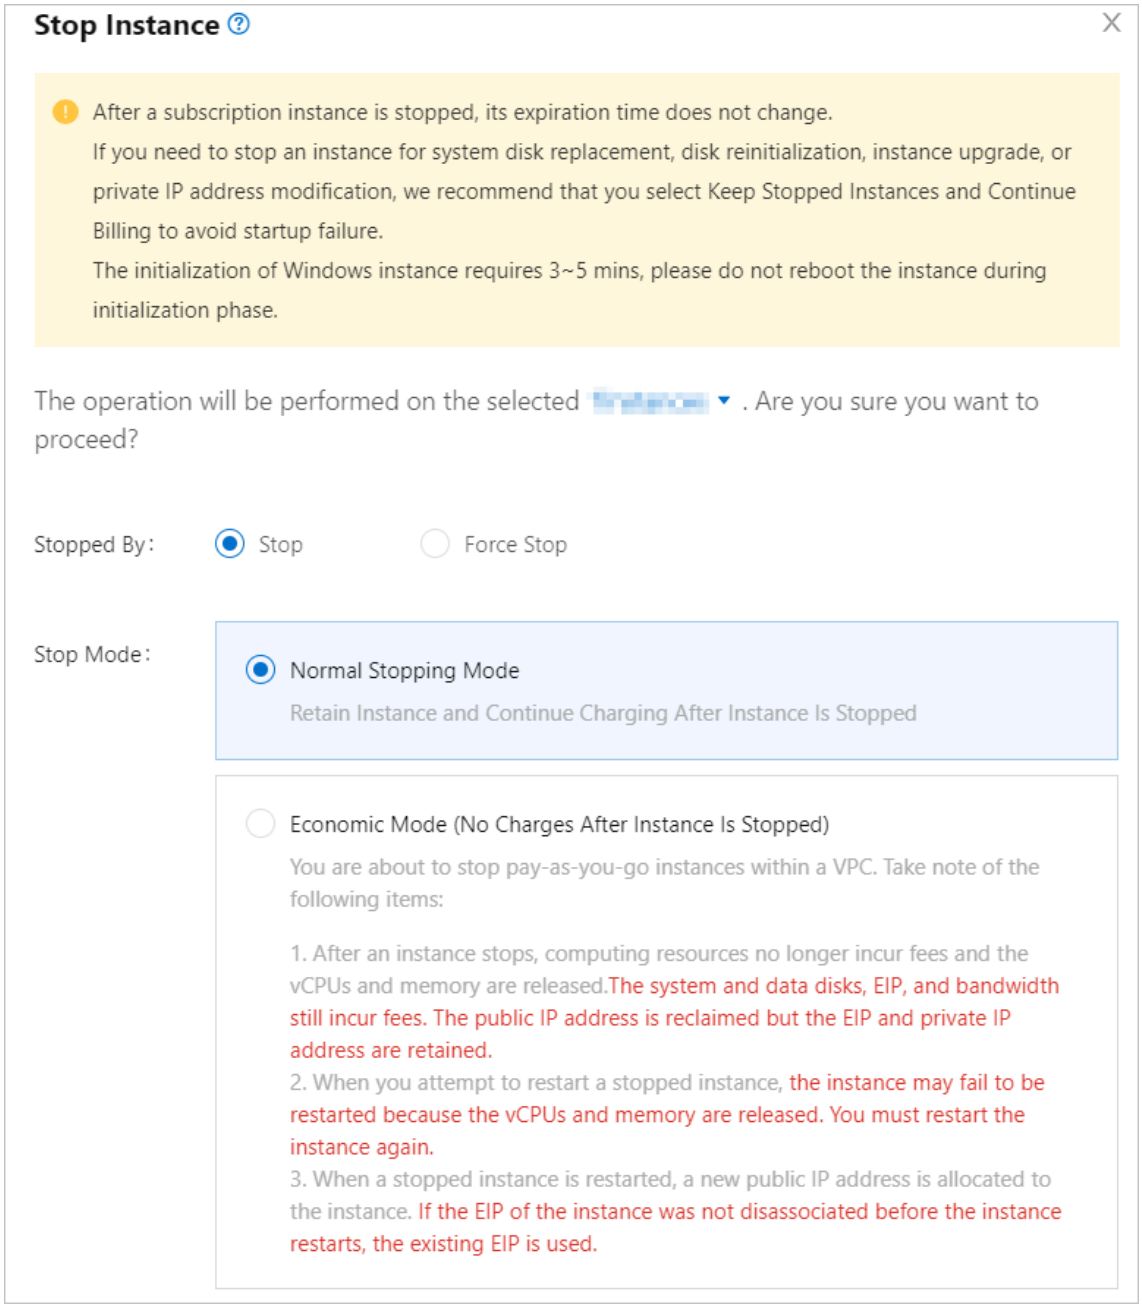

Result

The instance enters the **Stopped** state when it is stopped.

