Virtual Private Cloud

User Guide (Paris Region)

Issue 01

Date 2022-10-31





Copyright © Huawei Technologies Co., Ltd. 2024. 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.

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 Virtual Private Cloud?	1
1.2 Application Scenarios	2
1.3 Functions	2
1.4 VPC Connectivity	2
1.5 VPC and Other Services	3
1.6 User Permissions	4
1.7 Billing	4
1.8 Basic Concepts	11
1.8.1 Subnet	11
1.8.2 Elastic IP	11
1.8.3 Route Table	12
1.8.4 SNAT	14
1.8.5 Security Group	14
1.8.6 VPC Peering Connection	15
1.8.7 Network ACL	15
1.8.8 Virtual IP Address	18
1.8.9 Region and AZ	20
2 Getting Started	22
2.1 Quick Start	22
2.2 Typical Application Scenarios	22
2.3 Configuring a VPC for ECSs That Do Not Require Internet Access	23
2.3.1 Overview	23
2.3.2 Step 1: Create a VPC	24
2.3.3 Step 2: Create a Subnet for the VPC	27
2.3.4 Step 3: Create a Security Group	29
2.3.5 Step 4: Add a Security Group Rule	30
2.4 Configuring a VPC for ECSs That Access the Internet Using EIPs	33
2.4.1 Overview	36
2.4.2 Step 1: Create a VPC	38
2.4.3 Step 2: Create a Subnet for the VPC	41
2.4.4 Step 3: Assign an EIP and Bind It to an ECS	43
2.4.5 Step 4: Create a Security Group	45

2.4.6 Step 5: Add a Security Group Rule	46
2.5 Setting up an IPv6 Network	49
3 VPC and Subnet	58
3.1 VPC and Subnet Planning Suggestions	58
3.2 VPC	61
3.2.1 Creating a VPC	61
3.2.2 Modifying a VPC	64
3.2.3 Adding a Secondary IPv4 CIDR Block to a VPC	66
3.2.4 Deleting a Secondary IPv4 CIDR Block from a VPC	67
3.2.5 Deleting a VPC	68
3.2.6 Exporting VPC List	69
3.2.7 Obtaining a VPC ID	69
3.2.8 Viewing a VPC Topology	69
3.3 Subnet	70
3.3.1 Creating a Subnet for the VPC	70
3.3.2 Modifying a Subnet	72
3.3.3 Managing Subnet Tags	73
3.3.4 Exporting Subnet List	75
3.3.5 Viewing and Deleting Resources in a Subnet	75
3.3.6 Viewing IP Addresses in a Subnet	
3.3.7 Deleting a Subnet	
3.4 IPv4 and IPv6 Dual-Stack Network	78
4 Access Control	83
4.1 Differences Between Security Groups and Network ACLs	83
4.2 Security Group	
4.2.1 Security Groups and Security Group Rules	84
4.2.2 Default Security Group	88
4.2.3 Security Group Configuration Examples	
4.2.4 Managing a Security Group	
4.2.4.1 Creating a Security Group	
4.2.4.2 Cloning a Security Group	
4.2.4.3 Modifying a Security Group	
4.2.4.4 Deleting a Security Group	
4.2.5 Managing Security Group Rules	
4.2.5.1 Adding a Security Group Rule	
4.2.5.2 Fast-Adding Security Group Rules	
4.2.5.3 Modifying a Security Group Rule	
4.2.5.4 Replicating a Security Group Rule	
4.2.5.5 Importing and Exporting Security Group Rules	
4.2.5.6 Deleting a Security Group Rule	
4.2.6 Managing Instances Associated with a Security Group	
4.2.6.1 Adding an Instance to or Removing an Instance from a Security Group	10 <i>F</i>

4.2.6.2 Viewing the Security Group of an ECS	107
4.2.6.3 Changing the Security Group of an ECS	108
4.3 Network ACL	109
4.3.1 Network ACL Overview	109
4.3.2 Network ACL Configuration Examples	112
4.3.3 Managing Network ACLs	114
4.3.3.1 Creating a Network ACL	115
4.3.3.2 Modifying a Network ACL	115
4.3.3.3 Enabling or Disabling a Network ACL	116
4.3.3.4 Viewing a Network ACL	116
4.3.3.5 Deleting a Network ACL	117
4.3.4 Management Network ACL Rules	117
4.3.4.1 Adding a Network ACL Rule	117
4.3.4.2 Modifying a Network ACL Rule	119
4.3.4.3 Changing the Sequence of a Network ACL Rule	121
4.3.4.4 Enabling or Disabling a Network ACL Rule	122
4.3.4.5 Exporting and Importing Network ACL Rules	122
4.3.4.6 Deleting a Network ACL Rule	123
4.3.5 Managing Subnets Associated with a Network ACL	123
4.3.5.1 Associating Subnets with a Network ACL	124
4.3.5.2 Disassociating Subnets from a Network ACL	124
5 Elastic IP	126
5.1 EIP Overview	126
5.2 Assigning an EIP and Binding It to an ECS	127
5.3 Unbinding an EIP from an ECS and Releasing the EIP	129
5.4 Modifying an EIP Bandwidth	130
5.5 Exporting EIP Information	131
5.6 Managing EIP Tags	131
6 Shared Bandwidth	133
6.1 Shared Bandwidth Overview	133
6.2 Assigning a Shared Bandwidth	133
6.3 Adding EIPs to a Shared Bandwidth	
6.4 Removing EIPs from a Shared Bandwidth	
6.5 Modifying a Shared Bandwidth	
6.6 Deleting a Shared Bandwidth	
7 Route Tables	137
7.1 Route Tables and Routes	
7.2 Managing Route Tables	
7.2.1 Creating a Custom Route Table	
7.2.2 Associating a Route Table with a Subnet	
	141
7.2.3 Changing the Route Table Associated with a Subnet	

7.2.4 Viewing the Route Table Associated with a Subnet	143
7.2.5 Viewing Route Table Information	143
7.2.6 Exporting Route Table Information	144
7.2.7 Deleting a Route Table	144
7.3 Managing Routes	145
7.3.1 Adding a Custom Route	145
7.3.2 Modifying a Route	146
7.3.3 Replicating a Route	147
7.3.4 Deleting a Route	148
7.4 Configuring an SNAT Server	150
8 VPC Peering Connection	153
8.1 VPC Peering Connection Overview	153
8.2 VPC Peering Connection Usage Examples	155
8.3 Creating a VPC Peering Connection with Another VPC in Your Account	165
8.4 Creating a VPC Peering Connection with a VPC in Another Account	170
8.5 Obtaining the Peer Project ID of a VPC Peering Connection	175
8.6 Modifying a VPC Peering Connection	176
8.7 Viewing VPC Peering Connections	176
8.8 Deleting a VPC Peering Connection	177
8.9 Modifying Routes Configured for a VPC Peering Connection	177
8.10 Viewing Routes Configured for a VPC Peering Connection	179
8.11 Deleting Routes Configured for a VPC Peering Connection	180
9 VPC Flow Log	183
9.1 VPC Flow Log Overview	183
9.2 Creating a VPC Flow Log	184
9.3 Viewing a VPC Flow Log	186
9.4 Enabling or Disabling VPC Flow Log	189
9.5 Deleting a VPC Flow Log	189
10 Virtual IP Address	191
10.1 Virtual IP Address Overview	191
10.2 Assigning a Virtual IP Address	193
10.3 Binding a Virtual IP Address to an EIP or ECS	194
10.4 Binding a Virtual IP Address to an EIP	200
10.5 Unbinding a Virtual IP Address from an Instance	200
10.6 Unbinding a Virtual IP Address from an EIP	201
10.7 Releasing a Virtual IP Address	201
10.8 Disabling IP Forwarding on the Standby ECS	202
10.9 Disabling Source/Destination Check for an ECS NIC	203
11 Interconnecting with CTS	205
11.1 Supported VPC Operations	
11.2 Viewing Traces	208

12 Monitoring	209
12.1 Supported Metrics	209
12.2 Viewing Metrics	211
12.3 Creating an Alarm Rule	211
13	213
13.1	213
13.2 VPC Custom Policies	213
14 FAQ	215
14.1 General Questions	215
14.1.1 What Is a Quota?	215
14.2 Billing and Payments	215
14.2.1 Will I Be Billed for Using the VPC Service?	216
14.2.2 How Is an EIP Billed?	216
14.2.3 How Do I Change a Pay-per-Use EIP from Billing By Bandwidth to Traffic or from Billing to Bandwidth?	
14.2.4 Why Is My VPC Still Being Billed After It Was Deleted?	221
14.3 VPCs and Subnets	221
14.3.1 What Is Virtual Private Cloud?	221
14.3.2 Which CIDR Blocks Are Available for the VPC Service?	222
14.3.3 Can Subnets Communicate with Each Other?	222
14.3.4 What Subnet CIDR Blocks Are Available?	223
14.3.5 How Many Subnets Can I Create?	223
14.3.6 Why Can't I Delete My VPCs and Subnets?	223
14.3.7 Can I Change the VPC of an ECS?	227
14.4 EIPs	227
14.4.1 What Are the Differences Between EIP, Private IP Address, and Virtual IP Address?	227
14.4.2 How Do I Access the Internet Using an EIP Bound to an Extension NIC?	229
14.4.3 What Are the Differences Between the Primary and Extension NICs of ECSs?	230
14.4.4 Can an EIP That Uses Dedicated Bandwidth Be Changed to Use Shared Bandwidth?	231
14.4.5 Can I Bind an EIP to Multiple ECSs?	231
14.4.6 How Do I Access an ECS with an EIP Bound from the Internet?	231
14.4.7 Can I Bind an EIP of an ECS to Another ECS?	231
14.4.8 How Do I Unbind an EIP from an Instance and Bind a New EIP to the Instance?	232
14.4.9 Can I Bind an EIP to a Cloud Resource in Another Region?	233
14.4.10 Can I Change the Region of My EIP?	233
14.5 VPC Peering Connections	234
14.5.1 How Many VPC Peering Connections Can I Create in an Account?	234
14.5.2 Can a VPC Peering Connection Connect VPCs in Different Regions?	234
14.5.3 Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Co	
14.6 Virtual IP Addresses	241
14.6.1 Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?	241

14.6.2 Why is the Network Disconnected Between Servers Using a Virtual IP Address After an Active,	/
Standby Switchover?	245
14.7 Bandwidth	246
14.7.1 What Are Inbound Bandwidth and Outbound Bandwidth?	
14.7.2 How Do I Know If My EIP Bandwidth Limit Has Been Exceeded?	246
14.7.3 What Are the Differences Between Public Bandwidth and Private Bandwidth?	248
14.7.4 What Is the Bandwidth Size Range?	
14.7.5 What Bandwidth Types Are Available?	249
14.7.6 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth?	249
14.7.7 Is There a Limit to the Number of EIPs That Can Be Added to Each Shared Bandwidth?	249
14.7.8 What Is the Relationship Between Bandwidth and Upload/Download Rate?	249
14.8 Connectivity	250
14.8.1 Does a VPN Allow Communication Between Two VPCs?	250
14.8.2 Why Are Internet or Internal Domain Names in the Cloud Inaccessible Through Domain Name When My ECS Has Multiple NICs?	
14.8.3 What Are the Priorities of the Custom Route and EIP If Both Are Configured for an ECS to Ena the ECS to Access the Internet?	
14.8.4 Why Are There Intermittent Interruptions When a Local Host Accesses a Website Built on an E	
14.8.5 Why Do ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communication?	251
14.8.6 Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur Whe They Communicate?	
14.8.7 Why Can't My ECS Use Cloud-init?	255
14.8.8 Why Can't My ECS Access the Internet Even After an EIP Is Bound?	259
14.8.9 Why Does My ECS Fail to Obtain an IP Address?	262
14.8.10 How Do I Handle a VPN or Direct Connect Connection Network Failure?	264
14.8.11 Why Can My Server Be Accessed from the Internet But Cannot Access the Internet?	266
14.8.12 Why Can't I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?.	267
14.8.13 Why Does My ECS Fail to Communicate with Other After It Has Firewall Installed?	268
14.9 Routing	269
14.9.1 How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?	269
14.9.2 Can a Route Table Span Multiple VPCs?	270
14.9.3 How Many Routes Can a Route Table Contain?	270
14.9.4 Are There Any Restrictions on Using a Route Table?	271
14.9.5 Do the Same Routing Priorities Apply to Direct Connect Connections and Custom Routes in th Same VPC?	
14.9.6 Are There Different Routing Priorities of the VPN and Custom Routes in the Same VPC?	271
14.10 Security	271
14.10.1 Are the Security Group Rules Considered the Same If All Parameters Except Their Description the Same?	
14.10.2 How Do I Know the Instances Associated with a Security Group?	271
14.10.3 Why Can't I Delete a Security Group?	272
14.10.4 Can I Change the Security Group of an ECS?	272

A Change History	278
14.10.9 Why Do My Security Group Rules Not Take Effect?	274
14.10.8 Why Is Access from a Specific IP Address Still Allowed After a Network ACL Rule Access from the IP Address Has Been Added?	
14.10.7 Which Security Group Rule Has a High Priority When Multiple Security Group Ru	
14.10.6 Does a Modified Security Group Rule or a Network ACL Rule Take Effect Immedia Connections?	
14.10.5 How Do I Configure a Security Group for Multi-Channel Protocols?	272

Service Overview

1.1 What Is Virtual Private Cloud?

Overview

The Virtual Private Cloud (VPC) service enables you to provision logically isolated virtual networks for Elastic Cloud Servers (ECSs), improving cloud resource security and simplifying network deployment. You can configure and manage the virtual networks as required.

Within your own VPC, you can create security groups and VPNs, configure IP address ranges, specify bandwidth sizes, manage the networks in the VPC, and make changes to these networks as needed, quickly and securely. You can also define rules to control communications between ECSs in the same security group or in different security groups.

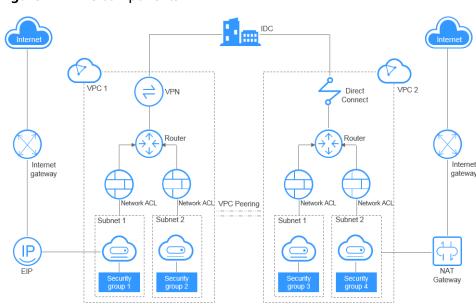


Figure 1-1 VPC components

Accessing the VPC Service

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

• Management console

You can use the console to directly perform operations on VPC resources. To access the VPC service, log in to the management console and select **Virtual Private Cloud** from the console homepage.

API

If you need to integrate a VPC into a third-party system for secondary development, you can use APIs to access the VPC service. For details, see the .

1.2 Application Scenarios

Hosting web applications

You can host web applications and websites in a VPC and use the VPC as a regular network. With EIPs, you can connect ECSs running your web applications to the Internet. A VPN gateway is used to establish a VPN tunnel between the web applications and the service system on the cloud, ensuring high-speed communication between the website and the service system.

Hosting services that demand high security

You can create a VPC and security groups to host multi-tier web applications in different security zones. You can associate web servers and database servers with different security groups and configure different access control rules for security groups. You can launch web servers in a publicly accessible subnet, and also run database servers in subnets that are not publicly accessible. In this way, you can ensure high security.

• Extending your corporate network into the cloud

You can establish a VPN connection between a VPC and a traditional data center to use the ECSs and block storage resources. Applications can be migrated to the cloud and additional web servers can be quickly deployed as needed when there is a spike in demand for computing resources. This way, less money has to be spent on IT and O&M and data is kept safer than in a traditional arrangement. A VPC can span multiple AZs, protecting from single points of failure and ensuring high availability for e-commerce systems.

1.3 Functions

lists common VPC functions.

Before using the VPC service, you should be familiar with the basic concepts, such as subnets, route tables, security groups, and EIPs. This will make it easier to understand VPC functions.

1.4 VPC Connectivity

You can use EIPs, load balancers, NAT gateways, VPN connections, and Direct Connect connections to access the Internet if required.

Use EIPs to Enable a Small Number of ECSs to Access the Internet

When only a few ECSs need to access the Internet, you can bind the EIPs to the ECSs. This will provide them with Internet access. You can also dynamically unbind the EIPs from the ECSs and bind them to NAT gateways and load balancers instead, which will also provide Internet access. The process is not complicated.

 Use NAT Gateways to Enable a Large Number of ECSs to Access the Internet

When a large number of ECSs need to access the Internet, the public cloud provides NAT gateways for the ECSs. With NAT gateways, you do not need to assign an EIP to each ECS, which reduces management costs incurred by an excessive number of EIPs. A NAT gateway offers both the SNAT and DNAT functions. SNAT allows multiple ECSs in the same VPC to share one or more EIPs to access the Internet. SNAT prevents the EIPs of ECSs from being exposed to the Internet. SNAT supports up to 1 million concurrent connections and 30,000 new connections. DNAT can implement port-level data forwarding. It maps EIP ports to ECS ports so that the ECSs in a VPC can share the same EIP and bandwidth to provide Internet-accessible services.

 Use ELB to Connect to the Internet If There Are a Large Number of Concurrent Requests

In high-concurrency scenarios, such as e-commerce, you can use load balancers provided by the ELB service to evenly distribute incoming traffic across multiple ECSs, allowing a large number of users to concurrently access your business system or application. ELB is deployed in the cluster mode. It provides fault tolerance for your applications by automatically balancing traffic across multiple AZs. You can also take advantage of deep integration with Auto Scaling (AS), which enables automatic scaling based on service traffic and ensures service stability and reliability.

• Use VPN or Direct Connect to Extend Your On-premises Data Center into the Cloud over the Internet

For customers with equipment rooms in their on-premises data centers, not all businesses of the customers will be migrated to the cloud because the customers want to reuse their legacy devices and require smooth business evolution. Then, you can use VPN or Direct Connect to interconnect your VPC and on-premises data center. A VPN connection routes traffic through the Internet, which allows you to use a private network with the price of the public network. A Direct Connect connection is a dedicated, private network connection that provides you with more efficient data transmission and more consistent network experience than Internet-based connections.

1.5 VPC and Other Services

FCS

The VPC service provides an isolated virtual network for ECSs. You can configure and manage the network as required. There are multiple connectivity options for ECSs to access the Internet. You can also define rules for communication between ECSs in the same security group or in different security groups.

ELB

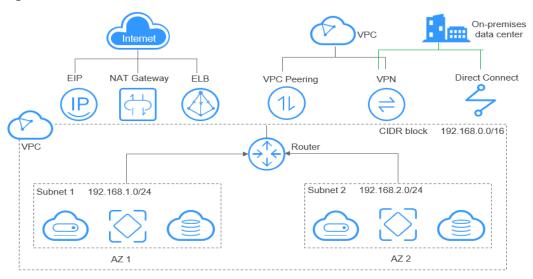
ELB uses the EIPs and bandwidths associated with the VPC service.

Cloud Eye

You can use Cloud Eye to monitor the status of your VPCs without adding plug-ins.

Figure 1-2 shows the relationship between VPC and other services.

Figure 1-2 VPC and other services



1.6 User Permissions

The cloud system provides two types of user permissions by default: user management and resource management. User management refers to the management of users, user groups, and user group rights. Resource management refers to the control operations that can be performed by users on cloud service resources.

1.7 Billing

Billing Items

The VPC service is free of charge.

Table 1-1 Billing items

Billing Item	Description
EIP	EIPs are required if your resources need to access the Internet.

The EIP service provides multiple billing modes.

- **EIP Billing Modes**
- Which Billing Option Is Right for Me?
- How Will I Be Billed If I Change My Bandwidth Size?
- How Do I Change the EIP Billing Mode?

EIP Billing Modes

EIPs can be billed on a pay-per-use basis.

- Figure 1-3
- Table 1-2

Figure 1-3 EIP billing

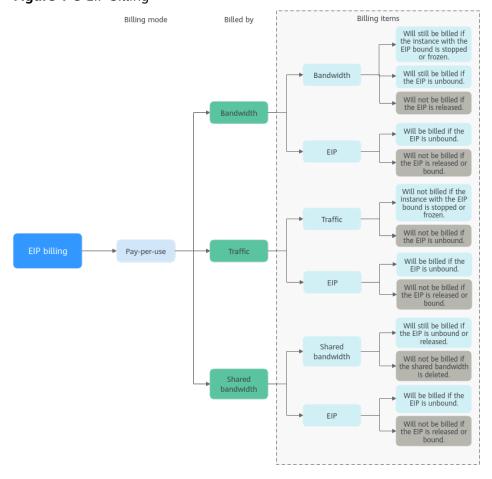


Table 1-2 EIP billing description

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
Pay- per- use	Bandwidth	Bandwid th EIP retention	If a pay-per-use EIP is billed by bandwidth: Bandwidth: You are billed based on your specified bandwidth size and usage duration. There is no limit on how much traffic you can use. After the EIP is purchased, you can change your specified bandwidth you use will not exceed the bandwidth you specified. EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: If the EIP is not bound to any instance, both the EIP and its bandwidth will be billed. If the EIP is bound to an instance, only the bandwidth will be billed. The bandwidth will be billed regardless of if the instance bound to the EIP is running or not. After the EIP is unbound from an instance, the bandwidth will continue to be billed. Unless it is released, the EIP will still be billed. If the EIP is released, both the EIP and its bandwidth will not be billed.

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
	Traffic	Traffic EIP retention n	If a pay-per-use EIP is billed by traffic: Traffic: You are billed based on your EIP type and the total amount of traffic going out of the cloud. The bandwidth size you set is only used to limit the maximum data transfer rate. To prevent high fees caused by burst traffic, specify a proper bandwidth size when you buy an EIP. If an EIP billed by traffic uses a dedicated bandwidth, only the bandwidth used in the outbound direction will be billed. EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: If the EIP is not bound to an instance, you will be billed for the EIP itself, but not for traffic. If the EIP is bound to an instance, only the used traffic will be billed. If the instance bound to the EIP stops running and there is no traffic generated, there will be no traffic or EIP fees. After the EIP is unbound from an instance, the traffic will not be billed but the EIP will still be billed. If the EIP is released, the EIP will not be billed.

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
	Shared bandw idth	Shared bandwid th EIP retention n	If a pay-per-use EIP is added to a shared bandwidth: • Share bandwidth: Only the shared bandwidth will be billed. There will be no additional bandwidth or traffic costs for EIPs added to the shared bandwidth. • EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: Shared bandwidth No operations on the EIP will affect the billing of a shared bandwidth. For example, if you have released the EIP but have not deleted the shared bandwidth will still be billed. After a shared bandwidth is deleted, it will no longer be billed. After a shared bandwidth is deleted, it will no longer be billed. If the EIP is not bound to an instance, the EIP will still be billed. If the EIP is unbound from an instance, the EIP will be billed to keep it allocated to your account unless it is released. If the EIP is released or bound to an instance, the EIP will not be billed.

To save money, you can add multiple EIPs in the same region to a shared bandwidth. A shared bandwidth can be billed on a pay-per-use basis. For details, see **Table 1-3**. Currently, only pay-per-use EIPs can be added to a shared bandwidth.

- You can add an EIP to a shared bandwidth when buying the EIP.
- You can also add an existing EIP to a shared bandwidth. After the EIP is added to a shared bandwidth, there will be no additional bandwidth or traffic cost. You will only be billed for the shared bandwidth.

Table 1-3 Shared bandwidth	billina	aetails
-----------------------------------	---------	---------

Billing Mode	Billed By	Billing Item	Billing Item Description
Pay- per-use	Bandwidt h	Bandwidth	You are billed based on your specified bandwidth size and usage duration. There is no limit on how much traffic you can use.
			After a shared bandwidth is purchased, you can change your specified bandwidth size. The bandwidth you use will not exceed the bandwidth you specified.

MOTE

- The price of bandwidth, traffic, and EIP depends on the region.
- The EIP bandwidth is the outbound 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, that consumes outbound bandwidth. Only the outbound bandwidth will be billed.
 - If your purchased or modified bandwidth is no more than 10 Mbit/s, the inbound bandwidth will be 10 Mbit/s, and the outbound bandwidth will be the same as the purchased or modified bandwidth.
 - If your purchased or modified bandwidth is more than 10 Mbit/s, both the bandwidths in inbound and outbound directions will be the same as the purchased or modified bandwidth.

Which Billing Option Is Right for Me?

EIPs can be billed by bandwidth or traffic. **Table 1-4** shows the application scenarios of different billing options.

Cloud Eye monitors your network metrics, such as bandwidth and traffic. Based on the bandwidth usage, you can determine which billing option (by bandwidth or by traffic) is more cost-effective. Here are some suggestions for your reference:

- If you need less than 5 Mbit/s of bandwidth for a short time and the traffic is light, set your EIP to be billed by traffic.
- If you need more than 5 Mbit/s of bandwidth and the bandwidth usage is greater than 20%, set your EIP to be billed by bandwidth.

•	•	3 .	
Billing Mode	Billed By	Scenario	
Pay-per- use	Bandwidth	Heavy or stable traffic	
	Traffic	Light or sharply fluctuating traffic	
	Shared bandwidth	Staggered traffic	

Table 1-4 Application scenarios of EIP billing options

How Will I Be Billed If I Change My Bandwidth Size?

If an EIP is not added to a shared bandwidth, the EIP uses the dedicated bandwidth regardless of it is billed by bandwidth or traffic. After an EIP is added to a shared bandwidth, only the shared bandwidth is billed.

When you change the bandwidth size, the bandwidth price and effective time depend on the billing mode, which applies to both dedicated and shared bandwidths. For details, see **Table 1-5**.

□ NOTE

Decreasing bandwidths may cause packet loss.

Table 1-5 Impact on billing after bandwidth size change

Billing Mode	Billed By	Change	Impact
use idth	Bandw idth	Increase or decrease the bandwidth	The change will take effect immediately.
	Traffic	Increase or decrease the bandwidth	The change will take effect immediately. The bandwidth size you set is only used to limit the maximum data transfer rate.

How Do I Change the EIP Billing Mode?

The EIP service has multiple billing modes you can choose from. You can change your EIP billing mode during the EIP usage period if necessary.

• Table 1-6

Ⅲ NOTE

Changing the billing mode does not change EIPs or interrupt their use.

Table 1 6 Em biding mode change description				
Change	Description			
 From billing by traffic (pay-per-use) to billing by bandwidth (pay-per- 	An EIP billed by traffic on a pay-per-use basis can be directly changed to be billed by bandwidth on a pay-per-use basis.			
use)From billing by bandwidth (pay-per-use) to billing by traffic (pay-	 An EIP billed by bandwidth on a pay-per-use basis can be directly changed to be billed by traffic on a pay-per-use basis. The new billing mode takes effect immediately. 			
per-use)	The new billing mode takes effect immediately.			

Table 1-6 EIP billing mode change description

If you want to change a pay-per-use EIP from billing by bandwidth to billing by traffic, refer to How Do I Change a Pay-per-Use EIP from Billing By Bandwidth to Traffic or from Billing By Traffic to Bandwidth?.

1.8 Basic Concepts

1.8.1 Subnet

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.

• After a subnet is created, its CIDR block cannot be modified. Subnets in the same VPC cannot overlap.

For example, if the CIDR block of VPC-A is 10.0.0.0/16, you can specify 10.0.0.0/24 for subnet A01, 10.0.1.0/24 for subnet A02, and 10.0.3.0/24 for subnet A03.

Ⅲ NOTE

By default, you can create a maximum of 100 subnets in each region. If this cannot meet your service requirements, request a quota increase by referring to **What Is a Quota?**

1.8.2 Elastic IP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers.

Each EIP can be used by only one cloud resource at a time.

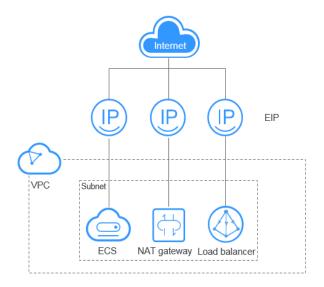


Figure 1-4 Accessing the Internet using an EIP

1.8.3 Route Table

Route Tables

A route table contains a set of routes that are used to determine where network traffic from your subnets in a VPC is directed. Each subnet must be associated with a route table. A subnet can only be associated with one route table, but you can associate multiple subnets with the same route table.

- Default route table: When you create a VPC, the system automatically
 generates a default route table for the VPC. If you create a subnet in the VPC,
 the subnet automatically associates with the default route table. The default
 route table ensures that subnets in a VPC can communicate with each other.
 - You can add routes to, delete routes from, and modify routes in the default route table, but cannot delete the table.
 - When you create a VPC endpoint, VPN or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.
- Custom route table: If you do not want to use the default route table, you can create a custom route table and associate it with the subnet. Custom route tables can be deleted if they are no longer required.

The custom route table associated with a subnet affects only the outbound traffic. The default route table of a subnet controls the inbound traffic.

□ NOTE

Route

You can add routes to default and custom route tables and configure the destination, next hop type, and next hop in the routes to determine where network traffic is directed. Routes are classified into system routes and custom routes.

• System routes: These routes are automatically added by the system and cannot be modified or deleted.

After a route table is created, the system automatically adds the following system routes to the route table, so that instances in a VPC can communicate with each other.

- Routes whose destination is 100.64.0.0/10 or 198.19.128.0/20.
- Routes whose destination is a subnet CIDR block.

In addition to the preceding system routes, the system automatically adds a route whose destination is 127.0.0.0/8. This is the local loopback address.

• Custom routes: These are routes that you can add, modify, and delete. The destination of a custom route cannot overlap with that of a system route.

You can add a custom route and configure the destination, next hop type, and next hop in the route to determine where network traffic is directed. **Table** 1-7 lists the supported types of next hops.

You cannot add two routes with the same destination to a VPC route table even if their next hop types are different. The route priority depends on the destination. According to the longest match routing rule, the destination with a higher matching degree is preferentially selected for packet forwarding.

Table 1-7 Next hop type

Next Hop Type	Description	Supported Route Table
Server	Traffic intended for the destination is forwarded to an ECS	Default route table
	in the VPC.	Custom route table
Extension NIC	Traffic intended for the destination is forwarded to the	 Default route table
	extension NIC of an ECS in the VPC.	 Custom route table
VPN gateway	Traffic intended for the destination is forwarded to a VPN gateway.	Custom route table
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.	Custom route table
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	Default route tableCustom route table

Next Hop Type	Description	Supported Route Table
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.	Default route tableCustom route table
Virtual IP address	Traffic intended for the destination is forwarded to a virtual IP address and then sent to active and standby ECSs to which the virtual IP address is bound.	Default route tableCustom route table

◯ NOTE

Currently, the route with the next hop type Direct Connect gateway cannot be configured. To configure it, submit a service ticket.

If you specify the destination when creating a resource, a system route is delivered. If you do not specify a destination when creating a resource, a custom route that can be modified or deleted is delivered.

For example, when you create a NAT gateway, the system automatically delivers a custom route without a specific destination (0.0.0.0/0 is used by default). In this case, you can change the destination. However, when you create a VPN gateway, you need to specify the remote subnet, that is, the destination of a route. In this case, the system delivers this system route. Do not modify the route destination on the **Route Tables** page. If you do, the destination will be inconsistent with the configured remote subnet. To modify the route destination, go to the specific resource page and modify the remote subnet, then the route destination will be changed accordingly.

1.8.4 **SNAT**

In addition to services provided by the system, some ECSs need to access the Internet to obtain information or download software. You can bind EIPs to virtual NICs (ports) of ECSs to enable the ECSs to access the Internet. However, assigning an EIP to each ECS consumes IPv4 addresses, incurs additional costs, and may increase the attack surface for a virtual environment. Therefore, SNAT is introduced to enable multiple ECSs to share one EIP.

On a public cloud, an EIP can be assigned to an ECS that serves as the SNAT router or gateway for other ECSs from the same subnet or VPC.

For details about how to configure SNAT, see **Configuring an SNAT Server**.

1.8.5 Security Group

A security group is a collection of access control rules for cloud resources, such as cloud servers, containers, and databases, that have the same security protection requirements and that are mutually trusted. After a security group is created, you can configure access rules that will apply to all cloud resources added to this security group.

Like whitelists, security group rules work as follows:

- Inbound rules control incoming traffic to instances in the security group.
 If an inbound request matches the source in an inbound security group rule with Action set to Allow, the request is allowed and other requests are denied.
 - By default, you do not need to configure deny rules in the inbound direction because requests that do not match allow rules will be denied.
- Outbound rules control outgoing traffic from instances in the security group.
 If the destination of an outbound security group rule with Action set to Allow is 0.0.0.0/0, all outbound requests are allowed.
 0.0.0.0/0 represents all IPv4 addresses.

1.8.6 VPC Peering Connection

A VPC peering connection is a networking connection that connects two VPCs for them to communicate using private IP addresses. The VPCs to be peered can be in the same account or different accounts, but must be in the same region.

You can use VPC peering connections to build networks in different scenarios.
 For details, see VPC Peering Connection Usage Examples.

Figure 1-5 shows an application scenario of VPC peering connections.

- There are two VPCs (VPC-A and VPC-B) in region A that are not connected.
- Service servers (ECS-A01 and ECS-A02) are in VPC-A, and database servers (RDS-B01 and RDS-B02) are in VPC-B. The service servers and database servers cannot communicate with each other.
- You need to create a VPC peering connection (peering-AB) between VPC-A and VPC-B so the service servers and database servers can communicate with each other.

VPC-A-Region A VPC-B-Region A 172.16.0.0/16 172.17.0.0/16 Security group (General-purpose web server) Security group (General-purpose web server) Subnet-A01-172.16.0.0/24 Subnet-B01-172.17.0.0/24 VPC peering connection VPC-A route table peering-AB VPC-B route table ECS-A01 Next Hop Destination Destination Next Hop 172.16.0.0/16 peering-AB 172.17.0.0/16 peering-AB ECS-A02

Figure 1-5 VPC peering connection network diagram

For details about VPC peering connections, see **VPC Peering Connection**.

1.8.7 Network ACL

A network ACL is an optional layer of security for your subnets. After you associate one or more subnets with a network ACL, you can control traffic in and out of the subnets.

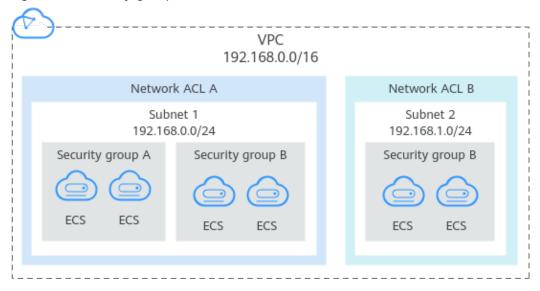


Figure 1-6 Security groups and network ACLs

Network ACL Basics

- Your VPC does not come with a network ACL, but you can create a network ACL and associate it with a VPC subnet if required. By default, each network ACL denies all inbound traffic to and outbound traffic from the associated subnet until you add rules.
- You can associate a network ACL with multiple subnets. However, a subnet can only be associated with one network ACL at a time.
- Each newly created network ACL is in the **Inactive** state until you associate subnets with it.
- Network ACLs use connection tracking to track traffic to and from instances.
 Changes to inbound and outbound rules do not take effect immediately for the existing traffic.

If you add, modify, or delete a network ACL rule, or associate or disassociate a subnet with or from a network ACL, all the inbound and outbound persistent connections will not be disconnected. New rules will only be applied for the new connections.

NOTICE

After a persistent connection is disconnected, new connections will not be established immediately until the timeout period of connection tracking expires. For example, after an ICMP persistent connection is disconnected, a new connection will be established and a new rule will apply when the timeout period (30s) expires.

- The timeout period of connection tracking varies by protocol. The timeout period of a TCP connection in the established state is 600s, and that of an ICMP connection is 30s. For other protocols, if packets are received in both inbound and outbound directions, the connection tracking timeout period is 180s. If packets are received only in one direction, the connection tracking timeout period is 30s.
- The timeout period of TCP connections varies by connection status. The timeout period of a TCP connection in the established state is 600s, and that of a TCP connection in the FIN-WAIT state is 30s.

Default Network ACL Rules

By default, each network ACL has preset rules that allow the following packets:

- Packets whose source and destination are in the same subnet.
- Broadcast packets with the destination 255.255.255.255/32, which is used to configure host startup information.
- Multicast packets with the destination 224.0.0.0/24, which is used by routing protocols.
- Metadata packets with the destination 169.254.169.254/32 and TCP port number 80, which is used to obtain metadata.
- Packets from CIDR blocks that are reserved for public services (for example, packets with the destination 100.125.0.0/16).
- A network ACL denies all traffic in and out of a subnet excepting the
 preceding packets. Table 1-8 shows the default rules. You cannot modify or
 delete the default rules.

Table 1-8 Default network ACL rules

Direction	Priorit y	Actio n	Protoco l	Sourc e	Destinatio n	Description
Inbound	*	Deny	All	0.0.0.0	0.0.0.0/0	Denies all inbound traffic.
Outboun d	*	Deny	All	0.0.0.0	0.0.0.0/0	Denies all outbound traffic.

How Traffic Matches Network ACL Rules

- Each network ACL rule has a priority value where a smaller value corresponds to a higher priority. Any time two rules conflict, the rule with the higher priority is the one that gets applied. The rule whose priority value is an asterisk (*) has the lowest priority.
- If multiple network ACL rules conflict, only the rule with the highest priority takes effect. If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule.

Application Scenarios

- If the application layer needs to provide services for users, traffic must be allowed to reach the application layer from all IP addresses. However, you also need to prevent illegal access from malicious users.
 - Solution: You can add network ACL rules to deny access from suspect IP addresses.
- How can I isolate ports with identified vulnerabilities? For example, how do I isolate port 445 that can be exploited by WannaCry worm?
 - Solution: You can add network ACL rules to deny access traffic from a specific port and protocol, for example, TCP port 445.
- No defense is required for the east-west traffic between subnets, but access control is required for north-south traffic.
 - Solution: You can add network ACL rules to protect north-south traffic.
- For frequently accessed applications, a security rule sequence may need to be adjusted to improve performance.
 - Solution: A network ACL allows you to adjust the rule sequence so that frequently used rules are applied before other rules.

1.8.8 Virtual IP Address

A virtual IP address can be shared among multiple ECSs. An ECS can have both private and virtual IP addresses, and you can access the ECS through either IP address. A virtual IP address has the same network access capabilities as a private IP address, including layer 2 and layer 3 communication in VPCs, access between VPCs using VPC peering connections, as well as access through EIPs, VPN connections, and Direct Connect connections.

You can bind ECSs deployed in active/standby mode with the same virtual IP address, and then bind an EIP to the virtual IP address. Virtual IP addresses can work together with Keepalived to ensure high availability and disaster recovery. If the active ECS is faulty, the standby ECS automatically takes over services from the active one.

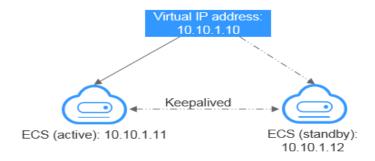
Networking

Virtual IP addresses are used for high availability and can work together with Keepalived to make active/standby ECS switchover possible. This way if one ECS goes down for some reason, the other one can take over and services continue uninterrupted. ECSs can be configured for HA or as load balancing clusters.

Networking mode 1: HA

If you want to improve service availability and avoid single points of failure, you can deploy ECSs in the active/standby mode or deploy one active ECS and multiple standby ECSs. In this arrangement, the ECSs all use the same virtual IP address. If the active ECS becomes faulty, a standby ECS takes over services from the active ECS and services continue uninterrupted.

