

Elastic IP

User Guide(ME-Abu Dhabi Region)

Issue	01
Date	2025-11-05



Copyright © Huawei Technologies Co., Ltd. 2025. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <https://www.huawei.com>

Email: support@huawei.com

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Service Overview..... 1

1.1 What Is Elastic IP?..... 1

1.2 Advantages..... 3

1.3 Application Scenarios..... 4

1.4 Functions..... 6

1.5 Notes and Constraints..... 7

1.6 Related Services..... 7

1.7 Billing..... 8

1.8 Basic Concepts..... 9

1.8.1 Dedicated Bandwidth..... 9

1.8.2 Shared Bandwidth..... 10

1.8.3 Region and AZ..... 10

2 Getting Started..... 12

2.1 Overview..... 12

2.2 Step 1: Create a VPC..... 13

2.3 Step 2: Create a Subnet for the VPC..... 19

2.4 Step 3: Assign an EIP and Bind It to an ECS..... 23

2.5 Step 4: Create a Security Group..... 26

2.6 Step 5: Add a Security Group Rule..... 31

3 Elastic IP..... 36

3.1 EIP Overview..... 36

3.2 Assigning an EIP..... 37

3.3 Modifying an EIP Bandwidth..... 40

3.4 Binding or Unbinding an EIP..... 41

3.5 Releasing an EIP..... 42

3.6 Exporting EIP Information..... 43

3.7 Managing EIP Tags..... 43

3.8 Assigning an EIP and Binding It to an ECS..... 44

3.9 EIP Configuration Examples..... 47

3.9.1 Changing an EIP for an Instance..... 47

4 Shared Bandwidth..... 51

4.1 Shared Bandwidth Overview..... 51

4.2 Assigning a Shared Bandwidth.....	52
4.3 Adding EIPs to or Removing EIPs from a Shared Bandwidth.....	53
4.4 Removing EIPs from a Shared Bandwidth.....	54
4.5 Modifying a Shared Bandwidth.....	55
4.6 Deleting a Shared Bandwidth.....	55
5 Cloud Eye Monitoring.....	57
5.1 Monitoring EIPs.....	57
5.2 Monitoring Metrics.....	57
5.3 Creating an Alarm Rule.....	60
6 Managing EIP Quotas.....	62
7 FAQs.....	63
7.1 Product Consultation.....	63
7.1.1 What Is the EIP Assignment Policy?.....	63
7.1.2 Why Is an EIP Newly Assigned the Same as the One I Released?.....	63
7.1.3 Can I Assign a Specific EIP?.....	63
7.1.4 Why Can't I Find My Assigned EIP on the Management Console?.....	64
7.1.5 Can a Bandwidth Be Used by Multiple Accounts?.....	64
7.1.6 How Many ECSs Can I Bind an EIP To?.....	64
7.1.7 How Can I Unbind an Existing EIP from an Instance and Bind Another EIP to the Instance?.....	64
7.2 EIP.....	66
7.2.1 What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?.....	66
7.2.2 How Do I Access an ECS with an EIP Bound from the Internet?.....	69
7.2.3 Can I Bind an EIP of an ECS to Another ECS?.....	69
7.2.4 Can I Bind an EIP to a Cloud Resource in Another Region?.....	69
7.2.5 Can Multiple EIPs Be Bound to an ECS?.....	69
7.2.6 How Do I Assign or Retrieve a Specific EIP?.....	71
7.3 Bandwidth.....	71
7.3.1 What Are Inbound Bandwidth and Outbound Bandwidth?.....	71
7.3.2 What Bandwidth Types Are Available?.....	72
7.3.3 What Is the Bandwidth Size Range?.....	72
7.3.4 How Do I Know If My EIP Bandwidth Has Been Exceeded?.....	72
7.3.5 How Do I Increase a Bandwidth to Be More Than 300 Mbit/s?.....	74
7.3.6 How Many EIPs Can I Add to Each Shared Bandwidth?.....	74
7.3.7 Can I Change the Dedicated Bandwidth Used by an EIP to a Shared Bandwidth?.....	75
7.3.8 What Is the Relationship Between Bandwidth and Upload/Download Rate?.....	75
7.3.9 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth? Can a Dedicated Bandwidth Be Changed to a Shared Bandwidth or the Other Way Around?.....	75
7.3.10 What Are the Differences Between Public Bandwidth and Private Bandwidth?.....	77
7.4 Connectivity.....	77
7.4.1 Why Can't My ECS Access the Internet Even After an EIP Is Bound?.....	77
7.4.2 What Should I Do If an EIP Cannot Be Pinged?.....	80
7.4.3 Why Does the Download Speed of My ECS Is Slow?.....	86

7.4.4 What Are the Priorities of the Custom Route and EIP If Both Are Configured for an ECS to Enable the ECS to Access the Internet?..... 86

1 Service Overview

1.1 What Is Elastic IP?

Introduction

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. A resource with an EIP can access the Internet directly, but a resource with only a private IP address cannot.


EIPs can only be bound to resources in the same region.

EIP

EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers.

Each EIP can be bound to only one cloud resource and both should be in the same region.

Figure 1-1 Connecting to the Internet using an EIP



EIP Bandwidth

Bandwidth refers to the maximum amount of data that can be transmitted in a given amount of time (generally one second). A larger bandwidth value indicates a stronger transmission capability. Bandwidth is classified into public bandwidth and private bandwidth.

Public bandwidth is the bandwidth consumed when data is transferred between the cloud platform and the Internet. Public bandwidth is classified into inbound bandwidth and outbound bandwidth. For details about the outbound bandwidth and inbound bandwidth, see [Figure 1-2](#).

- The metrics about outbound bandwidth on the Cloud Eye console are **Outbound Bandwidth** and **Outbound Traffic**.
- The metrics about inbound bandwidth on the Cloud Eye console are **Inbound Bandwidth** and **Inbound Traffic**.

Figure 1-2 Inbound bandwidth and outbound bandwidth




Table 1-1 Inbound bandwidth and outbound bandwidth

Type	Description
Outbound bandwidth	Bandwidth consumed when data is transferred from the cloud platform to the Internet. For example, the outbound bandwidth is consumed when ECSs provide services accessible from the Internet or when FTP clients download resources from the ECSs. The metrics about outbound bandwidth on the Cloud Eye console are Outbound Bandwidth and Outbound Traffic .
Inbound bandwidth	Bandwidth consumed when data is transferred from the Internet to the cloud platform. For example, the inbound bandwidth is consumed when downloading resources from the Internet to ECSs and when FTP clients upload resources to the ECSs. The metrics about inbound bandwidth on the Cloud Eye console are Inbound Bandwidth and Inbound Traffic .

Accessing EIP

You can access EIPs through the management console or using HTTPS-based APIs.

- Management console
Log in to the management console, select **Elastic IP** from the console homepage, and then perform operations on EIP resources.
- APIs
If you need to integrate the EIP service provided by the cloud system into a third-party system for secondary development, you can use an API to access the EIP service. For details, see the *Elastic IP API Reference*.

1.2 Advantages

- Flexibility
EIPs can be flexibly bound to or unbound from ECSs, BMSs, NAT gateways, load balancers, or virtual IP addresses. EIP bandwidths can be scaled according to service changes.
- Cost-effective
EIPs are available on a pay-per-use (billed by bandwidth or traffic) basis. You can use shared bandwidths to enjoy lower bandwidth costs.
- Ease of use
EIP binding, unbinding, and bandwidth adjustments take effect immediately.

1.3 Application Scenarios

EIP Application Scenarios

EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers.

Binding an EIP to an ECS


Scenario

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

Related Services

ECS, BMS, and VPC

Figure 1-3 EIP used by an ECS



Binding an EIP to a NAT Gateway

Scenario

After an EIP is bound to a NAT gateway and SNAT and DNAT rules are added, multiple cloud servers (such as ECSs and BMSs) can use the same EIP to access the Internet and provide services accessible from the Internet.


An SNAT rule allows servers in a specific VPC subnet to use the same EIP to access the Internet.

A DNAT rule enables servers in a VPC to provide services accessible from the Internet.

Related Services

NAT Gateway, ECS, BMS and VPC

Figure 1-4 EIP used by a NAT gateway



Binding an EIP to a Load Balancer

A high-availability load balancing network can be built with the help of EIP.

To handle a large number of concurrent requests from the Internet, you can deploy multiple ECSs in a VPC and use ELB to distribute requests across these servers to improve service stability and availability.


Scenario

After you bind an EIP to a load balancer, the load balancer can distribute requests from the Internet to backend servers.

Related Services

ELB, ECS, and VPC

Figure 1-5 EIP used by a load balancer



1.4 Functions


EIP provides various functions for you to flexibly configure services and build diversified networks.

EIP

The EIP service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths.

You can assign EIPs, bind them to or unbind them from cloud resources, release EIPs, modify EIP bandwidth, and upgrade static BGP EIPs to dynamic BGP EIPs.

Figure 1-6 Connecting to the Internet using an EIP



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

Shared Bandwidth

All ECSs, BMSs, and load balancers can share the same bandwidth if they reside in the same region and have EIPs bound.

You can assign, modify, delete a shared bandwidth, add EIPs to a shared bandwidth, and remove EIPs from a shared bandwidth.

For details, see [Shared Bandwidth Overview](#).

1.5 Notes and Constraints

Constraints on EIP Resources

Note the constraints on the following EIP resources before using them.

- [EIP Overview](#)
- [Shared Bandwidth Overview](#)

1.6 Related Services

EIP and Other Services

Table 1-2 EIP and other services

Service	EIP and Other Services	Reference
Elastic Cloud Server (ECS) Bare Metal Server (BMS)	Bind an EIP to a cloud resource to allow the resource to access the Internet.	Binding an EIP to a Cloud Resource for Internet Access
NAT Gateway	Use a public NAT gateway to enable servers to share one or more EIPs to access the Internet.	Using a Public NAT Gateway to Enable Servers to Share One or More EIPs to Access the Internet
Elastic Load Balance (ELB)	Use load balancers to distribute public traffic to multiple ECSs in a VPC based on domain names and paths.	Using ELB to Distribute Public Traffic Across Multiple Backend Servers in a VPC
Cloud Eye	Check metrics of EIPs on Cloud Eye, such as bandwidth and traffic usage.	Cloud Eye Monitoring

1.7 Billing

Billing Item

EIPs are billed on a pay-per-use basis. [Table 1-3](#) describes the EIP billing items.

Table 1-3 EIP billing

Billing Mode	Billed By	EIP Retention Fee	Bandwidth Price	Public Network Traffic Price
Pay-per-use	Bandwidth	- EIP retention fee is not included if the EIP is bound to an ECS, BMS, or load balancer. - EIP retention fee is included if the EIP is unbound but not released.	Included	Not included
	Traffic		Not included	Included

NOTE

- "Not included" indicates that the fee will not be included in the bill. "Included" indicates that the fee will be included in the bill.
- You can go to the product pricing details page to view details about the EIP pricing.

Billing Options

The public network bandwidth can be billed by fixed bandwidth or by traffic usage.

- By fixed bandwidth: You will be billed based on the bandwidth you specify. Your outbound bandwidth will not exceed the bandwidth specified.
- By traffic usage: You will be billed on a pay-per-use basis. Only traffic used in the outbound direction will be billed.

To prevent excessive fees due to sudden traffic spikes, you can set a peak value for the outbound bandwidth.

You can select the billing mode based on bandwidth usage.

EIPs are billed on a pay-per-use basis. For details, see [Table 1-4](#).

Table 1-4 EIP billing items

Billing Item	Description	Applicable Scenario
Bandwidth Price	The bandwidth is fixed and the traffic is not limited.	For heavy or stable traffic
Public network traffic	Specified maximum bandwidth, billed hourly by the amount of traffic used	For light or sharply fluctuating traffic
EIP retention fee	- If an EIP is bound to an instance, the EIP is free. - If an EIP is unbound from an instance but is not released, the EIP will be billed.	-

Configuration Changes

For a pay-per-use EIP, you can change the bandwidth name, size, and specify whether it is to be billed by bandwidth or traffic.

