

ALIBABA CLOUD

# Alibaba Cloud

专有宿主机  
Quick Start

Document Version: 20220216

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

# Table of Contents

1.Overview .....	05
2.Create a dedicated host .....	06
3.Create an ECS instance on a dedicated host .....	09

# 1. Overview

This topic describes how to create a subscription dedicated host (DDH) and Elastic Compute Service (ECS) instances on the dedicated host.

If you want to create a dedicated host and ECS instances, you must select the region, specifications, and billing method of the dedicated host. For more information, see [Billing overview](#).

For more information, see the following topics:

1. [Create a dedicated host](#)
2. [Create an ECS instance on a dedicated host](#)

## 2. Create a dedicated host

This topic describes how to create a subscription dedicated host in the Elastic Compute Service (ECS) console.

### Prerequisites

- An Alibaba Cloud account is created. For more information, see [Create an Alibaba Cloud account](#).
- A credit card or PayPal account is associated with your Alibaba Cloud account. For more information, see [Add a payment method](#).

### Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose **Instances & Images > Dedicated Hosts**.
3. In the top navigation bar, select a region.
4. On the Dedicated Hosts page, click **Create DDH**.
5. On the page that appears, set the following parameters:

i. **Billing Method**: Select **Subscription**.

ii. **Region**: Select a region and a zone where you want to create a dedicated host.

For example, to create a dedicated host in the China (Beijing) region, select **China (Beijing)** from the drop-down list. For more information, see [Regions and zones](#).

iii. **Dedicated Host Type**, **DDH Name**, and **Quantity**: Select a dedicated host type, enter a dedicated host name, and then specify the number of dedicated hosts that you want to purchase.

The dedicated host type determines the instance family and the maximum number of ECS instances that you can deploy on the dedicated host. The host types g6s, c6s, and r6s allow you to customize the vCPU-to-memory ratio. This allows you to flexibly allocate computing resources when you create ECS instances. For more information about dedicated host types, see .

 **Notice** ECS instances on local SSD i2 dedicated hosts do not support manual migration and automatic failover. If a local SSD i2 dedicated host fails, you can [submit a ticket](#) to apply for manual migration. However, the data on local disks will be lost after the migration.

iv. (Optional)Add tags.

You can assign tags to dedicated hosts and manage them in groups. For more information, see [Overview](#).

v. (Optional)Select a resource group.

You can separate dedicated hosts into different resource groups and set different access permissions for these dedicated hosts. For more information, see [Resource groups](#).

vi. In the **DDH Settings** section, set the following parameters.

Parameter	Description
<b>Allow Automatic Deployment</b>	<ul style="list-style-type: none"> <li>If you select Allow Automatic Deployment, ECS instances are automatically deployed on available dedicated hosts. For more information, see <a href="#">Automatic deployment</a>.</li> <li>If you do not select Allow Automatic Deployment, you must specify a dedicated host when you create an ECS instance.</li> </ul> <p>The Allow Automatic Deployment option is selected by default.</p>
<b>Automatic Instance Migration upon DDH Failure</b>	<ul style="list-style-type: none"> <li>If you select Automatic Instance Migration upon DDH Failure, ECS instances on a dedicated host are automatically migrated to another dedicated host if the original dedicated host fails.</li> <li>If you do not select this option, you must <a href="#">submit a ticket</a> to apply for a new dedicated host if the original dedicated host fails.</li> </ul> <p>The Automatic Instance Migration upon DDH Failure option is selected by default. You can modify this setting after the dedicated host is created. For more information, see <a href="#">Modify the settings of automatic instance migration for a dedicated host</a>.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> <b>Notice</b> This option is unavailable for local SSD i2 dedicated hosts.</p> </div>
<b>CPU Oversold Ratio</b>	<p>The CPU overcommit ratio affects the number of available vCPUs on a dedicated host. You can use the following formula to calculate the number of available vCPUs on a dedicated host: <math>\text{Number of available vCPUs} = \text{Number of physical CPU cores} \times 2 \times \text{CPU overcommit ratio}</math>. In scenarios where high CPU stability is not required, such as development and test environments, you can increase the number of available vCPUs on a dedicated host by increasing the CPU overcommit ratio. This way, you can deploy more ECS instances of the same specification on the dedicated host and reduce the unit deployment cost.</p> <div style="background-color: #e0f2f1; padding: 5px; border: 1px solid #ccc;"> <p> <b>Notice</b> You can configure CPU overcommit ratios only for the following dedicated host types: g6s, c6s, and r6s.</p> </div> <p>For example, the number of physical CPU cores on each g6s dedicated host is 52. If you set the CPU overcommit ratio of a g6s dedicated host to 4, the number of available vCPUs on the dedicated host is 416.</p>

vii. Set the **Duration** parameter and select the **Enable Auto-renewal** option based on your business requirements.

viii. Read and agree to the **Dedicated Host Terms of Service**.

ix. In the lower-right corner, confirm the cost next to **Total** and click **Preview**.