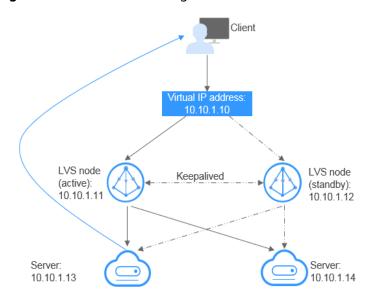
Figure 1-7 Networking diagram of the HA mode



- In this configuration, a single virtual IP address is bound to two ECSs in the same subnet.
- Keepalived is then used to configure the two ECSs to work in the active/ standby mode. Follow industry standards for configuring Keepalived. The details are not included here.
- **Networking mode 2**: HA load balancing cluster

If you want to build a high-availability load balancing cluster, use Keepalived and configure LVS nodes as direct routers.

Figure 1-8 HA load balancing cluster



- Bind a single virtual IP address to two ECSs.
- Configure the two ECSs as LVS nodes working as direct routers and use Keepalived to configure the nodes in the active/standby mode. The two ECSs will evenly forward requests to different backend servers.

- Configure two more ECSs as backend servers.
- Disable the source/destination check for the two backend servers.

Follow industry standards for configuring Keepalived. The details are not included here.

Application Scenarios

- Accessing the virtual IP address through an EIP
 If your application has high availability requirements and needs to provide services through the Internet, it is recommended that you bind an EIP to a virtual IP address.
- Using a VPN, Direct Connect, or VPC peering connection to access a virtual IP address

To ensure high availability and access to the Internet, use a VPN for security and Direct Connect for a stable connection. The VPC peering connection is needed so that the VPCs in the same region can communicate with each other.

1.8.9 Region and AZ

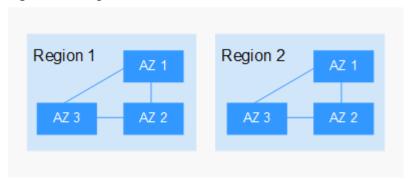
Concept

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 1-9 shows the relationship between regions and AZs.

Figure 1-9 Regions and AZs



Selecting a Region

Select a region closest to your target users for lower network latency and quick access.

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 Quick Start

This document describes how to prepare for and quickly create a VPC with an IPv4 or IPv6 CIDR block.

CIDR Block Types

IPv4: When you create a VPC and subnet, IPv4 CIDR block is used by default. Servers on the IPv4 network cannot access IPv6 services on the Internet or provide services accessible from users using an IPv6 client. For details about how to set up an IPv4 network, see .

IPv6: When you need to access the IPv6 services on the Internet or provide services accessible from users using an IPv6 client, you need to enable the IPv6 function. After the IPv6 function is enabled, you can provide services for users using an IPv4 or IPv6 client. For details about how to set up an IPv6 network, see .

2.2 Typical Application Scenarios

A VPC provides an isolated virtual network for ECSs. You can configure and manage the network as required.

- If any of your ECSs, for example, ECSs that function as the database of server nodes for website deployment, do not need to access the Internet, you can configure a VPC for the ECSs by following the instructions described in Configuring a VPC for ECSs That Do Not Require Internet Access.
- If your ECSs need to access the Internet, you can configure EIPs for them. For
 example, the ECSs functioning as the service nodes for deploying a website
 need to be accessed by users over the Internet. Then, you can configure a VPC
 for these ECSs by following the instructions provided in Configuring a VPC
 for ECSs That Access the Internet Using EIPs.
- When you need to access the IPv6 services on the Internet or provide services accessible from users using an IPv6 client, you need to enable the IPv6 function. After the IPv6 function is enabled, you can provide services for users using an IPv4 or IPv6 client.

2.3 Configuring a VPC for ECSs That Do Not Require Internet Access

2.3.1 Overview

If your ECSs do not require Internet access (for example, the ECSs functioning as the database nodes or server nodes for deploying a website), you can follow the procedure shown in **Figure 2-1** to configure a VPC for the ECSs.

Figure 2-1 Configuring the network

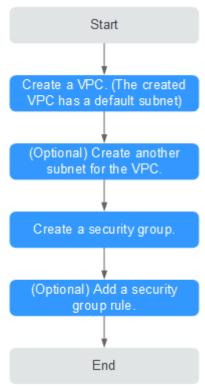


Table 2-1 describes the different tasks in the procedure for configuring the network.

Table 2-1 Configuration process description

Task	Description
Create a VPC.	This task is mandatory.
	After the VPC is created, you can create other required network resources in the VPC based on your service requirements.

Task	Description
Create another subnet for	This task is optional.
the VPC.	If the default subnet cannot meet your requirements, you can create one.
	The new subnet is used to assign IP addresses to NICs added to the ECS.
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.
Add a security group rule.	This task is optional.
	If the default rule meets your service requirements, you do not need to add rules to the security group.

2.3.2 Step 1: Create a VPC

Scenarios

A VPC provides an isolated virtual network for ECSs. You can configure and manage the network as required.

You can create a VPC by following the procedure provided in this section. Then, create subnets, security groups, and assign EIPs by following the procedure provided in subsequent sections based on your actual network requirements.

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 as prompted.

A default subnet will be created together with a VPC and you can also click **Add Subnet** to create more subnets for the VPC.

Table 2-2 VPC parameter descriptions

Parameter	Description	Example Value
Region	Select the region nearest to you to ensure the lowest latency possible.	-
Name	The VPC name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	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	192.168.0.0/16
Advanced Settings	Click the drop-down arrow to set advanced VPC parameters, including tags.	Retain the default settings.
Tag	The VPC tag, which consists of a key and value pair. You can add a maximum of 10 tags to each VPC.	Key: vpc_key1Value: vpc-01

Table 2-3 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	subnet-01
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings	Click the drop-down arrow to set advanced settings for the subnet, including Gateway and DNS Server Address .	Retain the default settings.
Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
DNS Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x

Parameter	Description	Example Value
Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet.	Key: subnet_key1Value: subnet-01

Table 2-4 VPC tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each VPC and can be the same for different VPCs. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	vpc_key1
Value	 Can contain a maximum of 43 characters. Can contain letters, digits, underscores (_), periods (.), and hyphens (-). 	vpc-01

Table 2-5 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

5. Confirm the current configuration and click **Create Now**.

2.3.3 Step 2: Create a Subnet for the VPC

Scenarios

A VPC comes with a default subnet. If the default subnet cannot meet your requirements, you can create one.

A subnet is configured with DHCP by default. When an ECS in this subnet starts, the ECS automatically obtains an IP address using DHCP.

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 2-6 Parameter descriptions

Parameter	Description	Example Value
VPC	The VPC for which you want to create a subnet.	-
Name	The subnet name.	Subnet
	The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable.	-
	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 cannot be disabled after the subnet is created.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings/ Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
Advanced Settings/DN S Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x

Parameter	Description	Example Value
Advanced Settings/Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet. The tag key and value must meet the requirements listed in Table 2-7.	Key: subnet_key1Value: subnet-01

Table 2-7 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

Precautions

When a subnet is created, there are five reserved IP addresses, which cannot be used. For example, in a subnet with CIDR block 192.168.0.0/24, the following IP addresses are reserved:

- 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: Gateway address.
- 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.

If you configured the default settings under **Advanced Settings** during subnet creation, the reserved IP addresses may be different from the default ones, but there will still be five of them. The specific addresses depend on your subnet settings.

2.3.4 Step 3: Create a Security Group

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The

cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. In the upper right corner, click **Create Security Group**. The **Create Security Group** page is displayed.
- 5. Configure the parameters as prompted.

Table 2-8 Parameter description

Parameter	Description	Example Value
Name	Mandatory	sg-AB
	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.	
Description	Optional	N/A
	Supplementary information about the security group. This parameter is optional.	
	The security group description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

2.3.5 Step 4: Add a Security Group Rule

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

Like whitelists, security group rules work as follows:

- Inbound rules control incoming traffic to instances in the security group.
 If an inbound request matches the source in an inbound security group rule with Action set to Allow, the request is allowed and other requests are denied.
 - By default, you do not need to configure deny rules in the inbound direction because requests that do not match allow rules will be denied.
- Outbound rules control outgoing traffic from instances in the security group. If the destination of an outbound security group rule with **Action** set to **Allow** is 0.0.0.0/0, all outbound requests are allowed.

0.0.0.0/0 represents all IPv4 addresses.

If the rules of the security group associated with your instance cannot meet your requirements, for example, you need to allow inbound traffic on a specific TCP port, you can add an inbound rule to allow traffic on the TCP port.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. Locate the row that contains the target security group, and click **Manage Rule** 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 2-9 Inbound rule parameter description

Param eter	Description	Example Value
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	TCP

Param eter	Description	Example Value
	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, or 22-30
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. For example:	0.0.0.0/0
	IP address:Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	If the source is a security group, this rule will apply to all instances associated with the selected security group.	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The inbound rule list is displayed.

8. On the **Outbound Rules** tab, click **Add Rule**.

The Add Outbound Rule dialog box is displayed.

9. Configure required parameters.

You can click + to add more outbound rules.

Table 2-10 Outbound rule parameter description

Param eter	Description	Example Value
Туре	Destination IP address version. You can select:IPv4IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	ТСР

Param eter	Description	Example Value
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.	22, or 22-30
	Outbound rules control outgoing traffic over specific ports from instances in the security group.	
Destina tion	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. For example:	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The outbound rule list is displayed.

2.4 Configuring a VPC for ECSs That Access the Internet Using EIPs

- 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 as prompted.

A default subnet will be created together with a VPC and you can also click **Add Subnet** to create more subnets for the VPC.

Table 2-11 VPC parameter descriptions

Parameter	Description	Example Value
Region	Select the region nearest to you to ensure the lowest latency possible.	-
Name	The VPC name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	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	192.168.0.0/16
Advanced Settings	Click the drop-down arrow to set advanced VPC parameters, including tags.	Retain the default settings.
Tag	The VPC tag, which consists of a key and value pair. You can add a maximum of 10 tags to each VPC.	Key: vpc_key1Value: vpc-01

Table 2-12 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	subnet-01
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings	Click the drop-down arrow to set advanced settings for the subnet, including Gateway and DNS Server Address .	Retain the default settings.
Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
DNS Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x

Parameter	Description	Example Value
Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet.	Key: subnet_key1Value: subnet-01

Table 2-13 VPC tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each VPC and can be the same for different VPCs. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	vpc_key1
Value	 Can contain a maximum of 43 characters. Can contain letters, digits, underscores (_), periods (.), and hyphens (-). 	vpc-01

Table 2-14 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

5. Confirm the current configuration and click **Create Now**.

2.4.1 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 2-2 to bind EIPs to the ECSs.

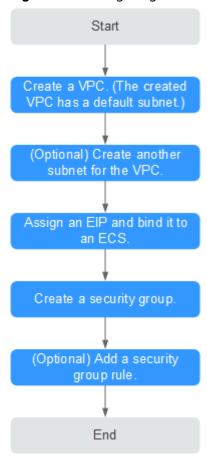


Figure 2-2 Configuring the network

Table 2-15 describes the different tasks in the procedure for configuring the network.

Table 2-15 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	This task is optional.
the VPC.	If the default subnet cannot meet your requirements, you can create one.
	The new subnet is used to assign IP addresses to NICs added to the ECS.

Task	Description
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.

2.4.2 Step 1: Create a VPC

Scenarios

A VPC provides an isolated virtual network for ECSs. You can configure and manage the network as required.

You can create a VPC by following the procedure provided in this section. Then, create subnets, security groups, and assign EIPs by following the procedure provided in subsequent sections based on your actual network requirements.

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 as prompted.

A default subnet will be created together with a VPC and you can also click **Add Subnet** to create more subnets for the VPC.

Table 2-16 VPC parameter descriptions

Parameter	Description	Example Value
Region	Select the region nearest to you to ensure the lowest latency possible.	-

Parameter	Description	Example Value
Name	The VPC name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	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).	192.168.0.0/16
	The following CIDR blocks are supported: • 10.0.0.0/8-24 • 172.16.0.0/12-24	
	172.16.0.0/12-24192.168.0.0/16-24	
Advanced Settings	Click the drop-down arrow to set advanced VPC parameters, including tags.	Retain the default settings.
Tag	The VPC tag, which consists of a key and value pair. You can add a maximum of 10 tags to each VPC.	Key: vpc_key1Value: vpc-01

Table 2-17 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	subnet-01
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24

Parameter	Description	Example Value
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable.	-
	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.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings	Click the drop-down arrow to set advanced settings for the subnet, including Gateway and DNS Server Address .	Retain the default settings.
Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
DNS Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x
Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet.	Key: subnet_key1Value: subnet-01

Table 2-18 VPC tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each VPC and can be the same for different VPCs. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	vpc_key1
Value	 Can contain a maximum of 43 characters. Can contain letters, digits, underscores (_), periods (.), and hyphens (-). 	vpc-01

Table 2-19 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, 	subnet_key1
	underscores (_), and hyphens (-).	
Value	 Can contain a maximum of 43 characters. 	subnet-01

5. Confirm the current configuration and click **Create Now**.

2.4.3 Step 2: Create a Subnet for the VPC

Scenarios

A VPC comes with a default subnet. If the default subnet cannot meet your requirements, you can create one.

A subnet is configured with DHCP by default. When an ECS in this subnet starts, the ECS automatically obtains an IP address using DHCP.

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 2-20 Parameter descriptions

Parameter	Description	Example Value
VPC	The VPC for which you want to create a subnet.	
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	Subnet
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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 cannot be disabled after the subnet is created.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings/ Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
Advanced Settings/DN S Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x
Advanced Settings/Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet. The tag key and value must meet the requirements listed in Table 2-21.	Key: subnet_key1Value: subnet-01

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

Table 2-21 Subnet tag key and value requirements

Precautions

When a subnet is created, there are five reserved IP addresses, which cannot be used. For example, in a subnet with CIDR block 192.168.0.0/24, the following IP addresses are reserved:

- 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: Gateway address.
- 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.

If you configured the default settings under **Advanced Settings** during subnet creation, the reserved IP addresses may be different from the default ones, but there will still be five of them. The specific addresses depend on your subnet settings.

2.4.4 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 an 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, click Assign EIP.
- 4. Set the parameters as prompted.

Table 2-22 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 you. The region selected for the EIP is its geographical location.	N/A
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 bandwidth types 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 traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. Shared Bandwidth: The bandwidth can be shared by multiple EIPs and is suitable for scenarios with staggered traffic. 	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
EIP Name	The name of the EIP.	eip-test
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 2-23.	Key: Ipv4_key1Value: 3005eip

Parameter	Description	Example Value
Quantity	The number of EIPs you want to purchase.	1

Table 2-23 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 (-). 	lpv4_key1
Value	 Can contain a maximum of 43 characters. Can contain letters, digits, underscores (_), periods (.), and hyphens (-). 	3005eip

5. Click Create Now.

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.

2.4.5 Step 4: Create a Security Group

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

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.

In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.

- 4. In the upper right corner, click **Create Security Group**. The **Create Security Group** page is displayed.
- 5. Configure the parameters as prompted.

Table 2-24 Parameter description

Parameter	Description	Example Value
Name	Mandatory	sg-AB
	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.	
Description	Optional	N/A
	Supplementary information about the security group. This parameter is optional.	
	The security group description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

2.4.6 Step 5: Add a Security Group Rule

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

Like whitelists, security group rules work as follows:

- Inbound rules control incoming traffic to instances in the security group.
 If an inbound request matches the source in an inbound security group rule with Action set to Allow, the request is allowed and other requests are denied.
 - By default, you do not need to configure deny rules in the inbound direction because requests that do not match allow rules will be denied.
- Outbound rules control outgoing traffic from instances in the security group.
 If the destination of an outbound security group rule with **Action** set to **Allow** is 0.0.0.0/0, all outbound requests are allowed.
 0.0.0.0/0 represents all IPv4 addresses.

If the rules of the security group associated with your instance cannot meet your requirements, for example, you need to allow inbound traffic on a specific TCP port, you can add an inbound rule to allow traffic on the TCP port.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. Locate the row that contains the target security group, and click **Manage Rule** 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 2-25 Inbound rule parameter description

Param eter	Description	Example Value
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	ТСР
	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, or 22-30

Param eter	Description	Example Value
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. For example:	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	If the source is a security group, this rule will apply to all instances associated with the selected security group.	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The inbound rule list is displayed.

8. On the **Outbound Rules** tab, click **Add Rule**.

The **Add Outbound Rule** dialog box is displayed.

9. Configure required parameters.

You can click + to add more outbound rules.

Table 2-26 Outbound rule parameter description

Param eter	Description	Example Value
Туре	Destination IP address version. You can select: • IPv4 • IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	ТСР
	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, or 22-30

Param eter	Description	Example Value
Destina tion	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. For example:	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The outbound rule list is displayed.

2.5 Setting up an IPv6 Network

Scenarios

This topic describes how to create a VPC with an IPv6 CIDR block and create an ECS with an IPv6 address in the VPC, so that the ECS can access the Internet using the IPv6 address. **Figure 2-3** shows the configuration process.

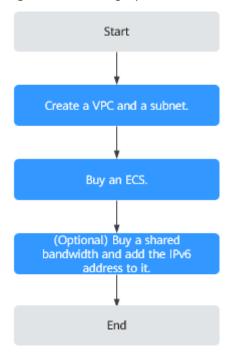


Figure 2-3 Setting up an IPv6 network

□ NOTE

If you already have a shared bandwidth, you can configure Internet access using an IPv6 address when purchasing an ECS.

Notes and Constraints

- The IPv4/IPv6 dual-stack function is currently free, but will be billed at a later date (price yet to be determined).
- Only certain ECS specifications support IPv6 networks and can use IPv4/IPv6 dual-stack networks. You need to select such ECSs in supported regions.
 To check which ECSs support IPv6:
 - On the ECS console: Click Create ECS. On the displayed page, view the ECS specifications.
 - If there is the **IPv6** parameter with the value of **Yes**, the ECS specifications support IPv6.

Application Scenarios of IPv4/IPv6 Dual Stack

Table 2-27 Application scenarios of IPv4/IPv6 dual stack

Applica tion Scenari o	Description	Subnet	ECS
Private commu nicatio n using IPv6 address es	Your applications deployed on ECSs need to communicate with other systems (such as databases) through private networks using IPv6 addresses.	IPv4 CIDR blockIPv6 CIDR block	 Private IPv4 address: used for private communication IPv6 address: used for private communication.
Public commu nicatio n using	Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.	IPv4 CIDR block IPv6	Private IPv4 address + IPv4 EIP: used for public network communication
IPv6 address es	Your applications deployed on ECSs need to both provide services accessible from the Internet and analyze the access request data using IPv6 addresses.	CIDR block	IPv6 address + shared bandwidth: used for public network communication

Step 1: Create a VPC

Before creating your VPCs, determine how many VPCs, the number of subnets, and what IP address ranges you will need.

Perform the following operations to create a VPC named **vpc-ipv6** and its default subnet named **subnet-ipv6**.

- 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.
- 4. Set the VPC and subnet parameters.

When configuring a subnet, select **Enable** for **IPv6 CIDR Block** so that the system will automatically allocate an IPv6 CIDR block to the subnet. IPv6 cannot be disabled after the subnet is created. Currently, customizing IPv6 CIDR block is not supported.

Table 2-28 VPC parameter descriptions

Parameter	Description	Example Value
Region	Select the region nearest to you to ensure the lowest latency possible.	-
Name	The VPC name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	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	192.168.0.0/16
Advanced Settings	Click the drop-down arrow to set advanced VPC parameters, including tags.	Retain the default settings.
Tag	The VPC tag, which consists of a key and value pair. You can add a maximum of 10 tags to each VPC.	Key: vpc_key1Value: vpc-01

Table 2-29 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	subnet-01
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings	Click the drop-down arrow to set advanced settings for the subnet, including Gateway and DNS Server Address .	Retain the default settings.
Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
DNS Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x

Parameter	Description	Example Value
Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet.	Key: subnet_key1Value: subnet-01

Click Create Now.

Step 2: Buy an ECS

On the management console, under **Compute**, click **Elastic Cloud Server**, and then click **Buy ECS**.

Configure the network for the ECS as follows:

- Network:
 - Select the created VPC vpc-ipv6.
 - Select the created subnet subnet-ipv6.
 - Select Self-assigned IPv6 address.

NOTICE

Select **Self-assigned IPv6 address** during ECS creation to assign an IPv6 address to the ECS. Otherwise, the IPv4/IPv6 dual-stack network cannot be used.

- Shared Bandwidth
 - If you select Do not configure, only IPv6 communication in a VPC is supported. If you want to enable Internet access, you need to perform operations in (Optional) Step 3: Buy a Shared Bandwidth and Add the IPv6 Address to It.
 - If you assign a shared bandwidth or select an existing shared bandwidth, the ECS can use the IPv6 address to access the Internet after the configuration is complete.
- Security Group: Select the default security group Sys-default. The default security group rule allows all outgoing IPv4 and IPv6 data packets and denies all inbound data packets. ECSs in the same security group can access each other without the need to add rules. You can also create a security group and add rules to it.
- EIP: Select **Not required**.

After the ECS is created, you can view the assigned IPv6 address on the ECS details page. You can also log in to the ECS and run the **ifconfig** command to view the assigned IPv6 address.

(Optional) Dynamically Assigning IPv6 Addresses

If an IPv6 address fails to be automatically assigned or the selected image does not support the function of automatic IPv6 address assignment, manually obtain the IPv6 address by referring to .

Ⅲ NOTE

If an ECS is created from a public image:

Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 is supported and then check whether dynamic IPv6 address assignment has been enabled. Currently, all Linux public images support IPv6, and dynamic IPv6 address assignment is enabled for Ubuntu 16 by default. You do not need to configure dynamic IPv6 address assignment for the Ubuntu 16 OS. For other Linux public images, you need to enable this function.

(Optional) Step 3: Buy a Shared Bandwidth and Add the IPv6 Address to It

By default, an IPv6 address can only be used for private network communication. If you want to use this IPv6 address to access the Internet or want it to be accessed by IPv6 clients on the Internet, you need to buy a shared bandwidth and add the IPv6 address to it.

If you already have a shared bandwidth, add the IPv6 address to the shared bandwidth.

Buying a Shared Bandwidth

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose Network > Elastic IP.
- In the navigation pane on the left, choose Elastic IP and Bandwidth > Shared Bandwidths.
- 4. In the upper right corner, click **Assign Shared Bandwidth**. On the displayed page, configure parameters as prompted.

Table 2-30 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 you.	N/A
Billed By	The billing method for the shared bandwidth.	Bandwidth
	You can specify a shared bandwidth to be billed by bandwidth or by traffic.	

Parameter	Description	Example Value
Bandwidth	The bandwidth size in Mbit/s. The minimum value is 5 Mbit/s. The maximum bandwidth can be 2000 Mbit/s.	10
Name	The name of the shared bandwidth.	Bandwidth-001

5. Click Create Now.

Adding the IPv6 Address to a Shared Bandwidth

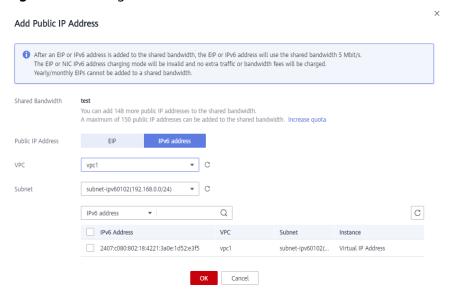
 On the Shared Bandwidths page, click Add Public IP Address in the Operation column.

Figure 2-4 Adding an IPv6 address to a shared bandwidth



2. Add the IPv6 address to the shared bandwidth.

Figure 2-5 Adding an IPv6 address to a shared bandwidth



3. Click OK.

Verifying the Result

Log in to the ECS and ping an IPv6 address on the Internet to verify network connectivity. Figure 2-6 shows an example command output.

Log in to the ECS using SSH or the RDP file through the EIP.

Figure 2-6 Verification

```
64 bytes from 2400:da00:2::29: icmp_seq=1 tt1=42 time=45.6 ms
64 bytes from 2400:da00:2::29: icmp_seq=2 tt1=42 time=45.1 ms
64 bytes from 2400:da00:2::29: icmp_seq=3 tt1=42 time=44.8 ms
64 bytes from 2400:da00:2::29: icmp_seq=4 tt1=42 time=45.1 ms
```

3 VPC and Subnet

3.1 VPC and Subnet Planning Suggestions

Before creating your VPCs, determine how many VPCs, the number of subnets, and what IP address ranges or connectivity options you will need.

- How Do I Determine How Many VPCs I Need?
- How Do I Plan Subnets?
- How Do I Plan Routing Policies?
- How Do I Connect to an On-Premises Data Center?
- How Do I Access the Internet?

How Do I Determine How Many VPCs I Need?

VPCs are region-specific. By default, networks in VPCs in different regions or even in the same region are not connected.

- One VPC
 - If your services do not require network isolation, a single VPC should be enough.
- Multiple VPCs

If you have multiple service systems in a region and each service system requires an isolated network, you can create a separate VPC for each service system.

MOTE

By default, you can create a maximum of five VPCs in each region. If this cannot meet your service requirements, request a quota increase by referring to **What Is a Quota?**

The following table lists the private CIDR blocks that you can specify when creating a VPC. Consider the following when selecting a VPC CIDR block:

- Number of IP addresses: Reserve sufficient IP addresses in case of business growth.
- IP address range: Avoid IP address conflicts if you need to connect a VPC to an on-premises data center or connect two VPCs.

Table 3-1 lists the supported VPC CIDR blocks.

Table 3-1 VPC CIDR blocks

VPC CIDR Block	IP Address Range	Maximum Number of IP Addresses
10.0.0.0/8-24	10.0.0.0-10.255.255.255	2^24-2=16777214
172.16.0.0/12-24	172.16.0.0-172.31.255.25 5	2^20-2=1048574
192.168.0.0/16-24	192.168.0.0-192.168.255. 255	2^16-2=65534

How Do I Plan Subnets?

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.

 After a subnet is created, its CIDR block cannot be modified. Subnets in the same VPC cannot overlap.

For example, if the CIDR block of VPC-A is 10.0.0.0/16, you can specify 10.0.0.0/24 for subnet A01, 10.0.1.0/24 for subnet A02, and 10.0.3.0/24 for subnet A03.

◯ NOTE

By default, you can create a maximum of 100 subnets in each region. If this cannot meet your service requirements, request a quota increase by referring to **What Is a Quota?**

When planning subnets, consider the following:

- You create different subnets for different modules in a VPC. For example, in VPC-A, you can create subnet A01 for web services, subnet A02 for management services, and subnet A03 for data services. You can leverage network ACLs to control access to each subnet.
- If your VPC needs to communicate with an on-premises data center through VPN or Direct Connect, ensure that the VPC subnet and the CIDR block used for communication in the data center do not overlap.

How Do I Plan Routing Policies?

When you create a VPC, the system automatically generates a default route table for the VPC. If you create a subnet in the VPC, the subnet automatically associates with the default route table. A route table contains a set of routes that are used to determine where network traffic from your subnets in a VPC is directed. The default route table ensures that subnets in a VPC can communicate with each other.

If you do not want to use the default route table, you can now create a custom route table and associate it with the subnets. The custom route table associated with a subnet affects only the outbound traffic. The default route table controls the inbound traffic.

You can add routes to default and custom route tables and configure the destination, next hop type, and next hop in the routes to determine where network traffic is directed. Routes are classified into system routes and custom routes.

- System routes: Routes that are automatically added by the system and cannot be modified or deleted. System routes allow instances in a VPC to communicate with each other.
- Custom routes: Routes that can be modified and deleted. The destination of a custom route cannot overlap with that of a system route.

You cannot add two routes with the same destination to a VPC route table even if their next hop types are different, because the destination determines the route priority. According to the longest match routing rule, the destination with a higher matching degree is preferentially selected for packet forwarding.

How Do I Connect to an On-Premises Data Center?

If you require interconnection between a VPC and an on-premises data center, ensure that the VPC does not have an overlapping IP address range with the on-premises data center to be connected.

As shown in **Figure 3-1**, you have VPC 1 in region A and VPC 2 and VPC 3 in region B. To connect to an on-premises data center, they can use a VPN, as VPC 1 does in Region A; or a Direct Connect connection, as VPC 2 does in Region B. VPC 2 connects to the data center through a Direct Connect connection, but to connect to another VPC in that region, like VPC 3, a VPC peering connection must be established.

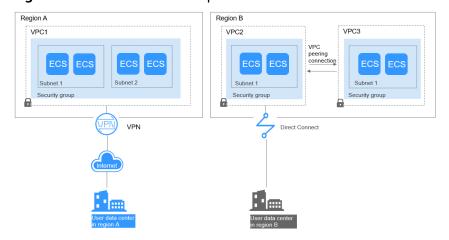


Figure 3-1 Connections to on-premises data centers

When planning CIDR blocks for VPC 1, VPC 2, and VPC 3:

- The CIDR block of VPC 1 cannot overlap with the CIDR block of the onpremises data center in Region A.
- The CIDR block of VPC 2 cannot overlap with the CIDR block of the onpremises data center in Region B.
- The CIDR blocks of VPC 2 and VPC 3 cannot overlap.

How Do I Access the Internet?

Use EIPs to enable a small number of ECSs to access the Internet.

When only a few ECSs need to access the Internet, you can bind the EIPs to the ECSs. This will provide them with Internet access. You can also dynamically unbind the EIPs from the ECSs and bind them to NAT gateways and load balancers instead, which will also provide Internet access. The process is not complicated.

Use a NAT gateway to enable a large number of ECSs to access the Internet.

When a large number of ECSs need to access the Internet, the public cloud provides NAT gateways for your ECSs. With NAT gateways, you do not need to assign an EIP to each ECS. NAT gateways reduce costs as you do not need so many EIPs. NAT gateways offer both source network address translation (SNAT) and destination network address translation (DNAT). SNAT allows multiple ECSs in the same VPC to share one or more EIPs to access the Internet. SNAT prevents the EIPs of ECSs from being exposed to the Internet. DNAT can implement portlevel data forwarding. It maps EIP ports to ECS ports so that the ECSs in a VPC can share the same EIP and bandwidth to provide Internet-accessible services.

Use ELB to access the Internet If there are a large number of concurrent requests.

In high-concurrency scenarios, such as e-commerce, you can use load balancers provided by the ELB service to evenly distribute incoming traffic across multiple ECSs, allowing a large number of users to concurrently access your business system or application. ELB is deployed in the cluster mode. It provides fault tolerance for your applications by automatically balancing traffic across multiple AZs. You can also take advantage of deep integration with Auto Scaling (AS), which enables automatic scaling based on service traffic and ensures service stability and reliability.

3.2 VPC

3.2.1 Creating a VPC

Scenarios

A VPC provides an isolated virtual network for ECSs. You can configure and manage the network as required.

You can create a VPC by following the procedure provided in this section. Then, create subnets, security groups, and assign EIPs by following the procedure provided in subsequent sections based on your actual network requirements.

Procedure

- 1. Log in to the management console.
- 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 as prompted.

A default subnet will be created together with a VPC and you can also click **Add Subnet** to create more subnets for the VPC.

Table 3-2 VPC parameter descriptions

Parameter	Description	Example Value
Region	Select the region nearest to you to ensure the lowest latency possible.	-
Name	The VPC name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	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	192.168.0.0/16
Advanced Settings	Click the drop-down arrow to set advanced VPC parameters, including tags.	Retain the default settings.
Tag	The VPC tag, which consists of a key and value pair. You can add a maximum of 10 tags to each VPC.	Key: vpc_key1Value: vpc-01

Table 3-3 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	subnet-01
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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.	
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings	Click the drop-down arrow to set advanced settings for the subnet, including Gateway and DNS Server Address .	Retain the default settings.
Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
DNS Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x

Parameter	Description	Example Value
Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet.	Key: subnet_key1Value: subnet-01

Table 3-4 VPC tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each VPC and can be the same for different VPCs. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	vpc_key1
Value	 Can contain a maximum of 43 characters. Can contain letters, digits, underscores (_), periods (.), and hyphens (-). 	vpc-01

Table 3-5 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

5. Confirm the current configuration and click **Create Now**.

3.2.2 Modifying a VPC

Scenarios

You can modify the following information about a VPC:

- Modifying the Name and Description of a VPC
- Modifying the CIDR Block of a VPC

NOTICE

If the **secondary IPv4 CIDR block** function is available in a region, the CIDR block of a VPC in this region cannot be modified through the console. You can call an API to modify VPC CIDR block. For details, see section "Updating VPC Information" in the *Virtual Private Cloud API Reference*.

Modifying the Name and Description of a VPC

- 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. Modify the name and description of a VPC using either of the following methods:
 - Method 1:
 - i. In the VPC list, click $\stackrel{\checkmark}{=}$ on the right of the VPC name.
 - ii. Enter the VPC name and click OK.
 - Method 2:
 - i. In the VPC list, click the VPC name with a hyperlink.The **Summary** page is displayed.
 - ii. Click $\stackrel{\checkmark}{=}$ on the right of the VPC name or description, enter the information, and click $\stackrel{\checkmark}{\sim}$.

Modifying the CIDR Block of a VPC

- 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 VPC list, locate the row that contains the VPC and click **Edit CIDR Block** in the **Operation** column.

The **Edit CIDR Block** dialog box is displayed.

4. Modify the VPC CIDR block as prompted.

NOTICE

A VPC CIDR block must be from 10.0.0.0/8-24, 172.16.0.0/12-24, or 192.168.0.0/16-24.

 If a VPC has no subnets, you can change both its network address and subnet mask.

X

Figure 3-2 Modifying network address and subnet mask



- If a VPC has subnets, you only can change its subnet mask.

Figure 3-3 Modifying subnet mask



5. Click OK.

3.2.3 Adding a Secondary IPv4 CIDR Block to a VPC

Scenarios

When you create a VPC, you specify a primary IPv4 CIDR block for the VPC, which cannot be changed. To extend the IP address range of your VPC, you can add a secondary CIDR block to the VPC.

If the **secondary IPv4 CIDR block** function is available in a region, the CIDR block of a VPC in this region cannot be modified through the console. You can call an API to modify VPC CIDR block. For details, see section "Updating VPC Information" in the *Virtual Private Cloud API Reference*.

Notes and Constraints

- You can allocate a subnet from either a primary or a secondary CIDR block of a VPC. A subnet cannot use both the primary and the secondary CIDR blocks.
 Subnets in the same VPC can communicate with each other by default, even if some subnets are allocated from the primary CIDR block and some are from the secondary CIDR block of a VPC.
- If a subnet in a secondary CIDR block of your VPC is the same as or overlaps with the destination of an existing route in the VPC route table, the existing route does not take effect.

If you create a subnet in a secondary CIDR block of your VPC, a route (the destination is the subnet CIDR block and the next hop is **Local**) is automatically added to your VPC route table. This route allows

communications within the VPC and has a higher priority than any other routes in the VPC route table. For example, if a VPC route table has a route with the VPC peering connection as the next hop and 100.20.0.0/24 as the destination, and a route for the subnet in the secondary CIDR block has a destination of 100.20.0.0/16, 100.20.0.0/16 and 100.20.0.0/24 overlaps and traffic will be forwarded through the route of the subnet.

- The allowed secondary CIDR block size is between a /28 netmask and /8 netmask.
- Table 3-6 lists the secondary CIDR blocks that are not supported.

Table 3-6 Restricted secondary CIDR blocks

Туре	CIDR Block (Not Supported)	
Reserved private CIDR blocks	 172.31.0.0/16 192.168.0.0/16 In-use primary CIDR blocks 	
Reserved system CIDR blocks	 100.64.0.0/10 214.0.0.0/7 198.18.0.0/15 169.254.0.0/16 	
Reserved public CIDR blocks	 0.0.0.0/8 127.0.0.0/8 240.0.0.0/4 255.255.255.255/32 	

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 VPC list, locate the row that contains the VPC and click **Edit CIDR Block** in the **Operation** column.
 - The **Edit CIDR Block** dialog box is displayed.
- 4. Click Add Secondary IPv4 CIDR Block.
- 5. Enter the secondary CIDR block and click **OK**.

3.2.4 Deleting a Secondary IPv4 CIDR Block from a VPC

Scenarios

If a secondary CIDR block of a VPC is no longer required, you can delete it.

 A secondary IPv4 CIDR block of a VPC can be deleted, but the primary CIDR block cannot be deleted. • If you want to delete a secondary CIDR block that contains subnets, you need to delete the subnets first.

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 VPC list, locate the row that contains the VPC and click **Edit CIDR Block** in the **Operation** column.

The **Edit CIDR Block** dialog box is displayed.

- 4. Locate the row that contains the secondary CIDR block to be deleted and click **Delete** in the **Operation** column.
- 5. Click **OK**.

3.2.5 Deleting a VPC

Scenarios

This section describes how to delete a VPC.

Notes and Constraints

If you want to delete a VPC that has subnets, custom routes, or other resources, you need to delete these resources as prompted on the console first and then delete the VPC.

You can refer to Why Can't I Delete My VPCs and Subnets?

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. On the **Virtual Private Cloud** page, locate the row that contains the VPC to be deleted and click **Delete** in the **Operation** column.

A confirmation dialog box is displayed.

4. Confirm the information and click Yes.

NOTICE

If a VPC cannot be deleted, a message will be displayed on the console. Delete the resources that are in the VPC by referring to **Why Can't I Delete My VPCs and Subnets?**

3.2.6 Exporting VPC List

Scenarios

Information about all VPCs under your account can be exported as an Excel file to a local directory.

This file records the names, ID, status, CIDR blocks, and the number of subnets of your VPCs.

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 upper right corner of the VPC list, click The system will automatically export information about all VPCs under your account in the current region as an Excel file to a local directory.

3.2.7 Obtaining a VPC ID

Scenarios

This section describes how to view and obtain a VPC ID.

If you create a VPC peering connection between two VPCs in different accounts, you need to obtain the project ID of the region that the peer VPC resides. You can recommend this section to the user of the peer VPC to obtain the project ID.

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. On the **Virtual Private Cloud** page, locate the VPC and click its name. The VPC details page is displayed.
- 4. In the **VPC Information** area, view the VPC ID.

Click next to ID to copy the VPC ID.

3.2.8 Viewing a VPC Topology

Scenarios

This section describes how to view the topology of a VPC. The topology displays the subnets in a VPC and the ECSs in the subnets.

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.

In the VPC list, click the name of the VPC for which the topology is to be viewed.

The VPC details page is displayed.

4. Click the **Topology** tab to view the VPC topology.

The topology displays the subnets in the VPC and the ECSs in the subnets. You can also perform the following operations on subnets and ECSs in the topology:

- Modify or delete a subnet.
- Add an ECS to a subnet, bind an EIP to the ECS, and change the security group of the ECS.

3.3 Subnet

3.3.1 Creating a Subnet for the VPC

Scenarios

A VPC comes with a default subnet. If the default subnet cannot meet your requirements, you can create one.

A subnet is configured with DHCP by default. When an ECS in this subnet starts, the ECS automatically obtains an IP address using DHCP.