Table 1-5 Impact on fees

Billing Mode	Change Scenario	Impact on Fees
Pay-per-use	Change the EIP to be billed by traffic or billed by bandwidth.	The change will take effect immediately.

1.8 Basic Concepts

1.8.1 Dedicated Bandwidth

If an EIP is not added to a shared bandwidth, the EIP uses a dedicated bandwidth no matter how it is billed.

You can modify the dedicated bandwidth as required. The modification is applied immediately.

A dedicated bandwidth can control how much data can be transferred using a single EIP. The EIP can be used by only one cloud resource, such as a NAT gateway, an ECS, or a load balancer. Generally, if all EIPs need to access the Internet simultaneously, dedicated bandwidths are recommended preferentially.

How Do I Use Dedicated Bandwidths?

- [Assigning an EIP](#)

- [Modifying an EIP Bandwidth](#)

1.8.2 Shared Bandwidth

A shared bandwidth can be shared by multiple EIPs to control the data transfer rate on these EIPs in a centralized manner.

If ECSs and load balancers with EIPs bound in the same region share the same bandwidth, VPCs and the region-level bandwidth can be managed in a unified manner, simplifying O&M statistics and network operations cost settlement.

NOTE

- A shared bandwidth cannot control how much data can be transferred using a single EIP. Data transfer rate on EIPs cannot be customized.

How Do I Use Shared Bandwidths?

- Assign a shared bandwidth and add your pay-per-use EIPs to the bandwidth.
 - [Assigning a Shared Bandwidth](#)
 - [Adding EIPs to a Shared Bandwidth](#)
- Assign a shared bandwidth, set **Billed By** to **Shared Bandwidth** and select the shared bandwidth when you assign EIPs.
 - [Assigning a Shared Bandwidth](#)
 - [Assigning an EIP](#)

1.8.3 Region and AZ


Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- A region is a physical data center, which is completely isolated to improve fault tolerance and stability. The region that is selected during resource creation cannot be changed after the resource is created.
- An AZ is a physical location where resources use independent power supplies and networks. A region contains one or more AZs that are physically isolated but interconnected through internal networks. Because AZs are isolated from each other, any fault that occurs in one AZ will not affect others.

[Figure 1-7](#) shows the relationship between regions and AZs.

Figure 1-7 Regions and AZs



Selecting a Region

You are advised to select a region close to you or your target users. This helps ensure low access latency.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see [Regions and Endpoints](#).

2 Getting Started

2.1 Overview

If your ECSs need to access the Internet (for example, the ECSs functioning as the service nodes for deploying a website), you can follow the procedure shown in [Figure 2-1](#) to bind EIPs to the ECSs.

Figure 2-1 Configuring the network

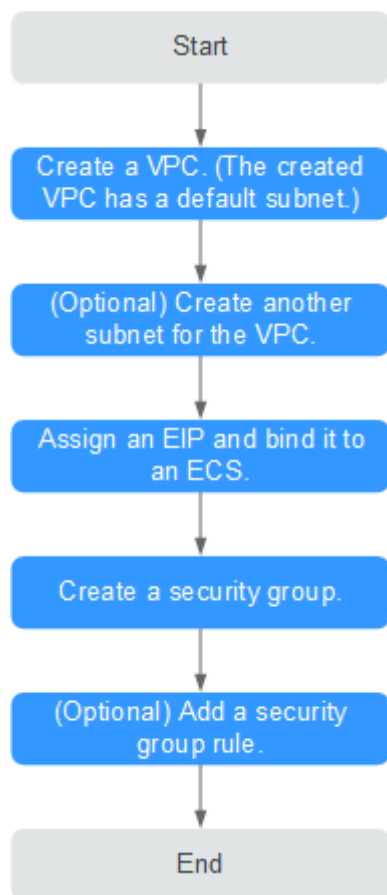


Table 2-1 describes the different tasks in the procedure for configuring the network.

Table 2-1 Configuration process description

Task	Description
Create a VPC.	This task is mandatory. A created VPC comes with a default subnet you specified. After the VPC is created, you can create other required network resources in the VPC based on your service requirements.
Create another subnet for the VPC.	This task is optional. If the default subnet cannot meet your requirements, you can create one. The new subnet is used to assign IP addresses to network interfaces attached to the ECS.
Assign an EIP and bind it to an ECS.	This task is mandatory. You can assign an EIP and bind it to an ECS for Internet access.
Create a security group.	This task is mandatory. You can create a security group and add ECSs in the VPC to the security group to improve ECS access security. After a security group is created, it has default rules, which allow all outgoing data packets. ECSs in a security group can access each other without the need to add rules.
Add a security group rule.	This task is optional. If the default rule does not meet your service requirements, you can add security group rules.

2.2 Step 1: Create a VPC

Scenarios

Virtual Private Cloud (VPC) allows you to provision logically isolated virtual private networks for cloud resources, such as cloud servers, containers, and databases.

You can create a VPC, specify a CIDR block, and create one or more subnets for the VPC. A VPC comes with a default route table that enables subnets in the VPC to communicate with each other.

Procedure

1. Log in to the management console.

2. Click  in the upper left corner and choose **Network > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. Click **Create VPC**.
The **Create VPC** page is displayed.
4. On the **Create VPC** page, set parameters for the VPC and subnets as prompted.

Table 2-2 VPC parameter description

Parameter	Description	Example Value
Region	Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the region nearest to where your services will be accessed.	-
Name	The VPC name. The name: • Can contain 1 to 64 characters. • Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	vpc-test
IPv4 CIDR Block	The CIDR block of the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset of the CIDR block for the VPC (for multiple subnets in the VPC). The following CIDR blocks are supported: • 10.0.0.0/8–24 • 172.16.0.0/12–24 • 192.168.0.0/16–24	10.0.0.0/8
Enterprise Project	The enterprise project to which the VPC belongs. An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is default . For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i> .	default

Parameter	Description	Example Value
Advanced Settings > Tag	The VPC tag. Click  to expand the configuration area and set this parameter. Add tags to help you quickly identify, classify, and search for your VPCs.	• Key: vpc_key1 • Value: vpc-01

Table 2-3 Subnet parameter descriptions

Parameter	Description	Example Value
AZ	An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network. Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.	AZ1
Name	The subnet name. The name: • Can contain 1 to 64 characters. • Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	subnet-01

Parameter	Description	Example Value
IPv4 CIDR Block	The IPv4 CIDR block of the subnet. A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets: Planning the CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to plan the CIDR block in advance based on the number of IP addresses required by your service. - The subnet CIDR block cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address. - The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks from the VPC available for new subnets, which can be a problem when you want to scale out services. Avoiding subnet CIDR block conflicts: If you need to connect two VPCs or connect a VPC to an on-premises data center, the CIDR blocks to be connected cannot be the same.	10.0.0.0/24
IPv6 CIDR Block	After the IPv6 function is enabled, the system automatically assigns an IPv6 CIDR block to the created subnet. Currently, the IPv6 CIDR block cannot be customized. IPv6 cannot be disabled after the subnet is created.	-

Parameter	Description	Example Value
Association Route Table	The default route table with which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. A default route table automatically comes with a VPC. Subnets in the VPC are automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.	-
Advanced Settings > Gateway	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. Retain the default value unless there are special requirements.	10.0.0.1
Advanced Settings > DNS Server Address	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. DNS server addresses allow ECSs in a VPC subnet to communicate with each other using private domain names. You can also directly access cloud services through private DNS servers. If you want to use other public DNS servers for resolution, you can change the default DNS server addresses. Changing the default DNS server addresses may cause communication failures in the subnet.	100.125.x.x

Parameter	Description	Example Value
Advanced Settings > DHCP Lease Time	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client. Unit: Day or hour After you change the DHCP lease time on the console, the change is applied automatically when the DHCP lease of an instance (such as ECS) is renewed. You can wait for the system to renew the lease or manually renew the lease. Renewing lease will not change the IP address used by the instance. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.	-
Advanced Settings > NTP Server Address	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. If you want to add NTP server addresses for a subnet, you can specify NTP Server Address. The IP addresses are added in addition to the default NTP server addresses. • If you add or change the NTP server addresses of a subnet, you need to renew the DHCP lease for or restart all the ECSs in the subnet to make the change take effect immediately. • If the NTP server addresses have been cleared out, restarting the ECSs will not help. You must renew the DHCP lease for all ECSs to make the change take effect immediately.	192.168.2.1
Advanced Settings > Tag	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. Add tags to help you quickly identify, classify, and search for your subnets.	• Key: subnet_key1 • Value: subnet-01

Parameter	Description	Example Value
Advanced Settings > Description	The gateway address of the subnet. Click  to expand the configuration area and set this parameter. Enter the description about the subnet in the text box as required. The subnet description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

5. Click **Create Now**.

Return to the VPC list and view the new VPC.

2.3 Step 2: Create a Subnet for the VPC

Scenarios

A subnet is a unique CIDR block with a range of IP addresses in a VPC. All resources in a VPC must be deployed on subnets.

When creating a VPC, you need to create at least one subnet. If one subnet cannot meet your requirements, you can create more subnets for the VPC.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and choose **Network > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.
4. Click **Create Subnet**.
The **Create Subnet** page is displayed.
5. Set the parameters as prompted.

Table 2-4 Subnet parameter descriptions

Parameter	Description	Example Value
VPC	The VPC for which you want to create a subnet.	vpc-test
Name	The subnet name. The name: • Can contain 1 to 64 characters. • Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	subnet-01

Parameter	Description	Example Value
AZ	An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network. Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.	AZ1
IPv4 CIDR Block	The IPv4 CIDR block of the subnet. A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets: • Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to plan the CIDR block in advance based on the number of IP addresses required by your service. – The subnet CIDR block cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. The first and last three addresses in a subnet are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address. – The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks from the VPC available for new subnets, which can be a problem when you want to scale out services. • Avoiding subnet CIDR block conflicts: Avoid CIDR block conflicts if you need to connect two VPCs or connect a VPC to an on-premises data center.	10.0.0.0/24

Parameter	Description	Example Value
IPv6 CIDR Block	If you select this option, the system automatically assigns an IPv6 CIDR block to the created subnet. Currently, the IPv6 CIDR block cannot be customized. IPv6 CIDR block cannot be disabled after the subnet is created.	-
Association Route Table	A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. A default route table automatically comes with a VPC. Subnets in the VPC are automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.	-
Advanced Settings > Gateway	Click  to expand the configuration area and set this parameter. The gateway address of the subnet. Retain the default value unless there are special requirements.	10.0.0.1
Advanced Settings > DNS Server Address	Click  to expand the configuration area and set this parameter. DNS server addresses allow ECSs in a VPC subnet to communicate with each other using private domain names. You can also directly access cloud services through private DNS servers. If you want to use other public DNS servers for resolution, you can change the default DNS server addresses. Changing the default DNS server addresses may cause communication failures in the subnet.	100.125.x.x

Parameter	Description	Example Value
Advanced Settings > NTP Server Address	Click  to expand the configuration area and set this parameter. If you want to add NTP server addresses for a subnet, you can specify NTP Server Address. The IP addresses here are added in addition to the default NTP server addresses. • If you add or change the NTP server addresses of a subnet, you need to renew the DHCP lease for or restart all the ECSs in the subnet to make the change take effect immediately. • If the NTP server addresses have been cleared out, restarting the ECSs will not help. You must renew the DHCP lease for all ECSs to make the change take effect immediately.	192.168.2.1
Advanced Settings > DHCP Lease Time	Click  to expand the configuration area and set this parameter. The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client. Unit: Day or hour If the time period is changed, the new lease time takes effect when the instance (such as an ECS) in the subnet is renewed next time. You can wait for the instance to be renewed automatically or manually modify the lease time. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.	-
Advanced Settings > Tag	Click  to expand the configuration area and set this parameter. Add tags to help you quickly identify, classify, and search for your subnets.	• Key: subnet_key 1 • Value: subnet-01

Parameter	Description	Example Value
Advanced Settings > Description	Click  to expand the configuration area and set this parameter. Enter the description about the subnet in the text box as required. The subnet description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

6. Click **Create Now**.

Return to the subnet list and view the new subnet.

