

Alibaba Cloud Elastic Compute Service

Network

Issue: 20190909

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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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 contact Alibaba Cloud directly if you discover any errors in this document.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus, page names, and other UI elements.	Click OK .
<code>Courier font</code>	It is used for commands.	Run the <code>cd / d C :/ windows</code> command to enter the Windows system folder.
<i>Italics</i>	It is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	It indicates that it is an optional value, and only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
<code>{}</code> or <code>{a b}</code>	It indicates that it is a required value, and only one item can be selected.	<code>swich {stand slave}</code>

Contents

Legal disclaimer.....	I
Generic conventions.....	I
1 Network types.....	1
2 Instance IP addresses.....	3
2.1 IP addresses of VPC-Connected ECS instances.....	3
2.2 IP addresses of a classic network-connected ECS instance.....	4
2.3 Intranet.....	6
3 Change IPv4 addresses.....	9
3.1 Change the private IP of an ECS instance.....	9
3.2 Change the public IP address of an ECS instance.....	9
3.3 Convert an ECS public IP address in a VPC to an Elastic IP Address.....	11
4 Elastic Network Interfaces.....	14
4.1 ENI overview.....	14
4.2 Create an ENI.....	16
4.3 Attach an ENI.....	18
4.4 Configure an ENI.....	20
4.5 Assign a secondary private IP address.....	24
4.6 Revoke a secondary private IP address.....	29
4.7 Modify an ENI.....	30
4.8 Detach an ENI from an instance.....	32
4.9 Delete an ENI.....	33
5 Multiqueue for NICs.....	35
6 Connect a classic network to a VPC.....	38
7 Network FAQ.....	40

1 Network types

Alibaba Cloud provides classic network and Virtual Private Cloud (VPC) network types.

Virtual Private Cloud (VPC)

VPCs are logically isolated networks established in Alibaba Cloud. You can customize the topology and IP addresses in a VPC. We recommend that the VPC network type is used if you have high network security requirements.

For more information about VPC, see [Virtual Private Cloud](#) documentation.

Classic network

A classic network is deployed in the public infrastructure of Alibaba Cloud, which is responsible for its planning and management. We recommend that the classic network type is used if your business requirements are high in terms of network usability.



Note:

If you purchased an ECS instance after 17:00 (UTC+8) on June 14, 2017, you cannot choose the classic network type.

VPC vs. Classic networks

The following table describes key network functions and indicates whether they are supported within VPCs and classic networks.

Items	VPC	Classic network
Two-layer logic isolation	Supported	Not supported
Custom private network blocks	Supported	Not supported
Private IP addresses	Unique within one VPC. Replicable between VPCs.	Unique in the global Classic network
Communicate within or between private networks	Able to communicate within a VPC, but isolated between VPCs	Able to communicate in one region and under one account
Tunneling	Supported	Not supported
Custom router	Supported	Not supported

Items	VPC	Classic network
Routing table	Supported	Not supported
Switches	Supported	Not supported
SDN	Supported	Not supported
Self-built NAT gateway	Supported	Not supported
Self-built VPN	Supported	Not supported

2 Instance IP addresses

2.1 IP addresses of VPC-Connected ECS instances

Each VPC-Connected ECS instance can communicate within an intranet by using a private IP address, or communicate over the Internet by using a public IP address.

Private IP addresses

Each VPC-Connected ECS instance is assigned a private IP address when it is created. That address is determined by the VPC and the CIDR block of the VSwitch to which the instance is connected.

Scenarios

A private IP address can be used in the following scenarios:

- Load balancing
- Communication among ECS instances within an intranet
- Communication between an ECS instance and other cloud products (such as OSS and RDS) within an intranet

For more information, see [#unique_7](#).

Modify a private IP address

You can modify the private IP address of a VPC-Connected ECS instance in the ECS console. For more information, see [Change the private IP of an ECS instance](#).

Public IP addresses

VPC-Connected ECS instances support two public IP address types:

- NatPublicIp, which is assigned to a VPC-Connected ECS instance, can be released only, and cannot be disassociated from the instance.
- Elastic public IP (EIP). For more information, see [What is an EIP address](#).

When a VPC-Connected ECS instance accesses the Internet, its public IP address is mapped to its private IP address through network address translation (NAT).

You cannot find a network interface for Internet access by running commands within the operating system.

Scenarios

NatPublicIp and EIP are applicable to different scenarios:

- **NatPublicIp:** If you want to assign a public IP address to a VPC-Connected ECS instance when creating the instance, and do not want to retain the public IP address when the instance is released, you can use a NatPublicIp address.
- **EIP:** If you want to keep a public IP address and associate it to any of your VPC-Connected ECS instances in the same region, you can use an EIP address.

Obtain a public IP address

- **NatPublicIp:** When creating a VPC-Connected ECS instance, if you select Assign a public IP, a NatPublicIp is assigned to the instance when it is created.
- **EIP:** You can apply for an EIP address and bind it to a VPC-Connected ECS instance. In this case, do not assign a NatPublicIp to an instance. For more information, see [Apply for an EIP address](#).

Release a public IP address

- **NatPublicIp:** When a NatPublicIp address is assigned to an instance, you can only release the IP address, but cannot disassociate it. Only a NatPublicIp address that is assigned to a Subscription instance can be released. For more information, see [#unique_11](#).
- **EIP:** If you do not need an EIP address, disassociate it from a VPC-Connected ECS instance and release it in the EIP console. For more information, see [Unbind and release an EIP address](#).

Billing

You are billed for outbound Internet traffic usage only. For more information, see [#unique_13](#).

2.2 IP addresses of a classic network-connected ECS instance

Currently, for ECS instances of the classic network type, IP addresses are distributed in a unified way and divided into public and private IP addresses. Private IP address

are mainly used for remote access to your instance or to the services deployed on your instance.

Intranet IP addresses

Each classic network-connected ECS instance is assigned a private, that is intranet, IP address.

Scenarios

Intranet IP addresses can be used in the following scenarios:

- Load balancing
- Mutual intranet access between ECS instances
- Mutual intranet access between ECS instances and other cloud services, such as OSS and RDS

Traffic generated through intranet IP addresses within an intranet is free of charge. For more information, see [Intranet](#).

Modify an intranet IP address

Once a classic network-connected ECS instance is created, you cannot change its intranet IP address.



Note:

Do not change an intranet IP address within a guest operating system. Otherwise, communication within an intranet is interrupted.

Public IP addresses

If you purchase bandwidth for Internet access, a public IP address is assigned to your classic network-connected ECS instance. You cannot change the public IP address once it is assigned.

Scenarios

A public IP address is used in the following scenarios:

- Mutual access between an ECS instance and the Internet
- Mutual Internet access between ECS instances and other Alibaba Cloud services

Assign a public IP address

When you create an ECS instance, a public IP address is assigned to it if Assign public IP is selected.

For a Subscription instance with no public IP address, you can use the [Upgrade Configuration](#) or the [Renew for Configuration Downgrade](#) feature to purchase public network bandwidth.



Note:

- For a Pay-As-You-Go classic network-connected ECS instance with no public IP address, you cannot assign a public IP address after the instance is created.
- For a classic network-connected ECS instance, you cannot disassociate or release its public IP address once the IP address is assigned. If you set the bandwidth to 0 Mbit/s when renewing an instance for configuration downgrade, in the next purchase cycle, the public IP address is retained, but the instance cannot access the Internet.

Billing

You are billed for usage of Internet outbound traffic only. For more information, see [Billing of network bandwidth](#).

Multicast and broadcast

Intranet IP addresses cannot be used for multicasting or broadcasting.

2.3 Intranet

If you need to transmit data between two ECS instances in the same region, use an intranet connection. Intranet connections can also be used to connect any combination of ECS, RDS, SLB, and OSS if they are deployed in the same region. However, the network speed is limited to one gigabit of shared bandwidth for non I/O optimized instances.

Alibaba Cloud instances can communicate over an intranet. The instances use one gigabit of shared bandwidth for non I/O optimized instances, and 10 gigabits of shared bandwidth for I/O optimized instances, with no special restrictions. However, because the intranet is a shared network, the bandwidth may fluctuate.

The following table describes how to enable intranet communication between ECS instances across different network types, depending on the number of accounts and whether the target regions and security groups are the same or different.

Network type	Accounts used	Regions	Security groups	How to enable intranet communication
VPC, same VPC	One account or multiple accounts	Same	Same	Enabled by default.
			Different	Authorize security groups for each other.
VPC, different VPCs	One account or multiple accounts	Same	Either the same or different	Use Express Connect. For more information, see Application scenarios from Product Introduction to Express Connect .
		Different	Different	
Classic	One account	Same	Same	Enabled by default.
	Multiple accounts		Either the same or different	Authorize security groups for each other. For more information, see Scenarios of security groups .

Private IP addresses are used for intranet communication. You cannot [change the private IP address](#) of an instance of the Classic network type, but you can change the private IP address of a VPC-Connected ECS instance. Private and public addresses of ECS instances do not support virtual IP (VIP) configuration.

By default, instances of different network types cannot communicate with one another in one intranet. However, VPC provides the [ClassicLink](#) function, which allows you to link an ECS instance in the classic network to cloud resources in a VPC through the intranet.