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 3-7 Parameter descriptions

Parameter	Description	Example Value
VPC	The VPC for which you want to create a subnet.	-
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	Subnet
IPv4 CIDR Block	The CIDR block for the subnet. This value must be within the VPC CIDR block.	192.168.0.0/24
IPv6 CIDR Block	Specifies whether to set IPv6 CIDR Block to Enable. 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 cannot be disabled after the subnet is created.	-
Associated Route Table	The default route table to which the subnet will be associated. You can change the route table to a custom route table on the Subnets page.	Default
Advanced Settings/ Gateway	The gateway address of the subnet. This IP address is used to communicate with other subnets.	192.168.0.1
Advanced Settings/DN S Server Address	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	100.125.x.x
Advanced Settings/Tag	The subnet tag, which consists of a key and value pair. You can add a maximum of 10 tags to each subnet. The tag key and value must meet the requirements listed in Table 3-8.	Key: subnet_key1Value: subnet-01

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, 	subnet_key1
	underscores (_), and hyphens (-).	
Value	Can contain a maximum of 43 characters.	subnet-01

Table 3-8 Subnet tag key and value requirements

6. Click OK.

Precautions

When a subnet is created, there are five reserved IP addresses, which cannot be used. For example, in a subnet with CIDR block 192.168.0.0/24, the following IP addresses are reserved:

- 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: Gateway address.
- 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.

If you configured the default settings under **Advanced Settings** during subnet creation, the reserved IP addresses may be different from the default ones, but there will still be five of them. The specific addresses depend on your subnet settings.

3.3.2 Modifying a Subnet

Scenarios

Modify the subnet name and DNS server address.

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.

In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.

- 4. In the subnet list, locate the target subnet and click its name. The subnet details page is displayed.
- 5. On the **Summary** tab, click on the right of the parameter to be modified and modify the parameter as prompted.

Table 3-9 Parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	Subnet
DNS Server Address	By default, two DNS server addresses are configured. You can change them as required. A maximum of two DNS server addresses are supported. Use commas (,) to separate every two addresses.	100.125.x.x
Description	Supplementary information about the subnet. This parameter is optional.	-
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

6. Click OK.

3.3.3 Managing Subnet Tags

Scenarios

You can add tags to subnets to help you identify and organize them.

You can add a tag to a subnet when creating the subnet, or you can add a tag to a created subnet on the subnet details page. A maximum of 10 tags can be added to each subnet.

A tag consists of a key and value pair. **Table 3-10** lists the tag key and value requirements.

Parameter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain letters, digits, underscores (_), and hyphens (-). 	subnet_key1
Value	Can contain a maximum of 43 characters.	subnet-01

Table 3-10 Subnet tag key and value requirements

Procedure

Search for subnets by tag key and value on the page showing the subnet list.

- 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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. In the search box above the subnet list, click the search box.

 Click the tag key and then the value as required. The system filters resources based on the tag you select.

Add, delete, edit, and view tags on the Tags tab of a subnet.

- 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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. In the subnet list, locate the target subnet and click its name.
- 5. On the subnet details page, 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 subnet, 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 key and value, and click **OK**.

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**.

3.3.4 Exporting Subnet List

Scenarios

Information about all subnets under your account can be exported as an Excel file to a local directory. This file records the name, ID, VPC, CIDR block, and associated route table of each subnet.

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. In the upper left corner of the subnet list, click **Export**.

 The system will automatically export information about all subnets under your account in a selected region as an Excel file to a local directory.

3.3.5 Viewing and Deleting Resources in a Subnet

Scenarios

VPC subnets have private IP addresses used by cloud resources. This section describes how to view resources that are using private IP addresses of subnets. If these resources are no longer required, you can delete them.

You can view resources, including ECSs, BMSs, load balancers, and NAT gateways.

NOTICE

After you delete all resources in a subnet by referring to this section, the message "Delete the resource that is using the subnet and then delete the subnet." is displayed when you delete the subnet, you can refer to Viewing IP Addresses in a Subnet.

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**. The **Subnets** page is displayed.
- 4. Locate the target subnet and click its name.
 - The subnet details page is displayed.
- 5. On the **Summary** page, view the resources in the subnet.
 - a. In the **VPC Resources** area, view the quantities of resources, such as ECSs, BMSs, network interfaces, and load balancers, in the subnet. Click the resource quantity with a hyperlink to view the resources in the subnet.
 - b. In the **Networking Components** area on the right of the page, view the NAT gateways in the subnet.
- 6. Delete resources from the subnet.

Table 3-11 Viewing and deleting resources in a subnet

Resource	Reference
Load balancer	You can directly switch to load balancers from the subnet details page.
	Click the load balancer quantity. The load balancer list is displayed.
	Locate the row that contains the load balancer and click Delete in the Operation column.

3.3.6 Viewing IP Addresses in a Subnet

Scenarios

A subnet is an IP address range in a VPC. This section describes how to view the used IP addresses in a subnet.

- Virtual IP addresses
- Private IP addresses
 - Used by the subnet itself, such as the gateway, system interface, and DHCP.
 - Used by cloud resources, such as ECSs, load balancers, and RDS instances.

Notes and Constraints

- A subnet cannot be deleted if its IP addresses are used by cloud resources.
- A subnet can be deleted if its IP addresses are used by itself.

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. Locate the target subnet and click its name.
 - The subnet details page is displayed.
- 5. Click the **IP Addresses** tab to view the IP addresses in the subnet.
 - a. In the virtual IP address list, you can view the virtual IP addresses assigned from the subnet.
 - b. In the private IP address list in the lower part of the page, you can view the private IP addresses and the resources that use the IP addresses of the subnet.

Follow-up Operations

If you want to view and delete the resources in a subnet, refer to Why Can't I Delete My VPCs and Subnets?

3.3.7 Deleting a Subnet

Scenarios

If your subnet is no longer required, you can delete it:

Notes and Constraints

If you want to delete a subnet that has custom routes, virtual IP addresses, or other resources, you need to delete these resources as prompted on the console first and then delete the subnet.

You can refer to Why Can't I Delete My VPCs and Subnets?

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. In the subnet list, locate the row that contains the subnet you want to delete and click **Delete** in the **Operation** column.
 - A confirmation dialog box is displayed.
- 5. Click Yes.

NOTICE

If a VPC cannot be deleted, a message will be displayed on the console. Delete the resources that are in the VPC by referring to **Why Can't I Delete My VPCs and Subnets?**

3.4 IPv4 and IPv6 Dual-Stack Network

What Is an IPv4/IPv6 Dual-Stack Network?

IPv4 and IPv6 dual-stack allows your resources, such as ECSs, to use both IPv4 and IPv6 addresses for private and public network communications. For example, if ECSs use the IPv4/IPv6 dual-stack network:

- ECSs can communicate with each other using private IPv4 addresses.
- ECSs can communicate with the Internet after they are bound with EIPs.
- ECSs can communicate with each other using IPv6 addresses.
- ECSs can communicate with the Internet after their IPv6 addresses are added to shared bandwidths.

□ NOTE

If you select **Enable** for **IPv6 CIDR Block** when creating a subnet, an IPv6 CIDR block will be automatically assigned to the subnet.

Basic operations on IPv4 and IPv6 dual-stack networks are the same as those on IPv4 networks, except some parameters. Check the console pages for details.

Notes and Constraints

- The IPv4/IPv6 dual-stack function is currently free, but will be billed at a later date (price yet to be determined).
- Only certain ECS specifications support IPv6 networks and can use IPv4/IPv6 dual-stack networks. You need to select such ECSs in supported regions.
 To check which ECSs support IPv6:
 - On the ECS console: Click Create ECS. On the displayed page, view the ECS specifications.

If there is the **IPv6** parameter with the value of **Yes**, the ECS specifications support IPv6.

IPv6 Application Scenarios

If your ECS supports IPv6, you can use the IPv4/IPv6 dual-stack network. **Table 3-12** shows the example application scenarios.

Table 3-12 Application scenarios of IPv4/IPv6 dual stack

Applica tion Scenari o	Description	Subnet	ECS	
Private commu nicatio n using IPv6 address es	Your applications deployed on ECSs need to communicate with other systems (such as databases) through private networks using IPv6 addresses.	IPv4 CIDR blockIPv6 CIDR block	 Private IPv4 address: used for private communication IPv6 address: used for private communication. 	
Public commu nicatio n using IPv6 address es	Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.	 IPv4 CIDR block IPv6 CIDR block 	CIDR block	 Private IPv4 address + IPv4 EIP: used for public network communication
	Your applications deployed on ECSs need to both provide services accessible from the Internet and analyze the access request data using IPv6 addresses.		IPv6 address + shared bandwidth: used for public network communication	

Basic Operations

Creating an IPv6 Subnet

Create an IPv6 subnet by following the instructions in **Creating a Subnet for the VPC**. Select **Enable** for **IPv6 CIDR Block**. An IPv6 CIDR block will be automatically assigned to the subnet. IPv6 cannot be disabled after the subnet is created.

Create Subnet * VPC vpc-03 IPv4 CIDR block: 192.168.3.0/24 The VPC already contains 1 subnets. **★** AZ AZ1 (?) subnet-8b9c * Name ★ IPv4 CIDR Block Available IP Addresses: 251 The CIDR block cannot be modified after the subnet has been created. IPv6 CIDR Block Enable ? Default (?) Associated Route Table Advanced Settings ▼ Gateway | DNS Server Address | NTP Server Address | DUCD Loses Time | Tag | Description Cancel

Figure 3-4 Creating an IPv6 subnet

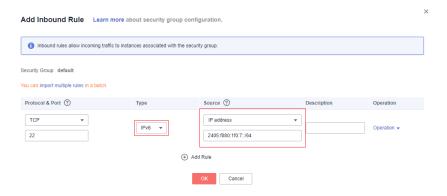
Viewing In-Use IPv6 Addresses

In the subnet list, click the subnet name. On the displayed page, view in-use IPv4 and IPv6 addresses on the **IP Addresses** tab.

Adding a Security Group Rule (IPv6)

Add a security group rule with **Type** set to **IPv6** and **Source** or **Destination** set to an IPv6 address or IPv6 CIDR block.

Figure 3-5 Adding a security group rule (IPv6)



Adding an IPv6 Network ACL Rule

Add a network ACL rule with **Type** set to **IPv6** and **Source** or **Destination** set to an IPv6 address or IPv6 CIDR block.

Figure 3-6 Adding an IPv6 network ACL rule



Adding an IPv6 EIP or Dual-Stack NIC IPv6 Address to a Shared Bandwidth

Add an IPv6 EIP or a dual-stack NIC IPv6 address to a shared bandwidth by following the instructions provided in **Adding EIPs to a Shared Bandwidth**.

Figure 3-7 Adding a dual-stack NIC IPv6 address to a shared bandwidth



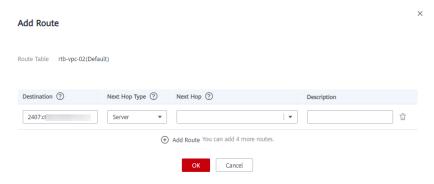
Adding a Route (IPv6)

Add a route with **Destination** and **Next Hop** set to an IPv4 or IPv6 CIDR block. For details about how to add a route, see **Adding a Custom Route**. If the destination is an IPv6 CIDR block, the next hop can only be an IP address in the same VPC as the IPv6 CIDR block.

Ⅲ NOTE

If the destination is an IPv6 CIDR block, the next hop type can only be an ECS, extension NIC, or virtual IP address. The next hop must also have IPv6 addresses.

Figure 3-8 Add Route



Assigning an IPv6 Virtual IP Address

Assign a virtual IPv4 or IPv6 address by referring to **Assigning a Virtual IP Address**.

Figure 3-9 Assigning a Virtual IP Address



∩ NOTE

Each virtual IPv6 address can only be bound to one dual-stack NIC.

Dynamically Assigning IPv6 Addresses

After an ECS is created successfully, you can view the assigned IPv6 address on the ECS details page. You can also log in to the ECS and run the **ifconfig** command to view the assigned IPv6 address.

If an IPv6 address fails to be automatically assigned or the selected image does not support the function of automatic IPv6 address assignment, manually obtain the IPv6 address by referring to "Dynamically Assigning IPv6 Addresses" in *Elastic Cloud Server User Guide*.

Ⅲ NOTE

If an ECS is created from a public image:

Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 is supported and then check whether dynamic IPv6 address assignment has been enabled. Currently, all Linux public images support IPv6, and dynamic IPv6 address assignment is enabled for Ubuntu 16 by default. You do not need to configure dynamic IPv6 address assignment for the Ubuntu 16 OS. For other Linux public images, you need to enable this function.

4 Access Control

4.1 Differences Between Security Groups and Network ACLs

You can configure network ACL and security group rules to protect the instances in your VPC, such as ECSs, CCI instances, and databases.

- A security group protects the instances in it.
- A network ACL protects associated subnets and all the resources in the subnets.

Figure 4-1 shows how security groups and network ACLs work. In **Figure 4-1**, security groups A and B protect the network security of ECSs. Network ACLs A and B add an additional layer of defense to subnets 1 and 2.

VPC 192.168.0.0/16 Network ACL A Network ACL B Subnet 1 Subnet 2 192.168.0.0/24 192.168.1.0/24 Security group A Security group B Security group B ECS **ECS ECS** ECS **ECS**

Figure 4-1 Security groups and network ACLs

Table 4-1 describes the differences between security groups and network ACLs.

Table 4-1 Differences between security groups and network ACLs

Category	Security Group	Network ACL
Protectio n Scope	Protects instances in a security group, such as ECSs, CCI instances, and databases.	Protects subnets and all the instances in the subnets.
Rules	Does not support Allow or Deny rules.	Supports both Allow and Deny rules.
Matching Order	If there are conflicting rules, they are combined and applied together.	If rules conflict, the rule with the highest priority takes effect.
Usage	 When creating an instance, such as an ECS, you must select a security group. If you do not have a security group, a default security group will be created for you. After creating an instance, you can: Add or remove the instance to or from the security group on the security group console. Associate or disassociate a security group with or from the instance on the instance console. 	Selecting a network ACL is not allowed when you create a subnet. You must create a network ACL, add inbound and outbound rules, associate subnets with it, and enable network ACL. The network ACL then protects the associated subnets and instances in the subnets.
Packets	Packet filtering based on the 3-tuple (protocol, port, and source/destination) is supported.	Packet filtering based on the 5- tuple (protocol, source port, destination port, and source/ destination) is supported.

4.2 Security Group

4.2.1 Security Groups and Security Group Rules

Security Groups

A security group is a collection of access control rules for cloud resources, such as cloud servers, containers, and databases, that have the same security protection requirements and that are mutually trusted. After a security group is created, you can configure access rules that will apply to all cloud resources added to this security group.

If you have not created any security groups yet, the system automatically creates a default security group for you and associates it with the instance (such as an ECS) when you create it. For details about the default security group, see **Default Security Group**.

Security Group Basics

- Security groups are stateful. If you send a request from your instance and the
 outbound traffic is allowed, the response traffic for that request is allowed to
 flow in regardless of inbound security group rules. Similarly, if inbound traffic
 is allowed, responses to allowed inbound traffic are allowed to flow out,
 regardless of outbound rules.
- Security groups use connection tracking to track traffic to and from instances.
 If an inbound rule is modified, the modified rule immediately takes effect for the existing traffic. Changes to outbound security group rules do not affect existing persistent connections and take effect only for new connections.
 - If you add, modify, or delete a security group rule, or add or remove an instance to or from a security group, the inbound connections of all instances in the security group will be automatically cleared.
 - The existing inbound persistent connections will be disconnected. All the new connections will match the new rules.
 - The existing outbound persistent connections will not be disconnected, and the original rule will still be applied. All the new connections will match the new rules

NOTICE

After a persistent connection is disconnected, new connections will not be established immediately until the timeout period of connection tracking expires. For example, after an ICMP persistent connection is disconnected, a new connection will be established and a new rule will apply when the timeout period (30s) expires.

- The timeout period of connection tracking varies by protocol. The timeout period of a TCP connection in the established state is 600s, and that of an ICMP connection is 30s. For other protocols, if packets are received in both inbound and outbound directions, the connection tracking timeout period is 180s. If packets are received only in one direction, the connection tracking timeout period is 30s.
- The timeout period of TCP connections varies by connection status. The timeout period of a TCP connection in the established state is 600s, and that of a TCP connection in the FIN-WAIT state is 30s.

Security Group Rules

A security group has inbound and outbound rules to control traffic that's allowed to reach or leave the instances associated with the security group. You can specify protocol, port, source/destination for a security group rule. **Table 4-2** describes key information about a security group rule.

Table 4-2 Security group rule information

Parameter	Description
Protocol	The network protocol used to match traffic in a security group rule. Currently, the value can be All , TCP , UDP , ICMP , or more.
Port	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.
	Outbound rules control outgoing traffic over specific ports from instances in the security group.
Source (Inbound)	The source in an inbound rule is used to match the IP address or address range of an external request. The source can be: • IP address:
	- Example IPv4 address: 192.168.10.10/32
	 Example IPv4 address range: 192.168.52.0/24 All IPv4 addresses: 0.0.0.0/0
	• Security group: You can select another security group in the same region under the current account as the source. For example, instance A is in security group A and instance B is in security group B. If security group A has an inbound rule with Action set to Allow and Source set to security group B, access from instance B is allowed to instance A.
Destination (Outbound)	The destination in an outbound rule is used to match the IP address or address range of an internal request. The destination can be:
	IP address:
	- Example IPv4 address: 192.168.10.10/32
	Example IPv4 address range: 192.168.52.0/24 All IPv4 addresses: 0.0.0.0/0
	• Security group: You can select another security group in the same region under the current account as the destination. For example, instance A is in security group A and instance B is in security group B. If security group A has an outbound rule with Action set to Allow and Destination set to security group B, access from instance A is allowed to instance B.

Like whitelists, security group rules work as follows:

Inbound rules control incoming traffic to instances in the security group.
 If an inbound request matches the source in an inbound security group rule with Action set to Allow, the request is allowed and other requests are denied.

By default, you do not need to configure deny rules in the inbound direction because requests that do not match allow rules will be denied.

Outbound rules control outgoing traffic from instances in the security group.
 If the destination of an outbound security group rule with Action set to Allow is 0.0.0.0/0, all outbound requests are allowed.

0.0.0.0/0 represents all IPv4 addresses.

Table 4-3 uses custom security group sg-AB as an example to describe its inbound and outbound rules in detail.

Table 4-3 Rules in security group sg-AB

Directio n	Protoc ol & Port	Source/ Destination	Description
Inbound	All	Source: sg-AB	Allows ECSs in the security group to communicate with each other.
Inbound	TCP: 22	Source: 0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over port 22 (SSH) for remotely logging in to Linux ECSs.
Inbound	TCP: 3389	Source: 0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over port 3389 (RDP) for remotely logging in to Windows ECSs.
Inbound	TCP: 80	Source: 10.5.6.30/32	Allows IP address 10.5.6.30 to access ECSs in the security group over port 80.
Outbou nd	All	Destination: 0.0.0.0/0	Allows access from ECSs in the security group to any IPv4 address over any port.

NOTICE

- After a port is enabled in a security group rule, ensure that the port in the instance is also enabled to ensure normal network communication.
- Generally, instances in the same security group can communicate with each other by default. If instances in the same security group cannot communicate with each other, the possible causes are as follows:
 - The inbound rule for communication between instances in the same security group is deleted.
 - Different VPCs cannot communicate with each other. The instances belong to the same security group but different VPCs.

You can use **VPC peering connections** to connect VPCs in different regions.

Security Group Constraints

By default, you can add up to 50 security group rules to a security group.

By default, you can add an ECS or extension NIC to up to five security groups.
 In such a case, the rules of all the selected security groups are aggregated to take effect.

4.2.2 Default Security Group

If you have not created any security groups yet, the system automatically creates a default security group for you and associates it with the instance when you create it. A default security group has the following rules:

- Inbound rules control incoming traffic to instances in a security group. Only instances in the same security group can communicate with each other, and all inbound requests are denied.
- Outbound rules allow all outbound traffic and response traffic to the outbound requests.

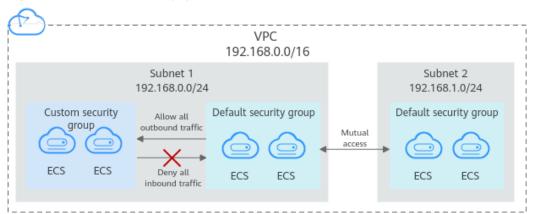


Figure 4-2 Default security group

∩ NOTE

- You cannot delete the default security group, but you can modify existing rules or add rules to the group.
- The default security group denies all external requests. To log in to an instance
 associated with this security group, add a security group rule by referring to Remotely
 Logging In to an ECS from a Local Server.

Table 4-4 describes the default rules for the default security group.

Table 4-4	Default s	ecurity	group	rules	5

Directi on	Protoc ol	Port/ Range	Source/ Destination	Description
Outbo und	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inboun d	All	All	Source: the current security group (for example, sg-xxxxx)	Allows communications among ECSs within the security group and denies all inbound traffic (incoming data packets).

4.2.3 Security Group Configuration Examples

Here are some common security group configuration examples for different scenarios, including remote login to ECSs, website access, and internal communication between instances in different security groups.

Generally, a security group denies all external requests by default. You need to add inbound rules to a security group based on the whitelist principle to allow specific external requests to access instances in the security group.

- Remotely Logging In to an ECS from a Local Server
- Remotely Connecting to an ECS from a Local Server to Upload or Download Files
- Setting Up a Website on an ECS to Provide Services Externally
- Using ping Command to Verify Network Connectivity
- Enabling ECSs In Different Security Groups to Communicate Through an Internal Network
- ECS Providing Database Access Service
- Allowing ECSs to Access Only Specific External Websites

By default, all outbound rules of a security group allow all requests from instances in the security group to access external networks. **Table 4-5** lists the rules.

Table 4-5 Default outbound rules in a security group

Directio n	Protoc ol & Port	Destination	Description
Outbou nd	All	0.0.0.0/0	This rule allows access from instances in the security group to any IPv4 address over any port.

Remotely Logging In to an ECS from a Local Server

A security group denies all external requests by default. To remotely log in to an ECS from a local server, add an inbound security group rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable the SSH (22) port. For details, see **Table 4-6**.
- To remotely log in to a Windows ECS using RDP, enable the RDP (3389) port. For details, see **Table 4-7**.

Table 4-6 Remotely logging in to a Linux ECS using SSH

Direction	Protocol & Port	Source
Inbound	TCP: 22	IP address: 0.0.0.0/0

Table 4-7 Remotely logging in to a Windows ECS using RDP

Direction	Protocol & Port	Source
Inbound	TCP: 3389	IP address: 0.0.0.0/0

NOTICE

If the source is set to 0.0.0.0/0, remotely logging in to the ECS through any IP address is allowed. To ensure security, set the source to a specific IP address based on service requirements. For details about the configuration example, see **Table 4-8**.

Table 4-8 Remotely logging in to an ECS using a specified IP address

ECS Type	Direction	Protocol & Port	Source
Linux ECS	Inbound	TCP: 22	IP address: 192.168.0.0/24
Windows ECS	Inbound	TCP: 3389	IP address: 10.10.0.0/24

Remotely Connecting to an ECS from a Local Server to Upload or Download Files

By default, a security group denies all external requests. If you need to remotely connect to an ECS from a local server to upload or download files, you need to enable FTP ports 20 and 21.

Table 4-9 Remotely connecting to an ECS from a local server to upload or download files

Direction	Protocol & Port	Source
Inbound	TCP: 20-21	IP address: 0.0.0.0/0

NOTICE

You must first install the FTP server program on the ECSs and check whether ports 20 and 21 are working properly.

Setting Up a Website on an ECS to Provide Services Externally

A security group denies all external requests by default. If you have set up a website on an ECS that can be accessed externally, you need to add an inbound rule to the ECS security group to allow access over specific ports, such as HTTP (80) and HTTPS (443).

Table 4-10 Setting up a website on an ECS to provide services externally

Direction	Protocol & Port	Source
Inbound	TCP: 80	IP address: 0.0.0.0/0
Inbound	TCP: 443	IP address: 0.0.0.0/0

Using ping Command to Verify Network Connectivity

By default, a security group denies all external requests. If you need to run the **ping** command on an ECS to verify network connectivity, add an inbound rule to the ECS security group to allow access over the ICMP port.

Table 4-11 Using **ping** command to verify network connectivity

Direction	Protocol & Port	Source
Inbound	ICMP: All	IP address: 0.0.0.0/0

Enabling ECSs In Different Security Groups to Communicate Through an Internal Network

ECSs in the same VPC but associated with different security groups cannot communicate with each other. If you want to share data between ECSs in a VPC, for example, ECSs in security group sg-A need to access MySQL databases in security group sg-B, you need to add an inbound rule to security group sg-B to allow access from ECSs in security group sg-A over MySQL port 3306.

Table 4-12 Enabling instances in different security groups to communicate through an internal network

Direction	Protocol & Port	Source
Inbound	TCP: 3306	Security group: sg-A

ECS Providing Database Access Service

A security group denies all external requests by default. If you have deployed the database service on an ECS and need to allow other ECSs to access the database service through an internal network, you need to add an inbound rule to the security group of the ECS with the database service deployed to allow access over

ports, for example, MySQL (3306), Oracle (1521), MS SQL (1433), PostgreSQL (5432) and Redis (6379).

Table 4-13 ECS providing database access service

Direction	Protocol & Port	Source	Description
Inbound	TCP: 3306	Security group: sg-A	This rule allows ECSs in security group sg-A to access the MySQL database service.
Inbound	TCP: 1521	Security group: sg-B	This rule allows ECSs in security group sg-B to access the Oracle database service.
Inbound	TCP: 1433	IP address: 172.16.3.21/32	This rule allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service.
Inbound	TCP: 5432	IP address: 192.168.0.0/24	This rule allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database service.

NOTICE

In this example, the source is for reference only. Set the source address based on actual requirements.

Allowing ECSs to Access Only Specific External Websites

By default, a security group allows all outbound traffic. **Table 4-15** lists the default rules. If you want to allow ECSs to access only specific websites, configure the security groups of the ECSs as follows:

1. First, add outbound rules to allow traffic over specific ports and to specific IP addresses.

Table 4-14 Enabling instances in different security groups to communicate through an internal network

Direction	Protocol & Port	Source
Outbound	TCP: 80	IP address: 132.15.XX.XX
Outbound	TCP: 443	IP address: 145.117.XX.XX

2. Then, delete the original outbound rules that allow all traffic shown in **Table** 4-15.

Directi on	Protoc ol & Port	Destination	Description
Outbou nd	All	0.0.0.0/0	This rule allows access from instances in the security group to any IPv4 address over any port.

Table 4-15 Default outbound rules in a security group

4.2.4 Managing a Security Group

4.2.4.1 Creating a Security Group

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

Notes and Constraints

If you have not created any security groups yet, the system automatically creates a default security group for you and associates it with the instance (such as an ECS) when you create it.

The default security group name is **default**. For details, see **Default Security Group**.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. In the upper right corner, click **Create Security Group**. The **Create Security Group** page is displayed.
- 5. Configure the parameters as prompted.

 Table 4-16 Parameter description

Parameter	Description	Example Value
Name	Mandatory	sg-AB
	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.	
Description	Description Optional	
	Supplementary information about the security group. This parameter is optional.	
	The security group description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

4.2.4.2 Cloning a Security Group

Scenarios

You can clone a security group from one region to another to quickly apply the security group rules to ECSs in another region.

You can clone a security group in the following scenarios:

- For example, you have security group sg-A in region A. If ECSs in region B require the same security group rules as those configured for security group sg-A, you can clone security group sg-A to region B, freeing you from creating a new security group in region B.
- If you need new security group rules, you can clone the original security group as a backup.
- Before you modify security group rules used by a service, you can clone the security group and modify the security group rules in the test environment to ensure that the modified rules work.

Notes and Constraints

- You can clone a security group from the same or a different region.
 - If you want to clone a security group from the same region, you can clone all rules in the security group.

- If you want to clone a security group from a different region, the system will clone only rules whose source and destination are IP addresses and rules whose source and destination is the current security group.
- Cloning a security group clones its security group rules, but not the instances associated with the security group.

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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. Locate the row that contains the security group, click **More** in the **Operation** column, and click **Clone**.
- 5. Select the region and name of the new security group as prompted.
- 6. Click **OK**.

You can then switch to the required region to view the cloned security group in the security group list.

4.2.4.3 Modifying a Security Group

Scenarios

After a security group is created, you can change its name and description.

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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. Locate the row that contains the security group, click **More** in the **Operation** column, and click **Modify**.

The **Modify Security Group** dialog box is displayed.

- 5. Modify the name and description of the security group as required.
- 6. Click **OK** to save the modification.

4.2.4.4 Deleting a Security Group

Scenarios

If your security group is no longer required, you can delete it.

Notes and Constraints

- The default security group is named **default** and cannot be deleted.
- If you want to delete a security group that is associated with instances, such as cloud servers, containers, and databases, you need to remove the instances from the security group first. For details, see Adding an Instance to or Removing an Instance from a Security Group.
- A security group cannot be deleted if it is used as the source or destination of a rule in another security group.

Delete or **modify** the rule and delete the security group again.

For example, if the source of a rule in security group **sg-B** is set to **sg-A**, you need to delete or modify the rule in **sg-B** before deleting **sg-A**.

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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. Locate the row that contains the target security group, click **More** in the **Operation** column, and click **Delete**.
 - A confirmation dialog box is displayed.
- 5. Confirm the information and click **Yes**.

4.2.5 Managing Security Group Rules

4.2.5.1 Adding a Security Group Rule

Scenarios

A security group is a collection of access control rules to control the traffic that is allowed to reach and leave the cloud resources that it is associated with. The cloud resources can be cloud servers, containers, databases, and more. A security group consists of inbound and outbound rules.

Like whitelists, security group rules work as follows:

- Inbound rules control incoming traffic to instances in the security group.
 If an inbound request matches the source in an inbound security group rule with Action set to Allow, the request is allowed and other requests are denied.
 - By default, you do not need to configure deny rules in the inbound direction because requests that do not match allow rules will be denied.
- Outbound rules control outgoing traffic from instances in the security group.
 If the destination of an outbound security group rule with Action set to Allow is 0.0.0.0/0, all outbound requests are allowed.

0.0.0.0/0 represents all IPv4 addresses.

If the rules of the security group associated with your instance cannot meet your requirements, for example, you need to allow inbound traffic on a specific TCP port, you can add an inbound rule to allow traffic on the TCP port.

Security Group Rule Configuration Examples

• Before configuring security group rules, you need to plan access policies for instances in the security group. For details about common security group rule configuration examples, see **Security Group Configuration Examples**.

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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. Locate the row that contains the target security group, and click **Manage Rule** 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 4-17 Inbound rule parameter description

Param eter	Description	Example Value
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	ТСР
	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, or 22-30

Param eter	Description	Example Value
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. For example:	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	If the source is a security group, this rule will apply to all instances associated with the selected security group.	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

7. Click **OK**.

The inbound rule list is displayed.

8. On the **Outbound Rules** tab, click **Add Rule**.

The **Add Outbound Rule** dialog box is displayed.

9. Configure required parameters.

You can click + to add more outbound rules.

Table 4-18 Outbound rule parameter description

Param eter	Description	Example Value
Туре	Destination IP address version. You can select:IPv4IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. Currently, the value can be All, TCP, UDP, ICMP, or more.	ТСР
	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, or 22-30

Param eter	Description	Example Value
Destina tion	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. For example:	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Descrip tion	Supplementary information about the security group rule. This parameter is optional.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

10. Click **OK**.

The outbound rule list is displayed.

Verifying Security Group Rules

After required security group rules are added, you can verify whether the rules take effect. 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 4-19** shows the rule.

Table 4-19 Security group rule

Direction	Protocol & Port	Source
Inbound	TCP: 80	0.0.0.0/0

Linux ECS

Check whether the security group rule takes effect on a Linux ECS:

- 1. Log in to the ECS.
- 2. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | grep 80

If information similar to Figure 4-3 is displayed, TCP port 80 is enabled.

Figure 4-3 Command output for the Linux ECS

tcp	0	0 0.0.0.0: <mark>80</mark>	0.0.0.0:*	LISTEN
-----	---	----------------------------	-----------	--------

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.

Windows ECS

To verify the security group rule on a Windows ECS:

- 1. Log in to the ECS.
- 2. Choose **Start** > **Run**. Type cmd to open the Command Prompt.
- 3. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | findstr 80

If information similar to **Figure 4-4** is displayed, TCP port 80 is enabled.

Figure 4-4 Command output for the Windows ECS



4. 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.

4.2.5.2 Fast-Adding Security Group Rules

Scenarios

The fast-adding rule function of security groups allows you to quickly add rules with common ports and protocols for remote login, ping tests, common web services, and database services.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. Locate the row that contains the target security group and click **Manage Rule** in the **Operation** column.

The page for configuring security group rules is displayed.

- 5. On the **Inbound Rules** tab, click **Fast-Add Rule**.
 - The Fast-Add Inbound Rule dialog box is displayed.
- 6. Configure required parameters.

Table 4-20 Inbound rule parameter description

Param eter	Description	Example Value
Protoco Is and Ports	 Common protocols and ports are provided for: Remote login and ping Web services Databases 	SSH (22)
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. You can specify: • Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6) • IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) • All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) • Security group: sg-abc 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
Descrip tion	(Optional) Supplementary information about the security group rule. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

7. Click **OK**.

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

- On the Outbound Rules tab, click Fast-Add Rule.
 The Fast-Add Outbound Rule dialog box is displayed.
- 9. Configure required parameters.

Table 4-21 Outbound rule parameter description

Param eter	Description	Example Value
Protoc ols and Ports	Common protocols and ports are provided for: Remote login and pingWeb servicesDatabases	SSH (22)
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Destin ation	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. You can specify:	0.0.0.0/0
	• Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6)	
	• IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
	• All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	Security group: sg-abc	
Descrip tion	(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.

4.2.5.3 Modifying a Security Group Rule

Scenarios

You can modify the port, protocol, and IP address of your security group rules as required to ensure the security of your instances.

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.

In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.

- 4. In the security group list, click the name of the security group.

 The security group details page is displayed.
- 5. Click the **Inbound Rules** or **Outbound Rules** tab as required. The security group rule list is displayed.
- 6. Locate the row that contains the rule and click **Modify** in the **Operation** column.
- 7. Modify the security group rule information as prompted and click **Confirm**.

4.2.5.4 Replicating a Security Group Rule

Scenarios

You can replicate an existing security group rule and modify it to quickly generate a new rule.

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 security group list, click the name of the security group. The security group details page is displayed.
- 4. Click the **Inbound Rules** or **Outbound Rules** tab as required. The security group rule list is displayed.
- 5. Locate the row that contains the rule and click **Replicate** in the **Operation** column.
 - The **Replicate Inbound Rule** dialog box is displayed.
- 6. Modify the security group rule information as prompted and click **OK**.

4.2.5.5 Importing and Exporting Security Group Rules

Scenarios

You can configure security group rules in an Excel file and import the rules to the security group. You can also export security group rules to an Excel file. You are advised to use this function in the following scenarios:

- If you want to quickly create or restore a security group rule, you can import your exported security group rule file to the security group.
- If you want to back up security group rules locally, you can export the rules to an Excel file.
- If you want to quickly apply the rules of one security group to another, or if you want to modify multiple rules of the current security group at once, you can import or export existing rules.

Notes and Constraints

- The security group rules to be imported must be configured based on the template. Do not add parameters or change existing parameters. Otherwise, the import will fail.
- If a security group rule to be imported is the same as an existing one, the security group rule cannot be imported. You can delete the rule and try again.

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 **Access Control** > **Security Groups**. The security group list is displayed.
- 4. On the security group list, click the name of the target security group. The security group details page is displayed.
- 5. Export and import security group rules.
 - Click to export all rules of the current security group to an Excel file.
 - Click to import security group rules from an Excel file into the current security group.

Table 4-22 describes the parameters in the template for importing rules.

Table 4-22 Template parameters

Param eter	Description	Example Value
Directi on	The direction in which the security group rule takes effect.	Inbound
	 Inbound: Inbound rules control incoming traffic to instances in the security group. 	
	 Outbound: Outbound rules control outgoing traffic from instances in the security group. 	
Protoc ol &	The network protocol used to match traffic in a security group rule.	ТСР
Port	Currently, the value can be All , TCP , UDP , ICMP , or more.	

Param eter	Description	Example Value
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.	22, or 22-30
	Inbound rules control incoming traffic over specific ports to instances in the security group.	
	Outbound rules control outgoing traffic over specific ports from instances in the security group.	
Туре	Source IP address version. You can select: IPv4 IPv6	IPv4
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. For example: • IP address:	sg- test[96a8a 93f-XXX- d7872990c 314]
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Destin ation	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. For example:	sg- test[96a8a 93f-XXX- d7872990c 314]
Descri ption	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Last Modifi ed	The time when the security group was modified.	-

4.2.5.6 Deleting a Security Group Rule

Scenarios

If your security group rule is no longer required, you can delete it.

Notes and Constraints

Security group rules use whitelists. Deleting a security group rule may result in ECS access failures. Security group rules work as follows:

- Inbound rule: If an inbound request matches the source in an inbound security group rule with **Action** set to **Allow**, the request is allowed.
- Outbound rule: If the destination of an outbound security group rule with **Action** set to **Allow** is 0.0.0.0/0, all outbound requests are allowed.

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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. In the security group list, click the name of the security group.

 The security group details page is displayed.
- 5. Click the **Inbound Rules** or **Outbound Rules** tab as required. The security group rule list is displayed.
- 6. In the security group rule list:
 - To delete a single security group rule, locate the row that contains the rule and click **Delete** in the **Operation** column.
 - To delete multiple security group rules, select multiple security group rules and click **Delete** in the upper left corner of the rule list.
- 7. Click Yes.

4.2.6 Managing Instances Associated with a Security Group

4.2.6.1 Adding an Instance to or Removing an Instance from a Security Group

Scenarios

When you create an instance, the system automatically adds the instance to a security group for protection.

- If one security group cannot meet your requirements, you can add an instance to multiple security groups.
- An instance must be added to at least one security group. If you want to change the security group for an instance, you can add the instance to a new security group and then remove the instance from the original security group.

Adding an Instance 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 Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. In the security group list, locate the row that contains the security group and click **Manage Instances** in the **Operation** column.

The **Associated Instances** tab is displayed.

5. Click an instance type.

The following operations use **Servers** as an example.

6. Click the **Servers** tab and click **Add**.

The **Add Server** dialog box is displayed.

7. In the server list, select one or more servers and click OK to add them to the current security group.

Removing an Instance from a Security Group

An instance must be added to at least one security group. If you want to remove an instance from a security group, the instance must be associated with at least two security groups now.

- 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.

- In the navigation pane on the left, choose Access Control > Security Groups.
 The security group list is displayed.
- 4. In the security group list, locate the row that contains the security group and click **Manage Instances** in the **Operation** column.

The **Associated Instances** tab is displayed.

5. Click an instance type.

The following operations use **Servers** as an example.

- 6. Click the **Servers** tab, select one or more servers, and click **Remove** in the upper left corner of the server list.
 - A confirmation dialog box is displayed.
- 7. Confirm the information and click Yes.

4.2.6.2 Viewing the Security Group of an ECS

Scenarios

View inbound and outbound rules of a security group used by an ECS.

Procedure

- 1. Log in to the management console.
- 2. In the upper left corner of the page, click =. In the service list, choose Computing > Elastic Cloud Server.

The ECS list is displayed.

- 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 and view information about the security group used by the ECS.