Constraints

- After a subnet is created, some IP addresses are reserved by the system and cannot be assigned to any instance. For example, in a subnet with CIDR block of 192.168.0.0/24, the following IP addresses are reserved by default:
 - 192.168.0.0: Network ID. This address is the beginning of the private IP address range and will not be assigned to any instance.
 - 192.168.0.1: The gateway address of the subnet.
 - 192.168.0.253: Reserved for the system interface. This IP address is used by the VPC for external communication.
 - 192.168.0.254: DHCP service address.
 - 192.168.0.255: Network broadcast address.

These IP addresses are for reference only. The system assigns reserved IP addresses based on your subnet settings. All other IP addresses in the subnet can be assigned to instances.

2.4 Step 3: Assign an EIP and Bind It to an ECS

Scenarios

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

Assigning a New EIP

- Log in to the management console.
- Click  in the upper left corner and select the desired region and project.
- Click  in the upper left corner and choose **Network > Elastic IP**.
- On the displayed page, click **Assign EIP**.
- Configure parameters as prompted.

Table 2-5 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
EIP Type	Dynamic BGP: Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.	Dynamic BGP
Billed By	The following options are available: • Bandwidth: You specify a maximum bandwidth and pay for the amount of time you use the bandwidth. This is suitable for scenarios with heavy or stable traffic. • Traffic: You specify a maximum bandwidth and pay for the total outbound traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. • Shared Bandwidth: A shared bandwidth can be shared by multiple EIPs. It controls the data transfer rate on these EIPs in a centralized manner. This is suitable for scenarios with staggered traffic.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth
Tag	The EIP tags. Each tag contains a key and value pair. The tag key and value must meet the requirements listed in Table 2-7 .	• Key: ipv4_key1 • Value: 3005eip

Parameter	Description	Example Value
Quantity	The number of EIPs to be assigned.	1
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default. For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i>.	default

Table 2-6 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth

Table 2-7 EIP tag requirements

Parameter	Requirement	Example Value
Key	- Cannot be left blank. - Must be unique for each EIP. - Can contain a maximum of 36 characters. - Can contain letters, digits, underscores (_), and hyphens (-).	Ipv4_key1

Parameter	Requirement	Example Value
Value	- Can contain a maximum of 43 characters. - Can contain letters, digits, underscores (_), periods (.), and hyphens (-).	eip-01

6. Click **Create Now**.
7. Click **Submit**.

Binding an EIP

1. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
2. Select the instance that you want to bind the EIP to.
3. Click **OK**.

2.5 Step 4: Create a Security Group

Scenarios

A security group consists of inbound and outbound rules to control the traffic that is allowed to flow into or out of instances (such as ECSs) in the security group. Security group rules are commonly used to allow or deny network traffic from specific sources or over specific protocols, block certain ports, and define specific access permissions for instances.

When creating an instance (for example, an ECS), you must associate it with a security group. If no security group has been created yet, a default security group will be created and associated with the instance. You can also create a security group and add inbound and outbound rules to allow specific traffic.

Security Group Templates

Several security group templates are provided for you to create a security group. A security group template has preconfigured inbound and outbound rules. You can select a template based on your service requirements. [Table 2-8](#) describes the security group templates.

Table 2-8 Security group rules

Template	Direction	Protocol/Port/Type	Source/Destination	Description	Scenario
General - purpose web server	Inbound	TCP: 22 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in the security group over port 22 (SSH) for remotely logging in to Linux instances.	- Remotely log in to an instance (such as an ECS) in a security group from an external network. - Enable external servers to ping the instances in a security group to verify network connectivity. - Use instances in a security group as web servers to provide website services accessible from the Internet.
		TCP: 3389 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 3389 (RDP) for remotely logging in to Windows instances.	
		TCP: 80 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 80 (HTTP) for visiting websites.	
		TCP: 443 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 443 (HTTPS) for visiting websites.	
		ICMP: All (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over any port for using the ping command to test connectivity.	
		All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network over any protocol and port.	

Template	Direction	Protocol/Port/Type	Source/Destination	Description	Scenario
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources over any protocol and port.	
All ports open	Inbound	All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network over any protocol and port.	Allowing any traffic to enter and leave a security group over any port may be risky.
		All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows any IP address to access the instances in a security group over any protocol and port.	
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources over any protocol and port.	
Fast-add rule	Inbound	All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network.	You can select protocols and ports that the inbound rule will apply to.
		Custom port and protocol	0.0.0.0/0	Allows all IP addresses to access ECSs in a security group over specified ports (TCP or ICMP) for different purposes.	

Template	Direction	Protocol/Port/Type	Source/Destination	Description	Scenario
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources using any protocol.	

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.
The security group list is displayed.
5. In the upper right corner, click **Create Security Group**.
The **Create Security Group** page is displayed.
6. Configure the parameters as prompted.

Figure 2-2 Create Security Group

Create Security Group

*

Name

sg-AB

*

Enterprise Project

default

?

Create Enterprise Project

*

Template

General-purpose web ser...

Description

The security group is for general-purpose web servers and includes default rules that allow all inbound ICMP traffic and inbound traffic on ports 22, 80, 443, and 3389. The security group is used for remote login, ping, and hosting a website on ECSs.

0/255

Hide Default Rule

Inbound

Outbound

Prio...	Action	Type	Protocol & Port	Source
1	Allow	IPv4	TCP: 22	0.0.0.0/0
1	Allow	IPv4	TCP: 3389	0.0.0.0/0
1	Allow	IPv4	TCP: 80	0.0.0.0/0
1	Allow	IPv4	TCP: 443	0.0.0.0/0
1	Allow	IPv4	ICMP: All	0.0.0.0/0
1	Allow	IPv4	All	sg-AB

OK

Cancel

Table 2-9 Parameter description

Parameter	Description	Example Value
Name	Mandatory Enter the security group name. The security group name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces. NOTE You can change the security group name after a security group is created. It is recommended that you give each security group a different name.	sg-AB
Enterprise Project	Mandatory When creating a security group, you can add the security group to an enabled enterprise project. An enterprise project lets you manage cloud resources and personnel by enterprise project. The default project is default. For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i>.	default
Template	Mandatory Several security group templates are provided for you to create a security group. A security group template has preconfigured inbound and outbound rules. You can select a template based on your service requirements.	General-purpose web server
Description	Optional Supplementary information about the security group. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

7. Confirm the inbound and outbound rules of the template and click **OK**.

2.6 Step 5: Add a Security Group Rule

Scenarios

A security group consists of inbound and outbound rules to control the traffic that is allowed to flow into or out of instances (such as ECSs) in the security group. Security group rules are commonly used to allow or deny network traffic from

specific sources or over specific protocols, block certain ports, and define specific access permissions for instances.

Adding Rules to a Security Group

1. Log in to the management console.
2. Click  in the upper left corner and choose **Network > Virtual Private Cloud**.
The VPC list page is displayed.
3. In the navigation pane on the left, choose **Access Control > Security Groups**.
The security group list is displayed.
4. Locate the target security group and click **Manage Rules** in the **Operation** column.
The page for configuring security group rules is displayed.
5. On the **Inbound Rules** tab, click **Add Rule**.
The **Add Inbound Rule** dialog box is displayed.
6. Configure required parameters.
You can click + to add more inbound rules.

Table 2-10 Inbound rule parameter description

Parameter	Description	Example Value
Type	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Inbound rules control incoming traffic over specific ports to instances in the security group.	22, 22-30

Parameter	Description	Example Value
Source	Source of the security group rule. The value can be IP address or Security group, to allow access from the IP addresses or the instances in the security group. If you select IP address for Source, you can enter multiple IP addresses in the IP address box. Each IP address represents a different security group rule. - IP address: - Single IP address: 192.168.10.10/32 - All IP addresses: 0.0.0.0/0 - IP address range: 192.168.1.0/24 If the source is a security group, this rule will apply to all instances associated with the selected security group.	0.0.0.0/0
Description	(Optional) Supplementary information about the security group rule. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

- Click **OK**.
The inbound rule list is displayed.
- On the **Outbound Rules** tab, click **Add Rule**.
The **Add Outbound Rule** dialog box is displayed.
- Configure required parameters.
You can click + to add more outbound rules.

Table 2-11 Outbound rule parameter description

Parameter	Description	Example Value
Type	Destination IP address version. You can select: IPv4 IPv6	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All, TCP, UDP, GRE, or ICMP.	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Outbound rules control outgoing traffic over specific ports from instances in the security group.	22, 22-30

Parameter	Description	Example Value
Destination	Destination of the security group rule. The value can be an IP address or a security group to allow access to IP addresses or instances in the security group. If you select IP address for Destination , you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. - IP address: - Single IP address: 192.168.10.10/32 - All IP addresses: 0.0.0.0/0 - IP address range: 192.168.1.0/24	0.0.0.0/0
Description	(Optional) Supplementary information about the security group rule. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

10. Click **OK**.

The outbound rule list is displayed and you can view your added rule.

Checking Whether a Port Is Enabled on an ECS

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

For example, if you have deployed a website on an ECS and want users to access your website through HTTP (80), you need to add an inbound rule to the ECS security group to allow access over the port. [Table 2-12](#) shows the rule.

Table 2-12 Security group rule

Direction	Type	Protocol & Port	Source
Inbound	IPv4	TCP: 80	IP address: 0.0.0.0/0

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

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

Figure 2-3 Command output for the Linux ECS

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

– **Checking a Windows ECS port**

- i. Choose **Start > Run**. Type **cmd** to open the Command Prompt.
- ii. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | findstr 80

If the following figure is displayed, TCP port 80 is enabled.

Figure 2-4 Command output for the Windows ECS

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

2. Enter **http://ECS EIP** in the address box of the browser and press **Enter**.
If the requested page can be accessed, the security group rule has taken effect.

3 Elastic IP

3.1 EIP Overview

EIP


The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. If a resource has an EIP bound, it can directly access the Internet. If a resource only has a private IP address, it cannot directly access the Internet.

EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers.

Each EIP can be bound to only one cloud resource and both should be in the same region.

You can use public NAT gateways to enable ECSs in the VPC to share an EIP to access or be accessed by the Internet. For details, see *NAT Gateway User Guide*.

Figure 3-1 Connecting to the Internet using an EIP



EIP Quotas

You can log in to the console to query your EIP quotas.

Binding an EIP to an Instance

Figure 3-2 Process for binding an EIP to an instance



Table 3-1 Process for binding an EIP to an instance

No.	Step	Description
1	Assigning an EIP	You can assign an EIP and bind it to cloud resources to allow them to access the Internet.
2	Binding an EIP to an Instance	• The procedure for binding an EIP varies depending on the target instance. • The EIP and the instance to be bound must be in the same region.

3.2 Assigning an EIP

Scenarios

You can assign an EIP and bind it to cloud resources to allow them to access the Internet. This section describes how to assign a new or specific EIP.

- By default, **new EIPs** are assigned at random.
 - If you assign a new EIP within 24 hours after releasing an EIP, the released EIP will be assigned first.
 - Other users can call APIs to assign the released EIP 24 hours after it is released.
- You can call APIs to **assign a specific EIP**.

Assigning a New EIP

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. On the displayed page, click **Assign EIP**.
5. Configure parameters as prompted.