Related information

- StopInstance

7.Start an instance

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. This topic describes how to start an instance in the ECS console.

Prerequisites

The instance that you want to start meets one of the following requirements:

- The instance is in the **Stopped** state.
- The instance is a pay-as-you-go instance that entered the **Expired** state due to an overdue payment. The overdue payment is settled but the instance cannot be automatically reactivated.

> Note

Context

If you have settled an overdue payment that caused a pay-as-you-go instance to stop but the instance cannot be automatically reactivated, the instance will still be released. You must manually reactivate the instance in a timely manner to prevent the instance from being released and affecting your business. For more information, see Pay-as-you-go.

After the overdue payment for a pay-as-you-go instance is settled, the system reactivates the instance. If the instance cannot be automatically reactivated, wait 10 minutes and check whether the instance is reactivated and enters the **Running** state. If the instance still cannot be automatically reactivated, manually reactivate it in a timely manner.

> **Note** After you reactivate a pay-as-you-go instance that was stopped due to an overdue payment, the instance begins to run again and resumes billing on a pay-as-you-go basis. Release instances that are no longer needed to avoid unnecessary costs.

The instance may fail to be manually reactivated if resources for the instance type are insufficient. Change the instance type. For more information, see Change the instance type of a pay-as-you-go instance. If the problem persists, submit a ticket.

Procedure

1.
2.
3.
4. Start instances.
   - To start a single instance, find the instance and choose **More > Instance Status > Start** in the **Actions** column.

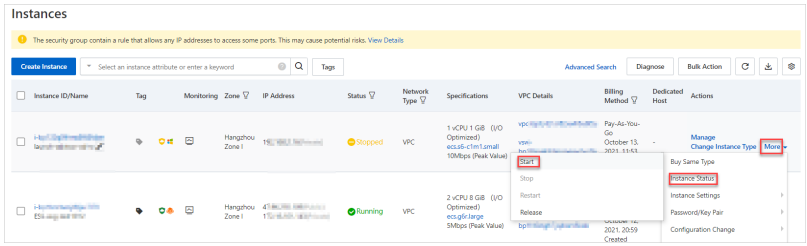

- To start multiple instances at a time, select the instances that you want to start and click **Start** in the lower part of the Instances page.

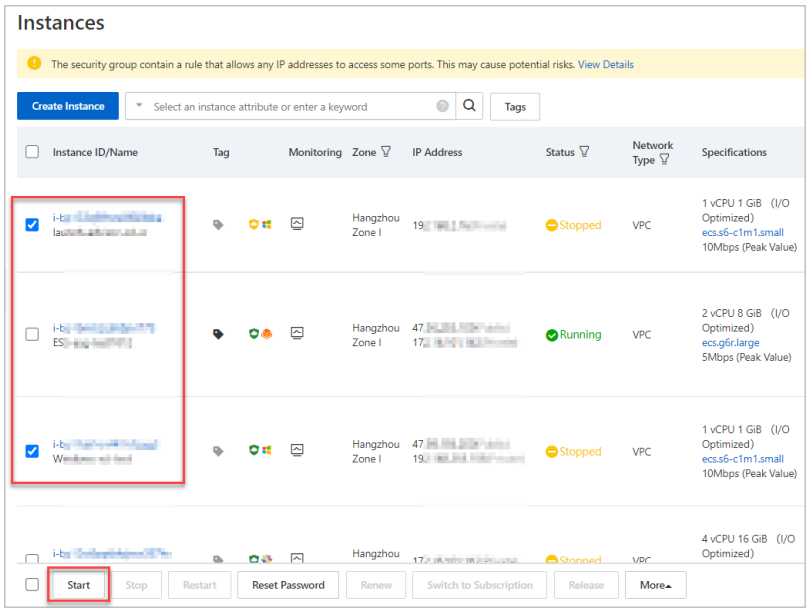

5. In the **Start Instance** dialog box, confirm the instance information and click **OK**.

Result

After the instance is started, it enters the **Running** state.