You can view the security groups associated with the ECS and the inbound and outbound rules.

4.2.6.3 Changing the Security Group of an ECS

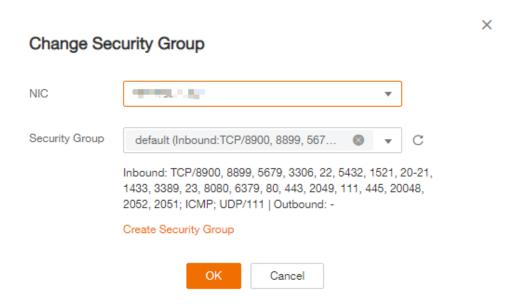
Scenarios

Change the security group associated with an ECS NIC.

Procedure

- 1. Log in to the management console.
- 2. Under Computing, click Elastic Cloud Server.
- In the ECS list, locate the row that contains the target ECS. Click More in the Operation column and select Manage Network > Change Security Group.
 The Change Security Group dialog box is displayed.

Figure 4-5 Change Security Group



4. Select the target NIC and security groups.

You can select multiple security groups. In such a case, the rules of all the selected security groups will be aggregated to apply on the ECS.

To create a security group, click Create Security Group.

◯ NOTE

Using multiple security groups may deteriorate ECS network performance. You are suggested to select no more than five security groups.

5. Click OK.

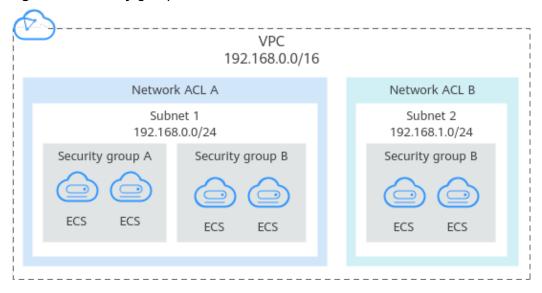
4.3 Network ACL

4.3.1 Network ACL Overview

A network ACL is an optional layer of security for your subnets. After you associate one or more subnets with a network ACL, you can control traffic in and out of the subnets.

Figure 4-6 shows how a network ACL works.

Figure 4-6 Security groups and network ACLs



Similar to security groups, network ACLs control access to subnets and add an additional layer of defense to your subnets. Security groups only have the "allow" rules, but network ACLs have both "allow" and "deny" rules. You can use network ACLs together with security groups to implement comprehensive and fine-grained access control.

Differences Between Security Groups and Network ACLs summarizes the basic differences between security groups and network ACLs.

Network ACL Basics

 Your VPC does not come with a network ACL, but you can create a network ACL and associate it with a VPC subnet if required. By default, each network ACL denies all inbound traffic to and outbound traffic from the associated subnet until you add rules.

- You can associate a network ACL with multiple subnets. However, a subnet can only be associated with one network ACL at a time.
- Each newly created network ACL is in the **Inactive** state until you associate subnets with it.
- Network ACLs use connection tracking to track traffic to and from instances.
 Changes to inbound and outbound rules do not take effect immediately for the existing traffic.

If you add, modify, or delete a network ACL rule, or associate or disassociate a subnet with or from a network ACL, all the inbound and outbound persistent connections will not be disconnected. New rules will only be applied for the new connections.

NOTICE

After a persistent connection is disconnected, new connections will not be established immediately until the timeout period of connection tracking expires. For example, after an ICMP persistent connection is disconnected, a new connection will be established and a new rule will apply when the timeout period (30s) expires.

- The timeout period of connection tracking varies by protocol. The timeout period of a TCP connection in the established state is 600s, and that of an ICMP connection is 30s. For other protocols, if packets are received in both inbound and outbound directions, the connection tracking timeout period is 180s. If packets are received only in one direction, the connection tracking timeout period is 30s.
- The timeout period of TCP connections varies by connection status. The timeout period of a TCP connection in the established state is 600s, and that of a TCP connection in the FIN-WAIT state is 30s.

Default Network ACL Rules

By default, each network ACL has preset rules that allow the following packets:

- Packets whose source and destination are in the same subnet.
- Broadcast packets with the destination 255.255.255.255/32, which is used to configure host startup information.
- Multicast packets with the destination 224.0.0.0/24, which is used by routing protocols.
- Metadata packets with the destination 169.254.169.254/32 and TCP port number 80, which is used to obtain metadata.
- Packets from CIDR blocks that are reserved for public services (for example, packets with the destination 100.125.0.0/16).
- A network ACL denies all traffic in and out of a subnet excepting the
 preceding packets. Table 4-23 shows the default rules. You cannot modify or
 delete the default rules.

Direction	Priorit y	Actio n	Protoco l	Sourc e	Destinatio n	Description
Inbound	*	Deny	All	0.0.0.0	0.0.0.0/0	Denies all inbound traffic.
Outboun d	*	Deny	All	0.0.0.0	0.0.0.0/0	Denies all outbound traffic.

Table 4-23 Default network ACL rules

How Traffic Matches Network ACL Rules

- Each network ACL rule has a priority value where a smaller value corresponds to a higher priority. Any time two rules conflict, the rule with the higher priority is the one that gets applied. The rule whose priority value is an asterisk (*) has the lowest priority.
- If multiple network ACL rules conflict, only the rule with the highest priority takes effect. If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule.

Application Scenarios

- If the application layer needs to provide services for users, traffic must be allowed to reach the application layer from all IP addresses. However, you also need to prevent illegal access from malicious users.
 - Solution: You can add network ACL rules to deny access from suspect IP addresses.
- How can I isolate ports with identified vulnerabilities? For example, how do I isolate port 445 that can be exploited by WannaCry worm?
 - Solution: You can add network ACL rules to deny access traffic from a specific port and protocol, for example, TCP port 445.
- No defense is required for the east-west traffic between subnets, but access control is required for north-south traffic.
 - Solution: You can add network ACL rules to protect north-south traffic.
- For frequently accessed applications, a security rule sequence may need to be adjusted to improve performance.
 - Solution: A network ACL allows you to adjust the rule sequence so that frequently used rules are applied before other rules.

Configuration Procedure

Figure 4-7 shows the procedure for configuring a network ACL.

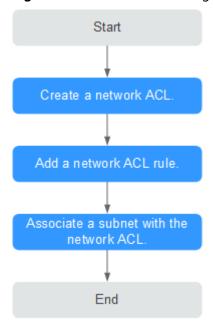


Figure 4-7 network ACL configuration procedure

- Create a network ACL by following the steps described in Creating a Network ACL.
- 2. Add network ACL rules by following the steps described in **Adding a Network ACL Rule**.
- 3. Associate subnets with the network ACL by following the steps described in Associating Subnets with a Network ACL. After subnets are associated with the network ACL, the subnets will be protected by the configured network ACL rules.

Notes and Constraints

- By default, each account can have up to 200 network ACLs in a region.
- A network ACL can contain no more than 20 rules in one direction, or performance will deteriorate.

4.3.2 Network ACL Configuration Examples

This section provides examples for configuring network ACLs.

- Denying Access from a Specific Port
- Allowing Access from Specific Ports and Protocols

Denying Access from a Specific Port

You might want to block TCP port 445 to protect against the WannaCry ransomware attacks. You can add a network ACL rule to deny all incoming traffic from TCP port 445.

Network ACL Configuration

Table 4-24 lists the inbound rules required.

Table 4-24 Network ACL rules

Dire ctio n	Ac ti on	Prot ocol	Sour ce	Source Port Range	Destination	Destina tion Port Range	Description
Inbo und	D en y	TCP	0.0.0. 0/0	1-6553 5	0.0.0.0/0	445	Denies inbound traffic from any IP address through TCP port 445.
Inbo und	All o w	All	0.0.0. 0/0	1-6553 5	0.0.0.0/0	All	Allows all inbound traffic.

□ NOTE

- By default, a network ACL denies all inbound traffic. You can add a rule to allow all inbound traffic if necessary.
- If you want a deny rule to be matched first, insert the deny rule above the allow rule. For details, see Changing the Sequence of a Network ACL Rule.

Allowing Access from Specific Ports and Protocols

In this example, an ECS in a subnet is used as the web server, and you need to allow inbound traffic from HTTP port 80 and HTTPS port 443 and allow all outbound traffic. You need to configure both the network ACL rules and security group rules to allow the traffic.

Network ACL Configuration

Table 4-25 lists the inbound and outbound rules required.

Table 4-25 Network ACL rules

Dire ctio n	Acti on	Protoc ol	Sourc e	Source Port Range	Desti natio n	Destina tion Port Range	Description
Inbo und	Allo w	ТСР	0.0.0.0 /0	1-65535	0.0.0. 0/0	80	Allows inbound HTTP traffic from any IP address to ECSs in the subnet through port 80.

Dire ctio n	Acti on	Protoc ol	Sourc e	Source Port Range	Desti natio n	Destina tion Port Range	Description
Inbo und	Allo w	ТСР	0.0.0.0	1-65535	0.0.0. 0/0	443	Allows inbound HTTPS traffic from any IP address to ECSs in the subnet through port 443.
Outb ound	Allo w	All	0.0.0.0	All	0.0.0. 0/0	All	Allows all outbound traffic from the subnet.

Security group configuration

Table 4-26 lists the inbound and outbound security group rules required.

Table 4-26 Security group rules

Direc tion	Protocol / Applicati on	Port	Source/ Destination	Description
Inbou nd	ТСР	80	Source: 0.0.0.0/0	Allows inbound HTTP traffic from any IP address to ECSs associated with the security group through port 80.
Inbou nd	ТСР	443	Source: 0.0.0.0/0	Allows inbound HTTPS traffic from any IP address to ECSs associated with the security group through port 443.
Outb ound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic from the security group.

A network ACL adds an additional layer of security. Even if the security group rules allow more traffic than that actually required, the network ACL rules allow only access from HTTP port 80 and HTTPS port 443 and deny other inbound traffic.

4.3.3 Managing Network ACLs

4.3.3.1 Creating a Network ACL

Scenarios

You can create a custom network ACL. By default, a newly created network ACL is disabled and has no inbound or outbound rules, or any subnets associated.

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 Access Control > Network ACLs.
- 4. In the right pane displayed, click Create Network ACL.
- 5. On the Create Network ACL page, configure parameters as prompted.

Table 4-27 Parameter descriptions

Parameter	Description	Example Value
Name	The network ACL name. This parameter is mandatory.	fw-92d3
	The name contains a maximum of 64 characters, which may consist of letters, digits, underscores (_), and hyphens (-). The name cannot contain spaces.	
Description	Supplementary information about the network ACL. This parameter is optional.	N/A
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

6. Click OK.

4.3.3.2 Modifying a Network ACL

Scenarios

Modify the name and description of a network ACL.

Procedure

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the displayed page, click on the right of **Name** and edit the network ACL name.
- 6. Click √ to save the new network ACL name.
- 7. Click on the right of **Description** and edit the network ACL description.
- 8. Click √ to save the new network ACL description.

4.3.3.3 Enabling or Disabling a Network ACL

Scenarios

After a network ACL is created, you may need to enable it based on network security requirements. You can also disable an enabled network ACL if needed. Before enabling a network ACL, ensure that subnets have been associated with the network ACL and that inbound and outbound rules have been added to the network ACL.

When a network ACL is disabled, custom rules will become invalid while default rules still take effect. Disabling a network ACL may interrupt network traffic. For information about the default network ACL rules, see **Default Network ACL Rules**.

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 Access Control > Network ACLs.
- 4. Locate the row that contains the network ACL, click **More** in the **Operation** column, and click **Enable** or **Disable**.
- 5. Click **Yes** in the displayed dialog box.

4.3.3.4 Viewing a Network ACL

Scenarios

View details about a network ACL.

Procedure

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- On the displayed page, click the Inbound Rules, Outbound Rules, and Associated Subnets tabs one by one to view details about inbound rules, outbound rules, and subnet associations.

4.3.3.5 Deleting a Network ACL

Scenarios

Delete a network ACL when it is no longer required.

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 **Access Control** > **Network ACLs**.
- Locate the network ACL, click More in the Operation column, and click Delete.
- 5. Click Yes.
 - □ NOTE

Deleting a network ACL will also disassociate its associated subnets and delete the network ACL rules.

4.3.4 Management Network ACL Rules

4.3.4.1 Adding a Network ACL Rule

Scenarios

Add an inbound or outbound rule based on your network security requirements.

Notes and Constraints

A network ACL can contain no more than 20 rules in one direction, or performance will deteriorate.

Procedure

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose **Access Control** > **Network ACLs**.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the **Inbound Rules** or **Outbound Rules** tab, click **Add Rule** to add an inbound or outbound rule.
 - Click + to add more rules.
 - Locate the row that contains the network ACL rule and click Replicate in the Operation column to replicate an existing rule.

Table 4-28 Parameter descriptions

Parameter	Description	Example Value
Туре	This parameter is available only after the IPv6 function is enabled.	IPv4
	The network ACL type. This parameter is mandatory. You can select a value from the drop-down list. Currently, only IPv4 and IPv6 are supported.	
Action	The action in the network ACL. This parameter is mandatory. You can select a value from the drop-down list. Currently, the value can be Allow or Deny .	Allow
Protocol	The protocol supported by the network ACL. This parameter is mandatory. You can select a protocol from the drop-down list. You can select TCP , UDP , ICMP , or All .	ТСР
Source	The source from which the traffic is allowed. The source can be an IP address or IP address range.	0.0.0.0/0
	IP address:	
	- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
Source Port Range	The source port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-). For example, 1-100 .	22, or 22-30

Parameter	Description	Example Value
Destination	The destination to which the traffic is allowed. The destination can be an IP address or IP address range.	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
Destination Port Range	The destination port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-). For example, 1-100 .	22, or 22-30
	You must specify this parameter if TCP or UDP is selected for Protocol .	
Description	Supplementary information about the network ACL rule. This parameter is optional.	N/A
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

6. Click **OK**.

4.3.4.2 Modifying a Network ACL Rule

Scenarios

Modify an inbound or outbound network ACL rule based on your network security requirements.

Procedure

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose **Access Control** > **Network ACLs**.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the **Inbound Rules** or **Outbound Rules** tab, locate the row that contains the target rule and click **Modify** in the **Operation** column. In the displayed

dialog box, configure parameters as prompted. **Table 4-29** lists the parameters to be configured.

Table 4-29 Parameter descriptions

Parameter	Description	Example Value
Туре	This parameter is available only after the IPv6 function is enabled. The network ACL type. This parameter is mandatory. You can select a value from	IPv4
	the drop-down list. Currently, only IPv4 and IPv6 are supported.	
Action	The action in the network ACL. This parameter is mandatory. You can select a value from the drop-down list. Currently, the value can be Allow or Deny .	Allow
Protocol	The protocol supported by the network ACL. This parameter is mandatory. You can select a protocol from the drop-down list. You can select TCP , UDP , ICMP , or All .	ТСР
Source	The source from which the traffic is allowed. The source can be an IP address or IP address range.	0.0.0.0/0
	IP address:	
	- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
Source Port Range	The source port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-). For example, 1-100 .	22, or 22-30

Parameter	Description	Example Value
Destination	The destination to which the traffic is allowed. The destination can be an IP address or IP address range.	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
Destination Port Range	The destination port number or port number range. The value ranges from 1 to 65535. For a port number range, enter two port numbers connected by a hyphen (-). For example, 1-100 .	22, or 22-30
	You must specify this parameter if TCP or UDP is selected for Protocol .	
Description	Supplementary information about the network ACL rule. This parameter is optional.	N/A
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

6. Click Confirm.

4.3.4.3 Changing the Sequence of a Network ACL Rule

Scenarios

If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule.

If multiple network ACL rules conflict, only the rule with the highest priority takes effect.

Procedure

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose **Access Control** > **Network ACLs**.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.

- 5. On the **Inbound Rules** or **Outbound Rules** tab, locate the target rule, click **More** in the **Operation** column, and select **Insert Rule Above** or **Insert Rule Below**.
- 6. In the displayed dialog box, configure required parameters and click **OK**. The rule is inserted. The procedure for inserting an outbound rule is the same as that for inserting an inbound rule.

4.3.4.4 Enabling or Disabling a Network ACL Rule

Scenarios

Enable or disable an inbound or outbound rule based on your network security requirements.

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.

- In the navigation pane on the left, choose Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the **Inbound Rules** or **Outbound Rules** tab, locate the row that contains the target rule, and click **More** and then **Enable** or **Disable** in the **Operation** column.
- 6. Click **Yes** in the displayed dialog box.

The rule is enabled or disabled. The procedure for enabling or disabling an outbound rule is the same as that for enabling or disabling an inbound rule.

4.3.4.5 Exporting and Importing Network ACL Rules

Scenarios

You can export inbound and outbound rules of a specific network ACL as an Excel file and then import these rules for another network ACL. Importing and exporting rules across regions are supported.

Notes and Constraints

- For optimal performance, import no more than 40 network ACL rules at a time.
- Importing rules will not delete existing rules.
- Importing duplicate rules will fail.

Exporting Network ACL Rules

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 Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. Click to export the inbound and outbound network ACL rules. The exported rules are stored in an Excel file. You need to download the file to a local directory.

Importing Network ACL Rules

- 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 Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. Click
- 6. Select the Excel file containing the exported network ACL rules and click **Upload** to import the rules.

4.3.4.6 Deleting a Network ACL Rule

Scenarios

Delete an inbound or outbound rule based on your network security requirements.

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 Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the **Inbound Rules** or **Outbound Rules** tab, locate the row that contains the target rule and click **Delete** in the **Operation** column.
- 6. Click **Yes** in the displayed dialog box.

4.3.5 Managing Subnets Associated with a Network ACL

4.3.5.1 Associating Subnets with a Network ACL

Scenarios

You can associate a network ACL with a subnet to protect resources in the subnet.

Notes and Constraints

- You can associate a network ACL with multiple subnets. However, a subnet can only be associated with one network ACL at a time.
- After a network ACL is associated with a subnet, the default network ACL rules deny all traffic to and from the subnet until you add custom rules to allow traffic. For details, see Adding a Network ACL Rule.

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 Access Control > Network ACLs.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the displayed page, click the **Associated Subnets** tab.
- 6. On the **Associated Subnets** tab, click **Associate**.
- 7. On the displayed page, select the subnets to be associated with the network ACL, and click **OK**.

MOTE

A subnet with a network ACL associated will not be displayed on the page for you to select. If you want to associate such a subnet with another network ACL, you must first disassociate the subnet from the original network ACL. One-click subnet association and disassociation are not supported currently. A subnet can only be associated with one network ACL.

4.3.5.2 Disassociating Subnets from a Network ACL

Scenarios

You can disassociate a subnet from its network ACL based on your network requirements.

Procedure

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.

- 3. In the navigation pane on the left, choose **Access Control** > **Network ACLs**.
- 4. Locate the target network ACL and click its name to switch to the page showing details of that particular network ACL.
- 5. On the displayed page, click the **Associated Subnets** tab.
- 6. On the **Associated Subnets** page, locate the row that contains the target subnet and click **Disassociate** in the **Operation** column.
- 7. Click **Yes** in the displayed dialog box.

5 Elastic IP

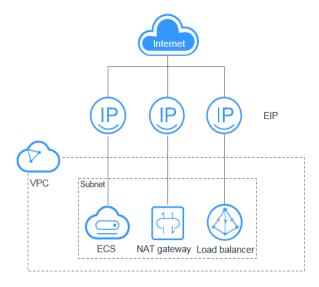
5.1 EIP Overview

EIP

The Elastic IP (EIP) service enables you to use static public IP addresses and scalable bandwidths to connect your cloud resources to the Internet. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers. Various billing modes are provided to meet diverse service requirements.

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

Figure 5-1 Accessing the Internet using an EIP



EIP Quotas

If you want to know the number of EIPs that can be assigned in a region, see **How Do I View My Quotas?**

Advantages

Flexibility

An EIP can be flexibly associated with or disassociated from the ECS, BMS, NAT gateway, load balancer, or virtual IP address. The bandwidth can be adjusted according to service changes.

Shared bandwidth

EIPs can use shared bandwidth to lower bandwidth costs.

Immediate use
 EIP binding, unbinding, and bandwidth adjustments take effect immediately.

5.2 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 an EIP

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. On the displayed page, click Assign EIP.
- 4. Set the parameters as prompted.

Table 5-1 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 you. The region selected for the EIP is its geographical location.	N/A

Parameter	Description	Example Value
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 bandwidth types 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 traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. Shared Bandwidth: The bandwidth can be shared by multiple EIPs and is suitable for scenarios with staggered traffic. 	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
EIP Name	The name of the EIP.	eip-test
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 5-2.	Key: lpv4_key1Value: 3005eip
Quantity	The number of EIPs you want to purchase.	1

Parameter Requirement **Example Value** Key Cannot be left blank. Ipv4 key1 • Must be unique for each EIP. • Can contain a maximum of 36 characters. • Can contain letters, digits, underscores (_), and hyphens (-). Value • Can contain a maximum of 43 3005eip characters. • Can contain letters, digits,

Table 5-2 EIP tag requirements

5. Click Create Now.

Binding an EIP

1. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.

underscores (_), periods (.), and

2. Select the instance that you want to bind the EIP to.

hyphens (-).

3. Click OK.

5.3 Unbinding an EIP from an ECS and Releasing the EIP

Scenarios

If you no longer need an EIP, unbind it from the ECS and release the EIP to avoid wasting network resources.

Notes and Constraints

- EIPs assigned and bound to load balancers in the ELB service are displayed in the EIP list of the VPC service, but you cannot unbind these EIPs from classic load balancers.
- Only EIPs with no instance bound can be released. If you want to release an EIP with an instance bound, you need to unbind EIP from the instance first.

Procedure

Unbinding a single EIP

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

4. Click **Yes** in the displayed dialog box.

Releasing a single EIP

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. On the displayed page, locate the row that contains the target EIP, click **More** and then **Release** in the **Operation** column.
- 4. Click **Yes** in the displayed dialog box.

Unbinding multiple EIPs at once

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. On the displayed page, select the EIPs to be unbound.
- 4. Click the **Unbind** button located above the EIP list.
- 5. Click **Yes** in the displayed dialog box.

Releasing multiple EIPs at once

- 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 released.
- 4. Click the **Release** button located above the EIP list.
- 5. Click **Yes** in the displayed dialog box.

5.4 Modifying an EIP Bandwidth

Scenarios

Modify the EIP bandwidth name, size, and billing option (by bandwidth or by traffic).

Procedure

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. Locate the row that contains the target EIP in the EIP list, click **More** in the **Operation** column, and select **Modify Bandwidth**.
- 4. Modify the bandwidth parameters as prompted.

You can modify the following parameters:

- Bandwidth Name
- Billing By: **Bandwidth** or **Traffic**
- Bandwidth (Mbit/s)
- 5. Click **Next**.
- 6. Click **Submit**.

You can also select multiple EIPs and click **Modify Bandwidth** above the list to modify their bandwidths in batches. Only dedicated bandwidths billed on a payper-use basis can be modified in batches.

5.5 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

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Network** > **Elastic IP**.
- 3. On the displayed page, click in the upper right corner of the EIP list. The system will automatically export all EIPs in the current region of your account to an Excel file and download the file to a local directory.

5.6 Managing EIP Tags

Scenarios

Tags can be added to EIPs to facilitate EIP identification and administration. 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 5-3** lists the tag key and value requirements.

Table 5-3 EIP tag requirements

Parameter	Requirement	Example Value
Key	Cannot be left blank.Must be unique for each EIP.	lpv4_key1
	Can contain a maximum of 36 characters.	
	 Can contain letters, digits, underscores (_), and hyphens (-). 	
Value	Can contain a maximum of 43 characters.	3005eip
	• Can contain letters, digits, underscores (_), periods (.), and hyphens (-).	

Procedure

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

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. Click the search box above the EIP list.
- 4. Select the tag key and value of the EIP.

You can add multiple tag keys and values to refine your search results. If you add more than one tag to search for EIPs, the system will display only the EIPs that contain all of the tags you specified.

5. Click OK.

The system displays the EIPs you are looking for based on the entered tag keys and values.

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

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. On the displayed page, locate the EIP whose tags you want to manage, and click the EIP name.
- 4. 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 taq.
 - 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**.

6 Shared Bandwidth

6.1 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, BMSs, and load balancers that have EIPs bound in the same region can share a bandwidth.

□ NOTE

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

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.

- 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 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.

6.2 Assigning a Shared Bandwidth

Scenarios

Assign a shared bandwidth for use with EIPs.

Procedure

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

Table 6-1 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 you.	N/A
Billed By	The billing method for the shared bandwidth.	Bandwidth
	You can specify a shared bandwidth to be billed by bandwidth or by traffic.	
Bandwidth	The bandwidth size in Mbit/s. The minimum value is 5 Mbit/s. The maximum bandwidth can be 2000 Mbit/s.	10
Name	The name of the shared bandwidth.	Bandwidth-001

5. Click Create Now.

6.3 Adding EIPs to a Shared Bandwidth

Scenarios

Add EIPs to a shared bandwidth and the EIPs can then share that bandwidth. You can add multiple EIPs to a shared bandwidth at the same time.

Notes and Constraints

- The type of EIPs must be the same as that of the shared bandwidth the EIPs to be added to.
- If it is a standard shared bandwidth, you can add dynamic BGP EIPs and IPv6 NICs to it. If it is a premium shared bandwidth, you can add premium BGP EIPs and IPv6 NICs to it.

Procedure

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Network** > **Elastic IP**.
- 3. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.
- 4. In the shared bandwidth list, locate the row that contains the shared bandwidth that you want to add EIPs to. In the **Operation** column, choose **Add EIP**, and select the EIPs to be added.

MOTE

- 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.
- 5. Click OK.

6.4 Removing EIPs from a Shared Bandwidth

Scenarios

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 choose **Network** > **Elastic IP**.
- In the navigation pane on the left, choose Elastic IP and Bandwidth > Shared Bandwidths.
- 4. In the shared bandwidth list, locate the row that contains the bandwidth from which EIPs are to be removed, choose **More** > **Remove EIP** in the **Operation** column, and select the EIPs to be removed in the displayed dialog box.
- 5. Click **OK**.

6.5 Modifying a Shared Bandwidth

Scenarios

This section describes how to change the name, billing mode, size, and billing option (by bandwidth or by traffic) of a shared bandwidth.

Procedure

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.

4. 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.

You can modify the following parameters:

- Bandwidth Name
- Billing By: **Bandwidth** or **Traffic**
- Bandwidth (Mbit/s)
- 5. Click **Next**.
- 6. Click **Submit**.

6.6 Deleting a Shared Bandwidth

Scenarios

Delete a shared bandwidth when it is no longer required.

Prerequisites

Before deleting a shared bandwidth, remove all the EIPs associated with it. For details, see **Removing EIPs from a Shared Bandwidth**.

Procedure

- 1. Log in to the management console.
- 2. Click \equiv in the upper left corner and choose **Network** > **Elastic IP**.
- 3. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.
- 4. 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**.
- 5. In the displayed dialog box, click **OK**.

7 Route Tables

7.1 Route Tables and Routes

Route Tables

A route table contains a set of routes that are used to determine where network traffic from your subnets in a VPC is directed. Each subnet must be associated with a route table. A subnet can only be associated with one route table, but you can associate multiple subnets with the same route table.

- Default route table: When you create a VPC, the system automatically generates a default route table for the VPC. If you create a subnet in the VPC, the subnet automatically associates with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.
 - You can add routes to, delete routes from, and modify routes in the default route table, but cannot delete the table.
 - When you create a VPC endpoint, VPN or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.
- Custom route table: If you do not want to use the default route table, you can create a custom route table and associate it with the subnet. Custom route tables can be deleted if they are no longer required.

The custom route table associated with a subnet affects only the outbound traffic. The default route table of a subnet controls the inbound traffic.

□ NOTE

Route

You can add routes to default and custom route tables and configure the destination, next hop type, and next hop in the routes to determine where network traffic is directed. Routes are classified into system routes and custom routes.

• System routes: These routes are automatically added by the system and cannot be modified or deleted.

After a route table is created, the system automatically adds the following system routes to the route table, so that instances in a VPC can communicate with each other.

- Routes whose destination is 100.64.0.0/10 or 198.19.128.0/20.
- Routes whose destination is a subnet CIDR block.

In addition to the preceding system routes, the system automatically adds a route whose destination is 127.0.0.0/8. This is the local loopback address.

• Custom routes: These are routes that you can add, modify, and delete. The destination of a custom route cannot overlap with that of a system route.

You can add a custom route and configure the destination, next hop type, and next hop in the route to determine where network traffic is directed. **Table 7-1** lists the supported types of next hops.

You cannot add two routes with the same destination to a VPC route table even if their next hop types are different. The route priority depends on the destination. According to the longest match routing rule, the destination with a higher matching degree is preferentially selected for packet forwarding.

Table 7-1 Next hop type

Next Hop Type	Description	Supported Route Table		
Server	Traffic intended for the destination is forwarded to an ECS in the VPC.	Default route tableCustom route table		
Extension NIC	Traffic intended for the destination is forwarded to the extension NIC of an ECS in the VPC.	Default route table		
VPN gateway	Traffic intended for the destination is forwarded to a VPN gateway.	Custom route table		
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.	Custom route table		
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	Default route tableCustom route table		
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection. • Default ro table • Custom ro table			

Next Hop Type	Description	Supported Route Table
Virtual IP address	Traffic intended for the destination is forwarded to a virtual IP address and then sent to active and standby ECSs to which the virtual IP address is bound.	Default route tableCustom route table

□ NOTE

Currently, the route with the next hop type Direct Connect gateway cannot be configured. To configure it, submit a service ticket.

If you specify the destination when creating a resource, a system route is delivered. If you do not specify a destination when creating a resource, a custom route that can be modified or deleted is delivered.

For example, when you create a NAT gateway, the system automatically delivers a custom route without a specific destination (0.0.0.0/0 is used by default). In this case, you can change the destination. However, when you create a VPN gateway, you need to specify the remote subnet, that is, the destination of a route. In this case, the system delivers this system route. Do not modify the route destination on the **Route Tables** page. If you do, the destination will be inconsistent with the configured remote subnet. To modify the route destination, go to the specific resource page and modify the remote subnet, then the route destination will be changed accordingly.

Custom Route Table Configuration Process

Figure 7-1 shows the process of creating and configuring a custom route table.

Create a route table.

Add a route.

Associate the route table with a subnet.

Figure 7-1 Route table configuration process

1. For details about how to create a custom route table, see **Creating a Custom Route Table**.

- 2. For details about how to add a custom route, see **Adding a Custom Route**.
- 3. For details about how to associate a subnet with a route table, see

 Associating a Route Table with a Subnet. After the association, the routes in the route table control the routing for the subnet.

7.2 Managing Route Tables

7.2.1 Creating a Custom Route Table

Scenarios

A VPC automatically comes with a default route table. If your default route table cannot meet your service requirements, you can create a custom route table.

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 > Route Tables**.
- 4. In the upper right corner, click **Create Route Table**. On the displayed page, configure parameters as prompted.

Table 7-2 Parameter descriptions

Parameter	Description	Example Value
Name	The name of the route table. This parameter is mandatory.	rtb-001
	The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	
VPC	The VPC that the route table belongs to. This parameter is mandatory.	vpc-001
Description	Supplementary information about the route table. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Parameter	Description	Example Value
Route Settings	The route information. This parameter is optional.	-
	You can add a route when creating the route table or after the route table is created. For details, see Adding a Custom Route. You can click + to add more routes.	

5. Click **OK**.

A message is displayed. You can determine whether to associate the route table with subnets immediately as prompted. If you want to associate immediately, perform the following operations:

- a. Click **Associate Subnet**. The route table details page is displayed.
- b. Click **Associate Subnet** and select the target subnets to be associated.
- c. Click **OK**.

7.2.2 Associating a Route Table with a Subnet

Scenarios

After a subnet is created, the system associates the subnet with the default route table of its VPC. If you want to use specific routes for a subnet, you can associate the subnet with a custom route table.

The custom route table associated with a subnet affects only the outbound traffic. The default route table determines the inbound traffic.

NOTICE

After a route table is associated with a subnet, the routes in the route table control the routing for the subnet and apply to all cloud resources in the subnet.

Notes and Constraints

- A subnet must have a route table associated and can only be associated with one route table.
- A route table can be associated with multiple subnets.

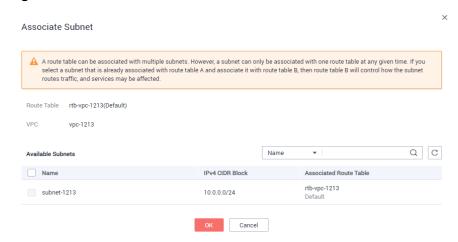
Procedure

- 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.

- In the navigation pane on the left, choose Virtual Private Cloud > Route Tables.
- 4. In the route table list, locate the row that contains the target route table and click **Associate Subnet** in the **Operation** column.
- 5. Select the subnet to be associated.

Figure 7-2 Associate Subnet



6. Click OK.

7.2.3 Changing the Route Table Associated with a Subnet

Scenarios

You can change the route table for a subnet. If the route table for a subnet is changed, routes in the new route table will apply to all cloud resources in the subnet.

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Route Tables.
- 4. Click the name of the target route table.
- 5. On the **Associated Subnets** tab page, click **Change Route Table** in the **Operation** column and select a new route table as prompted.
- 6. Click OK.

After the route table for a subnet is changed, routes in the new route table will apply to all cloud resources in the subnet.

7.2.4 Viewing the Route Table Associated with a Subnet

Scenarios

This section describes how to view the route table associated with a subnet.

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- Locate the target subnet and click its name.
 The subnet details page is displayed.

Figure 7-3 Viewing the route table associated with a subnet



- 5. In the right of the subnet details page, view the route table associated with the subnet.
- 6. Click the name of the route table.

The route table details page is displayed. You can further view the route information.

7.2.5 Viewing Route Table Information

Scenarios

This section describes how to view detailed information about a route table, including:

- Basic information, such as name, type (default or custom), and ID of the route table
- Routes, such as destination, next hop, and route type (system or custom)
- Associated subnets

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 > Route Tables**.
- 4. Click the name of the target route table.

The route table details page is displayed.

- a. On the **Summary** tab page, view the basic information and routes of the route table.
- b. On the **Associated Subnets** tab page, view the subnets associated with the route table.

7.2.6 Exporting Route Table Information

Scenarios

Information about all route tables under your account can be exported as an Excel file to a local directory. This file records the name, ID, VPC, type, and number of associated subnets of the route tables.

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 > Route Tables**.
- 4. On the displayed page, click in the upper right of the route table list.

 The system will automatically export information about all route tables under your account in the current region as an Excel file to a local directory.

7.2.7 Deleting a Route Table

Scenarios

This section describes how to delete a custom route table.

Notes and Constraints

- The default route table cannot be deleted.
- A custom route table with a subnet associated cannot be deleted directly.

If you want to delete such a route table, you can associate the subnet with another route table first by referring to **Changing the Route Table Associated with a Subnet**.

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 > Route Tables**.
- 4. Locate the row that contains the route table you want to delete and click **Delete** in the **Operation** column.
 - A confirmation dialog box is displayed.
- 5. Click Yes.

7.3 Managing Routes

7.3.1 Adding a Custom Route

Scenarios

Each route table contains a default system route, which indicates that ECSs in a VPC can communicate with each other. You can also add custom routes as required to forward the traffic destined for the destination to the specified next hop.

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 > Route Tables**.
- 4. In the route table list, click the name of the route table to which you want to add a route.
- 5. Click **Add Route** and set parameters as prompted.

You can click + to add more routes.

Parameter Description **Example Value** Destination Mandatory IPv4: 192.168.0.0/16 Enter the destination of the route. You can enter a single IP address or an IP address range in CIDR notation. The destination of each route in a route table must be unique. The destination cannot overlap with any subnet in the VPC. Next Hop VPC peering Mandatory connection Type Set the type of the next hop. When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to VPN gateway or Direct Connect gateway. Next Hop Mandatory peer-AB Set the next hop. The resources in the drop-down list box are displayed based

Table 7-3 Parameter descriptions

6. Click OK.

Description

7.3.2 Modifying a Route

Scenarios

This section describes how to modify a custom route in a route table.

on the selected next hop type.

text box as required.

Enter the description of the route in the

Optional

Notes and Constraints

- System routes cannot be modified.
- When you create a VPC endpoint, VPN or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.

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 > Route**Tables
- 4. In the route table list, click the name of the target route table.
- 5. Locate the row that contains the route to be modified and click **Modify** in the **Operation** column.
- 6. Modify the route information in the displayed dialog box.

Table 7-4 Parameter descriptions

Parameter	Description	Example Value
Destination	Mandatory Enter the destination of the route. You can enter a single IP address or an IP address range in CIDR notation. The destination of each route in a route table must be unique. The destination cannot overlap with any subnet in the VPC.	IPv4: 192.168.0.0/16
Next Hop Type	Mandatory Set the type of the next hop. NOTE When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to VPN gateway or Direct Connect gateway.	VPC peering connection
Next Hop	Mandatory Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.	peer-AB
Description	Optional Enter the description of the route in the text box as required.	-

7. Click **OK**.

7.3.3 Replicating a Route

Scenarios

This section describes how to replicate routes among all route tables of a VPC. VPC route tables include the default and custom route tables.

Notes and Constraints

Table 7-5 shows whether routes of different types can be replicated to default or custom route tables.

For example, if the next hop type of a route is a server, this route can be replicated to both default or custom route tables. If the next hop type of a route is a Direct Connect gateway, the route cannot be replicated to the default route table, but can be replicated to a custom route table.

Table 7-5 Route replication

Next Hop Type	Can Be Replicated to Default Route Table	Can Be Replicated to Custom Route Table
Local	No	No
Server	Yes	Yes
Extension NIC	Yes	Yes
VPN gateway	No	Yes
Direct Connect gateway	No	Yes
NAT gateway	Yes	Yes
VPC peering connection	Yes	Yes
Virtual IP address	Yes	Yes

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 > Route Tables**.
- 4. In the route table list, locate the row that contains the route table you want to replicate routes from and click **Replicate Route** in the **Operation** column.
- 5. Select the target route table that you want to replicate route to and the routes to be replicated as prompted.
 - The listed routes are those that do not exist in the target route table. You can select one or more routes to replicate to the target route table.
- 6. Click OK.

7.3.4 Deleting a Route

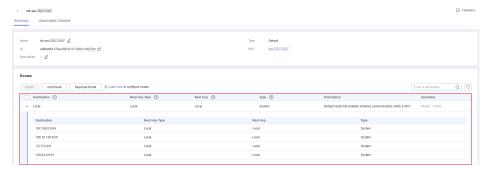
Scenarios

This section describes how to delete a custom route from a route table.

Notes and Constraints

• System routes cannot be deleted.

Figure 7-4 System routes



- The routes automatically delivered by VPN or Direct Connect to the default route table cannot be deleted. The next hop types of such routes are:
 - VPN gateway
 - Direct Connect gateway

The following figure shows a route with **VPN gateway** as **Next Hop Type**. If you want to delete such a route, click the next hop hyperlink to delete the corresponding resource.

Figure 7-5 Route delivered by VPN



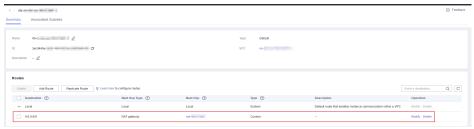
Procedure

- 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 > Route Tables**.
- 4. Locate the target route table and click its name. The route table details page is displayed.