Table 3-2 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
EIP Type	Dynamic BGP: Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.	Dynamic BGP
Billed By	The following options are available: • Bandwidth: You specify a maximum bandwidth and pay for the amount of time you use the bandwidth. This is suitable for scenarios with heavy or stable traffic. • Traffic: You specify a maximum bandwidth and pay for the total outbound traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. • Shared Bandwidth: A shared bandwidth can be shared by multiple EIPs. It controls the data transfer rate on these EIPs in a centralized manner. This is suitable for scenarios with staggered traffic.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth
Tag	The EIP tags. Each tag contains a key and value pair. The tag key and value must meet the requirements listed in Table 3-4 .	• Key: ipv4_key1 • Value: 3005eip

Parameter	Description	Example Value
Quantity	The number of EIPs to be assigned.	1
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default. For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i>.	default

Table 3-3 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth

Table 3-4 EIP tag requirements

Parameter	Requirement	Example Value
Key	- Cannot be left blank. - Must be unique for each EIP. - Can contain a maximum of 36 characters. - Can contain letters, digits, underscores (_), and hyphens (-).	Ipv4_key1

Parameter	Requirement	Example Value
Value	- Can contain a maximum of 43 characters. - Can contain letters, digits, underscores (_), periods (.), and hyphens (-).	eip-01

- Click **Create Now**.
- Click **Submit**.

Assigning a Specific EIP

If you want to retrieve an EIP that you have released within seven days (inclusive) or assign a specific EIP, you can use APIs.

You can set the value of **ip_address** to the one that you want to assign. For details, see section "Assigning an EIP" in *Elastic IP API Reference*.

- If the EIP has been assigned to another user, you will fail to assign your required EIP.
- The management console does not support assigning a specific EIP.

Why Can't I Find My Assigned EIP on the Management Console?

You can perform the following operations to locate an EIP if you cannot find it on the management console.

EIPs Not in the Current Region

- Log in to the management console.
- In the upper left corner of the console, select the region that the EIP to be queried belongs to.
- In the EIP list, view your EIPs.

3.3 Modifying an EIP Bandwidth

Scenarios

No matter which billing mode is used, if your EIP is not added to a shared bandwidth, it uses a dedicated bandwidth. A dedicated bandwidth can control how much data can be transferred using a single EIP.

This section describes how to increase or decrease the dedicated bandwidth size. Changing bandwidth size does not change the EIPs.

Procedure

- Log in to the management console.

2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. Locate the target EIP and choose **More > Modify Bandwidth** in the **Operation** column.
 - If it is a pay-per-use EIP, the **Modify Bandwidth** page is displayed.
5. Modify the bandwidth parameters as prompted.
6. Click **Next**.
7. Click **Submit**.

3.4 Binding or Unbinding an EIP

Binding an EIP to an Instance

Bind EIPs to resources such as ECSs, virtual IP addresses, and load balancers to allow them to access the Internet.

Binding an EIP to an ECS or Virtual IP Address

1. In the EIP list, locate the row that contains the EIP, and click **Bind**.
2. Select the instance.
3. Click **OK**.

NOTE

To bind an instance to an EIP:

- If the instance is an ECS:
 - The ECS must be in the running or stopped status.
 - The ECS must be in the same region as that of the EIP.
 - The ECS has no EIP bound to it.
- If the instance is a virtual IP address:
 - The virtual IP address must be in the same region as that of the EIP.
 - The virtual IP address must be in the available or assigned status.

Binding an EIP to a NAT Gateway

If you want to bind a NAT gateway to an EIP, the NAT gateway must be in the same region as that of the EIP. After an EIP is bound to a NAT gateway, ECSs associated with this gateway can share the EIP to access the Internet or provide services accessible from the Internet.

You can bind an EIP to a NAT gateway by configuring SNAT and DNAT rules for the gateway. Sections "Allowing a Private Network to Access the Internet Using SNAT" and "Allowing Internet Users to Access a Service in a Private Network Using DNAT" in the *NAT Gateway User Guide*.

Binding an EIP to a Load Balancer

If you want to bind a load balancer to an EIP, the load balancer must be in the same region as that of the EIP. Then, the load balancer can receive requests over

the Internet. For details, see "Binding or Unbinding an IPv4 EIP" in the *Elastic Load Balance User Guide*.

Unbinding an EIP from an Instance

If an EIP is no longer required, you can unbind it from your instance.

Unbinding an EIP from an ECS or Virtual IP Address

Unbinding a single EIP

1. Log in to the management console.
2. Click  in the upper left corner and choose **Network > Elastic IP**.
3. On the displayed page, locate the row that contains the target EIP, and click **Unbind** in the **Operation** column.

A confirmation dialog box is displayed.

4. Click **Yes** in the displayed dialog box.
In the EIP list, the target EIP has no associated instance.

Unbinding multiple EIPs at a time

1. Log in to the management console.
2. Click  in the upper left corner and choose **Network > Elastic IP**.
3. On the displayed page, select the EIPs to be unbound.
4. In the upper left corner of the EIP list, click **Unbind**.

A confirmation dialog box is displayed.

5. Click **Yes** in the displayed dialog box.
In the EIP list, the target EIPs have no associated instances.

3.5 Releasing an EIP

Scenarios

If an EIP is no longer required, you can unbind it from your instance and release it if it is a pay-per-use EIP. This section describes how to release an EIP.

Notes and Constraints

- An EIP that has been bound to an instance cannot be released.

Releasing a Pay-per-Use EIP

1. Log in to the management console.
2. Click  in the upper left corner and choose **Network > Elastic IP**.
3. In the EIP list, locate the row that contains the EIP and choose **More > Release** in the **Operation** column.

A confirmation dialog box is displayed.

- Click **Yes** in the displayed dialog box.
You can find that the EIP is not in the EIP list.

3.6 Exporting EIP Information

Scenarios

The information of all EIPs under your account can be exported in an Excel file to a local directory. The file records the ID, status, type, bandwidth name, and bandwidth size of EIPs.

Procedure

- Log in to the management console.
- Click  in the upper left corner and choose **Network > Elastic IP**.
- On the EIP list page, select one or more EIPs and click **Export** in the upper left corner.
The system will automatically export all EIPs to an Excel file and download the file to a local directory.

3.7 Managing EIP Tags

Scenarios

You can add tags to EIPs to help identify and manage them more easily. You can add a tag to an EIP when assigning the EIP. Alternatively, you can add a tag to an assigned EIP on the EIP details page. A maximum of 10 tags can be added to each EIP.

A tag consists of a key and value pair. [Table 3-5](#) lists the tag key and value requirements.

Table 3-5 EIP tag requirements

Parameter	Requirement	Example Value
Key	- Cannot be left blank. - Must be unique for each EIP. - Can contain a maximum of 36 characters. - Can contain letters, digits, underscores (_), and hyphens (-).	Ipv4_key1
Value	- Can contain a maximum of 43 characters. - Can contain letters, digits, underscores (_), periods (.), and hyphens (-).	eip-01

Procedure

Searching for EIPs by tag key and value on the EIP list page

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the search box above the EIP list, click anywhere in the box to set filters.
Select the tag key and then the value as required. The system filters resources based on the tag you select.

Adding, deleting, editing, and viewing tags on the Tags tab of an EIP

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. On the displayed page, locate the EIP whose tags you want to manage and click the EIP name.
5. On the page showing EIP details, click the **Tags** tab and perform desired operations on tags.
 - View tags.
On the **Tags** tab, you can view details about tags added to the current EIP, including the number of tags and the key and value of each tag.
 - Add a tag.
Click **Add Tag** in the upper left corner. In the displayed **Add Tag** dialog box, enter the tag key and value, and click **OK**.
 - Edit a tag.
Locate the row that contains the tag you want to edit, and click **Edit** in the **Operation** column. Enter the new tag value, and click **OK**.
The tag key cannot be modified.
 - Delete a tag.
Locate the row that contains the tag you want to delete, and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

3.8 Assigning an EIP and Binding It to an ECS

Scenarios

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

Assigning a New EIP

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.

3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. On the displayed page, click **Assign EIP**.
5. Configure parameters as prompted.

Table 3-6 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
EIP Type	Dynamic BGP: Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.	Dynamic BGP
Billed By	The following options are available: • Bandwidth: You specify a maximum bandwidth and pay for the amount of time you use the bandwidth. This is suitable for scenarios with heavy or stable traffic. • Traffic: You specify a maximum bandwidth and pay for the total outbound traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. • Shared Bandwidth: A shared bandwidth can be shared by multiple EIPs. It controls the data transfer rate on these EIPs in a centralized manner. This is suitable for scenarios with staggered traffic.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth

Parameter	Description	Example Value
Tag	The EIP tags. Each tag contains a key and value pair. The tag key and value must meet the requirements listed in Table 3-8 .	- Key: ipv4_key1 - Value: 3005eip
Quantity	The number of EIPs to be assigned.	1
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default. For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i>.	default

Table 3-7 Parameter descriptions

Parameter	Description	Example Value
Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-
Bandwidth	The bandwidth size in Mbit/s.	100
Bandwidth Name	The name of the bandwidth.	bandwidth

Table 3-8 EIP tag requirements

Parameter	Requirement	Example Value
Key	• Cannot be left blank. • Must be unique for each EIP. • Can contain a maximum of 36 characters. • Can contain letters, digits, underscores (_), and hyphens (-).	Ipv4_key1
Value	• Can contain a maximum of 43 characters. • Can contain letters, digits, underscores (_), periods (.), and hyphens (-).	eip-01

6. Click **Create Now**.
7. Click **Submit**.

Binding an EIP

1. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
2. Select the instance that you want to bind the EIP to.
3. Click **OK**.

3.9 EIP Configuration Examples

3.9.1 Changing an EIP for an Instance

Scenarios

If you want to change an EIP for an ECS, a load balancer, a NAT gateway, or other cloud resources, you need to unbind the current EIP from the cloud resource first. Then, you can bind a new EIP to the cloud resource to enable Internet access for it.

Changing an EIP for a Cloud Resource

Figure 3-3 Process description

Table 3-9 Process description

No.	Procedure	Description
1	Unbind an EIP	After an EIP is unbound from a cloud resource, the cloud resource can have a new EIP bound for Internet access.
2	Assign a new EIP	If you already have an EIP that you require, skip this step.
3	Bind a new EIP	After a cloud resource has a new EIP bound, it can access the Internet using the new EIP.
4	Release the EIP that has been unbound	• If an unbound EIP still needs to be used, skip this step. • If an unbound EIP is no longer required, you can release it.

Scenario 1: Unbinding an EIP from an ECS and Binding a New EIP to the ECS

1. Unbind an EIP.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and choose **Network > Elastic IP**.
 - c. On the displayed page, locate the row that contains the target EIP, and click **Unbind** in the **Operation** column.
A confirmation dialog box is displayed.
 - d. Click **Yes** in the displayed dialog box.
In the EIP list, the target EIP has no associated instance.

2. Assign an EIP.

NOTE

If you already have an EIP that you require, skip this step.

- a. Log in to the management console.
 - b. Click  in the upper left corner and choose **Network > Elastic IP**.
 - c. On the displayed page, click **Assign EIP**.
 - d. Configure parameters as prompted.
 - e. Click **Next**.
3. Bind the new EIP to the ECS.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and choose **Network > Elastic IP**.
 - c. On the **EIPs** page, locate the target EIP, and click **Bind** in the **Operation** column.
 - d. Select the desired ECS.
 - e. Click **OK**.

4. Release the EIP that is unbound.

 NOTE

If an unbound EIP is no longer required, you can release it.

- a. Log in to the management console.
- b. Click  in the upper left corner and choose **Network > Elastic IP**.
- c. In the EIP list, locate the row that contains the EIP and choose **More > Release** in the **Operation** column.
A confirmation dialog box is displayed.
- d. Click **Yes** in the displayed dialog box.
You can find that the EIP is not in the EIP list.

Scenario 2: Unbinding an EIP from a Load Balancer and Binding a New EIP to the Load Balancer

1. Assign an EIP by referring to [2](#).

 NOTE

If you already have an EIP that you require, skip this step.

2. Unbind an EIP from a load balancer and bind the new EIP to the load balancer.
 - a. On the displayed page, locate the row that contains the load balancer and click **More** in the **Operation** column.
 - i. Unbinding an IPv4 EIP
 - 1) Click **Unbind IPv4 EIP**.
 - 2) In the displayed dialog box, confirm the IPv4 EIP that you want to unbind and click **OK**.
 - ii. Binding an IPv4 EIP
 - 1) Click **Bind IPv4 EIP**.
 - 2) In the **Bind IPv4 EIP** dialog box, select the EIP you want to bind to the load balancer and click **OK**.
3. Release the EIP that was replaced. For details, see [4](#).