Related information

- StartInstance

8.Hibernate instances

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. If you do not need to use an ECS instance for a period of time but still want to retain the instance without performing operations such as configuration upgrade or downgrade on the instance, we recommend that you hibernate the instance. Hibernated instances are not the same as stopped instances. When hibernated instances are waked, they automatically restore their applications to the states the applications were in before hibernation. This allows the instances to resume services in a short time.

Context

When you hibernate an instance, the operating system of the instance saves data from the memory to the system disk of the instance. The saved data includes the applications that run in the operating system and the usage status of the applications. When you wake the instance, the operating system reads the data saved in the system disk, automatically restores the applications to the status before hibernation, and resumes the running state of the instance. In comparison, when you stop and restart an instance, the operating system restarts the backend services and applications.

> **Note** If the instance fails to be hibernated, the instance is automatically shut down. Data in the memory is not saved to the system disk. When the instance is started again, the operating system of the instance restarts the backend services and applications. The operating system cannot restore the applications to the status before hibernation.

Hibernation has different impacts on the billing of instances that use different billing methods:

- Subscription instance: The expiration time and billing of the hibernated instance are not affected.
- Pay-as-you-go instance: Whether the billing of the hibernated instance is affected is based on whether you select the **No Fees for Hibernated Instances** option when you hibernate the instance. The following table describes the billing details of resources.

Billing of resources on a hibernated instance

Resource	No Fees for Hibernated Instances	Retain Instance and Continue Charging After Instance Is Hibernated
Computing resource (vCPUs and memory)	Release and stop billing	Retain and continue billing
Disk (system disk and data disk)	Retain and continue billing	Retain and continue billing
Internal IP address	Retain and stop billing	Retain and stop billing
Public IP address	Release and stop billing. After the instance is started, a new public IP address is obtained.	Retain and stop billing
EIP	Retain and continue billing	Retain and continue billing
Bandwidth	Continue billing	Continue billing

Limits

- The instance hibernation feature is now available only in the US (Silicon Valley) and Germany (Frankfurt) regions, and will be gradually supported in other regions.
- Before you can hibernate an instance, the instance must meet the following requirements:
  - The instance hibernation feature is enabled when the instance is created.

    > **Note** The instance hibernation feature cannot be disabled after it is enabled. If you do not enable the instance hibernation feature when you create an instance, you cannot hibernate the instance.

  - The hibernation agent is installed on the instance.
- You can enable the instance hibernation feature only when you create an ECS instance by using an encrypted custom image. The following image versions are supported:
  - Windows Server 2016 or later
  - Ubuntu 18 or later
  - CentOS 7 or later
- If the instance hibernation feature is enabled for an ECS instance when the instance is created, you cannot perform the following operations on the instance:
  - Create custom images.
  - Create snapshots.
  - Change the instance type.
  - Change the operating system or system disk.
  - Change the bandwidth of subscription instances.
- If the instance hibernation feature is enabled for a preemptible instance, you can select only the No Fees for Hibernated Instances option when you hibernate the instance.
- You cannot hibernate ECS instances in scaling groups.

Step 1: Enable the instance hibernation feature

You must enable the instance hibernation feature when you create an ECS instance. Otherwise, you cannot hibernate the instance. When you create the instance, you must use an encrypted image.

1. Obtain an encrypted custom image.

   You can use one of the following methods to obtain an encrypted custom image:

   - Prepare an encrypted custom image that meets the hibernation requirements.
   - Copy an image and encrypt it at the same time. For more information, see Copy a custom image.

   > **Note** For more information about the limits on images, see Limits.

2. Create an ECS instance and enable the instance hibernation feature when you create the instance.

   For more information, see Create an instance by using the wizard. Take note of the parameters described in the following table.