Figure 7-6 Deleting a custom route



5. In the route list, locate the row that contains the route to be deleted and click **Delete** in the **Operation** column.

A confirmation dialog box is displayed.

6. Confirm the information and click Yes.

7.4 Configuring an SNAT Server

Scenarios

Together with VPC route tables, you can configure SNAT on an ECS to enable other ECSs that have no EIPs bound in the same VPC to access the Internet through this ECS.

The configured SNAT takes effect for all subnets in a VPC.

Prerequisites

- You have an ECS where SNAT is to be configured.
- The ECS where SNAT is to be configured runs Linux.
- The ECS where SNAT is to be configured has only one network interface card (NIC).

Procedure

- 1. Log in to the management console.
- 2. In the upper left corner of the page, click =. In the service list, choose **Computing > Elastic Cloud Server**.
- 3. On the displayed page, locate the target ECS in the ECS list and click the ECS name to switch to the page showing ECS details.
- 4. On the displayed ECS details page, click the **NICs** tab.
- 5. In the displayed area showing the NIC IP address details, disable **Source/ Destination Check**.

By default, the source/destination check is enabled. When this check is enabled, the system checks whether source IP addresses contained in the packets sent by ECSs are correct. If the IP addresses are incorrect, the system does not allow the ECSs to send the packets. This mechanism prevents packet spoofing, thereby improving system security. If the SNAT function is used, the SNAT server needs to forward packets. This mechanism prevents the packet sender from receiving returned packets. Therefore, you need to disable the source/destination check for SNAT servers.

- 6. Bind an EIP.
 - Bind an EIP to the private IP address of the ECS. For details, see
 Assigning an EIP and Binding It to an ECS.
 - Bind an EIP to the virtual IP address of the ECS. For details, see Binding a
 Virtual IP Address to an EIP or ECS.
- 7. On the ECS console, use the remote login function to log in to the ECS where you plan to configure SNAT.
- 8. Run the following command and enter the password of user **root** to switch to user **root**:

su - root

9. Run the following command to check whether the ECS can successfully connect to the Internet:

Before running the command, you must disable the response iptables rule on the ECS where SNAT is configured and configure security group rules.

ping www.google.com

The ECS can access the Internet if the following information is displayed:

[root@localhost ~]# ping www.google.com PING www.google.com (xxx.xxx.xxx.xxx) 56(84) bytes of data. 64 bytes from xxx.xxx.xxx.icmp_seq=1 ttl=51 time=9.34 ms 64 bytes from xxx.xxx.xxx.xxxx: icmp_seq=2 ttl=51 time=9.11 ms 64 bytes from xxx.xxx.xxxx.xxxx: icmp_seq=3 ttl=51 time=8.99 ms

10. Run the following command to check whether IP forwarding of the Linux OS is enabled:

cat /proc/sys/net/ipv4/ip_forward

In the command output, **1** indicates that IP forwarding is enabled, and **0** indicates that IP forwarding is disabled. The default value is **0**.

- If IP forwarding in Linux is enabled, go to step 13.
- If IP forwarding in Linux is disabled, go to 11 to enable IP forwarding in Linux.

Many OSs support packet routing. Before forwarding packets, OSs change source IP addresses in the packets to OS IP addresses. Therefore, the forwarded packets contain the IP address of the public sender so that the response packets can be sent back along the same path to the initial packet sender. This method is called SNAT. The OSs need to keep track of the packets where IP addresses have been changed to ensure that the destination IP addresses in the packets can be rewritten and that packets can be forwarded to the initial packet sender. To achieve these purposes, you need to enable the IP forwarding function and configure SNAT rules.

- 11. Use the vi editor to open the /etc/sysctl.conf file, change the value of net.ipv4.ip forward to 1, and enter :wq to save the change and exit.
- 12. Run the following command to make the change take effect:

sysctl -p /etc/sysctl.conf

13. Configure the SNAT function.

Run the following command to enable all ECSs on the network (for example, 192.168.1.0/24) to access the Internet using the SNAT function:

iptables -t nat -A POSTROUTING -o eth0 -s subnet -j SNAT --to nat-instance-ip

Figure 7-7 Configuring SNAT

```
| Troot@host-192-168-1-4 ~ | # vi /etc/sysctl.conf^C
| Troot@host-192-168-1-4 ~ | # ^C
| Troot@host-192-168-1-4 ~ | # iptables -t nat -A POSTROUTING -o eth0 -s 192.168.1.0
| V24 - j SNAT --to 192.168.1.4
```

MOTE

To ensure that the rule will not be lost after the restart, write the rule into the /etc/rc.local file.

- 1. Switch to the /etc/sysctl.conf file:
 - vi /etc/rc.local
- 2. Perform 13 to configure SNAT.
- 3. Save the configuration and exit:
 - :wa
- 4. Add the execution permissions for the rc.local file:
 - # chmod +x /etc/rc.local
- Check whether the configuration is successful. If information similar to Figure 7-8 (for example, 192.168.1.0/24) is displayed, the configuration was successful.

iptables -t nat --list

Figure 7-8 Verifying configuration

15. Add a route. For details, see section Adding a Custom Route.

Set the destination to **0.0.0.0/0**, and the next hop to the private or virtual IP address of the ECS where SNAT is deployed. For example, the next hop is **192.168.1.4**.

After these operations are complete, if the network communication still fails, check your security group and network ACL configuration to see whether required traffic is allowed.

8 VPC Peering Connection

8.1 VPC Peering Connection Overview

What Is a VPC Peering Connection?

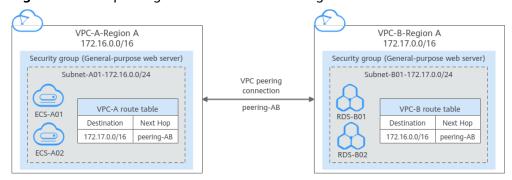
A VPC peering connection is a networking connection that connects two VPCs for them to communicate using private IP addresses. The VPCs to be peered can be in the same account or different accounts, but must be in the same region.

You can use VPC peering connections to build networks in different scenarios.
 For details, see VPC Peering Connection Usage Examples.

Figure 8-1 shows an application scenario of VPC peering connections.

- There are two VPCs (VPC-A and VPC-B) in region A that are not connected.
- Service servers (ECS-A01 and ECS-A02) are in VPC-A, and database servers (RDS-B01 and RDS-B02) are in VPC-B. The service servers and database servers cannot communicate with each other.
- You need to create a VPC peering connection (peering-AB) between VPC-A and VPC-B so the service servers and database servers can communicate with each other.

Figure 8-1 VPC peering connection network diagram



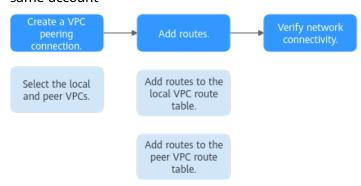
VPC Peering Connection Creation Process

A VPC peering connection can only connect VPCs in the same region.

• If two VPCs are in the same account, the process of creating a VPC peering connection is shown in **Figure 8-2**.

For details about how to create a VPC peering connection, see **Creating a VPC Peering Connection with Another VPC in Your Account**.

Figure 8-2 Process of creating a VPC peering connection between VPCs in the same account

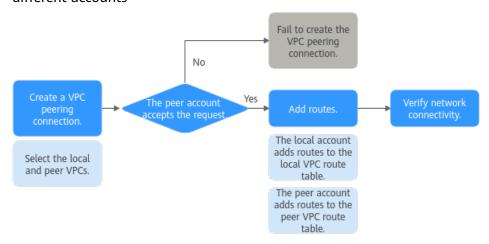


• If two VPCs are in different accounts, the process of creating a VPC peering connection is shown in **Figure 8-3**.

For details about how to create a VPC peering connection, see **Creating a VPC Peering Connection with a VPC in Another Account**.

If account A initiates a request to create a VPC peering connection with a VPC in account B, the VPC peering connection takes effect only after account B accepts the request.

Figure 8-3 Process of creating a VPC peering connection between VPCs in different accounts



Notes and Constraints

- A VPC peering connection can only connect VPCs in the same region.
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect.

8.2 VPC Peering Connection Usage Examples

A VPC peering connection is a networking connection between two VPCs in the same region and enables them to communicate. **Table 8-1** lists different scenarios of using VPC peering connections.

Table 8-1 VPC peering connection usage examples

Locati on	CIDR Block	Description	Usage Example
VPCs in the same region	 VPC CIDR blocks do not overlap. Subnet CIDR blocks of VPCs do not overlap. 	You can create VPC peering connections to connect entire CIDR blocks of VPCs. Then, all resources in the VPCs can communicate with each other.	 Peering Two or More VPCs Peering One Central VPC with Multiple VPCs
VPCs in the same region	 VPC CIDR blocks overlap. Some subnet CIDR blocks overlap. 	You can create VPC peering connections to connect specific subnets or ECSs from different VPCs. • To connect specific subnets from two VPCs, the subnet CIDR blocks cannot overlap. • To connect specific ECSs from two VPCs, each ECS must have a unique private IP address.	Peering Two VPCs with Overlapping CIDR Blocks Peering ECSs in a Central VPC with ECSs in Two Other VPCs
VPCs in the same region	 VPC CIDR blocks overlap. All subnet CIDR blocks overlap. 	VPC peering connections are not usable.	• Invalid VPC Peering Connections

Peering Two or More VPCs

 Two VPCs peered together: Figure 8-4 shows the networking diagram of a VPC peering connection that connects VPC-A and VPC-B.

VPC-A VPC-B 172.16.0.0/16 10.0.0.0/16 Subnet-A01 Subnet-B01 172.16.0.0/24 10.0.0.0/24 VPC peering connection ECS-A01 ECS-B01 Subnet-A02 Subnet-B02 172.16.1.0/24 10.0.1.0/24 ECS-A02 ECS-B02

Figure 8-4 Networking diagram (IPv4)

Table 8-2 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B

Table 8-3 VPC route tables (IPv4)

Rout e Tabl e	Destina tion	Next Hop	Rout e Type	Description
rtb- VPC- A	10.0.0.0/ 16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC- B	172.16.0 .0/16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

 Multiple VPCs peered together: Figure 8-5 shows the networking diagram of VPC peering connections that connect VPC-A, VPC-B, and VPC-C.

VPC-B VPC-C 10.0.0.0/16 192.168.0.0/16 Subnet-B02 Subnet-B01 Subnet-C01 Subnet-C02 192.168.1.0/24 10.0.0.0/24 10.0.1.0/24 192.168.0.0/24 ECS-B01 ECS-B02 ECS-C01 ECS-C02 VPC peering connections VPC-A 172.16.0.0/16 Subnet-A01 Subnet-A02 172.16.0.0/24 172.16.1.0/24 • ECS-A01 ECS-A02

Figure 8-5 Networking diagram (IPv4)

Table 8-4 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-B is peered with VPC-C.	Peering-BC	VPC-B	VPC-C

Table 8-5 VPC route tables (IPv4)

Rout e Tabl e	Destinat ion	Next Hop	Rout e Type	Description
rtb- VPC- A	10.0.0.0/ 16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
	192.168. 0.0/16	Peering -AC	Custo m	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.
rtb- VPC- B	172.16.0. 0/16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
	192.168. 0.0/16	Peering -BC	Custo m	Add a route with the CIDR block of VPC-C as the destination and Peering-BC as the next hop.
rtb- VPC- C	172.16.0. 0/16	Peering -AC	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.
	10.0.0.0/ 16	Peering -BC	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop.

Peering One Central VPC with Multiple VPCs

Figure 8-6 shows the networking diagram of VPC peering connections that connect VPC-B, VPC-C, VPC-D, VPC-E, VPC-F, VPC-G, and central VPC-A.

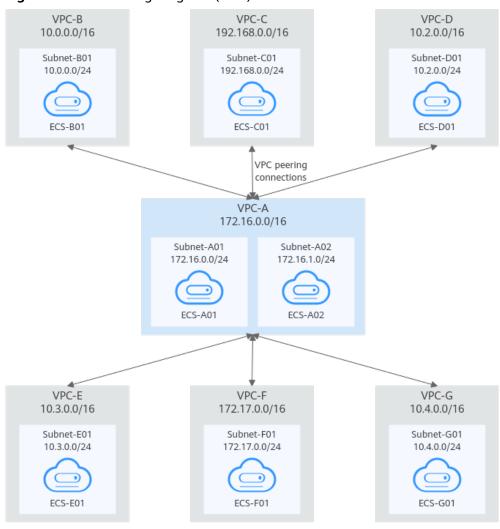


Figure 8-6 Networking diagram (IPv4)

Table 8-6 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-A is peered with VPC-D.	Peering-AD	VPC-A	VPC-D
VPC-A is peered with VPC-E.	Peering-AE	VPC-A	VPC-E
VPC-A is peered with VPC-F.	Peering-AF	VPC-A	VPC-F

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-G.	Peering-AG	VPC-A	VPC-G

Table 8-7 VPC route table details (IPv4)

Rout e Table	Destinati on	Next Hop	Route Type	Description
rtb- VPC- A	10.0.0.0/ 16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
	192.168.0 .0/16	Peering -AC	Custo m	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.
	10.2.0.0/ 16	Peering -AD	Custo m	Add a route with the CIDR block of VPC-D as the destination and Peering-AD as the next hop.
	10.3.0.0/ 16	Peering -AE	Custo m	Add a route with the CIDR block of VPC-E as the destination and Peering-AE as the next hop.
	172.17.0. 0/16	Peering -AF	Custo m	Add a route with the CIDR block of VPC-F as the destination and Peering-AF as the next hop.
	10.4.0.0/ 16	Peering -AG	Custo m	Add a route with the CIDR block of VPC-G as the destination and Peering-AG as the next hop.
rtb- VPC- B	172.16.0. 0/16	Peering -AB	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
rtb- VPC- C	172.16.0. 0/16	Peering -AC	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.
rtb- VPC- D	172.16.0. 0/16	Peering -AD	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop.
rtb- VPC-E	172.16.0. 0/16	Peering -AE	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.

Rout e Table	Destinati on	Next Hop	Route Type	Description
rtb- VPC-F	172.16.0. 0/16	Peering -AF	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.
rtb- VPC- G	172.16.0. 0/16	Peering -AG	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.

Peering Two VPCs with Overlapping CIDR Blocks

As shown in **Figure 8-7**, VPC-A and VPC-B have overlapping CIDR blocks, and their Subnet-A01 and Subnet-B01 also have overlapping CIDR blocks. In this case, a VPC peering connection can connect their Subnet-A02 and Subnet-B02 that do not overlap with each other.

VPC-B VPC-A 10.0.0.0/16 10.0.0.0/16 Subnet-A01 Subnet-B01 10.0.0.0/24 10.0.0.0/24 ECS-A01 ECS-B01 Subnet-A02 Subnet-B02 VPC 10.0.1.0/24 10.0.2.0/24 peering connection ECS-A02 ECS-B02

Figure 8-7 Networking diagram (IPv4)

Table 8-8 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B

Rout e Table	Destinat ion	Next Hop	Rout e Type	Description
rtb- VPC- A	10.0.2.0/ 24	Peering- AB	Custo m	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
rtb- VPC- B	10.0.1.0/ 24	Peering- AB	Custo m	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AB as the next hop.

Table 8-9 VPC route table details (IPv4)

Peering ECSs in a Central VPC with ECSs in Two Other VPCs

As shown in **Figure 8-8**, VPC-B and VPC-C have overlapping CIDR blocks, and their Subnet-B01 and Subnet-C01 have overlapping CIDR blocks. You can only create a VPC peering connection between ECSs.

- Use VPC peering connection Peering-AB to connect ECSs in Subnet-B01 and Subnet-A01
- Use VPC peering connection Peering-AC to connect ECSs in Subnet-C01 and Subnet-A01.

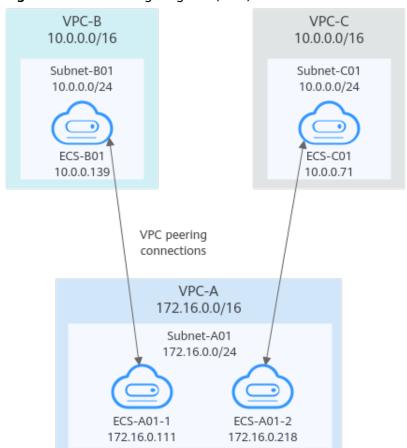


Figure 8-8 Networking diagram (IPv4)

Table 8-10 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
ECS-A01-1 in VPC-A is peered with ECS-B01 in VPC-B.	Peering-AB	VPC-A	VPC-B
ECS-A01-2 in VPC-A is peered with ECS-C01 in VPC-C.	Peering-AC	VPC-A	VPC-C

Table 8-11 VPC route table details (IPv4)

Rout e Table	Destinat ion	Next Hop	Route Type	Description
rtb- VPC- A	10.0.0.13 9/32	Peering -AB	Custo m	Add a route with the private IP address of ECS-B01 as the destination and Peering-AB as the next hop.
	10.0.0.71 /32	Peering -AC	Custo m	Add a route with the private IP address of ECS-C01 as the destination and Peering-AC as the next hop.
rtb- VPC- B	172.16.0. 111/32	Peering -AB	Custo m	Add a route with the private IP address of ECS-A01-1 as the destination and Peering-AB as the next hop.
rtb- VPC- C	172.16.0. 218/32	Peering -AC	Custo m	Add a route with the private IP address of ECS-A01-2 as the destination and Peering-AC as the next hop.

Invalid VPC Peering Connections

If VPCs with the same CIDR block also include subnets that overlap, VPC peering connections are not usable. VPC-A and VPC-B have the same CIDR block and their subnets have the same CIDR block. If a VPC peering connection is created between VPC-A and VPC-B, traffic cannot be routed between them because there are routes with the same destination.

In the rtb-VPC-A route table, the custom route for routing traffic from VPC-A to VPC-B and the local route have overlapping destinations. The local route has a higher priority and traffic will be forwarded within VPC-A and cannot reach VPC-B.

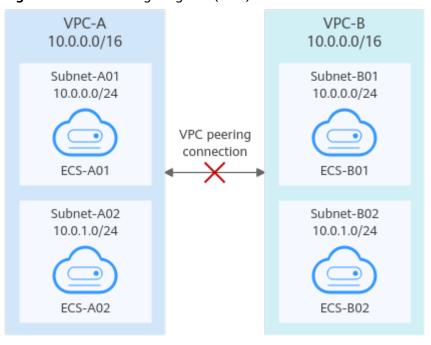


Figure 8-9 Networking diagram (IPv4)

Table 8-12 VPC route table details

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	10.0.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC-	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	10.0.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

8.3 Creating a VPC Peering Connection with Another VPC in Your Account

Scenarios

If two VPCs from the same region cannot communicate with each other, you can use a VPC peering connection. This section describes how to create a VPC peering connection between two VPCs in the same account.

This following describes how to create a VPC peering connection between VPC-A and VPC-B in account A to enable communications between ECS-A01 and RDS-B01.

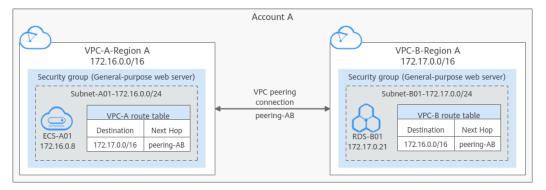
Procedure:

Step 1: Create a VPC Peering Connection

Step 2: Add Routes for the VPC Peering Connection

Step 3: Verify Network Connectivity

Figure 8-10 Networking diagram of a VPC peering connection between VPCs in the same account



Notes and Constraints

- Only one VPC peering connection can be created between two VPCs at the same time.
- A VPC peering connection can only connect VPCs in the same region.
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect.

Prerequisites

You have two VPCs from the same account in the same region. If you want to create one, see **Creating a VPC**.

Step 1: Create a VPC Peering Connection

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** > **VPC Peering Connections**.

The VPC peering connection list is displayed.

- 4. In the upper right corner of the page, click **Create VPC Peering Connection**. The **Create VPC Peering Connection** dialog box is displayed.
- 5. Configure the parameters as prompted. For details, see **Table 8-13**.

Table 8-13 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
VPC Peering Connection Name	Mandatory Enter a name for the VPC peering connection. The name can contain a maximum of 64 characters, including letters, digits, hyphens (-), and underscores (_).	peering-AB
Local VPC	Mandatory VPC at one end of the VPC peering connection. You can select one from the drop-down list.	VPC-A
Local VPC CIDR Block	CIDR block of the selected local VPC	172.16.0.0/16
Account	 Mandatory Options: My account and Another account Select My account. 	My account
Peer Project	The system fills in the corresponding project by default because My account is set to Account .	ab-cdef-1
	For example, if VPC-A and VPC-B are in account A and region A, the system fills in the correspond project of account A in region A by default.	

Parameter	Description	Example Value
Peer VPC	This parameter is mandatory if Account is set to My account .	VPC-B
	VPC at the other end of the VPC peering connection. You can select one from the dropdown list.	
Peer VPC CIDR Block	CIDR block of the selected peer VPC	172.17.0.0/16
	If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. For details, see VPC Peering Connection Usage Examples.	
Description	Optional Enter the description of the VPC peering connection in the text box as required.	peering-AB connects VPC-A and VPC-B.

6. Click **OK**.

A dialog box for adding routes is displayed.

- 7. Click Add Route or Add Later.
 - a. If you click **Add Route**, the **Local Routes** page is displayed. Then, go to **Step 2: Add Routes for the VPC Peering Connection**.
 - b. If you click **Add Later**, the VPC peering connection list is displayed.

Step 2: Add Routes for the VPC Peering Connection

To enable communications between VPCs connected by a VPC peering connection, you need to add forward and return routes to the route tables of the VPCs. For details, see VPC Peering Connection Usage Examples.

- 1. Add routes to the route table of the local VPC:
 - On the Local Routes tab of the VPC peering connection, click the Route Tables hyperlink.

The **Summary** tab of the default route table for the local VPC is displayed.

b. Click Add Route.

Table 8-14 describes the route parameters.

Table 8-14 Parameter description

Parameter	Description	Example Value
Destination	The peer VPC CIDR block, subnet CIDR block, or ECS IP address. For details, see VPC Peering Connection Usage Examples.	VPC-B CIDR block: 172.17.0.0/16
Next Hop Type	The next hop type. Select VPC peering connection .	VPC peering connection
Next Hop	The next hop address. Select the name of the current VPC peering connection.	peering-AB
Description	Supplementary information about the route. This parameter is optional.	-
	The route description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

c. Click **OK**.

You can view the route in the route list.

- 2. Add routes to the route table of the peer VPC:
 - a. On the **Peer Routes** tab of the VPC peering connection, click the **Route Tables** hyperlink.

The **Summary** tab of the default route table for the peer VPC is displayed.

b. Click Add Route.

Table 8-15 describes the route parameters.

Table 8-15 Parameter description

Parameter	Description	Example Value
Destination	The local VPC CIDR block, subnet CIDR block, or ECS IP address. For details, see VPC Peering Connection Usage Examples.	VPC-A CIDR block: 172.16.0.0/16
Next Hop Type	The next hop type. Select VPC peering connection.	VPC peering connection
Next Hop	The next hop address. Select the name of the current VPC peering connection.	peering-AB

Parameter	Description	Example Value
Description	Supplementary information about the route. This parameter is optional.	-
	The route description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

c. Click **OK**.

You can view the route in the route list.

Step 3: Verify Network Connectivity

After you add routes for the VPC peering connection, verify the communication between the local and peer VPCs.

- 1. Log in to ECS-A01 in the local VPC.
- 2. Check whether ECS-A01 can communicate with RDS-B01.

ping IP address of RDS-B01

Example command:

ping 172.17.0.21

If information similar to the following is displayed, ECS-A01 and RDS-B01 can communicate with each other, and the VPC peering connection between VPC-A and VPC-B is successfully created.

```
[root@ecs-A02 ~]# ping 172.17.0.21
PING 172.17.0.21 (172.17.0.21) 56(84) bytes of data.
64 bytes from 172.17.0.21: icmp_seq=1 ttl=64 time=0.849 ms
64 bytes from 172.17.0.21: icmp_seq=2 ttl=64 time=0.455 ms
64 bytes from 172.17.0.21: icmp_seq=3 ttl=64 time=0.385 ms
64 bytes from 172.17.0.21: icmp_seq=4 ttl=64 time=0.372 ms
...
--- 172.17.0.21 ping statistics ---
```

NOTICE

- In this example, ECS-A01 and RDS-B01 are in the same security group. If
 the instances in different security groups, you need to add inbound rules to
 allow access from the peer security group. For details, see Enabling ECSs
 In Different Security Groups to Communicate Through an Internal
 Network.
- If VPCs connected by a VPC peering connection cannot communicate with each other, refer to Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?.

8.4 Creating a VPC Peering Connection with a VPC in Another Account

Scenarios

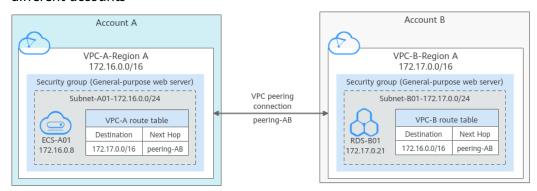
If two VPCs from the same region cannot communicate with each other, you can use a VPC peering connection. This section describes how to create a VPC peering connection between two VPCs in different accounts.

This following describes how to create a VPC peering connection between VPC-A in account A and VPC-B in account B to enable communications between ECS-A01 and RDS-B01.

Procedure:

- **Step 1: Create a VPC Peering Connection**
- Step 2: Peer Account Accepts the VPC Peering Connection Request
- Step 3: Add Routes for the VPC Peering Connection
- **Step 4: Verify Network Connectivity**

Figure 8-11 Networking diagram of a VPC peering connection between VPCs in different accounts



Notes and Constraints

- Only one VPC peering connection can be created between two VPCs at the same time.
- A VPC peering connection can only connect VPCs in the same region.
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect.
- For a VPC peering connection between VPCs in different accounts:
 - If account A initiates a request to create a VPC peering connection with a VPC in account B, the VPC peering connection takes effect only after account B accepts the request.
 - To ensure network security, do not accept VPC peering connections from unknown accounts.

Prerequisites

You have two VPCs in the same region, but they are from different accounts. If you want to create one, see **Creating a VPC**.

Step 1: Create a VPC Peering Connection

- 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 > VPC Peering Connections**.

The VPC peering connection list is displayed.

- 4. In the upper right corner of the page, click **Create VPC Peering Connection**. The **Create VPC Peering Connection** dialog box is displayed.
- 5. Configure the parameters as prompted. For details, see **Table 8-16**.

Table 8-16 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
VPC Peering Connection Name	Mandatory Enter a name for the VPC peering connection.	peering-AB
	The name can contain a maximum of 64 characters, including letters, digits, hyphens (-), and underscores (_).	
Local VPC	Mandatory VPC at one end of the VPC peering connection. You can select one from the drop-down list.	VPC-A
Local VPC CIDR Block	CIDR block of the selected local VPC	172.16.0.0/16
Account	 Mandatory Options: My account and Another account Select Another account. 	Another account

Parameter	Description	Example Value	
Peer Project ID	This parameter is mandatory because Account is set to Another account .	Project ID of VPC-B in region A: 067cf8aecf3XXX08322f	
	The project ID of the region that the peer VPC resides. For details about how to obtain the project ID, see Obtaining the Peer Project ID of a VPC Peering Connection.	13b	
Peer VPC ID	This parameter is mandatory because Account is set to Another account . ID of the VPC at the other end of the VPC peering connection. For details about how to obtain the ID, see Obtaining a VPC ID .	VPC-B ID: 17cd7278- XXX-530c952dcf35	
Description	Optional Enter the description of the VPC peering connection in the text box as required. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	peering-AB connects VPC-A and VPC-B.	

6. Click **OK**.

- If the message "Invalid VPC ID and project ID." is displayed, check whether the project ID and VPC ID are correct.
 - Peer Project ID: The value must be the project ID of the region where the peer VPC resides.
 - The local and peer VPCs must be in the same region.
- If the status of the created VPC peering connection is Awaiting acceptance, go to Step 2: Peer Account Accepts the VPC Peering Connection Request.

Step 2: Peer Account Accepts the VPC Peering Connection Request

After you create a VPC peering connection with a VPC in another account, you need to contact the peer account to accept the VPC peering connection request. In this example, account A notifies account B to accept the request. Account B needs to:

- 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 > VPC Peering Connections.

The VPC peering connection list is displayed.

- 4. In the VPC peering connection list, locate the VPC peering connection request to be accepted.
- 5. Locate the row that contains the target VPC peering connection and click **Accept Request** in the **Operation** column.
 - After the status of the VPC peering connection changes to **Accepted**, the VPC peering connection is created.
- 6. Go to Step 3: Add Routes for the VPC Peering Connection.

Step 3: Add Routes for the VPC Peering Connection

To enable communications between VPCs connected by a VPC peering connection, you need to add forward and return routes to the route tables of the VPCs. For details, see VPC Peering Connection Usage Examples.

Both accounts need to add a route to the route table of their VPC. In this example, account A adds a route to the route table of VPC-A, and account B adds a route to the route table of VPC-B.

- 1. Add routes to the route table of the local VPC:
 - a. In the VPC peering connection list of the local account, click the name of the target VPC peering connection.
 - The **Basic Information** tab of the VPC peering connection is displayed.
 - On the Local Routes tab of the VPC peering connection, click the Route Tables hyperlink.
 - The **Summary** tab of the default route table for the local VPC is displayed.
 - c. Click Add Route.

Table 8-17 describes the route parameters.

Table 8-17 Parameter description

Parameter	Description	Example Value
Destination	The peer VPC CIDR block, subnet CIDR block, or ECS IP address. For details, see VPC Peering Connection Usage Examples.	VPC-B CIDR block: 172.17.0.0/16
Next Hop Type	The next hop type. Select VPC peering connection .	VPC peering connection
Next Hop	The next hop address. Select the name of the current VPC peering connection.	peering-AB

Parameter	Description	Example Value
Description	Supplementary information about the route. This parameter is optional.	-
	The route description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

d. Click **OK**.

You can view the route in the route list.

- 2. Add routes to the route table of the peer VPC:
 - a. In the VPC peering connection list of the peer account, click the name of the target VPC peering connection.
 - The **Basic Information** tab of the VPC peering connection is displayed.
 - b. On the **Local Routes** tab of the VPC peering connection, click the **Route Tables** hyperlink.

The **Summary** tab of the default route table for the peer VPC is displayed.

c. Click **Add Route**.

Table 8-18 describes the route parameters.

Table 8-18 Parameter description

Parameter	Description	Example Value
Destination	The local VPC CIDR block, subnet CIDR block, or ECS IP address. For details, see VPC Peering Connection Usage Examples.	VPC-A CIDR block: 172.16.0.0/16
Next Hop Type	The next hop type. Select VPC peering connection .	VPC peering connection
Next Hop	The next hop address. Select the name of the current VPC peering connection.	peering-AB
Description	Supplementary information about the route. This parameter is optional. The route description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

d. Click **OK**.

You can view the route in the route list.

Step 4: Verify Network Connectivity

After you add routes for the VPC peering connection, verify the communication between the local and peer VPCs.

- 1. Log in to ECS-A01 in the local VPC.
- 2. Check whether ECS-A01 can communicate with RDS-B01.

ping IP address of RDS-B01

Example command:

ping 172.17.0.21

If information similar to the following is displayed, ECS-A01 and RDS-B01 can communicate with each other, and the VPC peering connection between VPC-A and VPC-B is successfully created.

```
[root@ecs-A02 ~]# ping 172.17.0.21
PING 172.17.0.21 (172.17.0.21) 56(84) bytes of data.
64 bytes from 172.17.0.21: icmp_seq=1 ttl=64 time=0.849 ms
64 bytes from 172.17.0.21: icmp_seq=2 ttl=64 time=0.455 ms
64 bytes from 172.17.0.21: icmp_seq=3 ttl=64 time=0.385 ms
64 bytes from 172.17.0.21: icmp_seq=4 ttl=64 time=0.372 ms
...
--- 172.17.0.21 ping statistics ---
```

NOTICE

- In this example, ECS-A01 and RDS-B01 are in the same security group. If
 the instances in different security groups, you need to add inbound rules to
 allow access from the peer security group. For details, see Enabling ECSs
 In Different Security Groups to Communicate Through an Internal
 Network.
- If VPCs connected by a VPC peering connection cannot communicate with each other, refer to Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?.

8.5 Obtaining the Peer Project ID of a VPC Peering Connection

Scenarios

If you create a VPC peering connection between two VPCs in different accounts, you can refer to this section to obtain the project ID of the region that the peer VPC resides.

Procedure

- Log in to the management console.
 The owner of the peer account logs in to the management console.
- 2. In the upper right corner of the page, select **My Credentials** from the username drop-down list.

3. In the project list, obtain the project ID.

8.6 Modifying a VPC Peering Connection

Scenarios

This section describes how to modify the name of a VPC peering connection.

Either owner of a VPC in a peering connection can modify the VPC peering connection in any state.

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 > VPC Peering Connections**.

The VPC peering connection list is displayed.

- 4. In the VPC peering connection list, locate the row that contains the target VPC peering connection and click **Modify** in the **Operation** column.
 - The **Modify VPC Peering Connection** dialog box is displayed.
- 5. Modify the VPC peering connection information and click **OK**.

8.7 Viewing VPC Peering Connections

Scenarios

This section describes how to view basic information about a VPC peering connection, including the connection name, status, and information about the local and peer VPCs.

If a VPC peering connection is created between two VPCs in different accounts, both the local and peer accounts can view information about the VPC peering connection.

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.

 In the navigation pane on the left, choose Virtual Private Cloud > VPC Peering Connections.

The VPC peering connection list is displayed.

4. In the VPC peering connection list, click the name of the target VPC peering connection

On the displayed page, view details about the VPC peering connection.

8.8 Deleting a VPC Peering Connection

Scenarios

This section describes how to delete a VPC peering connection.

Either owner of a VPC in a peering connection can delete the VPC peering connection in any state.

Notes and Constraints

The owner of either VPC in a peering connection can delete the VPC peering connection at any time. Deleting a VPC peering connection will also delete all information about this connection, including the routes in the local and peer VPC route tables added for the connection.

Procedure

- 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 > VPC Peering Connections**.

The VPC peering connection list is displayed.

- In the VPC peering connection list, locate the row that contains the target VPC peering connection and click **Delete** in the **Operation** column.
 A confirmation dialog box is displayed.
- 5. Click Yes.

8.9 Modifying Routes Configured for a VPC Peering Connection

Scenarios

This section describes how to modify the routes added for a VPC peering connection in the route tables of the local and peer VPCs.

- Modifying Routes of a VPC Peering Connection Between VPCs in the Same Account
- Modifying Routes of a VPC Peering Connection Between VPCs in Different Accounts

You can follow the instructions provided in this section to modify routes based on your requirements.

Modifying Routes of a VPC Peering Connection Between VPCs in the Same Account

- 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 > VPC Peering Connections**.

The VPC peering connection list is displayed.

4. In the VPC peering connection list, click the name of the target VPC peering connection.

The page showing the VPC peering connection details is displayed.

- 5. Modify the route added to the route table of the local VPC:
 - a. Click the Local Routes tab and then click the Route Tables hyperlink.
 The Summary tab of the default route table for the local VPC is displayed.
 - b. Locate the row that contains the route to be modified and click **Modify** in the **Operation** column.

The **Modify Route** dialog box is displayed.

- Modify the route and click OK.
- 6. Modify the route added to the route table of the peer VPC:
 - a. Click the **Peer Routes** tab and then click the **Route Tables** hyperlink.
 The **Summary** tab of the default route table for the peer VPC is displayed.
 - b. Locate the row that contains the route to be modified and click **Modify** in the **Operation** column.
 - The **Modify Route** dialog box is displayed.
 - c. Modify the route and click **OK**.

Modifying Routes of a VPC Peering Connection Between VPCs in Different Accounts

Only the account owner of a VPC can modify the routes added for the connection.

- 1. Log in to the management console using the account of the local VPC and modify the route of the local VPC:
 - a. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.
 - The Virtual Private Cloud page is displayed.
 - In the navigation pane on the left, choose Virtual Private Cloud > VPC
 Peering Connections.

The VPC peering connection list is displayed.

c. In the VPC peering connection list, click the name of the target VPC peering connection.

The page showing the VPC peering connection details is displayed.

- d. Modify the route added to the route table of the local VPC:
 - i. Click the **Local Routes** tab and then click the **Route Tables** hyperlink.
 - The **Summary** tab of the default route table for the local VPC is displayed.
 - ii. Locate the row that contains the route to be modified and click **Modify** in the **Operation** column.
 - The **Modify Route** dialog box is displayed.
 - iii. Modify the route and click **OK**.
- 2. Log in to the management console using the account of the peer VPC and modify the route of the peer VPC by referring to 1.

8.10 Viewing Routes Configured for a VPC Peering Connection

Scenarios

This section describes how to view the routes added to the route tables of local and peer VPCs of a VPC peering connection.

- Viewing Routes of a VPC Peering Connection Between VPCs in the Same Account
- Viewing Routes of a VPC Peering Connection Between VPCs in Different Accounts

If two VPCs cannot communicate through a VPC peering connection, you can check the routes added for the local and peer VPCs by following the instructions provided in this section.

Viewing Routes of a VPC Peering Connection Between VPCs in the Same Account

- 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 > VPC Peering Connections.

The VPC peering connection list is displayed.

4. In the VPC peering connection list, click the name of the target VPC peering connection.

The page showing the VPC peering connection details is displayed.

- 5. View the routes added for the VPC peering connection:
 - a. Click the **Local Routes** tab to view the local route added for the VPC peering connection.
 - b. Click the **Peer Routes** tab to view the peer route added for the VPC peering connection.

Viewing Routes of a VPC Peering Connection Between VPCs in Different Accounts

Only the account owner of a VPC in a VPC peering connection can view the routes added for the connection.

- 1. Log in to the management console using the account of the local VPC and view the route of the local VPC:
 - a. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.
 - The Virtual Private Cloud page is displayed.
 - b. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.
 - The VPC peering connection list is displayed.
 - c. In the VPC peering connection list, click the name of the target VPC peering connection.
 - The page showing the VPC peering connection details is displayed.
 - d. Click the **Local Routes** tab to view the local route added for the VPC peering connection.
- 2. Log in to the management console using the account of the peer VPC and view the route of the peer VPC by referring to 1.

8.11 Deleting Routes Configured for a VPC Peering Connection

Scenarios

This section describes how to delete routes from the route tables of the local and peer VPCs connected by a VPC peering connection.

- Deleting Routes of a VPC Peering Connection Between VPCs in the Same Account
- Deleting Routes of a VPC Peering Connection Between VPCs in Different Accounts

Deleting Routes of a VPC Peering Connection Between VPCs in the Same Account

- 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 > VPC Peering Connections**.

The VPC peering connection list is displayed.

4. In the VPC peering connection list, click the name of the target VPC peering connection.

The page showing the VPC peering connection details is displayed.

- 5. Delete the route added to the route table of the local VPC:
 - a. Click the Local Routes tab and then click the Route Tables hyperlink.
 The Summary tab of the default route table for the local VPC is displayed.
 - b. Locate the row that contains the route to be deleted and click **Delete** in the **Operation** column.

A confirmation dialog box is displayed.