 NOTE

If an unbound EIP is no longer required, you can release it.

Scenario 3: Unbinding an EIP from a NAT Gateway and Binding a New EIP to the NAT Gateway

1. Assign an EIP by referring to [2](#).

 NOTE

If you already have an EIP that you require, skip this step.

2. Modify an SNAT rule.

For details, see section "Modifying an SNAT Rule" of a public NAT gateway in *NAT Gateway User Guide*. In the EIP list, select the new EIP and deselect the existing EIP.

3. Modify a DNAT rule.

For details, see section "Modifying a DNAT Rule" of a public NAT gateway in the *NAT Gateway User Guide*. Select the newly assigned EIP.

4. Release the EIP that was replaced. For details, see [4](#).

 NOTE

If an unbound EIP is no longer required, you can release it.

4 Shared Bandwidth

4.1 Shared Bandwidth Overview

A shared bandwidth can be shared by multiple EIPs and controls the data transfer rate on these EIPs in a centralized manner. All ECSs and load balancers that have EIPs bound in the same region can share a bandwidth.

When you host a large number of applications on the cloud, if each EIP uses a bandwidth, a lot of bandwidths are required, increasing O&M workload. If all EIPs share the same bandwidth, VPCs and the region-level bandwidth can be managed in a unified manner, simplifying O&M statistics and network operations cost settlement.

NOTE

- A shared bandwidth cannot control how much data can be transferred using a single EIP. Data transfer rate on EIPs cannot be customized.

Advantages

- Lowered Bandwidth Costs
Region-level bandwidth sharing and multiplexing reduce bandwidth usage and O&M costs.
- Easy to Manage
Region-level bandwidth sharing and multiplexing simplify O&M statistics, management, and operations cost settlement.
- Flexible Operations
You can add pay-per-use EIPs (except for **5_gray** EIPs of dedicated load balancers) to or remove them from a shared bandwidth regardless of the type of instances that they are bound to.

Methods of Using a Shared Bandwidth

You can use the shared bandwidth in either of the two ways shown in the following table.

Table 4-1 Methods of using a shared bandwidth

Description	Step
Method 1: Assign a shared bandwidth and add your pay-per-use EIPs to the bandwidth.	1. Assigning a Shared Bandwidth 2. Adding EIPs to or Removing EIPs from a Shared Bandwidth
Method 2: Assign a shared bandwidth, set Billed By to Shared Bandwidth and select the shared bandwidth when you assign pay-per-use EIPs.	1. Assigning a Shared Bandwidth 2. Assigning an EIP

4.2 Assigning a Shared Bandwidth

Scenarios

When you host a large number of applications on the cloud, if each EIP uses a dedicated bandwidth, a lot of bandwidths are required, which incurs high costs. If all EIPs share the same bandwidth, your network operation costs will be lowered and your system O&M as well as resource statistics will be simplified.

Assign a shared bandwidth for use with EIPs.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the upper right corner, click **Assign Shared Bandwidth**. On the displayed page, configure parameters as prompted.
6. Set the parameters as prompted.

Table 4-2 Parameter descriptions

Parameter	Description	Example Value
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across different regions. For lower network latency and faster access to your resources, select the region nearest to you.	N/A
Billed By	The billing method for the shared bandwidth. You can specify a shared bandwidth to be billed by bandwidth.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s. The maximum bandwidth can be 300 Mbit/s.	10
Name	The name of the shared bandwidth.	Bandwidth-001
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is default . For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i> .	default

7. Click **Create Now**.

4.3 Adding EIPs to or Removing EIPs from a Shared Bandwidth

Scenarios

You can add multiple EIPs to a shared bandwidth or remove EIPs that are no longer required from a shared bandwidth.

You can add multiple EIPs to a shared bandwidth at the same time.

Constraints

- If it is a premium shared bandwidth, you can add premium BGP EIPs and IPv6 NICs to it.

Adding EIPs to a Shared Bandwidth

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the target shared bandwidth that you want to add EIPs to. In the **Operation** column, choose **Add Public IP Address**, and select the EIPs to be added.

NOTE

- After an EIP is added to a shared bandwidth, the dedicated bandwidth used by the EIP will become invalid and the EIP will start to use the shared bandwidth. The EIP's dedicated bandwidth will be deleted and will no longer be billed.
6. Click **OK**.

Removing EIPs from a Shared Bandwidth

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the row that contains the bandwidth from which EIPs are to be removed, choose **More > Remove Public IP Address** in the **Operation** column.
6. On the **Remove Public IP Address** page, select the EIPs to be removed.
7. Click **OK**.

4.4 Removing EIPs from a Shared Bandwidth

Scenarios

You can remove EIPs that are no longer required from a shared bandwidth if needed.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.

5. In the shared bandwidth list, locate the row that contains the bandwidth from which EIPs are to be removed, choose **More > Remove Public IP Address** in the **Operation** column, and select the EIPs to be removed in the displayed dialog box.
6. Click **OK**.

4.5 Modifying a Shared Bandwidth

Scenarios

You can modify the name and size of a shared bandwidth as required.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the row that contains the shared bandwidth you want to modify, click **Modify Bandwidth** in the **Operation** column, and modify the bandwidth settings.
6. Click **Next**.
7. Click **Submit**.

4.6 Deleting a Shared Bandwidth

Scenarios

Delete a shared bandwidth when it is no longer required.

Notes and Constraints

If you want to delete a shared bandwidth with EIPs added, you have to [remove the EIPs from the shared bandwidth](#) first.

Deleting a Pay-per-Use Shared Bandwidth

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Network > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the row that contains the shared bandwidth you want to delete, click **More** in the **Operation** column, and then click **Delete**.

6. Click **OK**.

5 Cloud Eye Monitoring

5.1 Monitoring EIPs

Scenario

Cloud Eye is a multi-dimensional resource monitoring service that you can use to monitor EIP and bandwidths in real time, set alarm rules, identify resource exceptions, and quickly respond to resource changes.

Viewing Metrics

1. Log in to the management console.
2. In the upper left corner of the page, click  to open the service list and choose **Management & Deployment > Cloud Eye**.
3. Click **Cloud Service Monitoring** on the left navigation pane, and choose **Elastic IP and Bandwidth**.
4. Locate the target metric and click **View Metric** in the **Operation** column to check detailed information.

5.2 Monitoring Metrics

Overview

This section describes the namespace, list, and measurement dimensions of metrics of EIPs and bandwidths that you can check on Cloud Eye. You can use APIs or the Cloud Eye console to query the metrics of the monitored metrics and generated alarms.

 NOTE

- Cloud Eye can monitor dimensions nested to a maximum depth of four levels (levels 0 to 3). 3 is the deepest level. For example, if the dimension information of a monitoring metric is "bandwidth_id,publicip_id", the dimension of "bandwidth_id" is numbered 0 and that of "publicip_id" is 1.
- If a bandwidth is increased or decreased, there is a delay of 5 to 10 minutes for the monitoring metrics to update for the new bandwidth.

Namespace

Namespace of EIPs and bandwidths: SYS.VPC

Monitoring Metrics

Table 5-1 EIP and bandwidth metrics

ID	Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
upstream_bandwidth	Outbound Bandwidth	Network rate of outbound traffic (Previously called "Upstream Bandwidth")	≥ 0	bit/s	1000 (SI)	bandwidth_id,publicip_id	1 minute
downstream_bandwidth	Inbound Bandwidth	Network rate of inbound traffic (Previously called "Downstream Bandwidth")	≥ 0	bit/s	1000 (SI)	bandwidth_id,publicip_id	1 minute

ID	Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
upstream_bandwidth_usage	Outbound Bandwidth Usage	Usage of outbound bandwidth in the unit of percent. Outbound bandwidth usage = Outbound bandwidth/ Purchased bandwidth	0-100	%	N/A	bandwidth_id,publicip_id	1 minute
downstream_bandwidth_usage	Inbound Bandwidth Usage	Usage of inbound bandwidth in the unit of percent. Inbound bandwidth usage = Inbound bandwidth/ Purchased bandwidth	0-100	%	N/A	bandwidth_id,publicip_id	1 minute
upstream	Outbound Traffic	Network traffic going out of the cloud platform (Previously called "Upstream Traffic")	≥ 0	Byte	1000 (SI)	bandwidth_id,publicip_id	1 minute
downstream	Inbound Traffic	Network traffic going into the cloud platform (Previously called "Downstream Traffic")	≥ 0	Byte	1000 (SI)	bandwidth_id,publicip_id	1 minute

If an object is in a hierarchical system, specify the monitored dimension in hierarchical form when you use an API to query the metrics of this object.

- To query a single metric by calling an API, the **mount_point** dimension is used as follows:
`dim.0=bandwidth_id,530cd6b0-86d7-4818-837f-935f6a27414d&dim.1=publicip_id,3773b058-5b4f-4366-9035-9bbd9964714a`
530cd6b0-86d7-4818-837f-935f6a27414d and **3773b058-5b4f-4366-9035-9bbd9964714a** are the dimension values of **bandwidth_id** and **publicip_id**, respectively. For details about how to obtain the values, see [Dimensions](#).
- To query multiple metrics by calling an API, the **mount_point** dimension is used as follows:

```
"dimensions": [  
  {  
    "name": "bandwidth_id",  
    "value": "530cd6b0-86d7-4818-837f-935f6a27414d"  
  },  
  {  
    "name": "publicip_id",  
    "value": "3773b058-5b4f-4366-9035-9bbd9964714a"  
  }  
]
```


530cd6b0-86d7-4818-837f-935f6a27414d and **3773b058-5b4f-4366-9035-9bbd9964714a** are the dimension values of **bandwidth_id** and **publicip_id**, respectively. For details about how to obtain the values, see [Dimensions](#).

Dimensions

Key	Value
publicip_id	EIP ID
bandwidth_id	Bandwidth ID

5.3 Creating an Alarm Rule

Scenarios

Cloud Eye allows you to use alarm templates to create alarm rules to monitor cloud resource usage and key operations. After an alarm rule is created, if a metric reaches the specified threshold or there is a specified event, Cloud Eye immediately informs you of the exception through Simple Message Notification (SMN).

This section describes how to create alarm rules to monitor **metrics**.

Creating an Alarm Rule

1. Log in to the management console.
2. In the upper left corner of the page, click  to open the service list and choose **Management & Deployment > Cloud Eye**.
3. In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
4. On the **Alarm Rules** page, click **Create Alarm Rule** or modify an existing alarm rule.
5. Configure the parameters and click **Create**.
After the alarm rule is configured, the system notifies you when an alarm is triggered.

NOTE

For more information about alarm rules, see *Cloud Eye User Guide*.

6 Managing EIP Quotas

What Is a Quota?

A quota limits the quantity of a resource available to users, thereby preventing spikes in the usage of the resource. For example, an EIP quota limits the number of EIPs that can be assigned.

You can also request for an increased quota if your existing quota cannot meet your service requirements.

How Do I View My Quotas?

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, click  .
The **Quotas** page is displayed.
4. View the used and total quota of each type of resources on the displayed page.
If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

1. Log in to the management console.
2. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Quotas** page is displayed.
3. Click **Increase Quota** in the upper right corner of the page.
4. On the **Create Service Ticket** page, configure parameters as required.
In the **Problem Description** area, fill in the content and reason for adjustment.
5. After all necessary parameters are configured, select **I have read and agree to the Ticket Service Protocol and Privacy Statement** and click **Submit**.

7 FAQs

7.1 Product Consultation

7.1.1 What Is the EIP Assignment Policy?

By default, EIPs are assigned randomly.

If an EIP is released by mistake, the system will preferentially assign you an EIP that you have released in the last 24 hours.

If you want a specific EIP that you released more than 24 hours ago, see [How Do I Assign or Retrieve a Specific EIP?](#)

If you do not want an EIP that you have released, it is recommended that you assign another EIP first and then release the one that you do not need.