6. In the **Preview** dialog box, confirm the configurations and click **Create Order**.

7. On the page that appears, follow the instructions to complete the payment.

## Result

You can view the created dedicated host on the Dedicated Hosts page. If the dedicated host is in the **Running** state, you can use the dedicated host. If the created dedicated host does not appear on the Dedicated Hosts page, wait for a while and refresh the page.

## What's next

You can perform the following operations:

- [Create an ECS instance on a dedicated host](#)
- [Migrate an ECS instance from a shared host to a dedicated host](#)

## 3. Create an ECS instance on a dedicated host

This topic describes how to create and configure an Elastic Compute Service (ECS) instance on a dedicated host.

### Prerequisites

- A dedicated host is created. For more information, see [Create a dedicated host](#).
- The real-name verification is complete if you need to create an ECS instance in the Chinese mainland. For more information, see [Real-name Registration FAQs](#).
- An IPv4 virtual private cloud (VPC) is created in the region where the dedicated host resides. For more information, see [Create an IPv4 VPC](#).
- (Optional) A security group is created in the same region as the dedicated host and rules are added to the security group if you do not use the default security group. For more information, see [Create a security group](#) and [Add security group rules](#).
- (Optional) An SSH key pair is created in the region where the dedicated host resides if you need to associate the SSH key pair with a Linux ECS instance. For more information, see [Create an SSH key pair](#).
- (Optional) The user data of ECS instances are prepared if you need to configure the user data. For more information, see [Overview of ECS instance user data](#).
- (Optional) A Resource Access Management (RAM) role is created for an ECS instance and authorized to access the ECS instance if you allow the instance to assume a RAM role. For more information, see [Attach an instance RAM role](#).

### Context

Only VPC-connected ECS instances can be created on a dedicated host. The ECS instances on a dedicated host have different features from the ECS instances on a shared host. For more information, see [Differences between ECS instances on a dedicated host and those on a shared host](#).

 **Note** We recommend that you create a pay-as-you-go ECS instance on a dedicated host and select a combination of billing methods to optimize your costs without compromising flexibility. For more information, see [Resource billing for ECS instances on a dedicated host](#).

### Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose **Instances & Images > Dedicated Hosts**.
3. In the top navigation bar, select a region.
4. Find the dedicated host on which you want to create an ECS instance. Click **Create Instance** in the **Actions** column.
5. Perform the following operations in the **Basic Configurations** step:
  - i. Specify **Dedicated Host**.

By default, the dedicated host that you previously selected on the Dedicated Hosts page is used. You can select another dedicated host.

- ii. Select or clear **Associate with DDH**. This option specifies whether the instance is still deployed on the current dedicated host if the instance is stopped, released, or restarted.
  - If you select **Associate with DDH**, the instance is deployed on the current dedicated host. If the resources of the dedicated host are insufficient, the instance fails to be restarted.
  - If you clear **Associate with DDH**, the instance is deployed on another dedicated host of your Alibaba Cloud account when the resources of the current dedicated host are insufficient. The dedicated host must be available for automatic instance deployment. For more information, see [Features](#).

iii. Specify **Billing Method**.

Select a billing method for the ECS instances based on the billing method of the dedicated host. You can select **Subscription** or **Pay-As-You-Go** as the billing method of ECS instances that run on a subscription dedicated host.

iv. Specify the type and quantity of the ECS instances.

The region and zone of the ECS instance must be the same as those of the dedicated host. The available instance types depend on the type and available computing resources of the dedicated host. For more information, see [Features](#).

- If the dedicated host supports only predefined instance types, click **Predefined Instance Type**. Then, select the instance type in the Instance Type section.