- c. Click **Yes**.
- 6. Delete the route added to the route table of the peer VPC:
 - a. Click the **Peer Routes** tab and then click the **Route Tables** hyperlink.
 The **Summary** tab of the default route table for the peer VPC is displayed.
 - b. Locate the row that contains the route to be deleted and click **Delete** in the **Operation** column.
 - A confirmation dialog box is displayed.
 - c. Click Yes.

Deleting Routes of a VPC Peering Connection Between VPCs in Different Accounts

Only the account owner of a VPC in a VPC peering connection can delete the routes added for the connection.

- 1. Log in to the management console using the account of the local VPC and delete the route of the local VPC:
 - a. Click in the upper left corner and choose **Network** > **Virtual Private Cloud**.

The Virtual Private Cloud page is displayed.

b. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.

The VPC peering connection list is displayed.

- c. In the VPC peering connection list, click the name of the target VPC peering connection.
 - The page showing the VPC peering connection details is displayed.
- d. Delete the route added to the route table of the local VPC:
 - i. Click the **Local Routes** tab and then click the **Route Tables** hyperlink.

- The **Summary** tab of the default route table for the local VPC is displayed.
- ii. Locate the row that contains the route to be deleted and click **Delete** in the **Operation** column.
 - A confirmation dialog box is displayed.
- iii. Click Yes.
- 2. Log in to the management console using the account of the peer VPC and delete the route of the peer VPC by referring to 1.

9 VPC Flow Log

9.1 VPC Flow Log Overview

What Is a VPC Flow Log?

A VPC flow log records information about the traffic going to and from a VPC. VPC flow logs help you monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification.

VPC flow logs must be used together with the Log Tank Service (LTS). Before you create a VPC flow log, you need to create a log group and a log topic in LTS.

Figure 9-1 shows the process for configuring VPC flow logs.

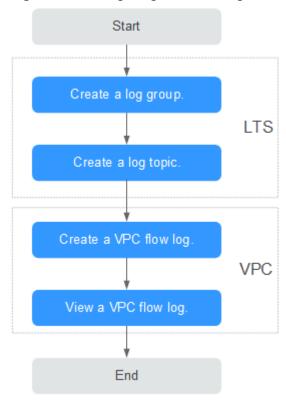


Figure 9-1 Configuring VPC flow logs

Notes and Constraints

- The following lists ECS types that support VPC flow logs in each region.
 - eu-west-1: C3, CC3, S3, and I3
 - eu-west-0: C3, CC3, S3, I3, P2, H1, and M2
- Each account can have up to 10 VPC flow logs in a region.

9.2 Creating a VPC Flow Log

Scenarios

A VPC flow log records information about the traffic going to and from a VPC.

Prerequisites

Ensure that the following operations have been performed on the LTS console:

- Create a log group.
- Create a log topic.

For more information about the LTS service, see the Log Tank Service User Guide.

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 **VPC Flow Logs**.
- 4. In the upper right corner, click **Create VPC Flow Log**. On the displayed page, configure parameters as prompted.

Table 9-1 Parameter descriptions

Parameter	Description	Example Value
Name	The VPC flow log name. The name can contain a maximum of 64 characters, which may consist of letters, digits, underscores (_), hyphens (-), and periods (.). The name cannot contain spaces.	flowlog-495d
Resource Type	The type of resources whose traffic is to be logged. You can select NIC .	NIC
Resource	The specific NIC whose traffic is to be logged. NOTE We recommend that you select an ECS that is in the running state. If an ECS in the stopped state is selected, restart the ECS after creating the VPC flow log for accurately recording the information about the traffic going to and from the ECS NIC.	N/A
Filter	 All traffic: specifies that both accepted and rejected traffic of the specified resource will be logged. Accepted traffic: specifies that only accepted traffic of the specified resource will be logged. Accepted traffic refers to the traffic permitted by the security group or network ACL. Rejected traffic: specifies that only rejected traffic of the specified resource will be logged. Rejected traffic refers to the traffic denied by the network ACL. 	All
Log Group	The log group created in LTS.	lts-group-abc
Log Stream	The log stream created in LTS.	lts-topic-abc

Parameter	Description	Example Value
Description	Supplementary information about the VPC flow log. This parameter is optional. The VPC flow log description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

Ⅲ NOTE

Only two flow logs, each with a different filter, can be created for a single resource under the same log group and log topic. Each VPC flow log must be unique.

5. Click **OK**.

9.3 Viewing a VPC Flow Log

Scenarios

View information about your flow log record.

The capture window is approximately 10 minutes, which indicates that a flow log record will be generated every 10 minutes. After creating a VPC flow log, you need to wait about 10 minutes before you can view the flow log record.

□ NOTE

If an ECS is in the stopped state, its flow log records will not be displayed.

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 **VPC Flow Logs**.
- 4. Locate the target VPC flow log and click **View Log Record** in the **Operation** column to view information about the flow log record in LTS.

The flow log record is in the following format:

Example 1: The following is an example of a flow log record in which data was recorded during the capture window:

1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983-bd6512aed4bd 192.168.0.154 192.168.3.25 38929 53 17 1 96 1548752136 1548752736 ACCEPT OK

Value 1 indicates the VPC flow log version. Traffic with a size of 96 bytes to NIC 1d515d18-1b36-47dc-a983-bd6512aed4bd during the past 10 minutes (from 16:55:36 to 17:05:36 on January 29, 2019) was allowed. A data packet

was transmitted over the UDP protocol from source IP address **192.168.0.154** and port **38929** to destination IP address **192.168.3.25** and port **53**.

Example 2: The following is an example of a flow log record in which no data was recorded during the capture window:

1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983-bd6512aed4bd - - - - - - 1431280876 1431280934 - NODATA

Example 3: The following is an example of a flow log record in which data was skipped during the capture window:

1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983-bd6512aed4bd - - - - - - 1431280876 1431280934 - SKIPDATA

Table 9-2 describes the fields of a flow log record.

Table 9-2 Log field description

Field	Description	Example Value
version	The VPC flow log version.	1
project-id	The project ID.	5f67944957444bd6bb4fe3b3 67de8f3d
interface-id	The ID of the NIC for which the traffic is recorded.	1d515d18-1b36-47dc-a983- bd6512aed4bd
srcaddr	The source IP address.	192.168.0.154
dstaddr	The destination IP address.	192.168.3.25
srcport	The source port.	38929
dstport	The destination port.	53
protocol	The Internet Assigned Numbers Authority (IANA) protocol number of the traffic. For details, see Assigned Internet Protocol Numbers.	17
packets	The number of packets transferred during the capture window.	1
bytes	The number of bytes transferred during the capture window.	96
start	The time, in Unix seconds, of the start of the capture window.	1548752136
end	The time, in Unix seconds, of the end of the capture window.	1548752736

Field	Description	Example Value
action	The action associated with the traffic: • ACCEPT: The recorded traffic was allowed by the security groups or network ACLs. • REJECT: The recorded traffic was denied by the security groups or network ACLs.	ACCEPT
log-status	 The logging status of the VPC flow log: OK: Data is logging normally to the chosen destinations. NODATA: There was no traffic of the Filter setting to or from the NIC during the capture window. SKIPDATA: Some flow log records were skipped during the capture window. This may be caused by an internal capacity constraint or an internal error. Example: When Filter is set to Accepted traffic, if there is accepted traffic, the value of log-status is OK. If there is no accepted traffic, the value of log-status is NODATA regardless of whether there is rejected traffic. If some accepted traffic is abnormally skipped, the value of log-status is SKIPDATA. 	OK

You can enter a keyword on the log topic details page on the LTS console to search for flow log records.

9.4 Enabling or Disabling VPC Flow Log

Scenarios

After a VPC flow log is created, the VPC flow log is automatically enabled. If you do not need to record flow log data, you can disable the corresponding VPC flow log. A disabled VPC flow log can be enabled again.

Notes and Constraints

- After a VPC flow log is enabled, the system starts to collect flow logs in the next log collection period.
- After a VPC flow log is disabled, the system stops collecting flow logs in the next log collection period. Generated flow logs will still be reported.

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 **VPC Flow Logs**.
- 4. Locate the VPC flow log to be enabled or disabled, and click **Enable** or **Disable** in the **Operation** column.
- 5. Click Yes.

9.5 Deleting a VPC Flow Log

Scenarios

Delete a VPC flow log that is not required. Deleting a VPC flow log will not delete the existing flow log records in LTS.

■ NOTE

If a NIC that uses a VPC flow log is deleted, the flow log will be automatically deleted. However, the flow log records are not deleted.

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 **VPC Flow Logs**.
- 4. Locate the row that contains the VPC flow log to be deleted and click **Delete** in the **Operation** column.

5. Click **Yes** in the displayed dialog box.

10 Virtual IP Address

10.1 Virtual IP Address Overview

What Is a Virtual IP Address?

A virtual IP address can be shared among multiple ECSs. An ECS can have both private and virtual IP addresses, and you can access the ECS through either IP address. A virtual IP address has the same network access capabilities as a private IP address, including layer 2 and layer 3 communication in VPCs, access between VPCs using VPC peering connections, as well as access through EIPs, VPN connections, and Direct Connect connections.

You can bind ECSs deployed in active/standby mode with the same virtual IP address, and then bind an EIP to the virtual IP address. Virtual IP addresses can work together with Keepalived to ensure high availability and disaster recovery. If the active ECS is faulty, the standby ECS automatically takes over services from the active one.

Networking

Virtual IP addresses are used for high availability and can work together with Keepalived to make active/standby ECS switchover possible. This way if one ECS goes down for some reason, the other one can take over and services continue uninterrupted. ECSs can be configured for HA or as load balancing clusters.

• Networking mode 1: HA

If you want to improve service availability and avoid single points of failure, you can deploy ECSs in the active/standby mode or deploy one active ECS and multiple standby ECSs. In this arrangement, the ECSs all use the same virtual IP address. If the active ECS becomes faulty, a standby ECS takes over services from the active ECS and services continue uninterrupted.

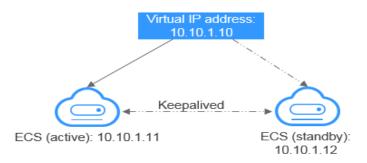


Figure 10-1 Networking diagram of the HA mode

- In this configuration, a single virtual IP address is bound to two ECSs in the same subnet.
- Keepalived is then used to configure the two ECSs to work in the active/ standby mode. Follow industry standards for configuring Keepalived. The details are not included here.
- **Networking mode 2**: HA load balancing cluster

If you want to build a high-availability load balancing cluster, use Keepalived and configure LVS nodes as direct routers.

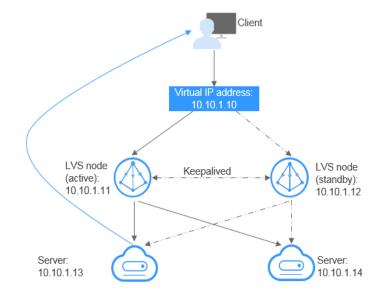


Figure 10-2 HA load balancing cluster

- Bind a single virtual IP address to two ECSs.
- Configure the two ECSs as LVS nodes working as direct routers and use Keepalived to configure the nodes in the active/standby mode. The two ECSs will evenly forward requests to different backend servers.
- Configure two more ECSs as backend servers.
- Disable the source/destination check for the two backend servers.

Follow industry standards for configuring Keepalived. The details are not included here.

Application Scenarios

- Accessing the virtual IP address through an EIP
 If your application has high availability requirements and needs to provide services through the Internet, it is recommended that you bind an EIP to a virtual IP address.
- Using a VPN, Direct Connect, or VPC peering connection to access a virtual IP address

To ensure high availability and access to the Internet, use a VPN for security and Direct Connect for a stable connection. The VPC peering connection is needed so that the VPCs in the same region can communicate with each other.

Notes and Constraints

- Virtual IP addresses are not recommended when multiple NICs in the same subnet are configured on an ECS. It is too easy for there to be route conflicts on the ECS, which would cause communication failure using the virtual IP address.
- If a virtual IP address is used in an active/standby scenario, disable IP forwarding on the standby ECS. For details, see Disabling IP Forwarding on the Standby ECS.

10.2 Assigning a Virtual IP Address

Scenarios

If an ECS requires a virtual IP address or if a virtual IP address needs to be reserved, you can assign a virtual IP address from the subnet.

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. In the subnet list, click the name of the subnet where a virtual IP address is to be assigned.
- 5. Click the **IP Addresses** tab and click **Assign Virtual IP Address**.
- 6. Select an IP address type. This parameter is available only in regions supporting IPv6.
 - IPv4
 - IPv6
- 7. Select a virtual IP address assignment mode.
 - Automatic: The system assigns an IP address automatically.
 - Manual: You can specify an IP address.

- 8. Select Manual and enter a virtual IP address.
- 9. Click OK.

You can then query the assigned virtual IP address in the IP address list.

10.3 Binding a Virtual IP Address to an EIP or ECS

Scenarios

You can use a virtual IP address and an EIP together.

If you bind a virtual IP address to ECSs that work in active/standby pairs and bind an EIP to the virtual IP address, you can access the ECSs over the Internet.

Notes and Constraints

- A virtual IP address can only be bound to one EIP.
- Do not bind more than eight virtual IP addresses to an ECS.
- A virtual IP address can be bound to a maximum of 10 ECSs.

If a virtual IP address is bound to an ECS, the virtual IP address is also associated with the security group of the ECS. A virtual IP address can be associated with up to 10 security groups.

Binding a Virtual IP Address to an EIP or ECS on the Console

- 1. Log in to the management console.
- Click = in the upper left corner and choose Network > Virtual Private Cloud.

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. Click the name with a hyperlink of the subnet that the virtual IP address belongs to.

The subnet details page is displayed.

- 5. On the **IP Addresses** tab, bind an EIP to the virtual IP address:
 - a. Locate the row that contains the virtual IP address and click **Bind to EIP** in the **Operation** column.

The **Bind to EIP** dialog box is displayed.

- b. Select an EIP and click **OK**.
 - In the virtual IP address list, you can view that the virtual IP address has an EIP bound.
- 6. On the **IP Addresses** tab, bind an instance to the virtual IP address:
 - a. Locate the row that contains the virtual IP address and click **Bind to Server** in the **Operation** column.

The **Bind to Server** dialog box is displayed.

b. Select an ECS and click OK.

In the virtual IP address list, you can view that the virtual IP address has an ECS bound.

NOTICE

- After a virtual IP address is bound to an ECS NIC, you need to manually configure the virtual IP address on the ECS. For details, see Configuring a Virtual IP Address for an ECS.
- If an ECS has multiple NICs, bind the virtual IP address to the primary NIC.
- An ECS NIC can have multiple virtual IP addresses bound.

Configuring a Virtual IP Address for an ECS

Manually configure the virtual IP address bound to an ECS.

The following OSs are used as examples here. For other OSs, see the help documents on their official websites.

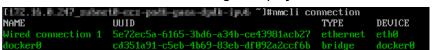
- Linux: CentOS 7.2 64bit and Ubuntu 22.04 server 64bit
- Windows: Windows Server

Linux (CentOS 7.2 64bit is used as an example.)

1. Obtain the NIC that the virtual IP address is to be bound and the connection of the NIC:

nmcli connection

Information similar to the following is displayed:



The command output in this example is described as follows:

- eth0 in the DEVICE column indicates the NIC that the virtual IP address is to be bound.
- Wired connection 1 in the NAME column indicates the connection of the NIC
- 2. Add the virtual IP address for the connection:

nmcli connection modify "Connection name of the N/C" +ipv4.addresses Virtual IP address

Configure the parameters as follows:

- Connection name of the NIC: The connection name of the NIC obtained in 1. In this example, the connection name is Wired connection 1.
- Virtual IP address. Enter the virtual IP address to be added. If you add multiple virtual IP addresses at a time, separate every two with a comma (,).

Example commands:

Adding a single virtual IP address: nmcli connection modify "Wired connection 1" +ipv4.addresses 172.16.0.125

- Adding multiple virtual IP addresses: nmcli connection modify "Wired connection 1" +ipv4.addresses 172.16.0.125,172.16.0.126
- 3. Make the configuration in 2 take effect:

nmcli connection up "Connection name of the NIC"

In this example, run the following command:

nmcli connection up "Wired connection 1"

Information similar to the following is displayed:

4. Check whether the virtual IP address has been bound:

ip a

Information similar to the following is displayed. In the command output, the virtual IP address 172.16.0.125 is bound to NIC eth0.

◯ NOTE

To delete an added virtual IP address, perform the following steps:

1. Delete the virtual IP address from the connection of the NIC:

nmcli connection modify "Connection name of the NIC" -ipv4.addresses Virtual IP address

To delete multiple virtual IP addresses at a time, separate every two with a comma (,). Example commands are as follows:

- Deleting a single virtual IP address: nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125
- Deleting multiple virtual IP addresses: nmcli connection modify "Wired connection 1" -ipv4.addresses 172.16.0.125,172.16.0.126
- 2. Make the deletion take effect by referring to 3.

Linux (Ubuntu 22.04 server 64bit is used as an example.)

If an ECS runs Ubuntu 22 or Ubuntu 20, perform the following operations:

1. Obtain the NIC that the virtual IP address is to be bound:

ifconfig

Information similar to the following is displayed. In this example, the NIC bound to the virtual IP address is **eth0**.

```
root@ecs-X-ubantu:~# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 172.16.0.210 netmask 255.255.255.0 broadcast 172.16.0.255
inet6 fe80::f816:3eff:fe01:f1c3 prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:01:f1:c3 txqueuelen 1000 (Ethernet)
RX packets 43915 bytes 63606486 (63.6 MB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 3364 bytes 455617 (455.6 KB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

2. Switch to the /etc/netplan directory:

cd /etc/netplan

- 3. Add a virtual IP address to the NIC.
 - a. Open the configuration file **01-netcfg.yaml**:

vim 01-netcfg.yaml

- b. Press i to enter the editing mode.
- c. In the NIC configuration area, add a virtual IP address. In this example, add a virtual IP address for **eth0**:

addresses:

- 172.16.0.26/32

The file content is as follows:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
     eth0:
        dhcp4: true
        addresses:
        - 172.16.0.26/32
     eth1:
       dhcp4: true
     eth2:
       dhcp4: true
     eth3:
        dhcp4: true
     eth4:
        dhcp4: true
```

- d. Press **Esc**, enter :wq!, save the configuration, and exit.
- 4. Make the configuration in 3 take effect:

netplan apply

5. Check whether the virtual IP address has been bound:

ip a

Information similar to the following is displayed. In the command output, the virtual IP address 172.16.0.26 is bound to NIC eth0.

```
root@ecs-X-ubantu:/etc/netplan# ip a
...
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default
qlen 1000
    link/ether fa:16:3e:01:f1:c3 brd ff:ff:ff:ff:ff
    altname enp0s3
    altname ens3
    inet 172.16.0.26/32 scope global noprefixroute eth0
    valid_lft forever preferred_lft forever
inet 172.16.0.210/24 brd 172.16.0.255 scope global dynamic noprefixroute eth0
    valid_lft 107999971sec preferred_lft 107999971sec
inet6 fe80::f816:3eff:fe01:f1c3/64 scope link
    valid_lft forever preferred_lft forever
```

□ NOTE

To delete an added virtual IP address, perform the following steps:

- 1. Open the configuration file **01-netcfg.yaml** and delete the virtual IP address of the corresponding NIC by referring to **3**.
- 2. Make the deletion take effect by referring to 4.

Windows OS (Windows Server is used as an example here.)

- 1. In **Control Panel**, click **Network and Sharing Center**, and click the corresponding local connection.
- 2. On the displayed page, click **Properties**.
- 3. On the Network tab page, select Internet Protocol Version 4 (TCP/IPv4).
- 4. Click **Properties**.
- 5. Select **Use the following IP address** and set **IP address** to the private IP address of the ECS, for example, 10.0.0.101.

Figure 10-3 Configuring private IP address

X Internet Protocol Version 4 (TCP/IPv4) Properties General You can get IP settings assigned automatically if your network supports this capability. Otherwise, you need to ask your network administrator for the appropriate IP settings. Obtain an IP address automatically Use the following IP address: IP address: 10 . 0 . 0 . 101 Subnet mask: 255 . 255 . 255 . 0 Default gateway: 10 . 0 . 0 . 1 Obtain DNS server address automatically Use the following DNS server addresses: Preferred DNS server: 100 . 125 . 1 . 250 Alternate DNS server: 114 . 114 . 114 . 114 Validate settings upon exit Advanced... OK Cancel

- 6. Click Advanced.
- 7. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address, for example, 10.0.0.154.

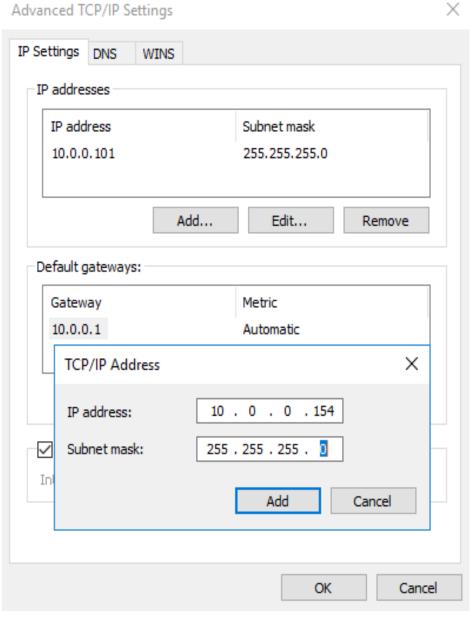


Figure 10-4 Configuring virtual IP address

- 8. Click **OK**.
- 9. In the **Start** menu, open the Windows command line window and run the following command to check whether the virtual IP address has been configured:

ipconfig /all

In the command output, **IPv4 Address** is the virtual IP address 10.0.0.154, indicating that the virtual IP address of the ECS NIC has been correctly configured.

10.4 Binding a Virtual IP Address to an EIP

Scenarios

This section describes how to bind a virtual IP address to an EIP.

Prerequisites

- You have configured the ECS networking based on **Networking** and ensure that the ECS has been bound with a virtual IP address.
- You have assigned an EIP.

Procedure

- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Network** > **Elastic IP**. The EIP list page is displayed.
- 3. Locate the row that contains the EIP to be bound to the virtual IP address, and click **Bind** in the **Operation** column.
- 4. In the **Bind EIP** dialog box, set **Instance Type** to **Virtual IP address**.
- 5. In the virtual IP address list, select the virtual IP address to be bound and click **OK**.

10.5 Unbinding a Virtual IP Address from an Instance

Scenarios

This section describes how to unbind a virtual IP address from an ECS.

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**. The **Subnets** page is displayed.
- 4. Click the name of the subnet that the virtual IP address belongs to.
 The **Summary** page is displayed.
- 5. Click the **IP Addresses** tab.
 - The virtual IP address list is displayed.
- 6. Locate the row that contains the virtual IP address, click **More** in the **Operation** column, and select **Unbind from Server**.
 - The **Bound Server** dialog box is displayed.

- 7. Unbind the virtual IP address from the instance.
 - a. Select the type of the instance bound to the virtual IP address.
 - b. Locate the row that contains the instance and click **Unbind** in the **Operation** column.
 - A confirmation dialog box is displayed.
 - c. Confirm the information and click Yes.

10.6 Unbinding a Virtual IP Address from an EIP

Scenarios

This section describes how to unbind a virtual IP address from an EIP.

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.

- In the navigation pane on the left, choose Virtual Private Cloud > Subnets.
 The Subnets page is displayed.
- 4. Click the name of the subnet that the virtual IP address belongs to.
 The **Summary** page is displayed.
- 5. Click the **IP Addresses** tab.
 - The virtual IP address list is displayed.
- 6. Locate the row that contains the virtual IP address, click **More** in the **Operation** column, and select **Unbind from EIP**.
 - A confirmation dialog box is displayed.
- 7. Confirm the information and click **Yes**.

10.7 Releasing a Virtual IP Address

Scenarios

If you no longer need a virtual IP address or a reserved virtual IP address, you can release it to avoid wasting resources.

Notes and Constraints

If you want to release a virtual IP address that is being used by a resource, refer to **Table 10-1**.

Prompts Cause Analysis and Solution This operation cannot be This virtual IP address is being used by an EIP or performed because the IP an ECS. address is bound to an Unbind the virtual IP address first. instance or an EIP. Unbind • EIP: Unbinding a Virtual IP Address from an the IP address and try again. • ECS: Unbinding a Virtual IP Address from an Instance Release the virtual IP address. This operation cannot be The virtual IP address is being used by an RDS DB performed because the IP instance. Delete the DB instance, which will also address is being used by a release its virtual IP address. system component.

Table 10-1 Releasing a virtual IP address that is being used by a resource

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 the name of the subnet that the virtual IP address belongs to.
- Click the IP Addresses tab, locate the row that contains the virtual IP address to be released, click More in the Operation column, and select Release.
 A confirmation dialog box is displayed.

Figure 10-5 Releasing a virtual IP address



6. Confirm the information and click Yes.

10.8 Disabling IP Forwarding on the Standby ECS

Scenarios

If a virtual IP address is used in an active/standby scenario, disable IP forwarding on the standby ECS.

Linux

- 1. Log in to the ECS.
- 2. Run the following command to switch to user **root**:

su root

3. Check whether IP forwarding is enabled:

cat /proc/sys/net/ipv4/ip_forward

In the command output, **1** indicates it is enabled, and **0** indicates it is disabled. The default value is **0**.

- If 1 is displayed, go to 4.
- If 0 is displayed, no further action is required.
- 4. Use either of the following methods to modify the configuration file:
 - Method 1: Use the vi editor to open the /etc/sysctl.conf file, change the value of net.ipv4.ip_forward to 0, and enter :wq to save the change and exit.
 - Method 2: Use the sed command. An example command is as follows:
 sed -i '/net.ipv4.ip forward/s/1/0/g' /etc/sysctl.conf
- 5. Make the modification take effect:

sysctl -p /etc/sysctl.conf

Windows

- 1. Log in to the ECS.
- 2. Open **Command Prompt** and run the following command:

ipconfig/all

In the command output, if the value of **IP Routing Enabled** is **No**, the IP forwarding function is disabled.

- 3. Press **Windows** and **R** keys together to open the **Run** box, and enter **regedit** to open the **Registry Editor**.
- 4. Set the value of IPEnableRouter under HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Services\Tcpip\Parameters to 0.
 - If the value is set to 0, IP forwarding will be disabled.
 - If the value is set to 1, IP forwarding will be enabled.

10.9 Disabling Source/Destination Check for an ECS NIC

Scenarios

If a virtual IP address is used in an HA load balancing cluster, you need to disable source/destination check for ECS NICs.

Procedure

1. Log in to the management console.

- 2. In the upper left corner of the page, click ≡. In the service list, choose Computing > Elastic Cloud Server.
- 3. In the ECS list, click the ECS name.
- 4. On the displayed ECS details page, click the **NICs** tab.
- 5. Check that **Source/Destination Check** is disabled.

11 Interconnecting with CTS

11.1 Supported VPC Operations

With CTS, you can record operations performed on the VPC service for further query, audit, and backtracking purposes.

Table 11-1 lists the VPC operations that can be recorded by CTS.

Table 11-1 VPC operations that can be recorded by CTS

Operation	Resource Type	Trace
Modifying a bandwidth	Bandwidth	modifyBandwidth
Assigning an EIP	EIP	createEip
Releasing an EIP	EIP	deleteEip
Binding an EIP	EIP	bindEip
Unbinding an EIP	EIP	unbindEip
Assigning a private IP address	Private IP address	createPrivateIp
Deleting a private IP address	Private IP address	deletePrivateIp
Creating a security group	security_groups	createSecurity-group
Updating a security group	security_groups	updateSecurity-group
Deleting a security group	security_groups	deleteSecurity-group
Creating a security group rule	security-group-rules	createSecurity-group-rule

Operation	Resource Type	Trace
Updating a security group rule	security-group-rules	updateSecurity-group-rule
Deleting a security group rule	security-group-rules	deleteSecurity-group-rule
Creating a subnet	Subnet	createSubnet
Deleting a subnet	Subnet	deleteSubnet
Modifying a subnet	Subnet	modifySubnet
Creating a VPC	VPC	createVpc
Deleting a VPC	VPC	deleteVpc
Modifying a VPC	VPC	modifyVpc
Creating a VPN	VPN	createVpn
Deleting a VPN	VPN	deleteVpn
Modifying a VPN	VPN	modifyVpn
Creating a router	routers	createRouter
Updating a router	routers	updateRouter
Adding an interface to a router	routers	addRouterInterface
Deleting an interface from a router	routers	removeRouterInterface
Creating a port	ports	createPort
Updating a port	ports	updatePort
Deleting a port	ports	deletePort
Creating a network	networks	createNetwork
Updating a network	networks	updateNetwork
Deleting a network	networks	deleteNetwork
Batch creating or deleting subnet tags	tag	batchUpdateTags
Batch creating or deleting VPC tags	tag	batchUpdateVpcTags
Creating a route table	routetables	createRouteTable
Updating a route table	routetables	updateRouteTable
Deleting a route table	routetables	deleteRouteTable

Operation	Resource Type	Trace
Creating a VPC peering connection	vpc-peerings	createVpcPeerings
Updating a VPC peering connection	vpc-peerings	updateVpcPeerings
Deleting a VPC peering connection	vpc-peerings	deleteVpcPeerings
Creating a network ACL group	firewall-groups	createFirewallGroup
Updating a network ACL group	firewall-groups	updateFirewallGroup
Deleting a network ACL group	firewall-groups	deleteFirewallGroup
Creating a network ACL policy	firewall-policies	createFirewallPolicy
Updating a network ACL policy	firewall-policies	updateFirewallPolicy
Deleting a network ACL policy	firewall-policies	deleteFirewallPolicy
Inserting a network ACL rule	firewall-policies	insertFirewallPolicyRule
Removing a network ACL rule	firewall-policies	removeFirewallPolicyRule
Creating a network ACL rule	firewall-rules	createFirewallRule
Updating a network ACL rule	firewall-rules	updateFirewallRule
Deleting a network ACL rule	firewall-rules	deleteFirewallRule
Creating an IP address group	address_group	createAddress_group
Updating an IP address group	address_group	updateAddress_group
Forcibly deleting an IP address group	address_group	force_deleteAddress_group
Deleting an IP address group	address_group	deleteAddress_group
Creating a flow log	flowlogs	createFlowLog

Operation	Resource Type	Trace
Updating a flow log	flowlogs	updateFlowLog
Deleting a flow log	flowlogs	deleteFlowLog

11.2 Viewing Traces

Scenarios

After CTS is enabled, CTS starts recording operations on cloud resources. The CTS management console stores the last seven days of operation records.

This section describes how to query or export the last seven days of operation records on the management console.

Procedure

- 1. Log in to the management console.
- 2. In the upper left corner of the page, click = to go to the service list. Under Management & Deployment, click Cloud Trace Service.
- 3. In the navigation pane on the left, choose **Trace List**.
- 4. Specify filters as needed. The following filters are available:
 - Trace Source, Resource Type, and Search By
 Select filters from the drop-down list.

 If you select Trace name for Search By, select a trace name.
 If you select Resource ID for Search By, select or enter a resource ID.
 If you select Resource name for Search By, select or enter a resource name.
 - Operator: Select a specific operator (a user other than an account).
 - Trace Status: Select All trace statuses, Normal, Warning, or Incident.
 - Search time range: In the upper right corner, choose Last 1 hour, Last 1 day, or Last 1 week, or specify a custom time range.
- 5. Click arrow on the left of the required trace to expand its details.
- Locate the required trace and click View Trace in the Operation column.
 A dialog box is displayed, showing the trace content.

12 Monitoring

12.1 Supported Metrics

Description

This section describes the namespace, list, and measurement dimensions of EIP and bandwidth metrics 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 alarms generated for EIPs and bandwidths.

Namespace

SYS.VPC

Monitoring Metrics

Table 12-1 EIP and bandwidth metrics

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
upstream _bandwid th	Outbo und Band width	Network rate of outbound traffic Unit: bit/s	≥ 0 bit/s	Bandwidth or EIP	1 minute
downstre am_band width	Inbou nd Band width	Network rate of inbound traffic Unit: bit/s	≥ 0 bit/s	Bandwidth or EIP	1 minute

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
upstream _bandwid th_usage	Outbo und Band width Usage	Usage of outbound bandwidth in the unit of percent. Outbound bandwidth usage = Outbound bandwidth/ Purchased bandwidth	0% to 100%	Bandwidth or EIP	1 minute
downstre am_band width_usa ge	Inbou nd Band width Usage	Usage of inbound bandwidth in the unit of percent. Inbound bandwidth usage = Inbound bandwidth/ Purchased bandwidth	0-100%	Bandwidth or EIP	1 minute
up_strea m	Outbo und Traffic	Network traffic going out of the cloud platform in a minute Unit: byte	≥ 0 bytes	Bandwidth or EIP	1 minute
down_str eam	Inbou nd Traffic	Network traffic going into the cloud platform in a minute Unit: byte	≥ 0 bytes	Bandwidth or EIP	1 minute

Dimensions

Кеу	Value
publicip_id	EIP ID
bandwidth_id	Bandwidth ID

If a monitored object has multiple dimensions, all dimensions are mandatory when you use APIs to query the metrics.

- Query a monitoring metric: dim.0=bandwidth_id,530cd6b0-86d7-4818-837f-935f6a27414d&dim.1=publici p id,3773b058-5b4f-4366-9035-9bbd9964714a
- Query monitoring metrics in batches:

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

12.2 Viewing Metrics

Scenarios

You can view the bandwidth and EIP usage.

You can view the inbound bandwidth, outbound bandwidth, inbound bandwidth usage, outbound bandwidth usage, inbound traffic, and outbound traffic in a specified period.

Procedure (Cloud Eye Console)

- 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.
- Click Cloud Service Monitoring on the left of the page, and choose Elastic IP and Bandwidth.
- 4. Locate the row that contains the target bandwidth or EIP and click **View**Metric in the Operation column to check the bandwidth or EIP monitoring information.

12.3 Creating an Alarm Rule

Scenarios

You can configure alarm rules to customize the monitored objects and notification policies. You can learn your resource statuses at any time.

Procedure

- 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 left navigation pane on the left, choose **Alarm Management > Alarm Rules**.
- 4. On the **Alarm Rules** page, click **Create Alarm Rule** and set required parameters, or modify an existing alarm rule.
- After the parameters are set, click Create.
 After the alarm rule is created, the system automatically notifies you if an alarm is triggered for the VPC service.

□ NOTE

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

13

13.1

If your cloud account meets your permissions requirements, you can skip this section.

Prerequisites

To grant permissions for other services, learn about all supported by IAM.

Process Flow

- 1. On the IAM console, create a user group and assign permissions to it (VPC ReadOnlyAccess as an example).
- 2. Create an IAM user and add it to the created user group.
- 3. Log in as the IAM user and verify permissions.

In the authorized region, perform the following operations:

 Choose Service List > Virtual Private Cloud. Then click Create VPC on the VPC console. If a message appears indicating that you have insufficient permissions to perform the operation, the VPCReadOnlyAccess policy is in effect.

13.2 VPC Custom Policies

Custom policies can be created to supplement the system-defined policies of VPC.

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

Example Custom Policies

Example 1: Allowing users to create and view VPCs

• Example 2: Denying VPC deletion

A deny policy must be used in conjunction with other policies to take effect. If the permissions assigned to a user contain both Allow and Deny actions, the Deny actions take precedence over the Allow actions.

The following method can be used if you need to assign permissions of the **VPC FullAccess** policy to a user but also forbid the user from deleting VPCs. Create a custom policy for denying VPC deletion, and assign both policies to the group the user belongs to. Then the user can perform all operations on VPC except deleting VPCs. The following is an example deny policy:

• Example 3: Defining permissions for multiple services in a policy

A custom policy can contain the actions of multiple services that are of the global or project-level type. The following is an example policy containing actions of multiple services:

14 FAQ

14.1 General Questions

14.1.1 What Is a Quota?

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, a VPC quota limits the number of VPCs that can be created.

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. In the upper right corner of the page, click The **Service Quota** page is displayed.
- 3. 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?

The system does not support online quota adjustment.

If you need to adjust a quota, contact the operations administrator.

14.2 Billing and Payments

14.2.1 Will I Be Billed for Using the VPC Service?

The VPC service is free, but EIP and bandwidth used together with a VPC will be billed based on standard pricing.

14.2.2 How Is an EIP Billed?

EIPs can be billed on a pay-per-use basis.

- Figure 14-1
- Table 14-1

Figure 14-1 EIP billing

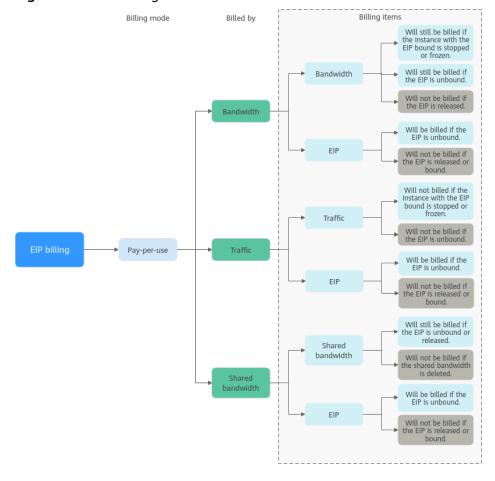


Table 14-1 EIP billing description

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
Pay- per- use	Bandwidth	Bandwid th EIP retention	If a pay-per-use EIP is billed by bandwidth: Bandwidth: You are billed based on your specified bandwidth size and usage duration. There is no limit on how much traffic you can use. After the EIP is purchased, you can change your specified bandwidth you use will not exceed the bandwidth you specified. EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: If the EIP is not bound to any instance, both the EIP and its bandwidth will be billed. If the EIP is bound to an instance, only the bandwidth will be billed. The bandwidth will be billed regardless of if the instance bound to the EIP is running or not. After the EIP is unbound from an instance, the bandwidth will continue to be billed. Unless it is released, the EIP will still be billed. If the EIP is released, both the EIP and its bandwidth will not be billed.

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
	Traffic	Traffic EIP retention n	If a pay-per-use EIP is billed by traffic: Traffic: You are billed based on your EIP type and the total amount of traffic going out of the cloud. The bandwidth size you set is only used to limit the maximum data transfer rate. To prevent high fees caused by burst traffic, specify a proper bandwidth size when you buy an EIP. If an EIP billed by traffic uses a dedicated bandwidth, only the bandwidth used in the outbound direction will be billed. EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: If the EIP is not bound to an instance, you will be billed for the EIP itself, but not for traffic. If the EIP is bound to an instance, only the used traffic will be billed. If the instance bound to the EIP stops running and there is no traffic generated, there will be no traffic or EIP fees. After the EIP is unbound from an instance, the traffic will not be billed but the EIP will still be billed. If the EIP is released, the EIP will not be billed.