Parameter	Description	Example
Instance	○ Instance Type: Select an I/O optimized instance type, except ECS Bare Metal Instance. ○ Memory: ■ Windows: Set the memory size to a value less than 16 GiB. ■ Linux: Set the memory size to a value less than 150 GiB.	ecs.g6e.large
Image	○ Select the encrypted custom image created in the previous step or an existing encrypted custom image that meets the hibernation requirements. ○ Select **Instance Hibernation** to enable the instance hibernation feature.	○ encrypted.windows2016 ○ Select **Instance Hibernation**.
Disk	○ System Disk: required. The system disk must meet the following requirements: ■ Category: ultra disk, standard SSD, or enhanced SSD (ESSD). ■ Capacity: The system disk capacity must be sufficient. We recommend that you set the system disk capacity to at least twice the memory size. This is because when the instance hibernation feature is enabled, the system disk reserves some space to store memory data. Therefore, the system disk capacity must be sufficient to ensure normal running of the operating system and applications when the system disk stores the memory data. ■ Encryption: By default, the system disk is encrypted if an encrypted image is used. ○ Data Disk: optional. To create data disks for an instance when you create the instance, you must select the disk categories and specify the sizes and quantity of the disks. You must also determine whether to encrypt the disks.	○ System Disk: Select Enhanced SSD (ESSD), set Disk Capacity to 60 GiB, select Disk Encryption, and then select **Default Service CMK** from the drop-down list. ○ Data Disk: Select Enhanced SSD (ESSD), set Disk Capacity to 40 GiB, and do not select Disk Encryption.

Parameter	Description	Example
Network	Select a virtual private cloud (VPC). **Note** ECS instances in the classic network do not support the instance hibernation feature.	[Default]vpc-bp1opxu1zkhn00g****

Step 2: Install the hibernation agent

After you enable the instance hibernation feature for the instance, you must install the hibernation agent on the instance before you can hibernate the instance.

1. Create and run one of the following commands to install the hibernation agent on the instance. For more information, see Use the immediate execution feature.
   - Windows instance:

     ```
     acs-plugin-manager.exe --exec --plugin ecs-hibernate-win --params "install"
     ```

   - Linux instance:

     ```
     acs-plugin-manager --exec --plugin ecs-hibernate-linux --params "install"
     ```

2. Restart the instance to make the hibernation agent take effect. For more information, see Restart an instance.

Step 3: Hibernate the instance

After the instance hibernation feature is enabled for the instance and the hibernation agent is installed on the instance, you can hibernate the instance in the Running state. You are unable to connect to the instance when the instance is hibernated.

1.
2.
3.
4. Find the instance that you want to hibernate and choose **More > Instance Status > Stop** in the **Actions** column.
5. In the **Stop Instance** dialog box, configure the parameters.
   i. Set **Stopped By** to **Hibernate**.

ii. Set **Stop Mode** to **Retain Instance and Continue Charging After Instance Is Hibernated** or **No Fees for Hibernated Instances** based on the billing method of the instance.

- You are charged for subscription instances even after the instances are hibernated. Select **Retain Instance and Continue Charging After Instance Is Hibernated**.
- For pay-as-you-go instances, you can select **Retain Instance and Continue Charging After Instance Is Hibernated** or **No Fees for Hibernated Instances**.

> **Note** For preemptible instances, select **No Fees for Hibernated Instances**.

For more information about the difference between **Retain Instance and Continue Charging After Instance Is Hibernated** and **No Fees for Hibernated Instances**, see Billing of resources on a hibernated instance in this topic.

iii. Click **OK**.

> **Note** The instance is stopped and enters the **Stopped** state. To start the instance, see Start an instance.

References

You can use Operation Orchestration Service (OOS) to hibernate and wake ECS instances at the scheduled time. This way, the hibernation and wake time of a large number of instances can be managed in an automated manner and the costs can be reduced by using the No Fees for Hibernated Instances feature. For more information, see Start and shut down ECS instances at the scheduled time.

9.Restart one or more instances

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. This topic describes how to restart one or more instances in the ECS console.

Prerequisites

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

Context

Restarting an instance will stop the instance. As a result, services provided by the instance are disrupted.

Procedure

1. Log on to the ECS console.
2. In the left-side navigation pane, choose **Instances & Images > Instances**.
3. Select the target region.
4. Find the target instances.
   - To restart a instance, choose **More > Instance Status > Restart** in the **Actions** column.
   - To restart multiple instances, select all required instances and then click **Restart** at the bottom of the instance list.
5. In the displayed **Restart Instance** dialog box, select a **Restart Mode**, and then click **OK**.

Related information

- RebootInstance

10.Release instances

From a resource management standpoint, FPGA-accelerated instances are considered as Elastic Compute Service (ECS) instances and are managed in the same way you manage other ECS instances. Only pay-as-you-go FPGA-accelerated instances (including preemptible instances) and expired subscription FPGA-accelerated instances can be released. This topic describes how to manually and automatically release pay-as-you-go instances.

Prerequisites

After an instance is released, its data is deleted and cannot be recovered. We recommend that you create snapshots to back up data before you release the instance. For more information, see Create a snapshot for a disk.