3 Change IPv4 addresses

3.1 Change the private IP of an ECS instance

After creating an ECS instance in a VPC network, you can change the private IP address and can change the VSwitch of the ECS instance.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, click Instances.
3. Select the target region.
4. In the Actions column, click More > Instance Status > Stop.
5. When the instance is stopped, click the instance ID to go to its Instance Details page.
6. In the Configuration Information panel, click More > Modify Private IP Address.
7. In Modify Private IP Address dialog, select a VSwitch, and then click Modify.

Make sure the current VSwitch and the selected VSwitch are in the same zone.



Note:

Enter a new IP address if you do not want to change the VSwitch of the ECS instance.

8. Go back to the instance page and, in the Actions column, click More > Instance Status > Restart to make the new private IP address take effect.

3.2 Change the public IP address of an ECS instance

If your ECS instance is assigned a public IP address, you can change the IP address within six hours after the instance is created regardless of whether the instance is in a classic network or in a VPC network.

Limits

- You can change the public IP address of an instance a maximum of three times.

- Changes to a public IP address must be made within six hours after the corresponding instance is created.

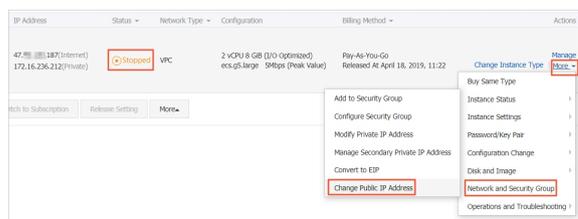
Prerequisites

- The instance must be in the Stopped state. For information about how to start or stop an instance, see [#unique_25](#).
- The instance must be assigned a public IP address. You can view the IP address in the IP Address column of the Instance List in the ECS console.

Procedure

To change the public IP address, follow these steps:

1. Find the target instance and then choose More > Network and Security Group > Change Public IP Address in the Actions column.



2. Click Start Now.

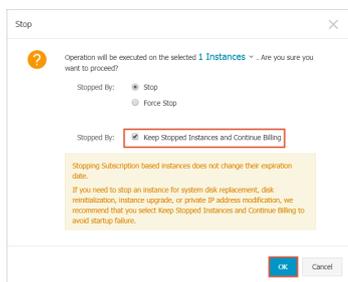
A new public IP address is displayed.

3. Click OK.

FAQ

- Can I change the public IP address of an instance if it was created more than six hours ago?
 - For instances in a VPC network, the public IP address can be converted into an Elastic IP Address (EIP). For information about how to convert the IP address, see [#unique_26](#).
 - For instances in a classic network, the public IP address cannot be changed if an instance has been created for more than six hours.

- Why is the Change Public IP Address option under Network and Security Group not displayed?
 - By default, the Change Public IP Address option is not displayed for instances that were created more than six hours ago.
 - If you enable the [no fees for stopped VPC instances](#) feature for an instance, make sure that this feature is disabled when you stop the instance. Otherwise, the instance will be temporarily released and the Change Public IP Address option is not displayed. You can disable this feature by selecting the Keep Stopped Instances and Continue Billing check box in the Stop dialog box.



- Can I change the private IP address of an instance?
 - This operation is allowed for instances in a VPC network. For information about how to change the private IP address, see [#unique_27](#).
 - This operation is not allowed for instances in a classic network.
- How do I obtain the public IP address of a created instance if the IP address is not assigned during instance creation?
 - For a Subscription instance, you can obtain the public IP address by upgrading or downgrading the network bandwidth configuration. For more information, see [#unique_28](#).
 - For a Pay-As-You-Go instance, you cannot obtain the public IP address after the instance is created. You can only [bind an EIP address](#).

3.3 Convert an ECS public IP address in a VPC to an Elastic IP Address

This topic describes how to convert the public IP address of an ECS instance in a VPC to an Elastic IP Address (EIP). After conversion, the public IP address is retained. You can keep the public IP address bound to the instance, or bind it to another instance. Converting the public IP address does not affect public access to your ECS instance, and does not cause transient traffic interruptions.

Prerequisites

Before converting a public IP address to an EIP, make sure the following requirements are met:

- The ECS instance has been assigned a public IP address.
- If the ECS instance is a subscription instance, you cannot convert the address within 24 hours prior to instance expiration.
- If the ECS instance is a pay-as-you-go instance, your account must not have any overdue payments.
- If the instance type of the ECS instance has been changed, wait until the change takes effect.
- You have stopped the ECS instance or the ECS instance is in the Running state. For more information, see [#unique_25](#).

Context

After the public IP address is converted to an EIP:

- The billing method of public bandwidth remains unchanged.
- The EIP is billed separately. For more information about EIP billing, see [EIP pricing](#). You can go to the Billing Management page, select [Usage Records](#), and select Elastic IP to export EIP usage records.

This section describes how to convert the public IP address of an ECS instance in a VPC to an EIP by using the ECS console. You can also convert the IP address by calling the [ConvertNatPublicIpToEip](#) operation. To call this operation, use SDK 4.3.0 or later. [Download](#) the latest SDK.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Instances & Images > Instances.
3. In the top navigation bar, select a region.
4. Find the target instance with the VPC network type, and then choose More > Network and Security Group > Convert to EIP from the Actions column.
5. In the dialog box that appears, confirm the information, and click OK.
6. Refresh the instance list.

Result

After the public IP address is converted to an EIP, the original IP address is followed by (EIP).

You can click the EIP to go to the [EIP console](#) to manage the EIP.

What's next

After the public IP address is converted to an EIP, you can:

- Unbind the EIP from the instance and bind it to another instance, or release it. For more information, see [#unique_12](#).
- [Add EIPs to an Internet Shared Bandwidth instance](#) to save costs. For more information, see [Select a product to gain access to the Internet](#) and [How to save Internet cost](#).

More information

[#unique_36](#)

4 Elastic Network Interfaces

4.1 ENI overview

An Elastic Network Interface (ENI) is a virtual network interface that can be attached to an ECS instance in a VPC. You can use an ENI to deploy a high-availability cluster, and perform low-cost failover and fine-grained network management in all Alibaba Cloud regions.

Scenarios

ENIs are suitable for:

- Deploying a high-availability cluster

An ENI is suitable for high-availability architecture for multiple network interfaces on a single instance.

- Providing a low-cost failover solution

You can detach an ENI from a failed ECS instance and then attach it to another ECS instance to quickly redirect traffic from the failed instance to a backup instance, thereby quickly restoring your services.

- Managing the network with refined controls

You can configure multiple ENIs for an instance in any Alibaba Cloud region. For example, you can use some ENIs for internal management and other ENIs for Internet business access, so as to isolate confidential data from business data. You can also configure specific security group rules for each ENI based on the source IP address, protocols, ports, and more to achieve secured traffic control.

ENI types

ENIs are classified into two types:

- Primary ENI

The ENI created by default upon the creation of an instance in a VPC. The life cycle of the primary ENI is the same as that of the instance, and you cannot remove the primary ENI from the instance.

- **Secondary ENI**

You can create a secondary ENI and attach it to an instance or detach it from the instance. The maximum number of ENIs that you can attach to one instance varies with the instance type. For more information, see [#unique_39](#).

ENI attributes

The following table describes ENI attributes.

Attribute	Quantity
Private IP address	Varies with instance types
MAC address	1
Security group	1 to 5
ENI name	1

Limits

ENIs have the following limits:

- There is an upper limit on the number of ENIs that can be created for one account in each region. For more information, see [Limits on ENIs](#).
- The ECS instance and its attached secondary ENI must be in the same zone and region, but can be in different VSwitches and security groups.
- The number of secondary ENIs that can be attached to an ECS instance depends on the instance type. For more information, see [#unique_39](#).
- Only I/O-optimized instance types support ENIs.
- ECS instances in a classic network do not support ENIs.
- The instance bandwidth varies with the instance type. You cannot increase the bandwidth of an ECS instance by attaching multiple ENIs to the instance.

Related operations

For images that cannot identify secondary ENIs, log on to the instance to [configure the ENI](#).

Console operations

In the ECS console, you can view information of an attached ENI. You can also perform the following actions with a secondary ENI only (a primary ENI is not supported):

- [#unique_42](#).
- [#unique_43](#).
- [Delete an ENI](#).
- [#unique_45](#).
- [#unique_46](#).

API operations

You can call [DescribeNetworkInterfaces](#) to query an ENI list, and call [DescribeInstances](#) to query the information of a specific ENI attached to an instance. Additionally, you can call the following API actions as needed for a secondary ENI only (a primary ENI is not supported):

- [CreateNetworkInterface](#)
- [DeleteNetworkInterface](#)
- [AttachNetworkInterface](#)
- [DetachNetworkInterface](#)
- [ModifyNetworkInterfaceAttribute](#)

4.2 Create an ENI

This topic describes how to create an elastic network interface (ENI) in the ECS console. You can use an ENI to deploy a high-availability cluster, and perform low-cost failover and fine-grained network management.

Background information

You can create an ENI by using either of the following two methods:

- [Attach an ENI when you create an instance](#). For more information, see [#unique_42](#). You can attach a maximum of two ENIs. One is the primary ENI and the other is the secondary ENI. A secondary ENI created in this way will be released with the instance if it is not detached from the instance. For information about how to detach an ENI, see [#unique_45](#).
- [Create a separate ENI](#). The created ENI can be attached to an instance. For more information, see [#unique_42](#). An ENI created in this way can only be used as a secondary ENI.

Limits

Before you create an ENI, note the following limits:

- Each ENI must be in a VSwitch of a VPC.
- Each ENI must belong to at least one security group.