7.1.2 Why Is an EIP Newly Assigned the Same as the One I Released?

If you have released EIPs in a region, the system preferentially assigns EIPs from the ones you released in the last 24 hours.

If you do not want an EIP that you have released, assign an EIP first and then release the one that you do not want.

You can assign a specific EIP by calling APIs. For details, see section "Assigning an EIP" in *Elastic IP API Reference*.

7.1.3 Can I Assign a Specific EIP?

By default, EIPs are assigned randomly.

- If you assign a new EIP within 24 hours after an EIP is released, the released EIP will be assigned first.
- Other users can call APIs to assign the released EIP 24 hours after it is released.

You can assign a specific EIP only by calling an API. For details, see section "Assigning an EIP" in the *Elastic IP API Reference*.

7.1.4 Why Can't I Find My Assigned EIP on the Management Console?

Symptom

After I logged in to the management console, I could not find my assigned EIP.

Possible Cause

Your EIP is not in the current region. For details, see [EIP Not in the Current Region](#).

EIP Not in the Current Region

Step 1 Log in to the management console.

Step 2 Locate the EIP.

1. In the upper left corner of the console, select the region to which the EIP to be queried belongs.
2. In the EIP list, view the assigned EIP.

----End

7.1.5 Can a Bandwidth Be Used by Multiple Accounts?

A bandwidth cannot be shared between different accounts. Each account can use and manage only its own EIP bandwidths.

7.1.6 How Many ECSs Can I Bind an EIP To?

An EIP can be bound to only one ECS.

An EIP cannot be shared by multiple ECSs, and the EIP and ECS must be in the same region. You can use public NAT gateways to enable the ECSs in the VPC to share an EIP to access or be accessed by the Internet.

For more information, see the *NAT Gateway User Guide*.

7.1.7 How Can I Unbind an Existing EIP from an Instance and Bind Another EIP to the Instance?

Scenario 1: Unbinding an EIP from an ECS and Binding a New EIP to the ECS

1. Unbind an EIP.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and choose **Network > Elastic IP**.
 - c. On the displayed page, locate the row that contains the target EIP, and click **Unbind**.

- d. Click **Yes**.
- 2. Assign an EIP.

 NOTE

If you already have an EIP that you require, skip this step.

- a. Log in to the management console.
- b. Click  in the upper left corner and choose **Network > Elastic IP**.
- c. On the displayed page, click **Assign EIP**.
- d. Set the parameters as prompted.
- e. Click **Next**.
- 3. Bind the new EIP to the ECS.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and choose **Network > Elastic IP**.
 - c. In the EIP list, locate the row that contains the EIP, and click **Bind**.
 - d. Select the desired ECS.
 - e. Click **OK**.
- 4. Release the EIP that is unbound.

 NOTE

If an unbound EIP is no longer required, you can release it.

- a. Log in to the management console.
- b. Click  in the upper left corner and choose **Network > Elastic IP**.
- c. In the EIP list, locate the row that contains the EIP, and choose **More > Release** in the **Operation** column.
- d. Click **Yes**.

Scenario 2: Unbinding an EIP from a Load Balancer and Binding a New EIP to the Load Balancer

- 1. Unbind an EIP.
 - a. Log in to the management console.
 - b. Click **Service List**. Under **Network**, click **Elastic Load Balance**.
 - c. In the load balancer list, locate the target load balancer and choose **More > Unbind EIP** in the **Operation** column.
 - d. Click **Yes**.
- 2. Assign an EIP by referring to [2](#).

 NOTE

If you already have an EIP that you require, skip this step.

- 3. Bind the new EIP to the load balancer.
 - a. Log in to the management console.
 - b. Click **Service List**. Under **Network**, click **Elastic Load Balance**.

- c. In the load balancer list, locate the target load balancer and choose **More** > **Bind EIP** in the **Operation** column.
- d. In the **Bind EIP** dialog box, select the EIP to be bound and click **OK**.
4. Release the EIP that was replaced. For details, see [4](#).

 NOTE

If an unbound EIP is no longer required, you can release it.

Scenario 3: Unbinding an EIP from a NAT Gateway and Binding a New EIP to the NAT Gateway

1. Assign an EIP by referring to [2](#).

 NOTE

If you already have an EIP that you require, skip this step.

2. Modify an SNAT rule.

For details, see section "Modifying an SNAT Rule" of a public NAT gateway in *NAT Gateway User Guide*. In the EIP list, select the new EIP and deselect the existing EIP.

3. Modify a DNAT rule.

For details, see section "Modifying a DNAT Rule" of a public NAT gateway in the *NAT Gateway User Guide*. Select the newly assigned EIP.

4. Release the EIP that was replaced. For details, see [4](#).

 NOTE

If an unbound EIP is no longer required, you can release it.

7.2 EIP

7.2.1 What Are the Differences Between EIPs, Private IP Addresses, and Virtual IP Addresses?

Different types of IP addresses have different functions.

Figure 7-1 IP address architecture




Table 7-1 Functions of different IP address types

IP Address Type	Description	Example Value
Private IP address	Private IP addresses come with your ECSs and belong to the VPC subnets of the ECSs. They are used for private communication on the cloud.	- Private IP address of ECS-A-01: 172.16.0.84 - Private IP address of ECS-B-01: 172.16.1.12

IP Address Type	Description	Example Value
Virtual IP address	A virtual IP address is a private IP address independently assigned from a VPC subnet. It can be released when no longer needed. You can: - Bind one or more virtual IP addresses to a cloud server so that you can use either the virtual or private IP address to access the server. If you have multiple services running on a cloud server, you can use different virtual IP addresses to access them. - Bind a virtual IP address to multiple cloud servers. You can use a virtual IP address and an HA software (such as Keepalived) to set up a high-availability active/standby cluster. If you want to improve service availability and avoid single points of failure, you can deploy cloud servers in the active/standby mode or deploy one active cloud server and multiple standby cloud servers. In this arrangement, the cloud servers all use the same virtual IP address. If the active cloud server goes down, the standby cloud server becomes the active server and continues to provide services. For more information about virtual IP addresses, see section "Virtual IP Address Overview" in the <i>Virtual Private Cloud User Guide</i>.	Bind virtual IP address (172.16.0.2) both ECS-A-01 and ECS-A-02. The active/standby switchover of ECS-A-01 and ECS-A-02 can be implemented by using Keepalived.
EIP	EIPs allow cloud resources to access the Internet. They can be flexibly bound to or unbound from instances. - You can bind an EIP to a virtual IP address to enable the ECSs with the virtual IP address bound to access the Internet. - You can also bind an EIP to the ECSs to enable them to access the Internet.	- Bind EIP (122.9.9.85) to virtual IP address (172.16.0.2) to enable ECS-A-01 and ECS-A-02 to access the Internet. - Bind EIP (122.9.9.87) to ECS-B-01 to enable ECS-B-01 to access the Internet.

7.2.2 How Do I Access an ECS with an EIP Bound from the Internet?

Each ECS is automatically added to a security group after being created to ensure its security. The security group denies access traffic from the Internet by default (except TCP traffic from port 22 through SSH to a Linux ECS and TCP traffic from port 3389 through RDP to a Windows ECS). To allow external access to ECSs in the security group, add an inbound rule to the security group.

You can set **Protocol** to **TCP**, **UDP**, **ICMP**, or **All** as required on the page for creating a security group rule.

- If your ECS needs to be accessible over the Internet and you know the IP address used to access the ECS, set **Source** to the IP address range containing the IP address.
- If your ECS needs to be accessible over the Internet but you do not know the IP address used to access the ECS, retain the default setting 0.0.0.0/0 for **Source**, and then set allowed ports to improve network security.

The default source **0.0.0.0/0** indicates that all IP addresses can access ECSs in the security group.

- Allocate ECSs that have different Internet access requirements to different security groups.

7.2.3 Can I Bind an EIP of an ECS to Another ECS?

Yes.

You can unbind the EIP from the original ECS. For details, see section "Unbinding an EIP from an ECS and Releasing the EIP" in the *Elastic IP User Guide*.

Then, bind the EIP to the target ECS. For details, see section "Assigning an EIP and Binding It to an ECS" in the *Elastic IP User Guide*.

7.2.4 Can I Bind an EIP to a Cloud Resource in Another Region?

An EIP cannot be bound to a cloud resource in another region.

The EIP and the cloud resource must be in the same region.

7.2.5 Can Multiple EIPs Be Bound to an ECS?

Scenarios

Multiple EIPs can be bound to an ECS, but this operation is not recommended.

If an ECS has multiple network interfaces attached and you want to bind multiple EIPs to this ECS, you need to configure policy-based routes for these network interfaces so that these extension network interfaces can communicate with external networks. For details, see [Configuration Example](#).

Configuration Example

[Table 7-2](#) lists ECS configurations.

Table 7-2 ECS configurations

Parameter	Configuration
Name	ecs_test
Image	CentOS 6.5 64bit
EIP	2
Primary NIC	eth0
Secondary NIC	eth1

Example 1:

If you intend to access public network 11.11.11.0/24 through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.
2. Run the following command to configure a route:
ip route add 11.11.11.0/24 dev eth1 via 192.168.2.1

In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

Example 2:

Based on example 1, if you intend to enable routing for default public network traffic through standby NIC **eth1**, perform the following operations to configure a route:

1. Log in to the ECS.
2. Run the following command to delete the default route:
ip route delete default

NOTICE

Exercise caution when deleting the default route because this operation will interrupt the network and result in SSH login failures.

3. Run the following command to configure a new default route:
ip route add 0.0.0.0/0 dev eth1 via 192.168.2.1
In the preceding command, **192.168.2.1** is the gateway IP address of standby NIC **eth1**.

7.2.6 How Do I Assign or Retrieve a Specific EIP?

If you want to retrieve an EIP that you have released or assign a specific EIP, you can use APIs by setting the value of **ip_address** to the one that you want to assign. For details, see section "Assigning an EIP" in *Elastic IP API Reference*.

NOTE

- If the EIP has been assigned to another user, you will fail to assign your required EIP.
- You cannot use the management console to assign a specific EIP.

7.3 Bandwidth

7.3.1 What Are Inbound Bandwidth and Outbound Bandwidth?

Bandwidth refers to the maximum amount of data that can be transmitted in a given amount of time (generally one second). A larger bandwidth value indicates a stronger transmission capability. Bandwidth is classified into public bandwidth and private bandwidth.

Public bandwidth is the bandwidth consumed when data is transferred between cloud instances and the Internet. Public bandwidth is classified into inbound bandwidth and outbound bandwidth. For details about the outbound bandwidth and inbound bandwidth, see [Table 7-3](#).

- **Outbound Bandwidth** means the same thing as **Upstream Bandwidth** or **Upstream Traffic** on the Cloud Eye console.
- **Inbound Bandwidth** means the same thing as **Downstream Bandwidth** and **Downstream Traffic** on the Cloud Eye console.

Figure 7-2 Inbound bandwidth and outbound bandwidth




Table 7-3 Inbound bandwidth and outbound bandwidth

Type	Description
Outbound bandwidth	Bandwidth consumed when data is transferred from cloud to the Internet. For example, the outbound bandwidth is used when ECSs provide services accessible from the Internet and FTP clients download resources from the ECSs. Outbound bandwidth means the same thing as upstream bandwidth on the Cloud Eye console.
Inbound bandwidth	Bandwidth consumed when data is transferred from the Internet to cloud. For example, the inbound bandwidth is used when resources are downloaded from the Internet to ECSs and FTP clients upload resources to the ECSs. Inbound bandwidth means the same thing as downstream bandwidth on the Cloud Eye console.

7.3.2 What Bandwidth Types Are Available?

There are dedicated or shared bandwidths.

If an EIP is not added to a shared bandwidth, the EIP uses the dedicated bandwidth no matter how it is billed.

- Dedicated bandwidths can be used by only one EIP.
- Shared bandwidths can be used by multiple EIPs.