> **Note** After an instance is released, snapshots and images that were manually created from the instance are not affected.

Context

- Subscription instance that have not expired cannot be released. Before a subscription instance that has not expired can be released, you must convert it into a pay-as-you-go instance. For more information, see Change the billing method of an instance from subscription to pay-as-you-go.
- You can manually release expired subscription instances. If you do not renew an expired instance within a specific period of time, the instance is automatically released.
- If economical mode is disabled for a pay-as-you-go instance, you continue to be charged for the instance until it is released.
- You can enable instance release protection for a pay-as-you-go instance to prevent irreversible data loss caused by accidental release operations. For more information, see Enable or disable release protection for ECS instances.
- If the Release Disk with Instance feature is disabled for a disk attached to an instance, the disk is automatically converted into a pay-as-you-go data disk and retained when the instance is released. For more information, see Release a disk.

Manually release instances

You can manually release pay-as-you-go instances in the ECS console.

1.
2.
3.
4. Release one or more pay-as-you-go instances at a time.
   - If you want to release a single pay-as-you-go instance at a time, find the instance that you want to release and choose **More > Instance Status > Release** in the **Actions** column.

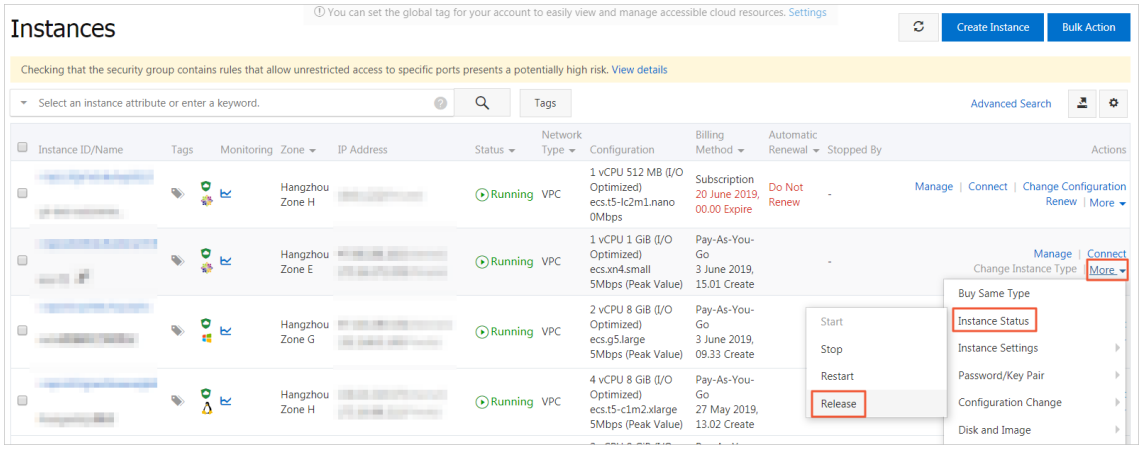

- If you want to release one or more pay-as-you-go instances at a time, click the Filter icon at the top of the **Billing Method** column and select Pay-As-You-Go from the drop-down list. In the displayed list of pay-as-you-go instances, select the instances that you want to release and then click **Release** below the instance list.

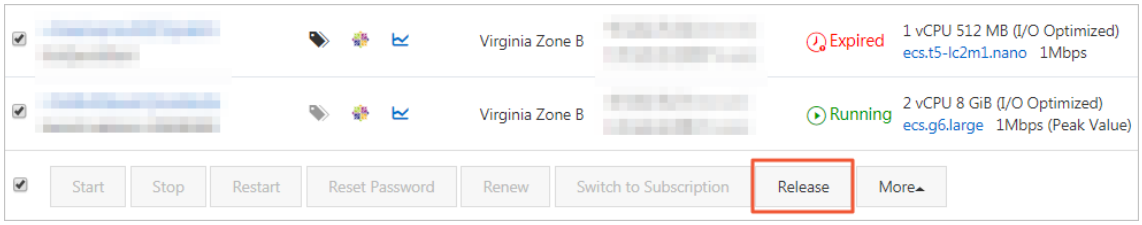

5. In the Release dialog box, select **Release Now**.
6. Click **Next**. Then, click **OK**.

Enable automatic release

You can enable automatic release for pay-as-you-go instances and set a time to automatically release the instances. If you set the automatic release time more than once, the most recent setting prevails.