Prerequisites

- A VPC and a VSwitch are created in the VPC.
- A security group is created in the same VPC.

Procedure

To create an ENI, follow these steps:

1. Click Create ENI.
2. In the displayed dialog box, complete the following configurations:
 - a. **Network Interface Name:** Enter a name for the ENI.
 - b. **VPC:** Select a VPC. When you attach an ENI to an instance, they must be in the same VPC.



Note:

After an ENI is created, you cannot change the VPC.

- c. **VSwitch:** Select a VSwitch. When you attach an ENI to an instance, they must be in the same zone, but they do not have to be in the same VSwitch.



Note:

After an ENI is created, you cannot change the VSwitch.

- d. **Primary Private IP:** Specify an IPv4 address as the private IP address of the ENI. The IPv4 address must be available in the CIDR block of the specified VSwitch. If you do not specify one, a private IP address is automatically assigned to your ENI after the ENI is created.
- e. **Security Group:** Select a security group in the selected VPC.
- f. **Description:** Optional. Enter a description for the ENI.
- g. Click OK.

On the Network Interfaces page, refresh the table. When the new ENI is in the Available state, it is created.

What to do next

After you create an ENI, you can:

- [Attach an ENI to an instance.](#)

- [Modify attributes of the ENI.](#)
- [Delete the ENI.](#)

4.3 Attach an ENI

This topic describes how to attach an Elastic Network Interface (ENI). Specifically, you either attach an ENI when you create an ECS instance, or you can alternatively create an ENI separately and then attach it to an ECS instance. Attaching an ENI allows you to build clusters with higher availability, perform failovers with lower costs, and manage your network with finer granularity.

Attach an ENI when you create an ECS instance

Limits

If you attach a secondary ENI, as opposed to a primary ENI, to an ECS instance and do not detach it from the ECS instance, the secondary ENI will be released when you release the ECS instance. For more information, see [Detach an ENI from an instance.](#)

Procedure

Before you begin, make sure that you have created an ECS instance. For the specific procedure, see [Step 2: Create an instance.](#)

When you attach an ENI to an ECS instance during the process of creating an ECS instance, configure the following parameters:

1. Basic configurations

- **Region:** ENIs are supported in all regions.
- **Instance type:** Select an I/O-optimized instance type that supports ENIs. For more information, see [Instance type families.](#)
- **Image:** The following image types support ENIs without any manual configuration required:
 - CentOS 7.3 64-bit
 - CentOS 6.8 64-bit
 - Windows Server 2016 Datacenter Edition 64-bit
 - Windows Server 2012 R2 Datacenter Edition 64-bit



Note:

For other image types, after you create an ECS instance, you must configure the ENI to enable the instance to support ENIs.

2. Networking

- **Network:** Select VPC, and then select a VPC and VSwitch that you created.
- **ENI:** Click Add ENI to attach the target ENI. The ENI and the instance must belong to the same VSwitch.



Note:

When you create an instance in the ECS console, you can attach up to two ENIs to the instance. One is the primary ENI, and the other is the secondary ENI. You can attach more secondary ENIs to the instance by using one of the following two methods:

- [Create an ENI](#) in the ECS console, and then [attach the ENI](#) to the instance.
- Call the API action [AttachNetworkInterface](#) to attach more ENIs to the instance.

Attach an ENI to an existing ECS instance

Limits

- The ENI can only be attached to the existing ECS instance as a secondary ENI, rather than a primary ENI.
- The ENI must be in the Available state.
- The ECS instance must be in the Stopped or Running state.
- The ENI can only be attached to a VPC ECS instance. The ENI and the instance must be in the same VPC.
- The VSwitch to which the ENI belongs must be in the same zone as the ECS instance to which the ENI is attached.
- The ENI can only be attached to an I/O-optimized instance.
- One ENI can be attached to only one VPC ECS instance, but one instance can be attached with multiple ENIs. For more information, see [Instance type families](#).

Prerequisites

- An ENI is created. For more information, see [Create an ENI](#).
- The ENI is in the Available state.

- The instance can be attached with secondary ENIs and is in the Stopped or Running state. For more information, see [Instance type families](#).

Procedure

1. Locate an available ENI, and then click Bind to Instance.
2. In the displayed dialog box, select the target instance, and then click OK.

Refresh the list. When the ENI is in the Bound state, the ENI is attached to the instance.



Notice:

If the last time your instance was started or restarted is earlier than April 1, 2018, then you must use the ECS console or call the API action [RebootInstance](#) to [Restart the instance](#), as opposed to logging on to the instance to restart it. Otherwise, the ENI cannot be attached to the instance.

What to do next

After you attach an ENI to an ECS instance, you can perform the following operations:

- [Detach the ENI from the instance](#) or [Delete the ENI](#).
- [Configure the ENI](#) if the image cannot identify the ENI.

4.4 Configure an ENI

This topic describes how to configure an ENI. For some images used by your instances, you may need to manually configure an ENI for these images so that the ENIs attached to your instances can be identified by the system.

Background information

If your instance is running one of the following images, ENIs are supported and you do not need to configure any ENIs manually.

- Centos 7.3 64-bit
- Centos 6.8 64-bit
- Windows Server 2008 R2 or later

If your instance is running an image not shown in the preceding list, and you want to attach an ENI to your instance, you must manually configure the ENI to be supported. This topic uses an instance running CentOS 7.2 64-bit as an example to describe how to configure an ENI manually.

Prerequisite

You have attached an ENI to an ECS instance.

Procedure

To configure an ENI manually, follow these steps:

1. Use the [#unique_47](#) interface or log on to the ECS console to obtain the following attributes of the ENI: the primary private IP address, subnet mask, the default route, and the MAC address. To obtain these attributes in the ECS console, follow these steps:
 - a. Log on to the [ECS console](#).
 - b. Find the target ENI and obtain its primary private IP address, subnet mask, default route, and MAC address. Example:

```
eth1  10 . 0 . 0 . 20 / 24  10 . 0 . 0 . 253  00 : 16 : 12
: E7 : 27
eth2  10 . 0 . 0 . 21 / 24  10 . 0 . 0 . 253  00 : 16 : 12
: 16 : EC
```

2. [Connect to the ECS instance](#).
3. Run the following command to generate the config file: `cat / etc / sysconfig / network - scripts / ifcfg -[network interface name in the OS]`.



Note:

- Pay attention to the relation between the network interface name in the OS and the MAC address.
- Pay attention to the relation between the network interface name in the OS and the MAC address. The default route must be set to `DEFROUTE = no`. Other editions must have the same configuration. Note that running the `ifup` command may change the active default route configuration after configuring the network interface.
- Example:

```
# cat / etc / sysconfig / network - scripts / ifcfg - eth1
DEVICE = eth1
BOOTPROTO = dhcp
ONBOOT = yes
TYPE = Ethernet
USERCTL = yes
PEERDNS = no
```

```
IPV6INIT = No
PERSISTENT_DHCLIENT = Yes
HWADDR = 00 : 16 : 3e : 12 : e7 : 27
DEFROUTE = noDefroute = No
```

4. To start the network interface, follow these steps:

- a. Run the `ifup [network interface name in the OS]` command to start the `dhclient` process, and initiate a DHCP request. Example:

```
# ifup eth1
# ifup eth2
```

- b. After a response is received, run the `ip a` command to check the IP allocation on the network interfaces, which must match with the information displayed on the ECS console. Example:

```
# ip a
1 : lo : mtu 65536 qdisc noqueue state UNKNOWN qlen 1
link / loopback 00 : 00 : 00 : 00 : 00 : 00 brd 00 : 00 : 00 : 00 : 00 : 00
inet 127 . 0 . 0 . 1 / 8 scope host loInet 125 . 0 . 0 . 1 / 8 Scope host Lo
valid_lft forever preferred_lft forever
2 : eth0 : mtu 1500 qdisc pfifo_fast state UP qlen 1000
10002 : eth0 : MTU 1500 qdisc glasstate up qlen 1000
link / ether 00 : 16 : 3e : 0e : 16 : 21 brd ff : ff : ff : ff : ff : ff
Inet 10 . 0 . 0 . 19 / 24 BRD glasscope Global Dynamic eth0
valid_lft 31506157se c preferred_lft 31506157se
cValid_lft 31506157se c preferred_lft 31506157se c
3 : eth1 : MTU 1500 qdisc glasstate up qlen 1000
link / ether 00 : 16 : 3e : 12 : e7 : 27 brd ff : ff : ff : ff : ff : ff
inet 10 . 0 . 0 . 20 / 24 brd 10 . 0 . 0 . 255 scope global dynamic eth1Inet 10 . 0 . 0 . 20 / 24 BRD glasscope Global Dynamic eth1
Valid_lft 31525994se c preferred_lft 31525994se c
4 : eth2 : MTU 1500 qdisc glasstate up qlen 1000
Link / ether 00 : 16 : Rye : 12 : 16 : ec brd ff : FF : FF : FF
inet 10 . 0 . 0 . 21 / 24 brd 10 . 0 . 0 . 255 scope global dynamic eth2
valid_lft 31526009se c preferred_lft 31526009se c
```

5. Set the metric for each network interface in the route table. In this example, set the metric parameters of `eth1` and `eth2` as follows.

```
eth1 : gw : 10 . 0 . 0 . 253 metric : 1001
```

```
eth2 : gw : 10 . 0 . 0 . 253 metric : 1002
```