7.3.3 What Is the Bandwidth Size Range?

The bandwidth range is from 1 Mbit/s to 300 Mbit/s.

7.3.4 How Do I Know If My EIP Bandwidth Has Been Exceeded?

Symptom

The bandwidth size configured when you assign a dedicated or shared bandwidth defines the maximum amount of outbound bandwidth supported. If an ECS running your web applications cannot be accessed smoothly from the Internet, check whether the bandwidth of the EIP bound to the ECS is greater than the configured bandwidth size.


NOTE

If the bandwidth exceeds the configured bandwidth size, there may be packet loss or remote login failure to an ECS. To prevent data loss, it is recommended that you monitor the bandwidth.

Troubleshooting

Troubleshoot the issue by following the procedure described below.

Figure 7-3 Troubleshooting procedure



Step 1 Check whether the EIP bandwidth usage has exceeded the configured size.

1. Check the monitoring data on the EIP bandwidth usage.
Check whether the inbound bandwidth and outbound bandwidth usage have exceeded the amount purchased.
2. Check **EIP bandwidth overflow** event.

If the bandwidth usage goes too high for a little while but it does not interrupt your services, ignore the problem. If the bandwidth usage goes too high many times or if the issue lasts for a long time, fix the problem as described in [Step 2](#).

Step 2 Fix the excessive bandwidth usage issue.

Traffic surges may cause the bandwidth to go beyond of the configured limit, causing packet loss.

Check whether the sudden traffic surge is normal.

1. If the traffic surge is normal, increase the bandwidth. For details, see .
2. If the traffic surge is not normal, for example, the surge was caused by attacks, refer to [System Processes Leading to High Bandwidth Usage](#).

Step 3 Check the alarm rule settings and EIP connectivity if the bandwidth usage has not exceeded the configured limit.

After doing the checks in [Step 1](#), if the bandwidth usage has not exceeded the configured limit or the purchased bandwidth:

- Check whether Cloud Eye alarm rules are configured properly.
If there are alarm rules for bandwidth usage on the Cloud Eye console, where the outbound bandwidth limit or the alarm period is set too small, the system may generate excessive alarms. You can refer to [Improper Cloud Eye Alarm Rules](#) to fix the problem.
- Check whether the EIP connection has failed.
If an ECS with an EIP bound cannot access the Internet, you can refer to section "Why Can't My ECS Access the Internet Even After an EIP Is Bound?" in the *Elastic IP User Guide*.

----End

System Processes Leading to High Bandwidth Usage

If some heavy-duty system processes or applications running on your ECS are causing the high bandwidth or CPU usage, your ECS will run slowly or may unexpectedly be inaccessible.

You can refer to the following to locate processes that have led to excessively high bandwidth or CPU usage, and optimize or stop the processes.

- Windows: Section "Why Is My Windows ECS Running Slowly?" in the *Elastic Cloud Server User Guide*
- Linux: Section "Why Is My Linux ECS Running Slowly?" in the *Elastic Cloud Server User Guide*.

Improper Cloud Eye Alarm Rules

If there are alarm rules for bandwidth usage on the Cloud Eye console, where the outbound bandwidth limit or the alarm period is set too small, the system may generate excessive alarms.

You need to set an appropriate alarm rule based on your assigned bandwidth. For example, if your purchased bandwidth is 5 Mbit/s, you can create an alarm rule to report an alarm if the outbound bandwidth reaches 4.8 Mbit/s for three consecutive periods. You can also increase your bandwidth. For details, see section "Modifying an EIP Bandwidth" in the *Elastic IP User Guide*. To create an alarm rule:

1. In the left navigation pane of the **Cloud Eye** console, choose **Alarm Management > Alarm Rules**.
2. Click **Create Alarm Rule** and configure an alarm rule to generate alarms when the bandwidth usage exceeds the configured limit.

7.3.5 How Do I Increase a Bandwidth to Be More Than 300 Mbit/s?

Symptom

The bandwidth of a pay-per-use EIP billed by traffic cannot be increased to be more than 300 Mbit/s.

Solution

The bandwidth of an EIP billed by traffic can be increased to a maximum of 300 Mbit/s. If a higher bandwidth is required, you need to change the EIP to be billed by bandwidth. Then, in certain regions, your bandwidth can be increased to a maximum of 2000 Mbit/s.

If your bandwidth usage is high, billing by bandwidth is more cost-effective than billing by traffic.

7.3.6 How Many EIPs Can I Add to Each Shared Bandwidth?

A shared bandwidth can be used by multiple EIPs.

By default, you can add a maximum of 20 EIPs to a shared bandwidth.

7.3.7 Can I Change the Dedicated Bandwidth Used by an EIP to a Shared Bandwidth?

Yes.

You cannot change the dedicated bandwidth used by an EIP to a shared bandwidth.

7.3.8 What Is the Relationship Between Bandwidth and Upload/Download Rate?

The bandwidth is measured in bit/s, indicating the number of binary bits transmitted per second. The download rate is measured in byte/s, indicating the number of bytes transmitted per second.

1 byte = 8 bits, that is, download rate = bandwidth/8

Due to various issues such as computer performance, network device quality, resource usage, and network peak hours, if the bandwidth is 1 Mbit/s, the actual upload or download rate is generally lower than 125 kByte/s (1 Mbit/s = 1,000 Kbit/s, upload or download rate = 1,000/8 = 125 kByte/s).

7.3.9 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth? Can a Dedicated Bandwidth Be Changed to a Shared Bandwidth or the Other Way Around?

You can select a dedicated or shared bandwidth based on your requirements by referring to [Table 7-4](#).

Table 7-4 Differences between dedicated and shared bandwidths

Item	Dedicated Bandwidth	Shared Bandwidth
Concepts	If you assign an EIP and do not add it to a shared bandwidth, the EIP uses a dedicated bandwidth by default no matter how the EIP is billed. A dedicated bandwidth can only be used by one EIP. Each EIP can only be bound to one cloud resource, such as an ECS, a NAT gateway, or a load balancer.	A shared bandwidth can be used by multiple EIPs. • The shared bandwidth is dynamically allocated to the EIPs based on the actual usage conditions. • Adding an EIP to or removing an EIP from a shared bandwidth does not affect your services.

Item	Dedicated Bandwidth	Shared Bandwidth
Features	- Stable performance: The bandwidth is dedicated, so your use of the bandwidth is not affected by other resources. This type of bandwidth is ideal for applications requiring high-performance networks. - Quality of service (QoS): Guaranteed stable bandwidths and low latency are suitable for real-time applications.	- Cost-effectiveness: Multiple EIPs sharing the same bandwidth can effectively reduce the overall costs. This type of bandwidth is suitable for users with limited budgets. - Flexibility: You can dynamically adjust the size of a shared bandwidth based on your requirements. - Performance fluctuation: When the bandwidth is used by multiple EIPs at the same time, the bandwidth allocated to each EIP is limited.
Applicable scenarios	- Bandwidth preemption needs to be avoided to ensure Internet access for all EIP at the same time. - High-performance and stable bandwidth is required, such as video streaming, online gaming, and financial transactions.	- Internet access needs to be scheduled at different times to optimize bandwidth usage. - There are no demanding requirements on bandwidth or multiple EIPs need to be used at the same time, such as web servers and test environments.
Changes between dedicated and shared bandwidths	A dedicated bandwidth cannot be changed to a shared bandwidth and vice versa. However, you can purchase a shared bandwidth for EIPs. - Add an EIP to a shared bandwidth and then the EIP will use the shared bandwidth. - Remove the EIP from the shared bandwidth and then the EIP will use the dedicated bandwidth.	

Bandwidth Preemption Descriptions

- Dedicated bandwidth:** Each EIP has a fixed bandwidth and is not affected by other EIPs.
For example, both EIP-A and EIP-B are allocated 20 Mbit/s of dedicated bandwidth.
If the bandwidth of EIP-A hits 30 Mbit/s, there will be packet loss due to the bandwidth limit, while the bandwidth of EIP-B remains idle.
- Shared bandwidth:** If the bandwidth usage of some EIPs is high, the idle bandwidth of other EIPs can be used.
For example, two EIPs (EIP-A and EIP-B) are added to a shared bandwidth of 40 Mbit/s.

- If EIP-A uses 30 Mbit/s and EIP-B uses 10 Mbit/s, the total bandwidth is 40 Mbit/s. EIP-A can use the idle bandwidth of EIP-B to increase its bandwidth and prevent packet loss.
- If EIP-A uses 30 Mbit/s and EIP-B uses 15 Mbit/s, the total bandwidth reaches 45 Mbit/s and exceeds the 40 Mbit/s limit. In this case, the flexibility of the shared bandwidth fails and there will be packet loss on both EIPs.

7.3.10 What Are the Differences Between Public Bandwidth and Private Bandwidth?

Public Bandwidth

Public bandwidth is the bandwidth consumed when data is transferred between cloud instances and the Internet. You can configure the public bandwidth when creating an ECS or bind an EIP to an ECS after the ECS is created.

Public bandwidth is classified into inbound bandwidth and outbound bandwidth.

Inbound bandwidth is the bandwidth consumed when data is transferred from the Internet to the cloud. For example, when resources are downloaded from the Internet to ECSs, that consumes inbound bandwidth.

Outbound bandwidth is the bandwidth consumed when data is transferred from the cloud to the Internet. For example, when ECSs provide services accessible from the Internet and external users download resources from the ECSs, this consumes outbound bandwidth.

Private Bandwidth

Private bandwidth is the bandwidth consumed when data is transferred between ECSs in the same region and on the same private network. ECSs can also be connected to cloud databases, load balancers, and OBS through private bandwidth. The private bandwidth size depends on the instance specifications.

For details, see section "ECS Types" in *Elastic Cloud Server Service Overview*.

7.4 Connectivity

7.4.1 Why Can't My ECS Access the Internet Even After an EIP Is Bound?


Symptom

An ECS with an EIP bound cannot access the Internet.

Troubleshooting


Figure 7-4 shows the networking diagram for an ECS to access the Internet using an EIP.

Figure 7-4 EIP network diagram



Locate the fault based on the following procedure.

Figure 7-5 Troubleshooting procedure



1. **Step 1: Check Whether the Is Running Properly**
2. **Step 2: Check Whether the Network Configuration of the Is Correct**
3. **Step 3: Check Whether an EIP Has Been Assigned and Bound to the**
4. **Step 4: Check Whether an EIP Is Bound to the Primary Network Interface of the**

5. [Step 5: Check Whether Required Security Group Rules Have Been Configured](#)
6. [Step 6: Check Whether Traffic from the Subnet Is Blocked](#)

Step 1: Check Whether the Is Running Properly

Check the status.

If the status is not **Running**, start or restart the .

Step 2: Check Whether the Network Configuration of the Is Correct

1. Check whether the 's network interface has an IP address assigned.
Log in to the , and run **ifconfig** or **ip address** to check the IP address of the ECS's network interface.
If the runs Windows, run **ipconfig**.
2. Check whether the ECS's network interface has a virtual IP address.
Log in to the , and run **ifconfig** or **ip address** to check whether the 's network interface has a virtual IP address. If the 's network interface has no virtual IP address, run the **ip addr add <virtual-IP-address> eth0** command to configure an IP address for the 's network interface.

Figure 7-6 Virtual IP address of a network interface

```
[root@demoserver ~]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether fa:16:3e:37:7b:62 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.30/24 brd 192.168.1.255 scope global dynamic eth0
        valid_lft 84950sec preferred_lft 84950sec
    inet 192.168.1.192/24 scope global secondary eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe37:7b62/64 scope link
        valid_lft forever preferred_lft forever
```

Check whether the ECS's network interface has a default route. If there is no default route, run **ip route add** to add one.

Figure 7-7 Default route

```
192.168.1.0/24 dev eth0 proto kernel scope link src 192.168.1.200
192.168.1.0/24 dev eth1 proto kernel scope link src 192.168.1.179
169.254.0.0/16 dev eth0 scope link metric 1002
default via 192.168.1.1 dev eth0 proto static
-bash-4.1#
```

Step 3: Check Whether an EIP Has Been Assigned and Bound to the

Check whether an EIP has been assigned and bound to the . If no EIP has been assigned, assign an EIP and bind it to the .