Billing Mode	Billed By	Billing Item	Billing Item Description	Impact of EIP Operations on Billing Items
	Shared bandw idth	Shared bandwid th EIP retention	If a pay-per-use EIP is added to a shared bandwidth: • Share bandwidth: Only the shared bandwidth will be billed. There will be no additional bandwidth or traffic costs for EIPs added to the shared bandwidth. • EIP retention: If an EIP is not released, it will continue to be billed even if it is not bound to an instance.	After an EIP is purchased: Shared bandwidth No operations on the EIP will affect the billing of a shared bandwidth. For example, if you have released the EIP but have not deleted the shared bandwidth will still be billed. After a shared bandwidth is deleted, it will no longer be billed. After a shared bandwidth is deleted, it will no longer be billed. EIP retention If the EIP is not bound to an instance, the EIP will still be billed. If the EIP is unbound from an instance, the EIP will be billed to keep it allocated to your account unless it is released. If the EIP is released or bound to an instance, the EIP will not be billed.

To save money, you can add multiple EIPs in the same region to a shared bandwidth. A shared bandwidth can be billed on a pay-per-use basis. For details, see **Table 14-2**. Currently, only pay-per-use EIPs can be added to a shared bandwidth.

- You can add an EIP to a shared bandwidth when buying the EIP.
- You can also add an existing EIP to a shared bandwidth. After the EIP is added to a shared bandwidth, there will be no additional bandwidth or traffic cost. You will only be billed for the shared bandwidth.

Table 14-2 Shared bandwidth billing details

Billing Mode	Billed By	Billing Item	Billing Item Description
Pay- per-use	Bandwidt h	Bandwidth	You are billed based on your specified bandwidth size and usage duration. There is no limit on how much traffic you can use.
			After a shared bandwidth is purchased, you can change your specified bandwidth size. The bandwidth you use will not exceed the bandwidth you specified.

14.2.3 How Do I Change a Pay-per-Use EIP from Billing By Bandwidth to Traffic or from Billing By Traffic to Bandwidth?

Table 14-3 EIP billing mode change description

Change	Description
From billing by traffic (payper-use) to billing by	A pay-per-use EIP billed by traffic can be directly changed to be billed by bandwidth.
bandwidth (pay-per-use)	After the change is successful, the new billing mode takes effect immediately.
From billing by bandwidth (pay-per-use) to billing by	A pay-per-use EIP billed by bandwidth can be directly changed to be billed by traffic.
traffic (pay-per-use)	After the change is successful, the new billing mode takes effect immediately.

Pay-per-Use EIPs: From Billing By Traffic to By Bandwidth

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

- 4. On the **Modify Bandwidth** page, change the billing option as prompted. You can also change the bandwidth name and size.
- 5. Click Next.
- 6. On the displayed page, confirm the configurations and click **Submit**.

□ NOTE

• Changing the billing options does not change EIPs or interrupt their use.

14.2.4 Why Is My VPC Still Being Billed After It Was Deleted?

Symptom

You deleted all VPCs under your account. However, there are still bills for VPC.

Reasons

VPCs are free, but you are still billed for the EIPs used together with a VPC.

- EIPs may be in use in other projects or regions. You can view all EIPs in the billing center, locate the EIP, and switch to the project or region where the EIP is located and release it.
- The information in the bill is from your previous settlement period. Generally, the consumption amount is not deducted from your account immediately after pay-per-use EIPs are released. Instead, bills are generated and the consumption amount is deducted from your account only after the settlement period ends.

14.3 VPCs and Subnets

14.3.1 What Is Virtual Private Cloud?

Within your own VPC, you can create security groups and VPNs, configure IP address ranges, specify bandwidth sizes, manage the networks in the VPC, and make changes to these networks as needed, quickly and securely. You can also define rules to control communications between ECSs in the same security group or in different security groups.

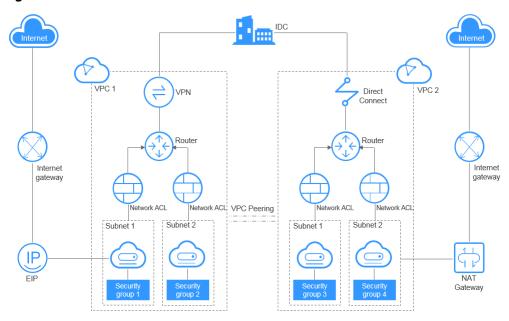


Figure 14-2 Product Architecture

14.3.2 Which CIDR Blocks Are Available for the VPC Service?

The following table lists the private CIDR blocks that you can specify when creating a VPC. Consider the following when selecting a VPC CIDR block:

- Number of IP addresses: Reserve sufficient IP addresses in case of business growth.
- IP address range: Avoid IP address conflicts if you need to connect a VPC to an on-premises data center or connect two VPCs.

Table 14-4 lists the supported VPC CIDR blocks.

Table 14-4 VPC CIDR blocks

VPC CIDR Block	IP Address Range	Maximum Number of IP Addresses
10.0.0.0/8-24	10.0.0.0-10.255.255.255	2^24-2=16777214
172.16.0.0/12-24	172.16.0.0-172.31.255.25 5	2^20-2=1048574
192.168.0.0/16-24	192.168.0.0-192.168.255. 255	2^16-2=65534

14.3.3 Can Subnets Communicate with Each Other?

- Subnets in the same VPC can communicate with each other by default.
- VPCs are isolated from each other. Subnets from different VPCs cannot communicate with each other. You can use a VPC peering connection to enable communication between VPCs in the same region.

□ NOTE

If subnets have network ACLs associated, network ACL rules should allow communication between the subnets.

14.3.4 What Subnet CIDR Blocks Are Available?

A subnet is an IP address range from a VPC. The VPC service supports CIDR blocks 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24.

Subnets must reside within your VPC, and the subnet masks used to define them can be between .

14.3.5 How Many Subnets Can I Create?

Each account can have a maximum of 100 subnets. If the number of subnets cannot meet your service requirements, request a quota increase. For details, see **What Is a Quota?**

14.3.6 Why Can't I Delete My VPCs and Subnets?

If VPCs and subnets are being used by other resources, you need to delete these resources first based on the prompts on the console before deleting the VPCs and subnets. This following provides detailed deletion prompts and corresponding deletion guide.

- Deleting Subnets
- Deleting VPCs

Deleting Subnets

You can refer to Table 14-5 to delete subnets.

Table 14-5 Deleting subnets

Prompts	Cause	Solution
You do not have permission to perform this operation.	Your account does not have permissions to delete subnets.	Contact the account administrator to grant permissions to your account and then delete the subnet.

Prompts	Cause	Solution
Delete custom routes from the associated route table of the subnet and then delete the subnet.	The route table has custom routes with the following as the next hop type: • Server • Extension NIC • Virtual IP address • NAT gateway	Delete the custom routes from the route table and then delete the subnet. 1. Viewing the Route Table Associated with a Subnet 2. Deleting a Route
Release any virtual IP addresses configured in the subnet and then delete the subnet.	The subnet has virtual IP addresses configured.	Release the virtual IP addresses from the subnet and then delete the subnet. Releasing a Virtual IP Address
Release any private IP addresses configured in the subnet and then delete the subnet.	The subnet has virtual IP addresses that are not used by any instance.	On the IP Addresses tab, release these private IP addresses that are not required and then delete the subnet. 1. Viewing IP Addresses in a Subnet 2. In the private IP address list, locate the IP address that is not being used and click Release in the Operation column. NOTICE If you want to release an in-use private IP address, you need to delete the resource that uses the IP address first.
Delete the resource (ECS or load balancer) that is using the subnet and then delete the subnet.	The subnet is being used by an ECS or a load balancer.	Delete the ECS or load balancer and then delete the subnet. Viewing and Deleting Resources in a Subnet

Prompts	Cause	Solution
Delete the load balancer that is using the subnet and then delete the subnet.	The subnet is being used by a load balancer.	Delete the load balancer and then delete the subnet. Viewing and Deleting Resources in a Subnet
Delete the NAT gateway that is using the subnet and then delete the subnet.	The subnet is being used by a NAT gateway.	Delete the NAT gateway and then delete the subnet. Viewing and Deleting Resources in a Subnet
Delete the resource that is using the subnet and then delete the subnet.	The subnet is being used by cloud resources.	On the IP Addresses tab, view the usage of the IP address, find the resource that is using the IP address, delete the resource, and delete the subnet. 1. Viewing IP Addresses in a Subnet 2. Locate resource based on the usage of the IP address. 3. Delete the resource and then delete the subnet.

Deleting VPCs

Before deleting a VPC, ensure that all subnets in the VPC have been deleted. You can refer to **Table 14-6** to delete VPCs.

Table 14-6 Deleting VPCs

Prompts	Cause	Solution
You do not have permission to perform this operation.	Your account does not have permissions to delete VPCs.	Contact the account administrator to grant permissions to your account and then delete the VPC.

Prompts	Cause	Solution
Delete the VPC endpoint service or the route configured for the service from the VPC route table and then delete the VPC.	The VPC route table has custom routes.	Delete the custom routes and then delete the VPC. 1. In the VPC list, locate the row that contains the VPC and click the number in the Route Tables column. The route table list is displayed. 2. Deleting a Route
	The VPC is being used by a VPC endpoint service.	Search for the VPC endpoint service on the VPC endpoint service console and delete it.
This VPC cannot be deleted because it has associated resources.	The VPC is being used by the following resources: Subnet VPC peering connection Custom route table	Click the resource name hyperlink as prompted to delete the resource. • Table 14-5 • Deleting a VPC Peering Connection • Deleting a Route Table
Delete the virtual gateway that is using the VPC and then delete the VPC.	The VPC is being used by a Direct Connect virtual gateway.	On the Direct Connect console, locate the virtual gateway and delete it.
Delete the VPN gateway that is using the VPC and then delete the VPC.	The VPC is being used by a VPN gateway.	On the VPN console, locate the VPN gateway and delete it.

Prompts	Cause	Solution
Delete all custom security groups in this region and then delete this last VPC.	In the current region, this is the last VPC and there are custom security groups. NOTICE You only need to delete the custom security groups. The default security group does not affect the deletion of VPCs.	Delete all custom security groups and then delete the VPC. Deleting a Security Group
Release all EIPs in this region and then delete this last VPC.	In the current region, this is the last VPC and there are EIPs.	Release all EIPs in this region and then delete this last VPC. Unbinding an EIP from an ECS and Releasing the EIP

14.3.7 Can I Change the VPC of an ECS?

Yes.

You can click **Change VPC** in the **Operation** column on the **Elastic Cloud Server** page.

For details, see section "Changing a VPC" in the Virtual Private Cloud User Guide.

14.4 EIPs

14.4.1 What Are the Differences Between EIP, Private IP Address, and Virtual IP Address?

Different types of IP addresses have different functions.

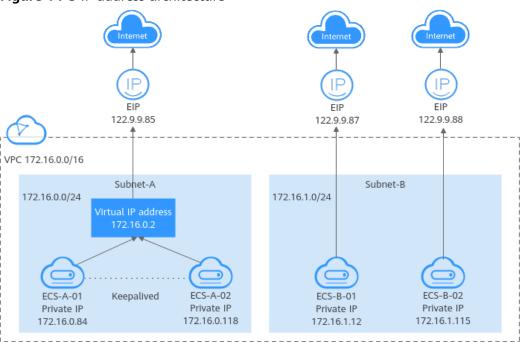


Figure 14-3 IP address architecture

Table 14-7 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
Virtual IP address	A virtual IP address can be shared among multiple ECSs. Two ECSs can work as an active and standby pair to achieve high availability by using a virtual IP address and Keepalived. If the active ECS is faulty, the virtual IP address can be dynamically switched to the standby ECS to continue providing services.	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.

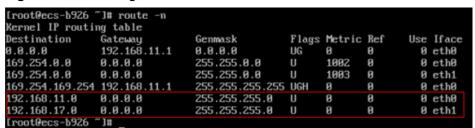
IP Address Type	Description	Example Value
EIP	 EIPs can be used by cloud resources for Internet access. You can bind an EIP to a virtual IP address to enable the ECSs with the virtual IP address bound to access the Internet. You can bind an EIP to an ECS to enable the ECS to access the Internet. Each EIP can be bound to only one ECS at a time. 	 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.

14.4.2 How Do I Access the Internet Using an EIP Bound to an Extension NIC?

1. After an EIP is bound to an extension NIC, log in to the ECS and use the **route** command to query the routes.

You can run **route --help** to learn more about the **route** command.

Figure 14-4 Viewing route information



2. Run the **ifconfig** command to view NIC information.

Figure 14-5 Viewing NIC information

```
[root@ecs-b926~]# ifconfig
eth8: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
          inet 192.168.11.42 netmask 255.255.255.0 broadcast 192.168.11.255
          inet6 fe80::f816:3eff:fef7:ic44 prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:f7:1c:44 txqueuelen 1000 (Ethernet)
          RX packets 127 bytes 21633 (21.1 KiB)
          RX errors 0 dropped 0 overruns 0 frame 0 TX packets 258 bytes 22412 (21.8 KiB)
          TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
          inet 192.168.17.191 netmask 255.255.26 broadcast 192.168.17.255
inet6 fe89::f816:3eff:fe1c:b57f prefixlen 64 scopeid 9x29<link>
ether fa:16:3e:1c:b5:7f txqueuelen 1898 (Ethernet)
          RX packets 11 bytes 1283 (1.2 KiB)
RX errors 0 dropped 0 overruns 0 frame 0
           TX packets 12 bytes 1388 (1.3 KiB)
          TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
           inet 127.0.0.1 netmask 255.0.0.0
          inet6 ::1 prefixlen 128 scopeid 0x10<host>
          loop txqueuelen 1 (Local Loopback)
RX packets 51 bytes 12018 (11.7 KiB)
          RX errors 0 dropped 0 overruns 0 frame 0
TX packets 51 bytes 12018 (11.7 KiB)
              errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

- 3. Enable access to the Internet through the extension NIC by default.
 - a. Run the following command to delete the default route of the primary NIC:

route del -net 0.0.0.0 gw 192.168.11.1 dev eth0

192.168.11.1 is the gateway of the subnet that the NIC works. You can view the gateway on the **Summary** tab page of the subnet on the management console.

Ⅲ NOTE

This operation will interrupt ECS communication. It is recommended that you perform the configuration by following step 4.

b. Run the following command to configure the default route for the extension NIC:

route add default gw 192.168.17.1

4. Configure Internet access from the extension NIC based on your destination address.

Run the following command to configure access to a specified CIDR block (for example, xx.xx.0.0/16) through the extension NIC:

You can configure the CIDR block as required.

route add -net xx.xx.0.0 netmask 255.255.0.0 gw 192.168.17.1

14.4.3 What Are the Differences Between the Primary and Extension NICs of ECSs?

The differences are as follows:

 Generally, the OS default routes preferentially use the primary NICs. If the OS default routes use the extension NICs, network communication will be interrupted. Then, you can check the route configuration to rectify the network communication error.

 Primary NICs can communicate with the public service zone (zone where PaaS and DNS services are deployed). Extension NICs cannot communicate this zone.

14.4.4 Can an EIP That Uses Dedicated Bandwidth Be Changed to Use Shared Bandwidth?

Yes. An EIP that uses a dedicated bandwidth can be changed to use a shared bandwidth.

14.4.5 Can I Bind an EIP to Multiple ECSs?

Each EIP can be bound to only one ECS at a time.

14.4.6 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. 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.

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

Yes.

You can unbind the EIP from the original ECS. For details, see **Unbinding an EIP** from an ECS and Releasing the EIP.

Then, bind the EIP to the target ECS. For details, see **Assigning an EIP and Binding It to an ECS**.

14.4.8 How Do I Unbind an EIP from an Instance and Bind a New 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. On the console homepage, under **Network**, click **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.

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

- a. Log in to the management console.
- b. On the console homepage, under **Network**, click **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. On the console homepage, under **Network**, click **Elastic IP**.
 - c. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
 - d. Select the desired ECS.
 - e. Click **OK**.
- 4. Release the EIP that is unbound.

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

- a. Log in to the management console.
- b. On the console homepage, under **Network**, click **Elastic IP**.
- 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 Networking, 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 Networking, 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.

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*.

4. Release the EIP that was replaced. For details, see 4.

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

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

No. EIPs and their associated cloud resources must be in the same region.

14.4.10 Can I Change the Region of My EIP?

The region of an EIP cannot be changed.

If you assigned an EIP in region A but need an EIP in region B, you cannot directly change the region of the assigned EIP from A to B. Instead, you have to assign an EIP in region B.

14.5 VPC Peering Connections

14.5.1 How Many VPC Peering Connections Can I Create in an Account?

If you use a VPC peering connection to connect VPCs in the same region, you can log in to the management console to view your VPC peering connection quota. For details, see **What Is a Quota?**

- Number of VPC peering connections that you can create in each region between VPCs in the same account: subject to the actual quota
- Number of VPC peering connections that you can create in each region between VPCs in different accounts: Accepted VPC peering connections use the quotas of both accounts. To-be-accepted VPC peering connections only use the quotas of accounts that request the connections.

An account can create VPC peering connections with different accounts if the account has enough quota.

14.5.2 Can a VPC Peering Connection Connect VPCs in Different Regions?

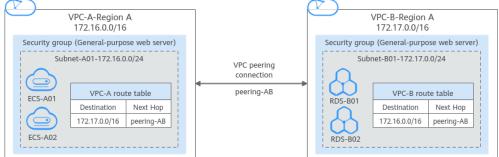
A VPC peering connection only can connect VPCs in the same region.

Figure 14-6 shows an application scenario of VPC peering connections.

- There are two VPCs (VPC-A and VPC-B) in region A that are not connected.
- Service servers (ECS-A01 and ECS-A02) are in VPC-A, and database servers (RDS-B01 and RDS-B02) are in VPC-B. The service servers and database servers cannot communicate with each other.
- You need to create a VPC peering connection (peering-AB) between VPC-A and VPC-B so the service servers and database servers can communicate with each other.

Figure 14-6 VPC peering connection network diagram

VPC-A-Region A



14.5.3 Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?

Symptom

After a VPC peering connection is created, the local and peer VPCs cannot communicate with each other.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Table 14-8 Possible causes and solutions

No.	Possible Cause	Solution
1	 Overlapping CIDR blocks of local and peer VPCs All their subnet CIDR blocks overlap. Some of their subnet CIDR blocks overlap. 	Refer to Overlapping CIDR Blocks of Local and Peer VPCs.
2	 Incorrect route configuration for the local and peer VPCs No routes are added. Incorrect routes are added. Destinations of the routes overlap with that configured for Direct Connect or VPN connections. 	Refer to Incorrect Route Configuration for Local and Peer VPCs.
3	 Incorrect network configuration The security group rules of the ECSs that need to communicate deny inbound traffic from each other. The firewall of the ECS NIC blocks traffic. The network ACL rules of the subnets connected by the VPC peering connection deny inbound traffic. Check the policy-based routing configuration of an ECS with multiple NICs. 	Refer to Incorrect Network Configuration.
4	ECS network failure	Refer to ECS Network Failure.

Overlapping CIDR Blocks of Local and Peer VPCs

If the CIDR blocks of VPCs connected by a VPC peering connection overlap, the connection may not take effect due to route conflicts.

Table 14-9 Overlapping CIDR blocks of local and peer VPCs

Scenario	Description	Solution
VPCs with overlapping CIDR blocks also include subnets that overlap.	As shown in Figure 14-7, the CIDR blocks of VPC-A and VPC-B overlap, and all their subnets overlap. Overlapping CIDR blocks of VPC-A and VPC-B: 10.0.0.0/16 Overlapping CIDR blocks of Subnet-A01 in VPC-A and Subnet-B01 in VPC-B: 10.0.0.0/24 Overlapping CIDR blocks	VPC-A and VPC-B cannot be connected using a VPC peering connection. Replan the network.
	of Subnet-A02 in VPC-A and Subnet-B02 in VPC-B: 10.0.1.0/24	
Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.	As shown in Figure 14-8, the CIDR blocks of VPC-A and VPC-B overlap, and some of their subnets overlap. Overlapping CIDR blocks of VPC-A and VPC-B: 10.0.0.0/16 Overlapping CIDR blocks of Subnet-A01 in VPC-A and Subnet-B01 in VPC-B: 10.0.0.0/24 CIDR blocks of Subnet-A02 in VPC-B do not overlap.	 A VPC peering connection cannot connect the entire VPCs, VPC-A and VPC-B. A connection can connect their subnets (Subnet-A02 and Subnet-B02) that do not overlap. For details, see Figure 14-9.

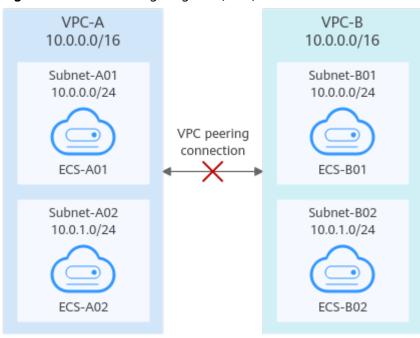
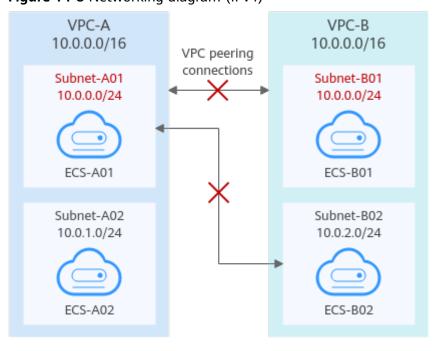


Figure 14-7 Networking diagram (IPv4)

Figure 14-8 Networking diagram (IPv4)



If CIDR blocks of VPCs overlap and some of their subnets overlap, you can create a VPC peering connection between their subnets with non-overlapping CIDR blocks. **Figure 14-9** shows the networking diagram of connecting Subnet-A02 and Subnet-B02. **Table 14-10** describes the routes required.

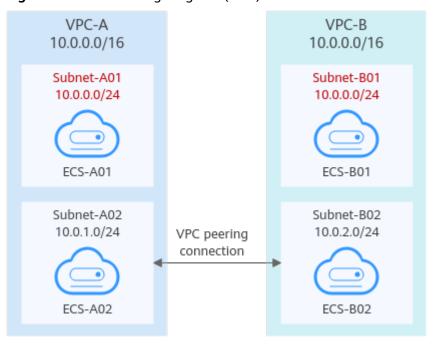


Figure 14-9 Networking diagram (IPv4)

Table 14-10 Routes required for the VPC peering connection between Subnet-A02 and Subnet-B02

Route Table	Destinat ion	Next Hop	Description
VPC-A route table	10.0.2.0/ 24	Peering- AB	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
VPC-B route table	10.0.1.0/ 24	Peering- AB	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AB as the next hop.

Incorrect Route Configuration for Local and Peer VPCs

Check the routes in the route tables of the local and peer VPCs by referring to **Viewing Routes Configured for a VPC Peering Connection**. **Table 14-11** lists the items that you need to check.

Table 14-11 Route check items

Item	Solution
Check whether routes are added to the route tables of the local	If routes are not added, add routes by referring to:
and peer VPCs.	 Creating a VPC Peering Connection with Another VPC in Your Account

Item	Solution
Check the destinations of routes added to the route tables of the local and peer VPCs.	If the route destination is incorrect, change it by referring to Modifying Routes Configured for a VPC Peering Connection .
 In the route table of the local VPC, check whether the route destination is the CIDR block, subnet CIDR block, or related private IP address of the peer VPC. 	
 In the route table of the peer VPC, check whether the route destination is the CIDR block, subnet CIDR block, or related private IP address of the local VPC. 	
Destinations of the routes overlap with that configured for Direct Connect or VPN connections.	Check whether any of the VPCs connected by the VPC peering connection also has a VPN or Direct Connect connection connected. If they do, check the destinations of their routes.
	If the destinations of the routes overlap, the VPC peering connection does not take effect. In this case, replan the network connection.

Incorrect Network Configuration

- Check whether security group rules of the ECSs that need to communicate allow inbound traffic from each other by referring to Viewing the Security Group of an ECS.
 - If the ECSs are associated with the same security group, you do not need to check their rules.
 - If the ECSs are associated with different security groups, add an inbound rule to allow access from each other by referring to Security Group Configuration Examples.
- Check whether the firewall of the ECS NIC blocks traffic.
 If the firewall blocks traffic, configure the firewall to allow inbound traffic.
- 3. Check whether network ACL rules of the subnets connected by the VPC peering connection deny inbound traffic.
 - If the network ACL rules deny inbound traffic, configure the rules to allow the traffic.
- 4. If an ECS has more than one NIC, check whether correct policy-based routing has been configured for the ECS and packets with different source IP addresses match their own routes from each NIC.
 - If an ECS has two NICs (eth0 and eth1):

- IP address of eth0: 192.168.1.10; Subnet gateway: 192.168.1.1
- IP address of eth1: 192.168.2.10; Subnet gateway: 192.168.2.1

Command format:

- ping -l /P address of eth0 Subnet gateway address of eth0
- ping -l IP address of eth1 Subnet gateway address of eth1

Run the following commands:

- ping -I 192.168.1.10 192.168.1.1
- ping -I 192.168.2.10 192.168.2.1

If the network communication is normal, the routes of the NICs are correctly configured.

Otherwise, you need to configure policy-based routing for the ECS with multiple NICs by referring to How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?

ECS Network Failure

- 1. Log in to the ECS.
- 2. Check whether the ECS NIC has an IP address assigned.
 - Linux ECS: Use the **ifconfig** or **ip address** command to view the IP address of the NIC.
 - Windows ECS: In the search box, enter cmd and press Enter. In the displayed command prompt, run the ipconfig command.

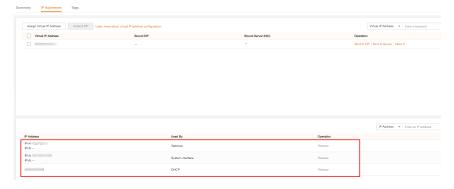
If the ECS NIC has no IP address assigned, see

- 3. Check whether the subnet gateway of the ECS can be pinged.
 - a. In the ECS list, click the ECS name.
 - The ECS details page is displayed.
 - b. On the ECS details page, click the hyperlink of VPC.
 - The Virtual Private Cloud page is displayed.
 - c. In the VPC list, locate the target VPC and click the number in the Subnets column.

The **Subnets** page is displayed.

- d. In the subnet list, click the subnet name.
 - The subnet details page is displayed.
- e. Click the **IP Addresses** tab and view the gateway address of the subnet.

Figure 14-10 Gateway address



f. Check whether the gateway communication is normal: ping Subnet gateway addressExample command: ping 172.17.0.1

14.6 Virtual IP Addresses

14.6.1 Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?

Symptom

After you bind a virtual IP address to an ECS NIC, you cannot ping the virtual IP address.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 14-11 Troubleshooting

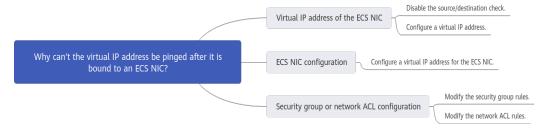


Table 14-12 Troubleshooting

Possible Cause	Solution
Virtual IP address of the ECS NIC	See Virtual IP Address of the ECS NIC
Virtual IP address of the internal NIC of the ECS	See Virtual IP Address of the Internal NIC of the ECS
Security group or network ACL configuration	See Security Group or Network ACL Configuration

Virtual IP Address of the ECS NIC

Check whether the source/destination check of the NIC is disabled and whether a virtual IP address is bound to the NIC.

1. Log in to the management console.

- 2. Click Service List and click Elastic Cloud Server under Computing.
- 3. In the ECS list, click the name of the ECS.
- 4. On the displayed ECS details page, click the **Network Interfaces** tab.
- 5. Ensure that **Source/Destination Check** is disabled.
- 6. Ensure that an IP address is displayed for **Virtual IP Address** on the NIC details page.

If there is no virtual IP address, click **Manage Virtual IP Address**. On the displayed **IP Addresses** tab, click **Assign Virtual IP Address**.

■ NOTE

To check whether a virtual IP address has been configured, **ifconfig** will not work. Use **ip address** instead. For more information, see .

Virtual IP Address of the Internal NIC of the ECS

The following uses Linux and Windows ECSs as examples to describe how to check whether an ECS NIC has a virtual IP address.

For a Linux ECS:

 Check if there is a NIC eth X:X: ifconfig

Figure 14-12 Checking for NIC ethX:X

```
[root@scy ~] # ifconfig
eth0: flags=4163<UP, BROADCAST, RUNNING, MULTICAST> mtu 1500
    inet 192.168.1.2 netmask 255.255.255.0 broadcast 192.168.1.255
    inet6 fe80::f816:3eff:fe4d:5b98 prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
    RX packets 77399 bytes 5101164 (4.8 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 68798 bytes 8090922 (7.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth0:1: flags=4163<UP, BROADCAST, RUNNING, MULTICAST> mtu 1500
    inet 192.168.1.137 netmask 255.255.255.0 broadcast 192.168.1.255
    ether fa:10:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
```

The command output in the preceding figure contains a NIC **eth***X:X.* **192.168.1.137** is its virtual IP address.

- If the NIC **eth***X:X* is there, the ECS NIC is correctly configured.
- If the NIC **eth** *X:X* cannot be found, perform the following operations:
- If the command output does not contain a NIC eth X:X, switch to the /etc/ sysconfig/network-scripts directory:

cd /etc/sysconfig/network-scripts

3. Run the following command to create and then modify the **ifcfg-eth0:1** file:

vi ifcfq-eth0:1

Add the following NIC information to the file:

BOOTPROTO=static DEVICE=eth0:1 HWADDR=fa:16:3e:4d:5b:98

```
IPADDR=192.168.1.137
GATEWAY=192.168.1.1
NETMASK=255.255.255.0
ONBOOT=yes
ONPARENT=yes
```

- 4. Press **Esc**, enter :wq!, and save the file and exit.
- 5. Restart the ECS and run the **ifconfig** command to check whether the virtual IP address has been configured for the ECS.

For a Windows ECS:

1. In the **Start** menu, open the Windows command line window and run the following command to check whether the virtual IP address has been configured:

ipconfig /all

Figure 14-13 Checking whether the virtual IP address has been configured

```
\Users\Administrator>ipconfig /all
Windows IP Configuration
   Host Name .
                                    . . . : dst-win
   Primary Dns Suffix . . . . . .
  : Hybrid
   WINS Proxy Enabled. . . .
Ethernet adapter Ethernet 5:
   Connection-specific DNS Suffix .:
   Description . . . . . . . . : Red Hat VirtIO Ethernet Adapter #2
Physical Address . . . . . . . : FA-16-3E-83-B2-73
  DHCP Enabled. . . . . . . . . No
Autoconfiguration Enabled . . . : Yes
Link-local IPv6 Address . . . : fe80::6182:a265:10bc:134e%3(Preferred)
IPv4 Address . . . . : 192.168.10.41(Preferred)
                                            : 255.255.255.0
   Subnet Mask . . . . . . . . . . . . .
  IPv4 Address. . . . . . . . . . : 192.168.10.137(Preferred)
  Subnet Mask . . . . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.10.1
   DHCPv6 IAID .
                                              184161854
   DHCPv6 Client DUID. . . .
                                       . . : 00-01-00-01-21-9F-1A-85-52-54-00-A6-AD-AC
                                            : 100.125.1.250
   DNS Servers . . . .
                                               114.114.114.114
   NetBIOS over Tcpip.
```

In the preceding command output, check whether the value of **IPv4 Address** (192.168.10.137) is the IP address of the ECS NIC.

- If yes, the virtual IP address has been configured for the ECS NIC.
- If no, perform the following operations:
- Choose Control Panel > Network and Internet > Network Connections.
 Right-click the corresponding local connection and then click Properties.
- 3. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.
- 4. Click **Properties**.
- 5. Select **Use the following IP address**, and set **IP address** to the private IP address displayed in **Figure 14-13**. For example, 192.168.10.41.

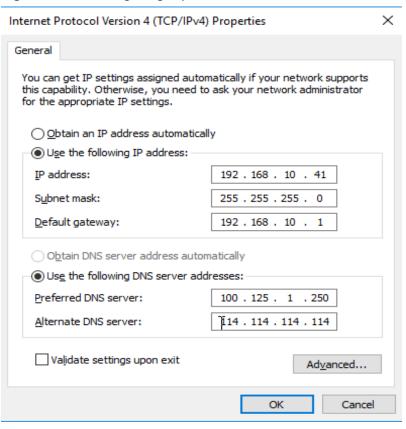


Figure 14-14 Configuring a private IP address

- 6. Click **Advanced**.
- 7. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address configured in **Figure 14-13**. For example, 192.168.10.137.

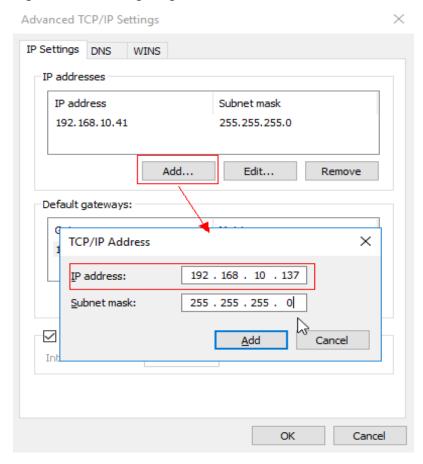


Figure 14-15 Configuring virtual IP address

Security Group or Network ACL Configuration

Check whether the ECS security groups and the network ACLs associated with the subnet used by the ECS NIC are blocking traffic.

- On the ECS details page, click the Security Groups tab and confirm that required security group rules have been configured for the virtual IP address. If the required security group rules have not been configured, click Change Security Group or Modify Security Group Rule to change the security group or modify the security group rules.
- Click Service List. Under Network, click Virtual Private Cloud. In the navigation pane on the left of the network console, click Network ACLs and check whether the network ACL rules associated with the subnet used by the ECS NIC are blocking access to the virtual IP address.

14.6.2 Why Is the Network Disconnected Between Servers Using a Virtual IP Address After an Active/Standby Switchover?

For an HA cluster using virtual IP addresses and Keepalived, if you find that the network between the client and the server is disconnected after an active/standby switchover, the possible cause is that the switchover is performed manually. As a result, the ARP table on the client is not updated, you can perform the following operations to update the ARP table:

- 1. Log in to the client.
- 2. Update the ARP table on the client.
 - Method 1: Trigger the client to learn the new MAC address corresponding to the virtual IP address:

ping Virtual IP address

Example command: ping 192.168.3.22

 Method 2: Clear the residual entries in the ARP table of the virtual IP address to trigger the client to learn the new ARP table:

arp -d Virtual IP address

Command example: arp -d 192.168.3.22

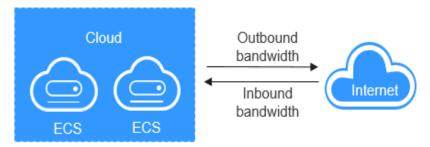
14.7 Bandwidth

14.7.1 What Are 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, that consumes outbound bandwidth.

Figure 14-16 Inbound bandwidth and outbound bandwidth



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

Symptom

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

■ NOTE

If the outbound bandwidth exceeds the configured bandwidth size, there may be packet loss. To prevent data loss, it is recommended that you monitor the bandwidth.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 14-17 Troubleshooting

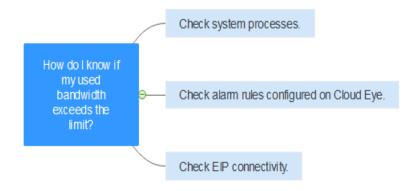


Table 14-13 Troubleshooting

Possible Cause	Description	Solution
System processes leading to high bandwidth	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.	See System Processes Leading to High Bandwidth Usage
Improper Cloud Eye alarm rules	If you have created 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.	See Improper Cloud Eye Alarm Rules

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 the processes that have led to excessively high bandwidth or CPU usage, and optimize or stop the processes.

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

Improper Cloud Eye Alarm Rules

If you have created 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 when the maximum outbound bandwidth reaches 4.8 Mbit/s three periods in a row. You can also **Modifying an EIP Bandwidth**.

- Log in to the management console, under Management & Deployment, click Cloud Eye. On 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 exceeds the configured limit.

14.7.3 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, that 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.

14.7.4 What Is the Bandwidth Size Range?

The bandwidth range is from 1 Mbit/s to 2000 Mbit/s.

If your current bandwidth cannot meet your requirements, you can increase it.

Procedure

Currently, a bandwidth that is greater than 2000 Mbit/s cannot be adjusted on the console.

If you need to adjust a quota, contact the operations administrator.

14.7.5 What Bandwidth Types Are Available?

There are dedicated bandwidths and shared bandwidths. A dedicated bandwidth can only be used by one EIP, but a shared bandwidth can be used by multiple EIPs.

14.7.6 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth?

A dedicated bandwidth can only be used by one EIP that is bound to one cloud resource, such as an ECS, a NAT gateway, or a load balancer.

A shared bandwidth can be shared by multiple EIPs. Adding an EIP to or removing an EIP from a shared bandwidth does not affect your services.

A dedicated bandwidth cannot be changed to a shared bandwidth or the other way around. You can purchase a shared bandwidth for your EIPs.

- After you add an EIP to a shared bandwidth, the EIP will use the shared bandwidth.
- After you remove an EIP from a shared bandwidth, the EIP will use the dedicated bandwidth.

14.7.7 Is There a Limit to the Number of EIPs That Can Be Added to Each Shared Bandwidth?

A maximum of 20 EIPs can be added to each shared bandwidth. If you want to add more EIPs to each shared bandwidth, request a quota increase. For details, see **What Is a Quota?**

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

The bandwidth is measured in bit/s, but the download rate is measured in byte/s.

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).

14.8 Connectivity

14.8.1 Does a VPN Allow Communication Between Two VPCs?

If the two VPCs are in the same region, you can use a VPC peering connection to enable communication between them.

If the two VPCs are in different regions, you can use a VPN to enable communication between the VPCs. The CIDR blocks of the two VPCs are the local and remote subnets, respectively.

14.8.2 Why Are Internet or Internal Domain Names in the Cloud Inaccessible Through Domain Names When My ECS Has Multiple NICs?

When an ECS has more than one NIC, if different DNS server addresses are configured for the subnets used by the NICs, the ECS cannot access the Internet or domain names in the cloud.

You can resolve this issue by configuring the same DNS server address for the subnets used by the same ECS. You can perform the following steps to modify DNS server addresses of subnets in a VPC:

- 1. Log in to the management console.
- 2. On the console homepage, under **Network**, click **Virtual Private Cloud**.
- 3. In the subnet list, locate the target subnet and click its name.
- 4. On the subnet details page, change the DNS server address of the subnet.
- 5. Click OK.

14.8.3 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.

14.8.4 Why Are There Intermittent Interruptions When a Local Host Accesses a Website Built on an ECS?

Symptom

After you build a website on an ECS, some users occasionally are unable to access the website through the local network.

Troubleshooting

- 1. Check the local network of the user.
 - If the local host communicates with the ECS using NAT, this problem may occur.
- 2. Run the following command to check whether **tcp_tw_recycle** is enabled on the ECS:

sysctl -a|grep tcp_tw_recycle

If the value of **tcp_tw_recycle** is **1**, the function is enabled.

3. Run the following command to check the number of lost packets of the ECS:

cat /proc/net/netstat | awk '/TcpExt/ { print \$21,\$22 }'

If the value of **ListenDrops** is not **0**, there is packet loss, that is, the network is faulty.

Procedure

This problem can be solved by modifying the kernel parameters of the ECS.

• Run the following command to temporarily modify the parameters (the parameters will change back after a restart):

sysctl -w net.ipv4.tcp_tw_recycle=0

- Perform the following operations to permanently modify the parameters:
 - a. Run the following command and modify the /etc/sysctl.conf file:

vi /etc/sysctl.conf

Add the following content to the file: net.ipv4.tcp_tw_recycle=0

- b. Press **Esc**, enter :wq!, and save the file and exit.
- c. Run the following command to make the modification take effect:sysctl -p

14.8.5 Why Do ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communication?

Symptom

Two ECSs (ecs01 and ecs02) are in the same subnet in a VPC. Their IP addresses are 192.168.1.141 and 192.168.1.40.

The **ecs01** can ping **ecs02** through a private IP address successfully, but **ecs02** cannot ping **ecs01** through a private IP address.

Troubleshooting

- Ping ecs01 from ecs02 through the EIP. If ecs01 can be pinged, the NIC of ecs01 is working properly.
- 2. Run the arp -n command on ecs02 to check whether the command output contains the MAC address of ecs01. If the command output does not contain the MAC address of ecs01, ecs02 fails to learn the MAC address of ecs01 when using the private IP address to ping ecs01.
- 3. Run the **ip a** command on **ecs01** to check the NIC configuration of **ecs01**. The following figure shows an example.

Figure 14-18 Viewing ecs01 NIC configuration

The IP address 192.168.1.40/32 should not be configured based on the command output. As a result, **ecs01** fails to send packets to **ecs02**.

Procedure

Modify the NIC configuration of **ecs01**. Run the following command to delete the redundant IP address, for example, 192.168.1.40/32, configured on the NIC **eth0**:

ip a del 192.168.1.40/32 dev eth0

14.8.6 Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur When They Communicate?

Symptom

Two ECSs in the same VPC cannot communicate with each other or there is packet loss when they communicate.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 14-19 Troubleshooting

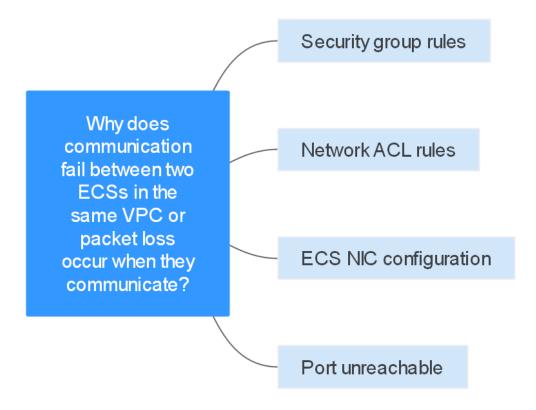


Table 14-14 Troubleshooting

Possible Cause	Solution
Security group rules	See Security Group Rules
Network ACL rules	See Network ACL Rules
ECS NIC configuration	See ECS NIC Configuration
Port unreachable	See Port Unreachable

Security Group Rules

Check whether the ECS NIC security group allows the outbound and inbound ICMP traffic.

Take the inbound direction as an example. The security group rules must contain at least one of the following rules.

Figure 14-20 Inbound security group rule



If packets of other protocols are tested, configure the security group rules to allow the corresponding protocol traffic. For example, if UDP packets are tested, check whether the security group allows the inbound UDP traffic.

Network ACL Rules

- Check whether the subnet of the ECS NIC has an associated network ACL.
- 2. Check the network ACL status in the network ACL list.
 - If **Disabled** is displayed in the **Status** column, the network ACL has been disabled. Go to 3.
 - If Enabled is displayed in the Status column, the network ACL has been enabled. Go to 4.
- 3. Click the network ACL name and configure rules on the **Inbound Rules** and **Outbound Rules** tabs to allow the ICMP traffic.
- 4. If the network ACL is disabled, all packets in the inbound and outbound directions are discarded by default. In this case, delete the network ACL or enable the network ACL and allow the ICMP traffic.

ECS NIC Configuration

The following procedure uses a Linux ECS as an example. For a Windows ECS, check the firewall restrictions.

- Check whether multiple NICs are configured for the ECS. If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policybased routing for the ECS. For details, see How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?
- Log in to the ECS and run the following command to check whether the NIC
 has been created and obtained a private IP address. If there is no NIC
 information or the private IP address cannot be obtained, contact technical
 support.

ifconfig

Figure 14-21 NIC IP address