- a. Run the following command to set the metric parameters.

```
# Ip - 4 route add default via glasdev eth1 metric
1001
# ip - 4 route add default via 10 . 0 . 0 . 253 dev
eth2 metric 1002
```

- b. Run the `route - n` command to check whether the configuration is successful. Example:

```
# route - n
Kernel IP routing table
Destinatio n Gateway Genmask Flags Metric Ref Use
Iface
0 . 0 . 0 . 0 10 . 0 . 0 . 253 0 . 0 . 0 . 0 UG 0 0 0
eth0
0 . 0 . 0 . 0 10 . 0 . 0 . 253 0 . 0 . 0 . 0 UG 1001 0
0 eth1
0 . 5 . 0 . 0 10 . 0 . 0 . 253 ug ub1002 0 0 eth2
10 . 0 . 0 . 0 0 . 5 . 0 . 0 255 . 25 . 25 . 0 u 0 0
0 eth0
10 . 0 . 0 . 0 0 . 0 . 0 . 0 255 . 255 . 255 . 0 U 0 0
0 eth1
10 . 0 . 0 . 0 0 . 5 . 0 . 0 255 . 25 . 25 . 0 u 0 0
0 eth2
169 . 254 . 0 . 0 0 . 0 . 0 . 0 255 . 0 . 0 U 1002 0 0
eth0
169 . 254 . 0 . 0 0 . 0 . 0 . 0 255 . 255 . 0 . 0 U 1003
0 0 eth1
169 . 254 . 0 . 0 0 . 0 . 0 . 0 255 . 255 . 0 . 0 U 1004
0 0 eth2
169 . 254 . 0 . 0 0 . 0 . 0 . 0 255 . 0 . 0 U
1004 0 0 eth2
```

6. To build a route table, follow these steps:



Note:

We recommend that you use the metric value as the route table name.

- a. Run the following command to build a route table.

```
# ip - 4 route add default via 10 . 0 . 0 . 253 dev
eth1 table 1001
# Ip - 4 route add default via glasdev eth2 table
1002
```

- b. Run the following command to check whether the route table is built successfully.

```
# ip route list table 1001
default via 10 . 0 . 0 . 253 dev eth1
# ip route list table 1002
```

```
default via 10.0.0.253 dev eth2
```

7. Configure the policy routing.

a. Run the following command to configure the policy routing.

```
# ip - 4 rule add from 10.0.0.20 lookup 1001
# ip - 4 rule add from 10.0.0.21 lookup 1002
```

b. Run `ip rule list` to view the routing rules.

```
# ip rule list
0 : from all lookup local
32764 : from 10.0.0.21 lookup 1002
32765 : from 10.0.0.20 lookup 1001
32766 : from all lookup main
32767 : from all lookup default
```

What to do next

After you have configured an ENI, you can perform the following operations:

- [Modify attributes of an ENI.](#)
- [Detach an ENI from an instance.](#)
- [Delete an ENI.](#)

4.5 Assign a secondary private IP address

This topic describes how to assign secondary private IP addresses to an Elastic Network Interface (ENI).

Scenarios

- **Optimize application usage**

If your ECS instance hosts multiple applications, you can assign multiple secondary private IP addresses to the corresponding ENI. In this way, each application uses a separate IP address for services, which optimizes the usage of the ECS instance.

- **Avoid service disruptions**

You can attach the ENI of an active ECS instance to another instance to direct traffic to the standby instance if the active instance fails, enabling service continuity.

Limits

- You can only attach an ENI to an ECS instance in a VPC. The ENI and the instance must be in the same VPC, VSwitch, and zone.

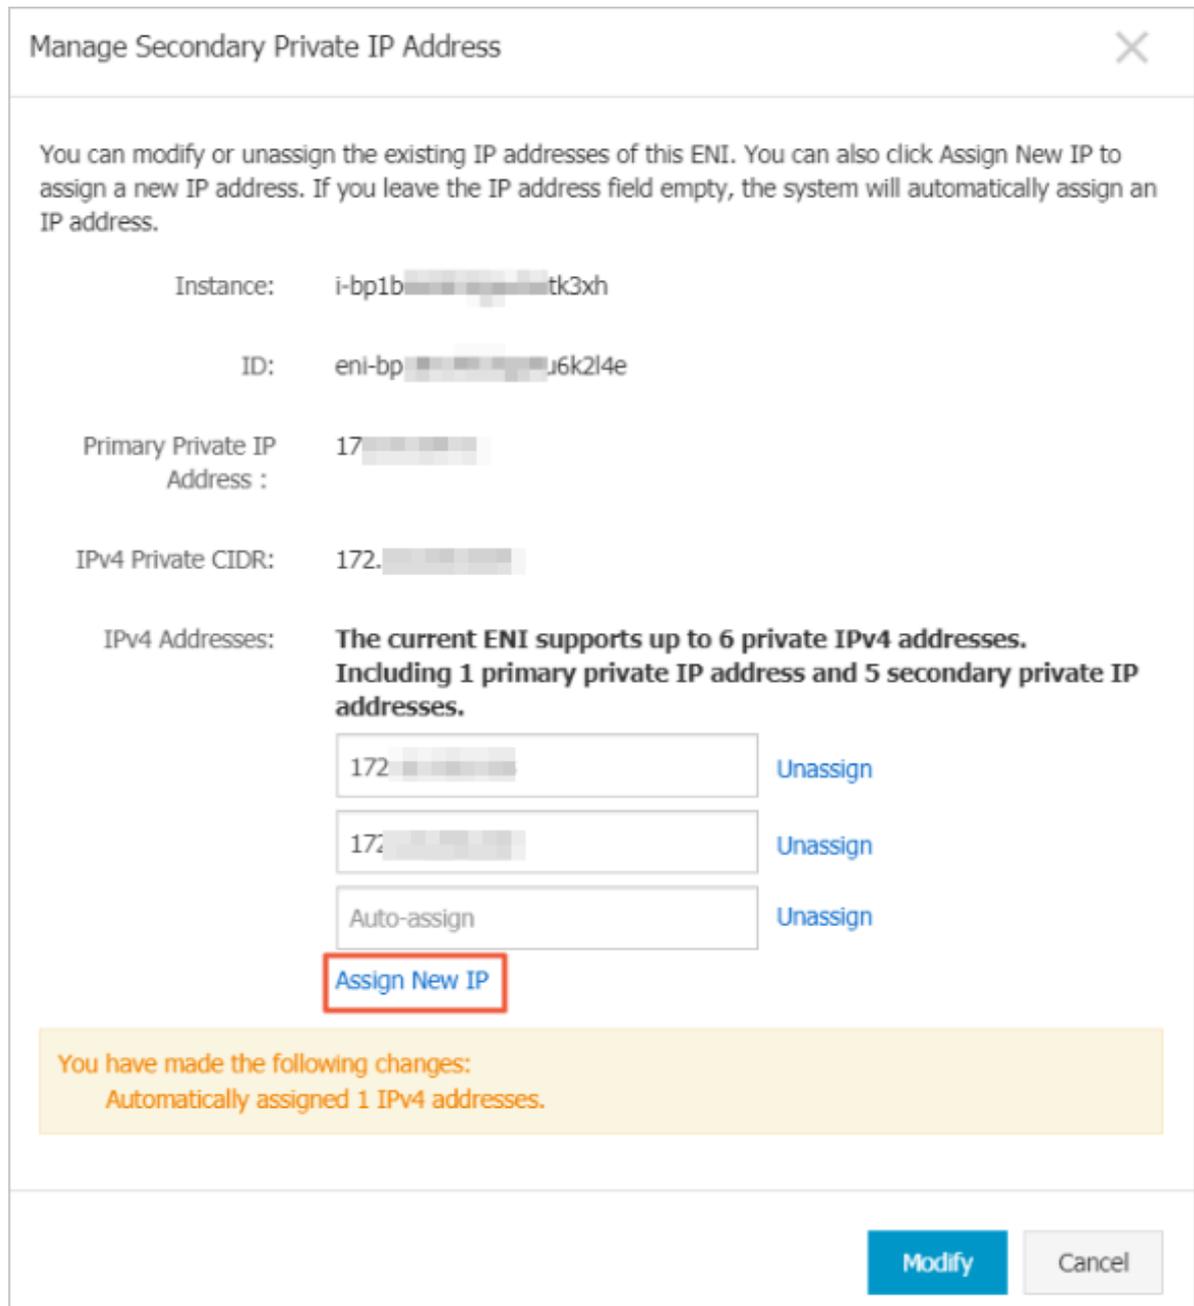
- Each VPC security group can contain a maximum of 2,000 private IP addresses, and the quota is shared among all corresponding primary and secondary ENIs.
- You can assign a maximum of 20 private IP addresses to an ENI.
 - If the target ENI is in the `Available` state, you can assign a maximum of 10 private IP addresses to the ENI.
 - If the target ENI is in the `InUse` state, the number of private IP addresses that you can assign to the ENI depends on the instance type. For more information, see [Instance type families](#).
- Your instance type must be able to support being assigned multiple secondary private IP addresses. For more information, see [Instance type families](#) or call the `DescribeInstanceTypes` API action.
- If you assign multiple secondary private IP addresses to a primary ENI, the instance to which the primary ENI is attached must be in the `Running` or `Stopped` state.

Assign a secondary private IP address to an ENI

1. On the Network Interfaces page, locate the target ENI, and then click **Manage Secondary Private IP Address** in the Actions column.