Step 4: Check Whether an EIP Is Bound to the Primary Network Interface of the

Check whether an EIP is bound to the primary network interface of the . If there is no EIP bound to the primary network interface of the , bind one.

You can view the network interface details by clicking the **Network Interfaces** tab on the details page. By default, the first record in the list is the primary network interface.

Step 5: Check Whether Required Security Group Rules Have Been Configured

For details about how to add security group rules, see [Adding a Security Group Rule](#).

If security group rules have not been configured, configure them based on your service requirements. (The remote IP address indicates the allowed IP address, and **0.0.0.0/0** indicates that all IP addresses are allowed.)

Step 6: Check Whether Traffic from the Subnet Is Blocked

Check whether the network ACL associated with the subnet of the ECS's network interface blocks traffic.

You can configure the network ACL on the VPC console. Make sure that the network ACL rules allow the traffic from the subnet.

7.4.2 What Should I Do If an EIP Cannot Be Pinged?

Symptom

After you purchase an EIP and bind it to an ECS, the local host or other cloud servers cannot ping the EIP of the ECS.

Fault Locating

The following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 7-8 Method of locating the failure to ping an EIP

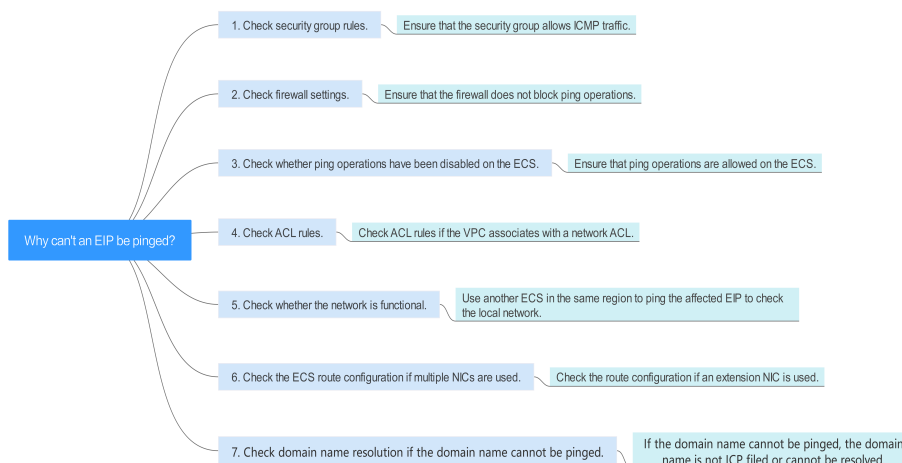


Table 7-5 Method of locating the failure to ping an EIP

Possible Causes	Solution
ICMP access rules are not added to the security group.	Add ICMP access rules to the security group. For details, see Checking Security Group Rules .
Ping operations are prohibited on the firewall.	Allow ping operations on the firewall. For details, see Checking Firewall Settings .
Ping operations are prohibited on the ECS.	Allow ping operations on the ECS. For details, see Checking Whether Ping Operations Have Been Disabled on the ECS .
Network ACL is associated.	If the VPC is associated with a network ACL, check the network ACL rules. For details, see Checking Network ACL Rules .
A network exception occurred.	Use another ECS in the same region to check whether the local network is functional. For details, see Checking Whether the Network Is Normal .
Routes are incorrectly configured if multiple NICs are used.	If the network is inaccessible due to an extension NIC, the fault is generally caused by incorrect route configurations. To resolve this issue, see Checking the ECS Route Configuration If Multiple NICs Are Used .
The domain name is not ICP licensed.	If the domain name cannot be pinged or cannot be resolved, see Checking Domain Name Resolution If the Domain Name Cannot Be Pinged to resolve this issue.

Checking Security Group Rules

ICMP is used for the ping command. Check whether the security group accommodating the ECS allows ICMP traffic.

1. Log in to the management console.
2. Under **Computing**, click **Elastic Cloud Server**.
3. On the **Elastic Cloud Server** page, click the name of the target ECS.
The page providing details about the ECS is displayed.
4. Click the **Security Groups** tab, expand the information of the security group, and view security group rules.
5. Click the security group ID.
The system automatically switches to the **Security Group** page.
6. On the **Outbound Rules** page, click **Add Rule**. In the displayed dialog box, set required parameters to add an outbound rule.

Figure 7-9 Adding an outbound rule

Add Outbound Rule

[Learn more](#) about security group configuration.

Security Group

default

You can import multiple rules in a batch.

Priority	Action	Protocol & Port	Type	Destination	Description	Operation
1-100	Allow	Protocols/ICMP All	IPv4	IP address 0.0.0.0/0		Replicate Delete

+

Add Rule

OK

Cancel

Table 7-6 Security group rules

Transfer Direction	Type	Protocol/Port Range	Destination
Outbound	IPv4	ICMP/Any	0.0.0.0/0 0.0.0.0/0 indicates all IP addresses.

7. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.

Figure 7-10 Adding an inbound rule

Add Inbound Rule

[Learn more](#) about security group configuration.

Security Group

default

You can import multiple rules in a batch.

Priority	Action	Protocol & Port	Type	Source	Description	Operation
1-100	Allow	Protocols/ICMP All	IPv4	IP address 0.0.0.0/0		Replicate Delete

+

Add Rule

OK

Cancel

Table 7-7 Security group rules

Transfer Direction	Type	Protocol/Port Range	Source
Inbound	IPv4	ICMP/Any	0.0.0.0/0 0.0.0.0/0 indicates all IP addresses.

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

Checking Firewall Settings

If a firewall is enabled on the ECS, check whether the firewall blocks the ping operations.

Linux

1. Consider CentOS 7 as an example. Run the following command to check the firewall status:

firewall-cmd --state

If **running** is displayed in the command output, the firewall has been enabled.

2. Check whether there is any ICMP rule blocking the ping operations.

iptables -L

If the command output shown in [Figure 7-11](#) is displayed, there is no ICMP rule blocking the ping operations.

Figure 7-11 Checking firewall rules

```
[root@ecs-3c4e ~]# iptables -L
Chain INPUT (policy ACCEPT)
target     prot opt source                destination           icmp echo-request

Chain FORWARD (policy ACCEPT)
target     prot opt source                destination

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination           icmp echo-reply
[root@ecs-3c4e ~]#
```

If the ping operations are blocked by an ICMP rule, run the following commands to modify the rule for unblocking:

iptables -A INPUT -p icmp --icmp-type echo-request -j ACCEPT

iptables -A OUTPUT -p icmp --icmp-type echo-reply -j ACCEPT

Windows

1. Log in to the Windows ECS, click the Windows icon in the lower left corner of the desktop, and choose **Control Panel > Windows Firewall**.
2. Click **Turn Windows Firewall on or off**.
View and set the firewall status.
3. If the firewall is **On**, go to step 4.
4. Check the ICMP rule statuses in the firewall.
 - a. In the navigation pane on the **Windows Firewall** page, click **Advanced settings**.
 - b. Enable the following rules:

Inbound Rules: File and Printer Sharing (Echo Request - ICMPv4-In)

Outbound Rules: File and Printer Sharing (Echo Request - ICMPv4-Out)

If IPv6 is enabled, enable the following rules:

Inbound Rules: File and Printer Sharing (Echo Request - ICMPv6-In) Outbound Rules: File and Printer Sharing (Echo Request - ICMPv6-Out)

Figure 7-12 Inbound Rules





Figure 7-13 Outbound Rules



Checking Whether Ping Operations Have Been Disabled on the ECS

Windows

Enable ping operations using the CLI.

1. Start the **Run** dialog box. Enter **cmd** and press **Enter**.
2. Run the following command to enable ping operations:
netsh firewall set icmpsetting 8

Linux

Check the ECS kernel parameters.

1. Check the **net.ipv4.icmp_echo_ignore_all** value in the **/etc/sysctl.conf** file. Value **0** indicates that ping operations are allowed, and value **1** indicates that ping operations are prohibited.
2. Allow ping operations.
 - Run the following command to temporarily allow the ping operations:
#echo 0 >/proc/sys/net/ipv4/icmp_echo_ignore_all

- Run the following command to permanently allow the ping operations:
`net.ipv4.icmp_echo_ignore_all=0`

Checking Network ACL Rules

By default, no ACL is configured for a VPC. If a network ACL is associated with a VPC, check the ACL rules.

1. Check whether the subnet of the ECS has been associated with a network ACL.
If an ACL name is displayed, the network ACL has been associated with the ECS.
2. Click the ACL name to view its status.
3. If the network ACL is enabled, add an ICMP rule to allow traffic.

NOTE

The default network ACL rule denies all incoming and outgoing packets. If a network ACL is disabled, the default rule is still effective.

Checking Whether the Network Is Normal

1. Use another ECS in the same region to check whether the local network is functional.
Use another ECS in the same region to ping the affected EIP. If the EIP can be pinged, the VPC is functional. In such a case, rectify the local network fault and ping the affected EIP again.
2. Check whether the link is accessible.
A ping failure is caused by packet loss or long delay, which may be caused by link congestion, link node faults, or heavy load on the ECS.

Checking the ECS Route Configuration If Multiple NICs Are Used

Generally, the default route of an OS will preferentially select the primary NIC. If an extension NIC is selected in a route and the network malfunctions, this issue is typically caused by incorrect route configuration.

- If the ECS has multiple NICs, check whether the default route is available.
 - a. Log in to the ECS and run the following command to check whether the default route is available:

ip route

Figure 7-14 Default route

```
[root@do-not-del-scy ~]# ip route
default via 192.168.2.1 dev eth0
169.254.0.0/16 dev eth0 scope link metric 1002
169.254.169.254 via 192.168.2.1 dev eth0 proto static
192.168.2.0/24 dev eth0 proto kernel scope link src 192.168.2.112
```

- b. If the route is unavailable, run the following command to add it:
ip route add default via XXXX dev eth0

 NOTE

In the preceding command, *XXXX* specifies a gateway IP address.

- If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policy routing on the ECS for network communication with the extension NIC.

Checking Domain Name Resolution If the Domain Name Cannot Be Pinged

If you can ping the EIP but not the domain name, the possible cause is that an error occurred in domain name resolution.

1. Check the domain name resolution.

If the domain name records are incorrectly configured, the domain name may fail to be resolved.

Switch to the DNS management console to view details about the domain name resolution.

2. Check the DNS server configuration.

If the system shows no server found after you ping a domain name, this issue may be caused by slow response from the DNS server.

7.4.3 Why Does the Download Speed of My ECS Is Slow?

Troubleshooting Process

If the download speed of an ECS is slow, check the following:

- Bandwidth limit exceeded: Your used bandwidth exceeds its limit and the limiting policy of the bandwidth takes effect, causing packet loss and slowing down the access. You can check the bandwidth usage or increase the bandwidth.
If your service traffic continues to be high, you can increase the bandwidth by referring to [Modifying a Shared Bandwidth](#).
- The memory usage of the ECS is higher than 80%.
For details, see section "Why Is My Linux ECS Running Slowly?" or "Why Is My Windows ECS Running Slowly?" in the *Elastic Cloud Server User Guide*
- Unstable carrier lines: The network between the local server and the cloud is unstable. Contact the carrier to check the network status.

7.4.4 What Are the Priorities of the Custom Route and EIP If Both Are Configured for an ECS to Enable the ECS to Access the Internet?

The priority of an EIP is higher than that of a custom route in a VPC route table. For example:

The VPC route table of an ECS has a custom route with 0.0.0.0/0 as the destination and NAT gateway as the next hop.

If an ECS in the VPC has an EIP bound, the VPC route table will have a policy-based route with 0.0.0.0/0 as the destination, which has a higher priority than its

custom route. In this case, traffic is forwarded to the EIP and cannot reach the NAT gateway.