1.
2.
3.
4. Configure automatic release for one or more pay-as-you-go instances at a time.
   - If you want to have a single pay-as-you-go instance automatically released at a time, find the instance that you want to release and choose **More > Instance Status > Release** in the **Actions** column.
   - If you want to have one or more pay-as-you-go instances automatically released at a time, click the Filter icon at the top of the **Billing Method** column and select Pay-As-You-Go from the drop-down list. In the displayed list of pay-as-you-go instances, select the instances that you want to release and then click **Release** below the instance list.
5. In the Release dialog box, select **Scheduled Release**.
6. Turn on **Automatic Release** and specify a date and time to release the selected instances.

> **Note** The automatic release time must be at least 30 minutes later than the current time and accurate to the minute.

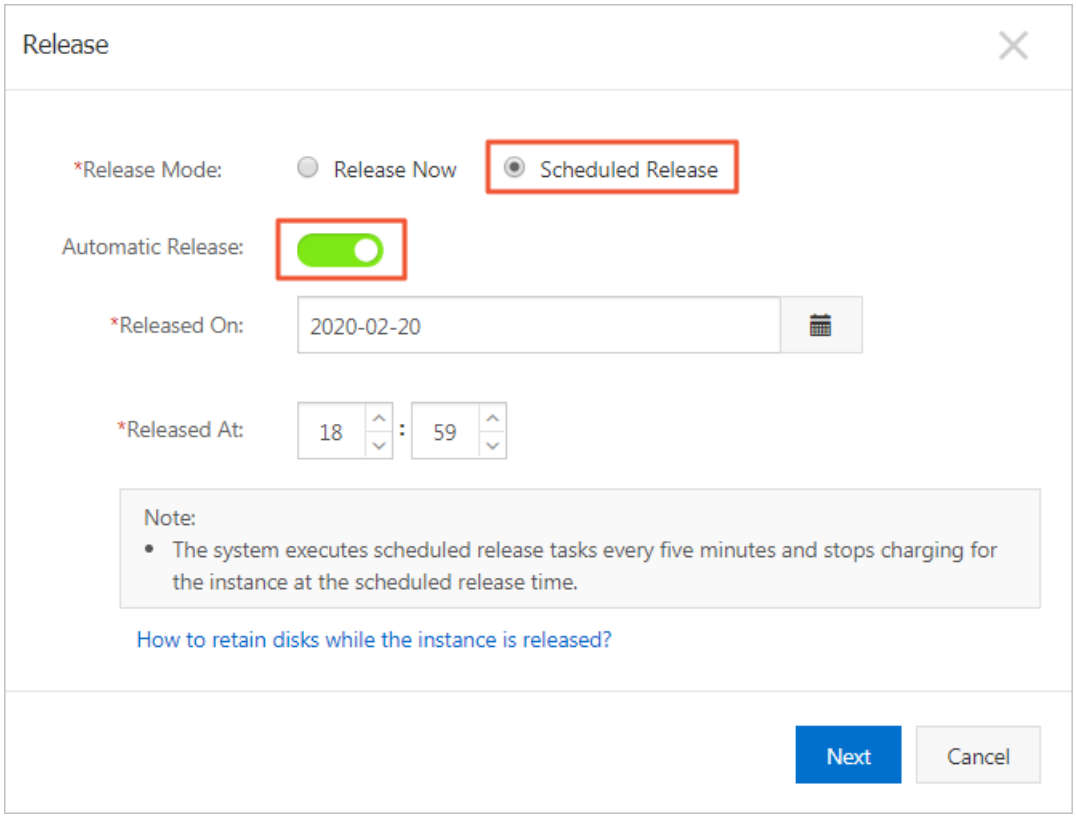

7. Click **Next**. Then, click **OK**.

Disable automatic release

1.
2.
3.
4. Disable automatic release for one or more pay-as-you-go instances at a time.
   - If you want to disable automatic release for a single pay-as-you-go instance at a time, find the instance for which you want to disable the automatic release feature and choose **More > Instance Status > Release** in the **Actions** column.
   - If you want to disable automatic release for one or more pay-as-you-go instances at a time, click the Filter icon at the top of the **Billing Method** column and select Pay-As-You-Go from the drop-down list. In the displayed list of pay-as-you-go instances, select the instances for which you want to disable automatic release and then click **Release** below the instance list.
5. In the Release dialog box, select **Scheduled Release**.
6. Turn off **Automatic Release**.
7. Click **Next**. Then, click **OK**.

Related information

References

- DeleteInstance
- ModifyInstanceAutoReleaseTime