- 2. In the displayed dialog box, click Assign New IP once or multiple times if additional IP addresses are needed.

You can also enter one or more secondary private IP addresses that are within the IPv4 Private CIDR. If you do not enter any secondary private IP address, the system randomly assigns IP addresses that are within the IPv4 Private CIDR.



- 3. Click Modify.
- 4. Optional. If you use automatic assignment of a secondary private IP address, click Manage Secondary Private IP Address in the Actions column to view the

assigned secondary private IP address, and then configure this IP address for an ECS instance.

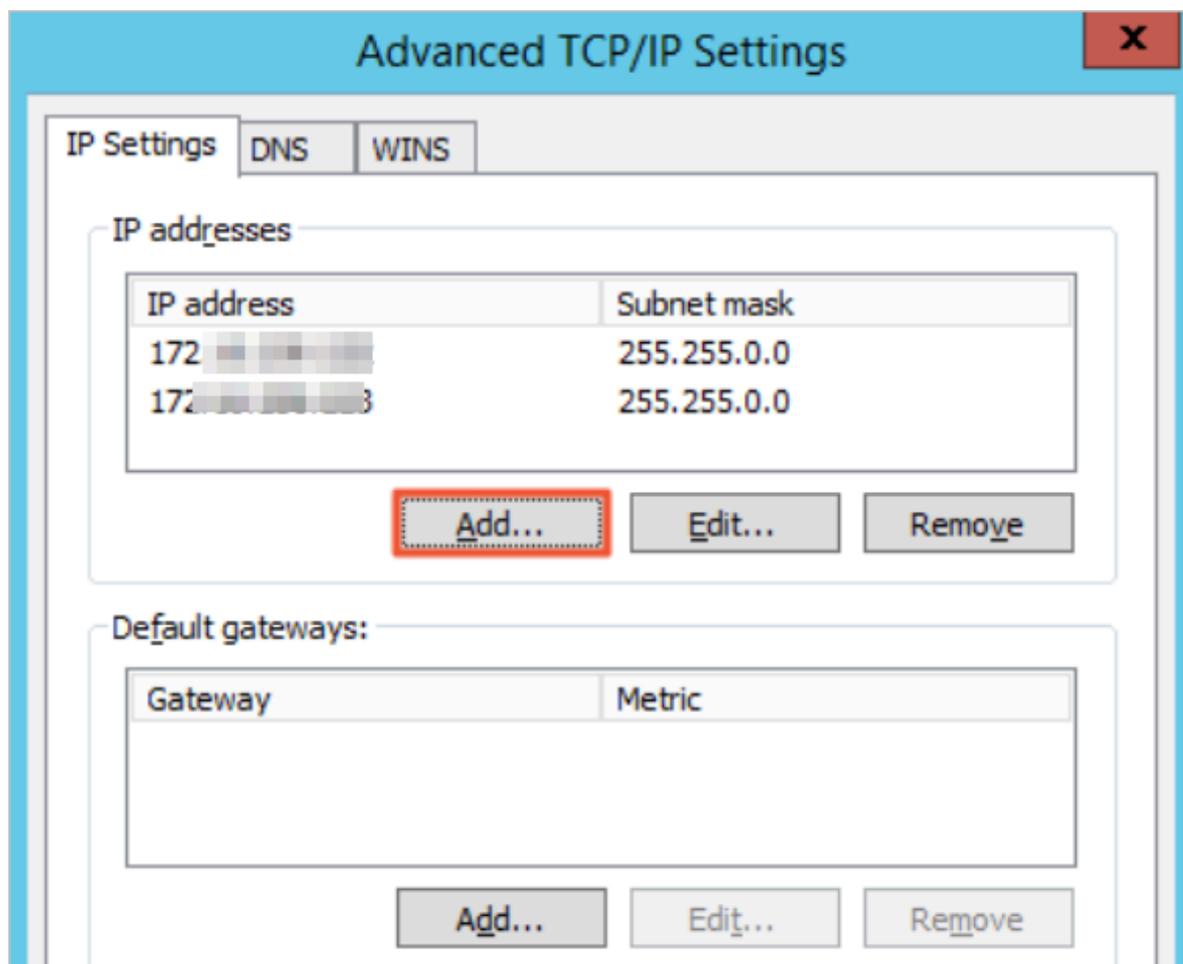
5. Optional. If the target ENI is not attached, attach it to an ECS instance. For more information, see [#unique_42](#).

Related API: [AssignPrivateIpAddresses](#)

Assign a secondary private IP address to a Windows instance

1. Connect to the target instance. For more information, see [#unique_73](#).
2. Open the Network and Sharing Center.
3. Click Change adapter settings.
4. Double-click the current network connection name, and then click Properties.
5. Double-click Internet Protocol Version 4 (TCP/IPv4).
6. Select Use the following IP address, and then click Advanced.
7. Click Add, and then enter the assigned IP Address and set the Subnet Mask.

You can add multiple IP addresses to the same adapter.



8. Click OK.

Assign a secondary private IP address to a Linux instance

1. Connect to the target instance. For more information, see [#unique_73](#).
2. Follow the instructions in the method that corresponds to the OS of your instance to assign a secondary private IP address.

In the following example, a primary ENI named `eth0` are used. If you use a secondary ENI, you must modify the ENI identifier as needed.

- **RHEL series: CentOS 6/7, Red Hat 6/7, and Aliyun Linux 17**

- a. Run the `vi / etc / sysconfig / network - scripts / ifcfg - eth0 : 0` command to open the network configuration file and add the following configuration items:

```
DEVICE = eth0 : 0
TYPE = Ethernet
BOOTPROTO = static
ONBOOT = yes
IPADDR =< IPv4 address 1 >
NETMASK =< IPv4 mask >
GATEWAY =< IPv4 gateway >
```

If you assign multiple IP addresses, run the `vi / etc / sysconfig / network - scripts / ifcfg - eth0 : 1` command to open the network configuration file and add the following configuration items:

```
DEVICE = eth0 : 1
TYPE = Ethernet
BOOTPROTO = static
ONBOOT = yes
IPADDR =< IPv4 address 2 >
NETMASK =< IPv4 mask >
GATEWAY =< IPv4 gateway >
```

- b. Run the `service network restart` or `systemctl restart network` command to restart the network service.

- **Debian series: Ubuntu 14/16 and Debian/8/9**

- a. Run the `vi / etc / network / interfaces` command to open the network configuration file and add the following configuration items:

```
auto eth0 : 0
iface eth0 : 0 inet static
address < IPv4 address 1 >
netmask < IPv4 mask >
gateway < IPv4 gateway >
```

```
auto eth0 : 1
iface eth0 : 1 inet static
address < IPv4 address 2 >
netmask < IPv4 mask >
gateway < IPv4 gateway >
```

- b. Run the `service networking restart` or `systemctl restart networking` command to restart the network service.

- SLES series: SUSE 11/12 and OpenSUSE 42

- a. Run the `vi / etc / sysconfig / network / ifcfg - eth0` command to open the network configuration file and add the following configuration items:

```
IPADDR_0 =< IPv4 address 1 >
NETMASK_0 =< subnet prefix length >
LABEL_0 = ' 0 '

IPADDR_1 =< IPv4 address 2 >
NETMASK_1 =< subnet prefix length >
LABEL_1 = ' 1 '
```

- b. Run the `service network restart` or `systemctl restart network` command to restart the network service.

What to do next

If you no longer require the current number of secondary private IP addresses, you can revoke one or more of them from the target ENI. For more information, see [#unique_74](#).

4.6 Revoke a secondary private IP address

This topic describes how to revoke a secondary private IP address from an Elastic Network Interface (ENI).

Limits

The primary private IP address cannot be revoked.

Prerequisites

- At least one secondary private IP addresses is assigned to the target ENI.
- The target ENI is in the `Available` or `InUse` state.
- If the secondary private IP addresses to be revoked is assigned to the primary ENI, the instance to which the primary ENI is attached must be in the `Running` or `Stopped` state.

Procedure

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Network & Security > ENI.
3. In the upper-left corner, select the target region.
4. On the Network Interfaces page, locate the target ENI, and then click Manage Secondary Private IP Address in the Actions column.
5. In the Manage Secondary Private IP Address dialog box, click Unassign once or multiple times if additional IP addresses need to be revoked.
6. Click Modify.

Related API: [UnassignPrivateIpAddresses](#)

What to do next

If your application requirements change, you can assign multiple secondary private IP address to an ENI. For more information, see [#unique_77](#).

4.7 Modify an ENI

This topic describes how to modify primary and secondary Elastic Network Interfaces (ENIs). You can only modify the primary ENI by configuring its associated instance with a different security group as needed, and you can modify a secondary ENI by changing its attributes (such as the name, associated security group, and description).

Limits

Before you can modify the security group to which an ENI belongs, the ENI and its associated ECS instance must meet the following limits:

- An ECS instance cannot be added to a basic security group and an advanced security group at the same time.
- An ENI cannot be added to a basic security group and an advanced security group at the same time.
- An ENI can be attached to an ECS instance only if they belong to the same type of security group.

For more information, see [#unique_79](#).

Modify a primary ENI

To modify a primary ENI, follow these steps:

**Note:**

The primary ENI and the secondary ENIs of an ECS instance can belong to different security groups. This means that if you associate the ECS instance with another security group, the primary ENI will also be associated with this security group, but the secondary ENIs will remain in the previous security group.