```
Iroot@ecs-acl Tlm ifconfig

Link encap:Ethernet HWaddr FA:16:3E:BC:B7:81
inet addr 192.168.72.289 Bcast:192.168.72.255 Mask:255.255.8
inet6 addr: fe88::f816:3eff:febc:b781/64 Scope:Link
UP BROADCAST RUNNING MULTICAST MTU:1580 Metric:1
RX packets:881 errors:8 dropped:8 overruns:8 frame:8
TX packets:547 errors:8 dropped:8 overruns:8 carrier:8
collisions:8 txqueuelen:1880
RX bytes:49684 (48.4 KiB) TX bytes:44454 (43.4 KiB)
Interrupt:46
```

3. If the CPU usage exceeds 80%, the ECS communication may be adversely affected. Run the following command to check whether the CPU usage of the ECS is too high:

top

4. Run the following command to check whether the ECS has any restrictions on security group rules:

iptables-save

5. Run the following command to check whether the **/etc/hosts.deny** file contains the IP addresses that limit communication:

vi /etc/hosts.deny

If the **hosts.deny** file contains the IP address of another ECS, delete the IP address from the **hosts.deny** file and save the file.

Port Unreachable

- 1. If a port of the ECS cannot be reached, check whether the security group rules and network ACL rules enable the port.
- 2. On the Linux ECS, run the following command to check whether the ECS listens on the port: If the ECS does not listen on the port, the ECS communication may be adversely affected.

netstat -na | grep < Port number>

14.8.7 Why Can't My ECS Use Cloud-init?

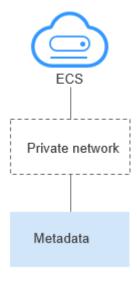
Symptom

An ECS is not able to use cloud-init.

Troubleshooting

Figure 14-22 shows the process for an ECS to obtain metadata using the cloud-init.

Figure 14-22 Process for obtaining metadata



Check the following possible causes.

The ECS has no IP address obtained.

Obtain an IP address.

Incorrect route for 169.254.169.254

Fail to obtain the ECS metadata.

Fail to log in to the ECS or create a non-root user after cloud-init is configured.

Check the format of the /etc/cloud/cloud.cfg configuration file.

Fail to use an obtained private key to log in to an ECS after the ECS starts (Fail to obtain the ECS login password).

Figure 14-23 Possible causes

Table 14-15 Possible causes

Possible Cause	Solution
The ECS has no IP address obtained.	See The ECS Has Not Obtained IP Address
Incorrect route for 169.254.169.254	See Incorrect Route for 169.254.169.254
Fail to obtain the ECS metadata.	See Did Not Obtain the ECS Metadata
Fail to log in to the ECS or create a non-root user after cloud-init is configured.	Check the format of the /etc/cloud/cloud.cfg configuration file. For details, see Cannot Log in to the ECS or Create a Non-root User After Cloud-init Is Configured.
Fail to use an obtained private key to log in to an ECS after the ECS starts (Fail to obtain the ECS login password).	Restart the ECS and try again.

The ECS Has Not Obtained IP Address

Check whether the ECS has obtained an IP address.

If no IP address is obtained, run the **dhclient** command to obtain the IP address (this command varies depending on the ECS OSs). Alternatively, you can run the **ifdown** *ethx* command to disable the network port and then run the **ifup** *ethx* command to enable it to allow the ECS NIC to automatically obtain an IP address again.

Figure 14-24 ECS IP address

```
-bash-4.1# ifconfig
eth0
          Link encap:Ethernet HWaddr FA:16:3E:BD:36:DD
          inet addr:192.168.1.200 Bcast:192.168.1.255 Mask:255.255.255.0
           inet6 addr: fe80::f816:3eff:febd:36dd/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
          RX packets:73008 errors:0 dropped:0 overruns:0 frame:0
          TX packets:26295 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
           RX bytes:4162713 (3.9 MiB) TX bytes:2336476 (2.2 MiB)
           Interrupt:35
eth1
          Link encap:Ethernet HWaddr FA:16:3E:A9:C7:1D
           inet addr:192.168.1.179 Bcast:192.168.1.255 Mask:255.255.255.0
           inet6 addr: fe80::f816:3eff:fea9:c71d/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST MTU:1500 Metric:1
          RX packets:45026 errors:0 dropped:0 overruns:0 frame:0
          TX packets:12244 errors:0 dropped:0 overruns:0 carrier:0 collisions:0 txqueuelen:1000
          RX bytes:1270534 (1.2 MiB) TX bytes:4178924 (3.9 MiB)
           Interrupt:34
lo
          Link encap:Local Loopback
           inet addr:127.0.0.1 Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
UP LOOPBACK RUNNING MTU:65536 Metric:1
RX packets:1 errors:0 dropped:0 overruns:0 frame:0
           TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:0
           RX bytes:28 (28.0 b) TX bytes:28 (28.0 b)
```

Incorrect Route for 169.254.169.254

Ping **169.254.169.254/32** from the ECS. If the IP address cannot be pinged, perform the following steps:

1. Check the exact route configured on the ECS for IP address **169.254.169.254/32**.

In most cases, the next hop of the exact route for IP address **169.254.169.254/32** is the same as that of the default route for the IP address.

Figure 14-25 Route for IP address 169.254.169.254/32

```
-bash-4.1# ip route
169.254.169.254 via 192.168.1.1 dev eth0 proto static
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#
```

2. If there is no exact route for IP address **169.254.169.254/32**, the cause is as follows:

Images with CentOS 5 OSs are not compatible with cloud-init. To use cloud-init, select a different OS.

- 3. If the next hop of the exact route for IP address **169.254.169.254/32** is different from that of the default route for the IP address:
 - If the ECS was created before cloud-init was enabled, run service network restart to obtain the correct route.

Did Not Obtain the ECS Metadata

Run the following command on the ECS to obtain the metadata:

curl http://169.254.169.254/openstack/latest/meta_data.json

If information similar to that shown in **Figure 14-26** is displayed, the ECS successfully obtains the metadata.

Figure 14-26 Command output

```
-bash-4.1# curl http://169.254.169.254/openstack/latest/meta_data.json
{"random_seed": "rTVrsD1Eh6AjUKLnvq51U8S0pH6xC78MFRTeWIOmunBNyqos6q/EsAEJondF8iJkMDGOTzbCTb815HNtS9X
kHu61u+yBfAeyUkxAj6OAv8KHMFgDv6xDfhKu6qyjCrjXn5hUFvqfZ/yaJ3LrAEjB8Nj59h1+umbFiB0Yc2WzYmTqWjXYRNvpmqJM
sIKYMOCLuFbwYo2aK1y27/WEVZDV8QlGpRkkuMwFaCN/rQQ/hHd+3UwSJbArsqVeoWCTp5oxixLiCJzSSHAKz41UiZiRxuYwmBgo
iTFtopvZTvmYEk1FmkZsy7h6FPOkgmjgpn+1kZf9qqht1vpyRr2pUAaFAeZa4z7QX1RtmwJ77MUyGlbea85/IPDUE1J/GJpoH1/+z
rDye1A09CsOG1VFuELadYDcrWA4k42f0o7dDmEyDmINnE8eeqa5r7EohbO4KTimzi+3nb10QjPq/S7J+mFM/UoZEJH0bZE4uw1Aj
2nhvy/pc6ho7fQKbx8C78fbiPh59CKyFOWB35nNJ/CZNHBTd3UdG25SQ701FnA+NtDbeo8+gB5iFLvWewwBG5BLcjmffjh9+mqot
45ae6ZcexUsIff scqm8jwCnCimthJIYGmbxu+6Fm9xpLDopDFrRtBUcRSNt1K67JprBSRppc+4sMyyiuKYIJOTUJYQYDBUZB7F30
=", "uuid": "53ebb737-ddc5-4303-9fac-aar72b0Db1D1a", "availability_zone": "eu-de-82", "hostname": "ec
s-gjm-55eb.novaloca1", "launch_index": 0, "meta": {"metering.image_id": "98721f93-722f-4386-a975-3cb
df1abf56d", "metering.imagetype": "gold", "metering.resourcespeccode": "c2.large.oracle", "metering.
cloudServiceType": "sys.service.type.ec2", "image_name": "AutoC_OTC_OEL_6.8", "metering.resourcetype
": "1", "os_bit": "64", "vpc_id": "120b71c7-94ac-45b8-8ed6-30aafc8fbdba", "os_type": "Linux", "charg
ing_mode": "8"}, "project_id": "efdf974f549b4eaab05c3903ddd2ab0e", "name": "ecs-gjm-55eb"}-bash-4.1#
```

Cannot Log in to the ECS or Create a Non-root User After Cloud-init Is Configured

Check whether the **/etc/cloud/cloud.cfg** configuration file format is correct. For details, see the file format requirements for different Linux distributions. The following figure shows an example **/etc/cloud/cloud.cfg** configuration file for Ubuntu.

Figure 14-27 Configuration file

```
# This will affect which distro class gets used
distro: rhel
# Default user name + that default users groups (if added/used)
default_user:
  name: Iinux // Specifies the username for login.
  lock_passwd: False // The value False indicates that the password login mode is enabled. For some OSs, the value 0 indicates that the password login mode is enabled.
  gecos: Cloud User
  groups: USerS //Specifies whether the user will be added to a group. This parameter is optional. The groups parameter value must be an existing group under /etc/group in the system
  passwd: $6$163DBVXK$Zh41chiJR7NuZvtJHsYBQJIg5RoQCRL51X2Hsgj2s5JwXI7KUO1we8WYcwbzea52VNpRnNo28vmxxCyU6LwoD0
  sudo: ["ALL=(ALL) NOPASSWD:ALL"] // Specifies that all permissions of user root will be granted to the use
  shell: /bin/bash // Specifies that the bash shell is used.
# Other config here will be given to the distro class and/or path classes
paths:
   cloud_dir: /var/lib/cloud/
   templates_dir: /etc/cloud/templates/
ssh_svcname: sshd
```

Obtained Private Key Cannot Be Used to Log in to an ECS After the ECS Starts (Failed to Obtain the ECS Login Password)

Restart the ECS to rectify the fault.

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

Symptom

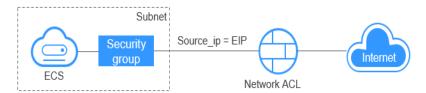
An ECS with an EIP bound cannot access the Internet.

Troubleshooting

Checking EIP Connectivity

Figure 14-28 shows the networking diagram for an ECS to access the Internet using an EIP.

Figure 14-28 EIP network diagram



Locate the fault based on the following procedure.



Figure 14-29 Troubleshooting procedure

- 1. Step 1: Check Whether the ECS Is Running Properly
- 2. Step 2: Check Whether the Network Configuration of the ECS Is Correct
- 3. Step 3: Check Whether an EIP Has Been Assigned and Bound to the ECS
- 4. Step 4: Check Whether an EIP Is Bound to the Primary NIC of the ECS
- 5. Step 5: Check Whether Required Security Group Rules Have Been Configured.
- 6. Step 6: Check Whether Traffic from the ECS Subnet Is Blocked

Step 1: Check Whether the ECS Is Running Properly

Check the ECS status.

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

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

1. Check whether the ECS NIC has an IP address assigned.

Log in to the ECS, and run **ifconfig** or **ip address** to check the ECS NIC IP address.

If the ECS runs Windows, run ipconfig.

2. Check whether the ECS NIC has a virtual IP address.

Log in to the ECS, and run **ifconfig** or **ip address** to check whether the ECS NIC has a virtual IP address. If the ECS NIC has no virtual IP address, run the **ip addr add** *virtual IP address* **eth0** command to configure an IP address for the ECS NIC.

Figure 14-30 Virtual IP address of a NIC

```
[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:37:75:62 brd 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 NIC has a default route. If there is no default route, run **ip route add** to add one.

Figure 14-31 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 ECS

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

Step 4: Check Whether an EIP Is Bound to the Primary NIC of the ECS

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

You can view the NIC details by clicking the **NICs** tab on the ECS details page. By default, the first record in the list is the primary NIC.

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 ECS Subnet Is Blocked

Check whether the network ACL of the NIC subnet blocks certain traffic from the subnet.

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

14.8.9 Why Does My ECS Fail to Obtain an IP Address?

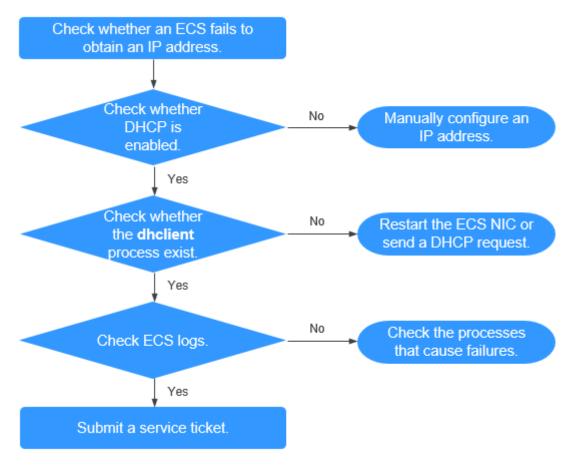
Symptom

The private IP address of the ECS fails to be obtained.

Troubleshooting

Locate the fault based on the following procedure.

Figure 14-32 Troubleshooting process



- 1. Checking Whether DHCP Is Enabled
- 2. Checking Whether the dhclient Process Exist
- 3. Checking ECS Logs

Checking Whether DHCP Is Enabled

Check whether the DHCP function of the subnet is enabled (enabled by default).

Switch to the subnet details page. If DHCP is disabled, you must manually configure a static IP address for the ECS by referring to step 3.

Checking Whether the dhclient Process Exist

1. Check whether the **dhclient** process exists:

ps -ef | grep dhclient

- 2. If the **dhclient** process does not exist, log in to the ECS and restart the ECS NIC or send a DHCP request.
 - Linux:

Run the **dhclient ethx** command. If **dhclient** commands are supported, run the **ifdown ethx;ifup ethx** command. In the command, *ethx* indicates the ECS NIC, for example, **eth0** and **eth1**.

- Windows:

Disconnect the network connection and connect it.

- 3. If the DHCP client does not send requests for a long time, for example, the fault occurs again after the NIC restarts, you can use the following method to configure the static IP address.
 - Linux:
 - i. Run the following command to open the /etc/sysconfig/networkscripts/ifcfg-eth0 file:
 - vi /etc/sysconfig/network-scripts/ifcfg-eth0
 - ii. Modify the following configuration items in the /etc/sysconfig/ network-scripts/ifcfg-eth0 file.

BOOTPROTO=static

IPADDR=192.168.1.100 #IP address

NETMASK=255.255.255.0 #Subnet mask

GATEWAY=192.168.1.1 #Gateway address

- iii. Run the following command to restart the network service: service network restart
- Windows:

On the Local Area Connection Status tab, click Properties. In the displayed area, select Internet Protocol Version 4 (TCP/IPv4) and click Properties. In the displayed area, enter the IP address, subnet mask, and the default gateway address.

Checking ECS Logs

Check the ECS messages log in the /var/log/messages directory.

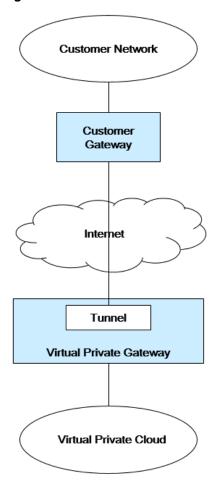
Search for the NIC MAC address and check whether there are any processes causing failures in obtaining IP addresses over DHCP.

14.8.10 How Do I Handle a VPN or Direct Connection Network Failure?

VPN Network

Figure 14-33 shows your network, the customer gateway, the VPN, and the VPC.

Figure 14-33 VPN network



Customer Self-Check Guidance

1. Provide your network information.

Obtain information listed in **Table 14-16**. This table lists example values. You can determine the actual values based on the example values. You must obtain all actual values of your project.

◯ NOTE

You can print this table and fill in your values.

Item Description Example Valu e VPC CIDR block Required for customer Example: N/A gateway configuration 10.0.0.0/16 VPC ID N/A N/A N/A CIDR block of subnet N/A Example: N/A 1 (can be the same 10.0.1.0/24 as the VPC CIDR block) **ECS ID** N/A N/A N/A N/A Customer gateway N/A N/A type (for example, Cisco) Public IP address N/A The value must N/A

Table 14-16 Network information

2. Provide your gateway configuration information.

used by the customer

gateway

You can check the gateway connectivity issues based on the following steps: You must take the IKE, IPsec, ACL rules, and route selection into consideration.

be static.

You can rectify the failure in any desired sequence. However, it is recommended that you check for the failure in the following sequence: IKE, IPsec, ACL rules, and route selection.

- a. Obtain the IKE policy used by your gateway device.
- b. Obtain the IPsec policy used by your gateway device.
- c. Obtain the ACL rule used by your gateway device.
- d. Check whether your gateway device can communicate with the gateway devices on the cloud.

□ NOTE

The commands used on different gateway devices are different. You can run the commands based on your gateway device (such as Cisco, H3C, AR, or Fortinet device) to obtain the preceding required information.

O&M Operations That Require Assistance

You must send communication requests from the ECSs to the remote device.

Method:

Log in to an ECS and ping an IP address in your on-premises data center.

14.8.11 Why Can My Server Be Accessed from the Internet But Cannot Access the Internet?

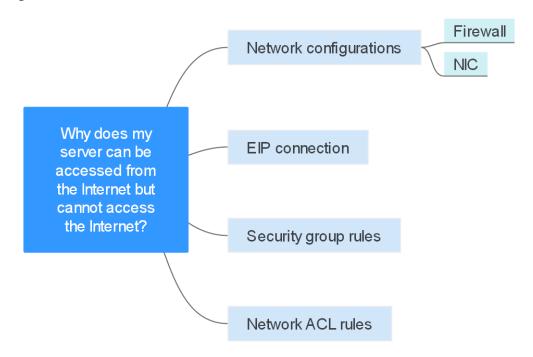
Symptom

The server can be accessed from, but cannot access the Internet.

Troubleshooting

Check the following possible causes.

Figure 14-34 Possible causes

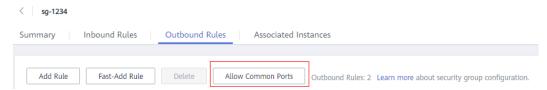


Security Group Rules

Check if there is a security group rule for the server denying the outbound traffic.

By default, a security group allows all outbound traffic. If the outbound traffic is denied, you can or click **Allow Common Ports**.

Figure 14-35 Allow Common Ports



Network ACL Rules

Check whether the network ACL of the subnet that the server belongs to denies the outbound traffic.

By default, a network ACL denies all outbound traffic. You need to add an outbound rule with **Action** set to **Allow** to the network ACL associated with the server.

14.8.12 Why Can't I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?

Symptom

You have enabled IPv4/IPv6 dual stack for an ECS, but the ECS cannot access websites using IPv6 addresses.

Troubleshooting

- Check whether the IPv4/IPv6 dual stack is correctly configured and whether the dual-stack NIC of the ECS has obtained an IPv6 address.
- Check whether the obtained IPv6 address of the dual-stack NIC has been added to a shared bandwidth.
- If the ECS has multiple NICs, check whether policy-based routes have been configured for these NICs.

Solution

 When you buy an ECS, select Automatically-assigned IPv6 address for Network.

□ NOTE

If an ECS is created from a public image:

Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 is supported and then check whether dynamic IPv6 address assignment has been enabled. Currently, all Linux public images support IPv6, and dynamic IPv6 address assignment is enabled for Ubuntu 16 by default. You do not need to configure dynamic IPv6 address assignment for the Ubuntu 16 OS. For other Linux public images, you need to enable this function.

- By default, IPv6 addresses can only be used for private network communication. If you want to use an IPv6 address to access the Internet or want it to be accessed by IPv6 clients on the Internet, you need to add the IPv6 address to a shared bandwidth.
 - If you already have a shared bandwidth, add the IPv6 address to it.
- If an ECS has multiple NICs, the primary NIC can communicate with external networks by default, but the extension NICs cannot. To enable extension NICs to communicate with external works either, you need to configure policybased routes for these NICs.

14.8.13 Why Does My ECS Fail to Communicate with Other After It Has Firewall Installed?

Symptom

An ECS has a single NIC and failed to communicate with others after the ECS has a firewall installed. An example scenario is as follows:

In a VPC, there are three ECSs. Services are deployed on ECS 1 and ECS 2, and a third-party firewall is installed on ECS X. Traffic from ECS 1 and ECS 2 needs to be filtered by the firewall of ECS X.

Fault Locating

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Table 14-17 Fault locating

Possible Cause	Solution
Security group rules	See Whether Security Group Rules Are Configured
Source/destination check	See Whether Source/Destination Check Is Disabled
VPC custom routes	See Whether VPC Custom Routes Are Added

Whether Security Group Rules Are Configured

Subnets in the same VPC can communicate with each other. If your service ECS cannot communicate with the ECS that has firewall installed, check whether they are in the same security group.

If the ECSs are in different security groups, you need to add rules to the security groups to allow access from each other.

Whether Source/Destination Check Is Disabled

Check whether the source/destination check function is disabled on the NIC of the ECS with firewall installed. If the function is not disabled, perform the following operations to disable it:

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- 3. Click **Service List** and choose **Compute** > **Elastic Cloud Server**.
- In the ECS list, click the target ECS name.
 The Summary tab page of the ECS is displayed.

Click the NICs tab, click to expand information about the primary NIC, and check whether Source/Destination Check is disabled.
 If it is not disabled, disable it.

Whether VPC Custom Routes Are Added

Check whether the subnet route table of the service VPC has a route pointing to the ECS with firewall installed.

If there is no such a route, add a custom route with next hop set to ECS and destination set to the ECS with the firewall installed.

14.9 Routing

14.9.1 How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?

Background

If an ECS has multiple NICs, the primary NIC can communicate with external networks by default, but the extension NICs cannot. To enable extension NICs to communicate with external works either, you need to configure policy-based routes for these NICs.

Scenarios

This example describes how to configure policy-based routes for an ECS with two NICs. Figure 14-36 shows the networking. The details are as follows:

- The primary and extension NICs on the source ECS are in different subnets of the same VPC.
- The source and destination ECSs are in different subnets of the same VPC and the two ECSs can communicate with each other through primary NICs without configuring policy-based routes.
- After policy-based routes are configured for the two NICs of the source ECS, both the primary and extension NICs can communicate with the destination ECS.

NOTICE

You can select a destination IP address based on service requirements. Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

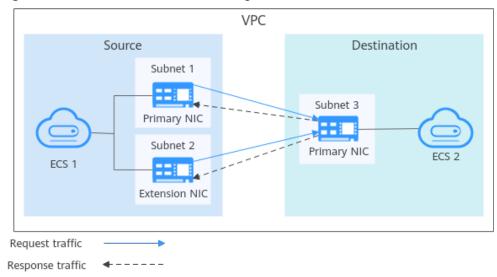


Figure 14-36 Dual-NIC ECS networking

Operation Guide

This document describes how to configure policy-based routes for Linux and Windows ECSs. For details, see **Table 14-18**.

Table 14-18 Operation instructions

OS Type	IP Address Version	Procedure
Linux	IPv4	Take an ECS running CentOS 8.0 (64-bit) as
	IPv6	an example.
Windows	IPv4	Take an ECS running Windows Server 2012
	IPv6	(64-bit) as an example.

14.9.2 Can a Route Table Span Multiple VPCs?

A route table cannot span multiple VPCs.

A route table contains a set of routes that are used to determine where network traffic from your subnets in a VPC is directed. A VPC has a default route table and can have multiple custom route tables.

Each subnet in a VPC must be associated with a route table. A subnet can only be associated with one route table at a time, but you can associate multiple subnets in a VPC with the same route table.

14.9.3 How Many Routes Can a Route Table Contain?

Currently, a route table can contain 100 routes.

14.9.4 Are There Any Restrictions on Using a Route Table?

- An ECS providing SNAT must have **Unbind IP from MAC** enabled.
- The destination of each route in a route table must be unique. The next hop
 must be a private IP address or a virtual IP address in the VPC. Otherwise, the
 route table will not take effect.
- If a virtual IP address is set to be the next hop in a route, EIPs bound with the virtual IP address in the VPC will become invalid.

14.9.5 Do the Same Routing Priorities Apply to Direct Connect Connections and Custom Routes in the Same VPC?

No. Direct Connect connections and custom routes are used in different scenarios, so the routing priorities are different.

14.9.6 Are There Different Routing Priorities of the VPN and Custom Routes in the Same VPC?

No. The routing priority of custom routes and that of VPNs are the same.

14.10 Security

14.10.1 Are the Security Group Rules Considered the Same If All Parameters Except Their Description Are the Same?

Yes. You cannot add or import a security group rule that has the same parameters but a different description than an existing rule in the security group.

14.10.2 How Do I Know the Instances Associated with a Security Group?

When you create an instance, such as ECS, cloud container, or database, you need to add the instance to a security group. To delete a security group, you must remove all instances from the security group first.

You can perform the following operations to view the instances associated with a security group:

- 1. In the security group list, locate the row that contains the target security group and click **Manage Instance** in the **Operation** column.
 - On the **Associated Instances** tab, you can view instances associated with the security group, such as servers and extension NICs.
 - If there is no instance associated with the security group on the **Associated Instances** tab, but the system still displays a message indicating that the security group has instances associated, perform the following operations:
- 2. Go to the console of the corresponding service, select the same region as the security group, and check whether the resources listed in **Table 14-19** exist.

Table 14-19 Check list

Product Category	Product/Instance
Databases	GaussDB
	DDS
	DDM
Middleware	Kafka
	RabbitMQ
	DMS (for RocketMQ)
	APIG
Big data	DataArts Studio
	DWS
	CSS

14.10.3 Why Can't I Delete a Security Group?

- The default security group is named **default** and cannot be deleted.
- If you want to delete a security group that is associated with instances, such as cloud servers, containers, and databases, you need to remove the instances from the security group first.
- A security group cannot be deleted if it is used as the source or destination of a rule in another security group.

You need to delete or modify the rule first and delete the security group.

For example, if the source of a rule in security group **sg-B** is set to **sg-A**, you need to delete or modify the rule in **sg-B** before deleting **sg-A**.

14.10.4 Can I Change the Security Group of an ECS?

Yes. Log in to the ECS console, switch to the page showing ECS details, and change the security group of the ECS.

14.10.5 How Do I Configure a Security Group for Multi-Channel Protocols?

ECS Configuration

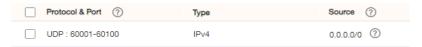
The TFTP daemon determines whether a configuration file specifies the port range. If you use a TFTP configuration file that allows the data channel ports to be configurable, it is a good practice to configure a small range of ports that are not listened on.

Security Group Configuration

You can configure port 69 and configure data channel ports used by TFTP for the security group. In RFC1350, the TFTP protocol specifies that ports available to data channels range from 0 to 65535. However, not all these ports are used by the TFTP daemon processes of different applications. You can configure a smaller range of ports for the TFTP daemon.

The following figure provides an example of the security group rule configuration if the ports used by data channels range from 60001 to 60100.

Figure 14-37 Security group rule



14.10.6 Does a Modified Security Group Rule or a Network ACL Rule Take Effect Immediately for Existing Connections?

- Security groups use connection tracking to track traffic to and from instances.
 If an inbound rule is modified, the modified rule immediately takes effect for the existing traffic. Changes to outbound security group rules do not affect existing persistent connections and take effect only for new connections.
 - If you add, modify, or delete a security group rule, or add or remove an instance to or from a security group, the inbound connections of all instances in the security group will be automatically cleared.
 - The existing inbound persistent connections will be disconnected. All the new connections will match the new rules.
 - The existing outbound persistent connections will not be disconnected, and the original rule will still be applied. All the new connections will match the new rules.
- Network ACLs use connection tracking to track traffic to and from instances.
 Changes to inbound and outbound rules do not take effect immediately for the existing traffic.
 - If you add, modify, or delete a network ACL rule, or associate or disassociate a subnet with or from a network ACL, all the inbound and outbound persistent connections will not be disconnected. New rules will only be applied for the new connections.

NOTICE

After a persistent connection is disconnected, new connections will not be established immediately until the timeout period of connection tracking expires. For example, after an ICMP persistent connection is disconnected, a new connection will be established and a new rule will apply when the timeout period (30s) expires.

- The timeout period of connection tracking varies by protocol. The timeout period of a TCP connection in the established state is 600s, and that of an ICMP connection is 30s. For other protocols, if packets are received in both inbound and outbound directions, the connection tracking timeout period is 180s. If packets are received only in one direction, the connection tracking timeout period is 30s.
- The timeout period of TCP connections varies by connection status. The timeout period of a TCP connection in the established state is 600s, and that of a TCP connection in the FIN-WAIT state is 30s.

14.10.7 Which Security Group Rule Has a High Priority When Multiple Security Group Rules Conflict?

Security group rules use the whitelist mechanism. If multiple security group rules conflict, the rules are aggregated to take effect.

14.10.8 Why Is Access from a Specific IP Address Still Allowed After a Network ACL Rule That Denies the Access from the IP Address Has Been Added?

Network ACL rules have priorities. A smaller priority value represents a higher priority. Each network ACL includes a default rule whose priority value is an asterisk (*). Default rules have the lowest priority.

If rules conflict, the rule with the highest priority takes effect.

If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule. For example, if the priority of rule A is 1 but you need rule B to take priority over rule A, insert rule B before rule A. Then, rule B will have a priority of 1 and rule A will be 2. Similarly, if rule B is less important than rule A, insert rule B after rule A.

When a rule that denies access from a specified IP address is added, insert the rules that allow access from all IP addresses at the end. Then, the rule that denies access from the specified IP address will take priority over the other rules and will be effective. For details, see .

14.10.9 Why Do My Security Group Rules Not Take Effect?

Symptom

The security group rules you have configured for an ECS have not taken effect.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 14-38 Troubleshooting

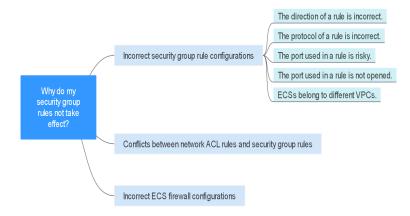


Table 14-20 Troubleshooting

Possible Cause	Solution
Incorrect Security Group Rule Configurations	See Incorrect Security Group Rule Configuration
Conflicts Between Network ACL Rules and Security Group Rules	See Conflicts Between Network ACL Rules and Security Group Rules

Incorrect Security Group Rule Configuration

If security group rules are incorrectly configured, ECSs cannot be protected. Check the security group rules based on the following causes:

- 1. The direction of a rule is incorrect.
- 2. The protocol of a rule is incorrect.
- 3. The port used in a rule is risky and cannot be accessed.
- 4. The port used in a rule is not opened. You can perform the following steps to check whether a port is being listened on the server.

For example, you have deployed a website on ECSs. Users need to access your website over TCP (port 80), and you have added the security group rule shown in **Table 14-21**.

Table 14-21 Security group rule

Direction	Protocol & Port	Source
Inbound	TCP: 80	0.0.0.0/0

Linux ECS

Check whether the security group rule takes effect on a Linux ECS:

- a. Log in to the ECS.
- b. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | grep 80

If information similar to **Figure 14-39** is displayed, TCP port 80 is enabled.

Figure 14-39 Command output for the Linux ECS



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.

Windows ECS

To verify the security group rule on a Windows ECS:

- a. Log in to the ECS.
- b. Choose **Start** > **Run**. Type cmd to open the Command Prompt.
- c. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | findstr 80

If information similar to **Figure 14-40** is displayed, TCP port 80 is enabled.

Figure 14-40 Command output for the Windows ECS



- d. 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.
- 5. ECSs belong to different VPCs. If two ECSs are in the same security group but in different VPCs, the ECSs cannot communicate with each other. To enable communications between the ECSs, use a VPC peering connection to connect the two VPCs. For details about VPC connectivity, see .

Conflicts Between Network ACL Rules and Security Group Rules

Security groups operate at the ECS level, whereas network ACLs operate at the subnet level.

For example, if you configure an inbound security group rule to allow access over port 80 and a network ACL rule to deny access over port 80, the security group rule will not take effect.

A Change History

Release Date	What's New	
2022-10-31	This release incorporates the following changes:	
	Added Can I Bind an EIP of an ECS to Another ECS?	
	 Added Can I Bind an EIP to a Cloud Resource in Another Region? 	
	Added What Bandwidth Types Are Available?	
	 Added What Is the Relationship Between Bandwidth and Upload/Download Rate? 	
	Added Can a Route Table Span Multiple VPCs?	
	 Added Are the Security Group Rules Considered the Same If All Parameters Except Their Description Are the Same? 	
2022-09-25	This release incorporates the following changes:	
	 Modified supported ECS flavors in section VPC Flow Log Overview. 	
2022-05-25	This release incorporates the following changes:	
	 Modified supported CIDR blocks in Adding a Secondary IPv4 CIDR Block to a VPC. 	
	 Added information about EIP billing in Assigning an EIP and Binding It to an ECS and Modifying an EIP Bandwidth. 	
	 Added information about shared bandwidth billing in Assigning a Shared Bandwidth and Modifying a Shared Bandwidth. 	
2022-01-25	This release incorporates the following change:	
	Added the supported netmask of a secondary CIDR block in Adding a Secondary IPv4 CIDR Block to a VPC.	

Release Date	What's New
2021-08-26	 This release incorporates the following changes: Added description about secondary IPv4 CIDR blocks in Modifying a VPC. Added sections Adding a Secondary IPv4 CIDR Block to a VPC and Deleting a Secondary IPv4 CIDR Block from
2021-06-30	 a VPC. This release incorporates the following changes: Added section Shared Bandwidth. Added IPv6-related description in VPC and Subnet and Access Control. Deleted the AZ parameter in Creating a VPC and
2021-04-20	 Creating a Subnet for the VPC. This release incorporates the following changes: Added supported ECS flavors in section VPC Flow Log Overview. Added procedure for manually configuring a virtual IP address in section Binding a Virtual IP Address to an EIP or ECS.
2020-12-03	 This release incorporates the following changes: Updated the subnet operation procedure based on console changes in section "VPC and Subnet" and adjusted its structure. Added section VPC and Subnet Planning Suggestions. Added section Exporting EIP Information. Added section VPC Flow Log. Modified the steps in section Changing the Security Group of an ECS. Deleted FAQ "Can a Route Table Span Multiple VPCs?" Updated the console strings. Updated the operation procedure based on GUI changes in section Virtual IP Address.

Release Date	What's New
2020-05-15	This release incorporates the following changes:
	 Added the example of allowing external access to a specified port in the section "Security Group Configuration Examples".
	Deleted section "Deleting a VPN".
	Optimized figure examples in this document.
	Optimized description in section "Network ACL Configuration Examples".
	Optimized description about default network ACL rules in section "Network ACL Overview".
	Added basic information to sections "Security Group Overview" and "Network ACL Overview".
	 Add rules when configuring Denying Access from a Specific Port.
	 Added FAQ "Does a Security Group Rule or a Network ACL Rule Immediately Take Effect for Its Original Traffic After It Is Modified?"
	 Modified FAQ "How Can I Delete a Subnet That Is Being Used by Other Resources?"
2018-09-30	This release incorporates the following changes:
	Added description about how to create multiple subnets at a time to section "Creating a VPC".
	 Added description about how to add multiple network ACL rules at a time and parameter Description to section "Adding a Network ACL Rule".
2018-08-15	This issue is the first official release.