Family	Instance Type	vCPUs	Memory	Clock Speed	Internal Network Bandwidth	Packet Forwarding Rate	IPv6-supported	Physical Processor
Compute Type c5	ecs.c5.large	2 vCPUs	4 GiB	2.5 GHz/2.7 GHz	1 Gbps	300,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.xlarge	4 vCPUs	8 GiB	2.5 GHz/2.7 GHz	1.5 Gbps	500,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.2xlarge	8 vCPUs	16 GiB	2.5 GHz/2.7 GHz	2.5 Gbps	800,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.3xlarge	12 vCPUs	24 GiB	2.5 GHz/2.7 GHz	4 Gbps	900,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.4xlarge	16 vCPUs	32 GiB	2.5 GHz/2.7 GHz	5 Gbps	1,000,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.6xlarge	24 vCPUs	48 GiB	2.5 GHz/2.7 GHz	7.5 Gbps	1,500,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.8xlarge	32 vCPUs	64 GiB	2.5 GHz/2.7 GHz	10 Gbps	2,000,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY
Compute Type c5	ecs.c5.16xlarge	64 vCPUs	128 GiB	2.5 GHz/2.7 GHz	20 Gbps	4,000,000 PPS	Yes	Intel Xeon(Skylake) Platinum 8163 / Intel Xeon(Cascade Lake) Platinum 8269CY

DDH Resources: vCPUs: 84 vCPUs available/86 vCPUs in total; Memory: 168 GiB available/172 GiB in total

Selected Instance Type: ecs.c5.large (2 vCPU 4 GiB, Compute Type c5)

Quantity: 1 Units. 2 vCPUs have been enabled, and 84 more vCPUs can be enabled. The selected instance type occupies 2 vCPUs. You can create a maximum of 42 more ECS instances.

- If the dedicated host supports custom instance types, click **Custom Instance Type**. Then, you can adjust the sliders to specify the number of vCPUs and the size of the memory. The minimum scaling step size of vCPUs is 1. If more than one vCPU is required, you must set this parameter to an even number, for example, 2 or 4. The minimum scaling step size of the memory is 1 GiB.

Instance Type: Predefined Instance Type | Custom Instance Type

vCPU: 1 Cores | 25 Cores | 50 Cores | 75 Cores | 104 Cores. Slider set to 2. Cores

Memory: 1 GiB | 92 GiB | 184 GiB | 276 GiB | 372 GiB. Slider set to 4. GiB

DDH Resources: vCPUs: 104 vCPUs available/104 vCPUs in total; Memory: 372 GiB available/372 GiB in total

Selected Instance Type: ecs.ddh6s.custom.c2m4 (2 vCPU 4 GiB)

Quantity: 1 Units. 0 vCPUs have been enabled, and 104 more vCPUs can be enabled. The selected instance type occupies 2 vCPUs. You can create a maximum of 52 more ECS instances.

v. Select an image.

You can select a public image, custom image, shared image, or an image that is purchased from Alibaba Cloud Marketplace. For more information, see [Select an image](#).

 Note

- To associate an SSH key pair with the ECS instance, you must select a Linux operating system from the drop-down list for the image.
- To configure user data of the instance, you must select an image that supports user data. For more information, see [Manage the user data of Linux instances](#).

vi. Configure the storage.

- **System Disk:** Required. The system disk is used to install the operating system. You must specify the type and capacity of the system disk.
  - **Disk Type:** All available types of disks in the current region are listed in this section.
  - **Capacity:** The default capacity of the system disk is 40 GiB. The maximum capacity is 500 GiB. If the size of the selected image is larger than 40 GiB, the image size is the default value. The minimum size of the system disk is related to the image. The actual size is displayed on the buy page.

Image	System disk size (GiB)
Linux (excluding CoreOS and Red Hat Enterprise Linux)	[max{20, Image size}, 500]
FreeBSD	[max {30, Image size}, 500]
CoreOS	[max {30, Image size}, 500]
Red Hat	[max {40, Image size}, 500]
Windows	[max {40, Image size}, 500]

- **Data Disk:** To add data disks, you must specify the disk type, capacity, and quantity. You must also specify whether to enable **disk encryption**. You can create an empty data disk or create a data disk from a snapshot.

 **Note** In this case, the data disk has the following features:

- The billing method of the data disk is the same as that of the instance.
- When you enable or disable the Release Disk with Instance feature for a system disk or a data disk, take note of the following information:
  - If the Release Disk with Instance feature is enabled for the disk, the disk is automatically released when its associated instance is released.
  - If the Release Disk with Instance feature is disabled for the disk, the disk is retained as a pay-as-you-go data disk 15 days after its associated instance expires, 15 days after a payment becomes overdue for the instance, or when the instance is manually released. If you create a disk in the Chinese mainland, you must complete real-name verification for your account to ensure that the disk can be retained.

 **Note** The retained disk is billed on a pay-as-you-go basis. You can log on to the Billing Management console and view consumption details by disk ID.

- A maximum of 16 data disks can be attached to a single instance. For more information, see the "Elastic Block Storage (EBS) limits" section in [Limits](#).