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Network & Security > Security Groups.
3. In the top navigation bar, select a region.
4. Find the target security group, and then click Manage Instances in the Actions column.
5. On the Instances in Security Group page, modify the security group with which the primary ENI is associated.
 - To add the primary ENI to a new security group, follow these steps:
 - a. In the upper-right corner of the Instances in Security Group page, click Add Instance.
 - b. In the displayed dialog box, select an instance ID to which the primary ENI is attached, and then click OK.

The primary ENI is added to the new security group along with the corresponding ECS instance.
 - To remove the primary ENI from the current security group, follow these steps:
 - a. On the Instances in Security Group page, select one or more instances, and then click Remove from Security Group.
 - b. In the displayed dialog box, click OK.

The primary ENI is removed from the current security group along with the corresponding ECS instance. Note that the primary ENI and the ECS instance must belong to at least one security group.
6. Go back to the Security Groups page and find the target primary ENI to verify that the settings have taken effect.

Related APIs:

- [JoinSecurityGroup](#)
- [LeaveSecurityGroup](#)

Modify a secondary ENI

To modify a secondary ENI, follow these steps:

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, choose Network & Security > ENI.
3. In the top navigation bar, select a region.
4. Find the target secondary ENI, and then click Modify in the Actions column.
5. In the displayed dialog box, modify the ENI attributes as follows:
 - **Network Interface Name:** Set a new ENI name according to the rules displayed under this field.
 - **Security Group:** Select a new security group for the ENI, or remove the ENI from a security group. Note that the ENI must be associated with at least one security group.
 - **Description:** Modify the description according to the rules displayed under this field.
6. Click OK.

Related API: [ModifyNetworkInterfaceAttribute](#)

4.8 Detach an ENI from an instance

You can only detach a secondary ENI from an instance. You cannot detach the primary ENI.

Limits

Before you detach a secondary ENI from an instance, note the following limits:

- The secondary ENI must be in the Bound state.
- The instance to which the ENI belongs must be in the Stopped or Running state.

Prerequisites

The secondary ENI **is attached to an instance**. Before you detach a secondary ENI from an instance, the instance must be in the Stopped or Running state.

Procedure

To detach a secondary ENI from an instance, follow these steps:

1. Find the target ENI, and in the Actions column, click Unbind.

2. In the displayed dialog box, confirm the information, and then click OK.

After, in the Network Interfaces page, refresh the table. When the selected ENI is in the Available state, it is detached from the instance.

What to do next

After an ENI is detached from an instance, you can:

- [Attach the ENI to another instance.](#)
- [Delete the ENI.](#)
- [Modify attributes of the ENI.](#)

4.9 Delete an ENI

You can only delete a secondary ENI. You cannot delete the primary ENI of an instance.

After a secondary ENI is deleted:

- The primary private IP address of the secondary ENI is released automatically.
- The deleted secondary ENI is automatically removed from all associated security groups.

If you release an instance, any attached ENIs will be deleted along with its release. You can choose to detach the ENI first and then release the corresponding instance separately.

Limits

You can only delete an ENI in the Available status.

Prerequisite

If an ENI is [attached to an instance](#), you must first [detach it from the instance](#) to delete it separately.

Procedure

To delete an ENI, follow these steps:

1. Log on to the [ECS console](#).
2. In the left-side navigation pane, select Networks and Security > ENI.
3. Select the target region.
4. Find the target ENI, and in the Actions column, click Delete.

5. Click OK.

In the Network Interfaces page, refresh the table. If the ENI is no longer displayed, it is deleted successfully.

5 Multiqueue for NICs

Multiqueued NICs route NIC interruptions in ECS instances to different CPUs. Results of network PPS and bandwidth tests show that a solution that uses two queues instead of one queue can enhance network performance by between 50% to 100%.

ECS instance types supporting multiqueue

See [#unique_86](#) to find instance types that support multiqueue and the number of queues that are supported.

Images supporting multi-queue

The following public images officially provided by Alibaba Cloud support multiqueue :



Note:

Whether an image supports multiqueue is not related to the memory address width of the operating system.

- CentOS 6.8/6.9/7.2/7.3/7.4
- Ubuntu 14.04/16.04
- Debian 8.9
- SUSE Linux Enterprise Server 12 SP1

Support for SUSE Linux Enterprise Server 12 SP2 edition is in development. Support for Windows 2012 R2 and Windows 2016 is by invitation.

Configure multi-queue support for NICs on a Linux ECS instance

We recommend that you use one of the latest Linux distributions, such as CentOS 7.2, to configure multi-queue for the NICs.

Here we take CentOS 7.2 as an example to illustrate how to configure multi-queue for the NIC. In this example, two queues are configured, and the NIC name is eth0.

- To check whether the NIC supports multi-queue, run the command: `ethtool -l eth0`.
- To enable multi-queue for the NIC, run the command: `ethtool -L eth0 combined 2`.

- If you are using more than one NIC, configure each NIC.

```
[ root @ localhost ~]# ethtool -l eth0
Channel parameters for eth0 :
Pre-set maximums :
RX : 0
TX : 0
Other : 0
Combined : 2 # This line indicates that a maximum
of two queues can be configured
Current hardware settings :
RX : 0
TX : 0
Other : 0
Combined : 1 # It indicates that one queue is
currently taking effect
[ root @ localhost ~]# ethtool -L eth0 combined 2 #
It sets eth0 to use two queues currently
```

- We recommend that you enable the `irqbalance` service so that the system can automatically adjust the allocation of the NIC interrupts on multiple CPU cores. Run the command: `systemctl start irqbalance` (this feature is enabled by default in CentOS 7.2).
- If the network performance is not improved as expected after the multi-queue feature is enabled, you can enable the RPS feature. To do so, see the following Shell script:

```
#!/ bin / bash
cpu_num=$(grep -c processor /proc/cpuinfo)
quotient=$((cpu_num / 8))
if [ $quotient -gt 2 ]; then
    quotient = 2
elif [ $quotient -lt 1 ]; then
    quotient = 1
fi
for i in $(seq $quotient)
do
    cpuset="${cpuset} f "
done
for rps_file in $(ls /sys/class/net/eth*/queues
/ rx -*/ rps_cpus )
do
    echo $cpuset > $rps_file
done
```

Configure multi-queue support for NICs on a Windows ECS instance



Note:

We are inviting Windows users to sign up and test multiqueue support for performance improvement. Note that the overall performance increase is not as great when compared with performance increase of Linux systems.

If you are using a Windows instance, you must install the driver to use the multiqueue feature for NICs.

To install the driver for Windows systems, follow these steps:

1. [Open a ticket](#) to request and download the driver installation package.
2. Unzip the driver installation package. For Windows 2012/2016 systems, use the driver in the Win8/amd64 folder.
3. Upgrade the NIC driver:
 - a. Select Device Manager > Network adapters.
 - b. Right click Red Hat VirtIO Ethernet Adapter and select Update Driver.
 - c. Select the Win8/admin64 directory of the driver directory that you have unzipped, and update the driver.
4. Restart the Windows system after the driver is upgraded for the multiqueue feature to take effect.

6 Connect a classic network to a VPC

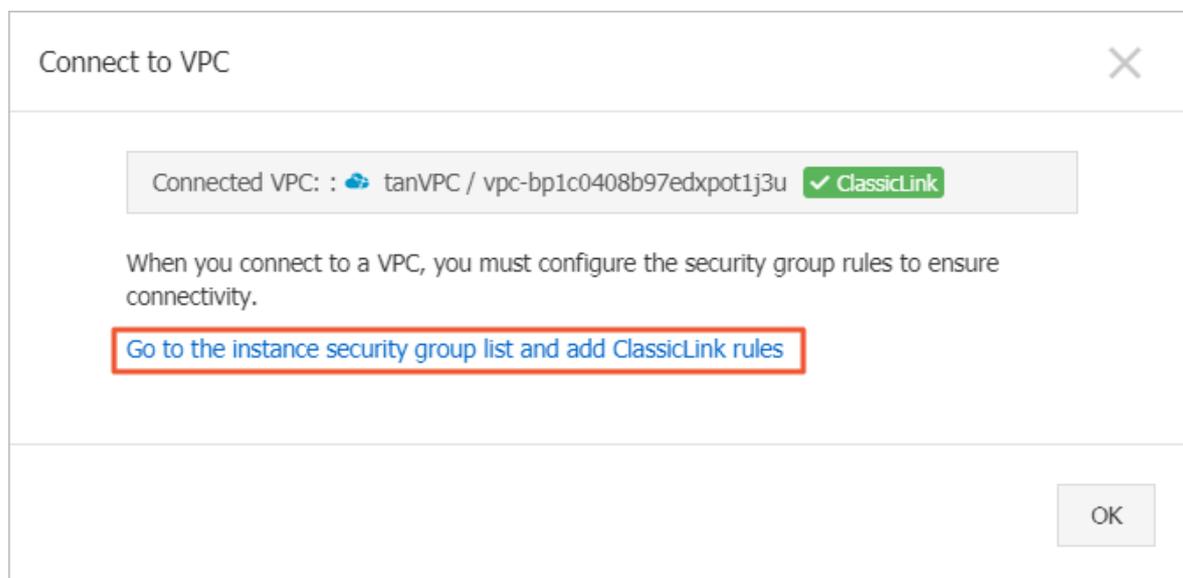
This topic describes how to connect a classic network to a VPC. You can set up a ClassicLink connection so that ECS instances of the classic network type can access cloud resources in a VPC through the intranet.

Prerequisites

Make sure that you are aware of the limits of ClassicLink. For more information, see [#unique_88](#).

Procedure

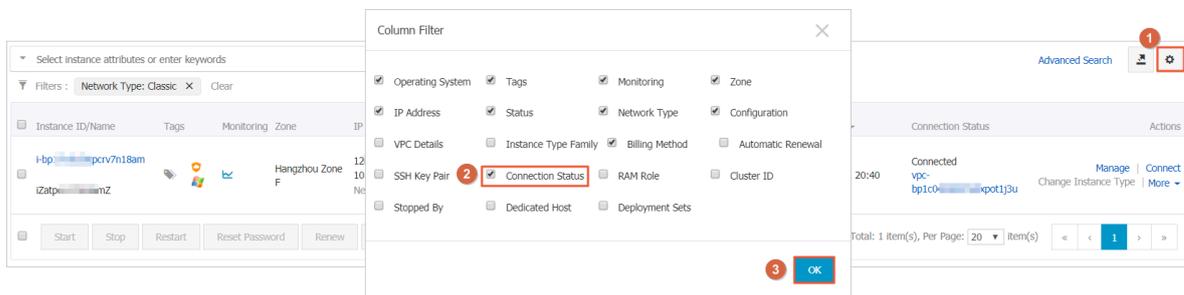
1. Log on to the [VPC console](#).
2. Select the region of the target VPC, and click the ID of the target VPC.
3. On the VPC Details page, click Enable ClassicLink. In the displayed dialog box, click OK.
4. Find the target ECS instance of the classic network type, and then choose More > Network and Security Group > Connect to VPC.
5. In the displayed dialog box, select the target VPC and click OK, and then click the security group configuration link.



6. Click Add ClassicLink Rules and configure the security rule according to the following information. Then, click OK.

Configuration	Description
Classic Security Group	Display the classic network security group.
Select VPC Security Group	Select a security group to use. Up to five security groups can be selected.
Mode	Select one of the following modes: <ul style="list-style-type: none"> Classic <=> VPC : The connected resources can access each other (recommended). Classic => VPC : Authorize the classic ECS instance to access cloud resources in the connected VPC. VPC => Classic : Authorize the cloud resources in the connected VPC to access the classic ECS instance.
Protocol Type and Port Range	Select the protocol and port used for the communication. The port must be in the form of xx/xx. For example, if port 80 is used, enter 80/80.
Priority	Set the priority for the rule. A smaller number represents a higher priority, for example, 1.
Description	Enter a description for the security rule.

7. Return to the [ECS console](#). On the Instance List page, click the Column Filter icon in the upper-right corner, and then select the Connection Status check box. Then, click OK.



If Connection Status is Connected, ECS instances of the classic network are connected to the VPC network.

7 Network FAQ

This topic lists FAQ related to ECS network.

- **Public bandwidth**
 - [What are the inbound and outbound bandwidths of ECS instances?](#)
 - [I bought 5 Mbit/s of public bandwidth for my ECS instance. What is the difference between the inbound bandwidth and outbound bandwidth of the instance?](#)
 - [Does my ECS instance use its public bandwidth exclusively or share the bandwidth with other instances?](#)
 - [How is the public bandwidth of ECS instances billed?](#)
 - [Why is 200 Kbit/s of inbound traffic already consumed on my newly created ECS instance?](#)
 - [How can I view the public traffic statistics of my ECS instance?](#)
 - [Why is the bandwidth usage of my ECS instance displayed in CloudMonitor different from that displayed in the ECS console?](#)
 - [My ECS instance has been stopped. Why am I still billed for outbound traffic from it on a pay-as-you-go basis?](#)
- **IP addresses**
 - [How can I query the IP addresses of my ECS instance?](#)
 - [How can I disable the public NIC of my ECS instance?](#)
- **Network access and traffic direction**
 - [Why can't I access a website hosted on an ECS instance?](#)
 - [An unusual logon to my ECS instance has been detected. What can I do?](#)
 - [What is traffic scrubbing?](#)
 - [How can I cancel traffic scrubbing for my ECS instance?](#)
 - [How can I request reverse lookup for my ECS instance?](#)
 - [Can an IP address point to multiple reverse lookup domain names?](#)

- **Public IP addresses**
 - [Can I change the public IP \(IPv4\) address of my ECS instance six hours after the instance is created? How?](#)
 - [Why can't I find the option to change the public IP address of my ECS instance in the ECS console?](#)
 - [Apart from the public IP address, can I change the private IP address of my ECS instance?](#)
 - [If no public IP \(IPv4\) address has been assigned to my ECS instance when the instance was created, how can I assign a public IP address to the instance?](#)
- **Network basics**
 - [What is a BGP data center?](#)
 - [What are WAN and LAN?](#)
 - [How can I express a subnet mask?](#)
 - [How can I plan subnets?](#)

What are the inbound and outbound bandwidths of ECS instances?

Bandwidth type	Description
Inbound bandwidth	<p>The bandwidth for inbound traffic of an ECS instance. For example:</p> <ul style="list-style-type: none"> · Traffic that occurs when you download external resources to your ECS instances · Traffic that occurs when you upload resources to your ECS instances by using an FTP client
Outbound bandwidth	<p>The bandwidth for outbound traffic of an ECS instance. For example:</p> <ul style="list-style-type: none"> · Traffic that occurs when your ECS instances provide external access · Traffic that occurs when you download resources from your ECS instances by using an FTP client

I bought 5 Mbit/s of public bandwidth for my ECS instance. What is the difference between the inbound bandwidth and outbound bandwidth of the instance?

The 5 Mbit/s you bought is the outbound bandwidth. The inbound bandwidth of your instance is capped at 100 Mbit/s.

- Outbound bandwidth is also called downstream bandwidth. The outbound bandwidth of an ECS instance is capped at 200 Mbit/s regardless of whether the instance resides in a VPC or classic network.
- Inbound bandwidth is also called upstream bandwidth. The maximum inbound bandwidth varies depending on the outbound bandwidth:
 - When the outbound bandwidth is less than or equal to 100 Mbit/s, the maximum inbound bandwidth is 100 Mbit/s.
 - When the outbound bandwidth is greater than 100 Mbit/s, the inbound bandwidth is the same as the outbound bandwidth.

Does my ECS instance use its public bandwidth exclusively or share the bandwidth with other instances?

Your ECS instance uses its public bandwidth exclusively.

How is the public bandwidth of ECS instances billed?

For details, see [#unique_13](#).

Why is 200 Kbit/s of inbound traffic already consumed on my newly created ECS instance?

This traffic was generated by Address Resolution Protocol (ARP) broadcast packets. Your new ECS instance is assigned to a large network segment. When an ARP request packet is sent to the gateway to request the IP address of an ECS instance within the network segment, the gateway broadcasts the ARP request packet to all the ECS instances within that network segment. Your new ECS instance receives the packet, and inbound traffic is generated. If the IP address of your new ECS instance is not requested, the instance does not reply with an ARP response packet.

How can I view the public traffic statistics of my ECS instance?

To view the public traffic statistics about your ECS instance, perform the following steps:

1. Log on to the [ECS console](#).
2. In the top navigation bar of the ECS console, choose Billing Management > Billing Management.
3. In the left-side navigation pane, choose Bill > Bill.
4. On the Bills page that appears, click the Bills tab. Specify a billing period. Then, click the filter icon to the right of Product Detail and select Elastic Compute Service (ECS) - Pay by quantity from the option list.

5. Click **Export Billing Overview (CSV)**. In the **Export Billing Overview (CSV)** dialog box, enter the captcha and click **OK**.
6. Open the exported CSV file to view the public traffic statistics about your ECS instance.

Why is the bandwidth usage of my ECS instance displayed in CloudMonitor different from that displayed in the ECS console?

ECS instances function as back-end servers of SLB instances and use the Layer 7 HTTP forwarding model. In this forwarding model, SLB instances forward client requests to ECS instances, and the ECS instances use their outbound bandwidth to return responses to the corresponding users. The bandwidth consumed by these responses is not displayed in the ECS console, but the traffic generated by the responses is counted towards the outbound traffic of the SLB instances and displayed in CloudMonitor. Therefore, the bandwidth usage of your ECS instance displayed in CloudMonitor is different from that displayed in the ECS console.

My ECS instance has been stopped. Why am I still billed for outbound traffic from it on a pay-as-you-go basis?

- **Problem description:** Your ECS instance is in the **Stopped** state when viewed from the ECS console, but is in the **Cleaning** state when viewed from the **Anti-DDoS basic** console. You are billed for outbound traffic from the instance on a pay-as-you-go basis every hour.
- **Cause:** HTTP flood protection is enabled for your ECS instance. After HTTP flood protection is enabled, the security mechanism sends probe packets to potential attack sources, which generates a large volume of outbound traffic.
- **Solution:** Disable HTTP flood protection for your ECS instance.

How can I query the IP addresses of my ECS instance?

- **Linux instance**

Run the `ifconfig` command to view NIC information. View the IP addresses, subnet masks, gateways, DNS servers, and MAC address in the command output.

- **Windows instance**

In the CLI, run the `ipconfig / all` command to view NIC information. View the IP addresses, subnet masks, gateways, DNS servers, and MAC address in the command output.