- If you create an instance from an instance family with local SSDs, such as the i2 type, the information of the local SSDs is displayed and cannot be modified. For more information about the local SSDs of an instance family, see [Instance family](#).

vii. Click **Next: Network and Security Group**.

6. Specify the parameters in the **Network and Security Group** step.

i. Specify **Network Type**.

You can select only **VPC** as the network type. Select a VPC and vSwitch from the drop-down lists. If no VPC or vSwitch is created, you can use the default VPC or vSwitch.

ii. Configure the public bandwidth.

- If you want to assign a public IP address to the instance, you must select **Assign Public IP Address**. Then, select **Pay-By-Traffic** or **Pay-By-Bandwidth** as the billing method of the public bandwidth, and specify the bandwidth. The public IP address cannot be detached from the instance. For more information about the billing method of public bandwidth, see [Public bandwidth](#).
- If your instance does not need to access the Internet or your VPC-type instance uses an elastic IP address (EIP) to access the Internet, you do not need to assign a public IP address. You can associate an EIP with or disassociate an EIP from an instance at any time.

iii. Specify **Security Group**.

Select a security group whose rules meet your business requirements.

 **Note** If you have not created a security group, you can use the default security group. For information about the rules of the default security group, see [Default security groups](#).

iv. Specify **Elastic Network Interface**.

If the selected instance type supports Elastic Network Interface (ENI), you can add an ENI and select a vSwitch for the ENI.

 **Note** By default, the ENI is released together with the instance. You can also call the [DetachNetworkInterface](#) operation to detach the ENI from the instance.

v. Click **Next: System Configurations**.7. (Optional) Specify the parameters in the **System Configurations** step.i. Select and configure **Logon Credentials**.

You can specify a key pair and password. You can also select **Set Later** to configure the logon credential after the instance is created. Select a logon credential based on the operating system of the image.

- Linux: You can use a key pair or password as the logon credential.
- Windows: You can use only a password as the logon credential.

ii. Specify **Instance Name** and **Host**.iii. Set the parameters in the **Advanced** section.

- RAM Role: You can assign a RAM role to the instance.
- User Data: You can customize the startup behavior of the instance or transfer data to the instance.

iv. Click **Next: Grouping**.8. (Optional) Specify the parameters in the **Grouping** step.

- If you have multiple instances, we recommend that you tag the instances for easy identification.
- If you are an enterprise user, and you have activated Resource Management and created resource groups, you can manage instances by using the resource groups.

## 9. Confirm the order and create the ECS instances.

i. In the **Configurations Selected** section, confirm the configurations of the instances.

You can click the



icon to modify the configurations.

- ii. Configure the release or renewal of the ECS instances based on the billing method.
  - If the billing method of the instances is **Subscription**, you must specify the duration, and specify whether to turn on the **Enable Auto-renewal** switch.

 **Note** The expiration date of the subscription instances cannot be later than that of the subscription dedicated host.

- If the billing method of the instances is **Pay-As-You-Go**, you can select **Automatic Release** and specify the time to automatically release the instances.
- iii. Confirm the **Total** cost in the lower-right corner of the page. The total cost includes the cost of the image, data disks, and public bandwidth.
  - iv. Read and accept the **ECS Terms of Service**.
  - v. Proceed based on the billing method of the ECS instances:
    - If the billing method of the ECS instances is **Subscription**, click **Create Order**.
    - If the billing method of the ECS instances is **Pay-As-You-Go**, click **Create Instance**.

## Result

After the instances are created, click **Console** in the upper-right corner of the page to go back to the ECS console. On the **Instances** page, select the region where the created instances reside. Then, you can view the instance ID, public IP address, and private IP address of the created instances. To view the information of the dedicated host where each ECS instance is created, perform the following operations: Click the **Column Filters** button. In the Column Filters dialog box, select **Dedicated Host** and click **OK**.

 **Note** For more information about **Column Filters**, see [Migrate ECS instances between different dedicated hosts in the DDH console](#).

## What's next

- You can build an FTP site on the ECS instance to upload files to the instance. For more information about how to deploy the FTP service, see [Build an FTP site on an ECS instance](#).
- After the ECS instances are created, we recommend that you check whether the operating systems are compliant and hardened.
  - [Harden operating system security for Linux](#)
  - [Harden operating system security for Windows](#)
- If you added data disks when you created the instance, you must format and partition the data disks before they can be used. For more information, see [Partition and format a data disk on a Linux instance](#) or [Partition and format a data disk on a Windows instance](#).

## Related information

- [RunInstances](#)