How can I disable the public NIC of my ECS instance?

- Linux instance
 1. Run the `ifconfig` command to view the public NIC name of your instance.
 2. Run the `Ifdown` command to disable the public NIC. For example, if the public NIC name is `eth1`, use the `ifdown eth1` command.



Note:

You can also run the `Ifup` command to re-enable the NIC. For example, if the public NIC name is `eth1`, use the `ifup eth1` command.

- Windows instance
 1. In the CLI, run the `ipconfig` command to view information about the public NIC.
 2. Open the Control Panel and choose Network and Internet > View network status and tasks. In the Network and Sharing Center window that appears, click Change adapter settings in the left-side navigation pane to disable the public NIC.

Why can't I access a website hosted on an ECS instance? A message similar to "Sorry, your access is blocked because the requested URL may pose a security threat to the website" is displayed.

- **Problem description:** When you access a website built on an ECS instance, you are prompted with a message similar to "Sorry, your access is blocked because the requested URL may pose a security threat to the website."
- **Cause:** Web Application Firewall (WAF) identifies your access to the requested URL as an attack and blocks your access.
- **Solution:** Add the public IP address, Elastic IP Address, or NAT IP address of the ECS instance to the WAF whitelist. For more information, see [Avoid Anti-DDoS Basic false positives by using a whitelist](#).

An unusual logon to my ECS instance has been detected. What can I do?

Perform the following steps to solve the problem:

1. Check the logon time to see whether this logon is performed by you or another administrator.

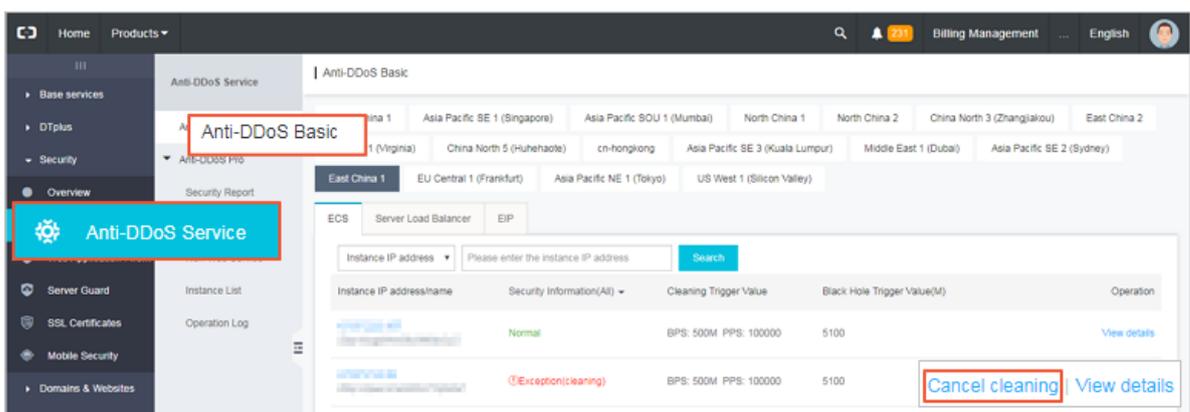
2. If no, it is an unauthorized logon. Perform the following steps:
 - a. [Reset the password](#).
 - b. Check whether the ECS instance is infected by viruses..
 - c. Configure security groups to [allow only specific IP addresses to log on](#).

What is traffic scrubbing?

The traffic scrubbing service monitors the inbound traffic to ECS instances in real time, and provides quick identification of unusual traffic, such as DDoS attacks. Anti-DDoS Basic, which includes the traffic scrubbing service, is enabled on ECS instances by default. When ECS instances are under attack, the traffic scrubbing service automatically detects the attacks and scrubs the traffic for ECS instances without affecting your services. When an anomaly is detected, the traffic scrubbing service redirects suspicious traffic from the network where the destination ECS instance resides to the scrubbing device. The scrubbing device identifies and removes malicious traffic and returns legitimate traffic to the network. This ensures only legitimate traffic is forwarded to the destination ECS instance.

How can I cancel traffic scrubbing for my ECS instance?

1. Log on to the [Alibaba Cloud Security Anti-DDoS Basic console](#).
2. Click the ECS tab. Then in the ECS instance list, find the IP address of your ECS instance that is in the cleaning state. In the Actions column, click View Details.
3. Click Cancel cleaning.



How can I request reverse lookup for my ECS instance?

Reverse lookup is used in mail services. It enables mail servers to reject all mails sent from the IP addresses that are mapped to unregistered domain names. Most of spammers use dynamic IP addresses or IP addresses that are mapped to unregistered domain names to send unwanted emails and escape tracking. After reverse lookup is

enabled on a mail server, the server rejects the mails that are sent from dynamic IP addresses and unregistered domains. This greatly reduces the number of spams.

You can submit a ticket in the [ticket system](#) to request reverse lookup for your ECS instance. We recommend that you specify the region, public IP address, and registered domain name of your ECS instance in the ticket to improve the ticket processing efficiency.

After your request is approved, you can use the `dig` command to check whether reverse lookup has taken effect for your instance. For example:

```
dig -x 121 . 196 . 255 .** + trace + nodnssec
```

If information similar to the following is displayed in the command output, reverse lookup has taken effect for your instance.

```
1 . 255 . 196 . 121 . in - addr . arpa . 3600 IN PTR ops .  
alidns . com .
```

Can an IP address point to multiple reverse lookup domain names?

No, each IP address can point to only one reverse lookup domain name. For example, you cannot configure an IP address such as 255.196.121.1 to be reversely resolved to multiple domain names such as mail.abc.com, mail.ospf.com, and mail.zebra.com.

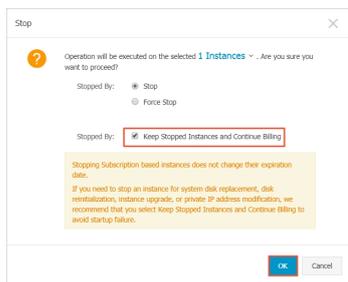
Can I change the public IP (IPv4) address of my ECS instance six hours after the instance is created? How?

- For instances in a VPC network, the public IP address can be converted into an Elastic IP Address (EIP). For information about how to convert the IP address, see [#unique_26](#).
- For instances in a classic network, the public IP address cannot be changed if an instance has been created for more than six hours.

Why can't I find the option to change the public IP address of my ECS instance in the ECS console?

- By default, the Change Public IP Address option is not displayed for instances that were created more than six hours ago.
- If you enable the [no fees for stopped VPC instances](#) feature for an instance, make sure that this feature is disabled when you stop the instance. Otherwise, the instance will be temporarily released and the Change Public IP Address option is

not displayed. You can disable this feature by selecting the **Keep Stopped Instances and Continue Billing** check box in the Stop dialog box.



Apart from the public IP address, can I change the private IP address of my ECS instance?

- This operation is allowed for instances in a VPC network. For information about how to change the private IP address, see [#unique_117](#).
- This operation is not allowed for instances in a classic network.

If no public IP (IPv4) address has been assigned to my ECS instance when the instance was created, how can I assign a public IP address to the instance?

- For a Subscription instance, you can obtain the public IP address by upgrading or downgrading the network bandwidth configuration. For more information, see [#unique_28](#).
- For a Pay-As-You-Go instance, you cannot obtain the public IP address after the instance is created. You can only [bind an EIP address](#).

What is a BGP data center?

Border Gateway Protocol (BGP) is primarily used for interconnection between Internet autonomous systems (AS). The main function of BGP is to control route propagation and select the best routes. A BGP data center is a data center that uses BGP to implement dual-line or multi-line interconnection.

China Netcom, China Telecom, China Railcom, and some large privately owned IDC carriers all have autonomous system numbers (ASNs). Most major network carriers in China use BGP to achieve multi-line interconnection with their own ASNs.

To achieve multi-line interconnection in this manner, an IDC must obtain a CIDR block and an ASN from the China Internet Network Information Center (CNNIC) or Asia-Pacific Network Information Center (APNIC), and then broadcast this CIDR block to the networks of other carriers through BGP. After networks are interconnected through BGP, the backbone routers of the network carriers will determine the

optimal routes to the CIDR block of the IDC to ensure high-speed access for users of different network carriers.

What are WAN and LAN?

- A wide area network (WAN) is also known as an external or public network. It is a telecommunications network that connects different smaller networks, including local area networks (LANs) and metro area networks (MANs). Each WAN extends over a large geographical area, such as across cities, states, or countries, and may cover continents to provide telecommunications services and form an international telecommunications network. WAN is not equal to Internet.
- A LAN is also known as an internal network. A LAN is a network that interconnects computers within a small area. Users can manage files, share application software and printers, schedule work for work groups, and communicate with each other such as sending emails or faxes within a LAN. A LAN is a closed network that can consist of two computers in an office or thousands of computers in a company. In Alibaba Cloud public cloud, ECS instances in the same region can be created in the same type of networks and communicate with each other through internal networks. ECS instances in different regions are isolated from each other.

How can I express a subnet mask?

You can express a subnet mask in either of the following method:

- Dotted decimal notation. For example:

The default subnet mask of a class A network is 255.0.0.0.

- Append a forward slash (/) and a number ranging from 1 to 32 to the end of an IP address to define a subnet mask. The number indicates the length of the network identification bit in the subnet mask. For example:

192.168.0.3/24.

How can I plan subnets?

For the best practices of planning subnets, see [Plan and Design VPC](#).