

# Virtual Private Cloud

## User Guide

**Issue** 51  
**Date** 2024-06-05



**Copyright © Huawei Cloud Computing 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 Cloud Computing Technologies Co., Ltd.

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are the property 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 Cloud 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.

---

# Contents

---

<b>1 Permissions Management.....</b>	<b>1</b>
1.1 Creating a User and Granting VPC Permissions.....	1
1.2 VPC Custom Policies.....	2
<b>2 VPC and Subnet.....</b>	<b>5</b>
2.1 VPC and Subnet Planning Suggestions.....	5
2.2 VPC.....	9
2.2.1 Creating a VPC.....	9
2.2.2 Adding a Secondary IPv4 CIDR Block to a VPC.....	22
2.2.3 Obtaining a VPC ID.....	23
2.2.4 Modifying a VPC.....	24
2.2.5 Managing VPC Tags.....	25
2.2.6 Viewing a VPC Topology.....	27
2.2.7 Exporting VPC List.....	27
2.2.8 Deleting a Secondary IPv4 CIDR Block from a VPC.....	28
2.2.9 Deleting a VPC.....	28
2.3 Subnet.....	29
2.3.1 Creating a Subnet for the VPC.....	29
2.3.2 Modifying a Subnet.....	35
2.3.3 Managing Subnet Tags.....	40
2.3.4 Exporting Subnet List.....	42
2.3.5 Viewing and Deleting Resources in a Subnet.....	42
2.3.6 Viewing IP Addresses in a Subnet.....	45
2.3.7 Deleting a Subnet.....	46
<b>3 Route Tables and Routes.....</b>	<b>48</b>
3.1 Route Tables and Routes.....	48
3.2 Managing Route Tables.....	53
3.2.1 Creating a Custom Route Table.....	53
3.2.2 Associating a Route Table with a Subnet.....	54
3.2.3 Changing the Route Table Associated with a Subnet.....	55
3.2.4 Viewing the Route Table Associated with a Subnet.....	55
3.2.5 Viewing Route Table Information.....	56
3.2.6 Exporting Route Table Information.....	57

3.2.7 Deleting a Route Table.....	57
3.3 Managing Routes.....	58
3.3.1 Adding Routes to a Route Table.....	58
3.3.2 Modifying a Route.....	59
3.3.3 Replicating a Route.....	61
3.3.4 Deleting a Route.....	62
3.4 Route Configuration Examples.....	64
3.4.1 Configuring an SNAT Server to Enable ECSs to Share an EIP to Access the Internet.....	64
<b>4 Virtual IP Address.....</b>	<b>68</b>
4.1 Virtual IP Address Overview.....	68
4.2 Assigning a Virtual IP Address.....	70
4.3 Binding a Virtual IP Address to an EIP or ECS.....	71
4.4 Binding a Virtual IP Address to an EIP.....	78
4.5 Unbinding a Virtual IP Address from an Instance.....	78
4.6 Unbinding a Virtual IP Address from an EIP.....	79
4.7 Releasing a Virtual IP Address.....	80
4.8 Disabling IP Forwarding on the Standby ECS.....	81
4.9 Disabling Source/Destination Check for an ECS NIC.....	82
<b>5 Elastic Network Interface and Supplementary Network Interface.....</b>	<b>83</b>
5.1 Elastic Network Interface.....	83
5.1.1 Elastic Network Interface Overview.....	83
5.1.2 Creating a Network Interface.....	84
5.1.3 Viewing Basic Information About a Network Interface.....	85
5.1.4 Attaching a Network Interface to a Cloud Server.....	86
5.1.5 Binding an EIP to a Network Interface.....	86
5.1.6 Binding a Network Interface to a Virtual IP Address.....	87
5.1.7 Detaching a Network Interface from an Instance or Unbinding an EIP from a Network Interface....	87
5.1.8 Changing Security Groups That Are Associated with a Network Interface.....	88
5.1.9 Deleting a Network Interface.....	89
5.2 Supplementary Network Interfaces.....	90
5.2.1 Supplementary Network Interface Overview.....	90
5.2.2 Creating a Supplementary Network Interface.....	91
5.2.3 Viewing Basic Information About a Supplementary Network Interface.....	95
5.2.4 Binding or Unbinding an EIP to or from a Supplementary Network Interface.....	95
5.2.5 Changing Security Groups That Are Associated with a Supplementary Network Interface.....	96
5.2.6 Deleting a Supplementary Network Interface.....	98
<b>6 Access Control.....</b>	<b>99</b>
6.1 What Is Access Control?.....	99
6.2 Security Group.....	106
6.2.1 Security Groups and Security Group Rules.....	106
6.2.2 Default Security Groups.....	117

6.2.3 Security Group Examples.....	119
6.2.4 Common Ports Used by ECSs.....	124
6.2.5 Managing a Security Group.....	126
6.2.5.1 Creating a Security Group.....	127
6.2.5.2 Cloning a Security Group.....	131
6.2.5.3 Modifying a Security Group.....	132
6.2.5.4 Viewing the Details of a Security Group.....	133
6.2.5.5 Deleting a Security Group.....	133
6.2.5.6 Managing Security Group Tags.....	134
6.2.6 Managing Security Group Rules.....	136
6.2.6.1 Adding a Security Group Rule.....	136
6.2.6.2 Fast-Adding Security Group Rules.....	144
6.2.6.3 Allowing Common Ports with A Few Clicks.....	150
6.2.6.4 Modifying a Security Group Rule.....	152
6.2.6.5 Replicating a Security Group Rule.....	153
6.2.6.6 Importing and Exporting Security Group Rules.....	153
6.2.6.7 Deleting a Security Group Rule.....	158
6.2.7 Managing Instances Associated with a Security Group.....	159
6.2.7.1 Adding an Instance to or Removing an Instance from a Security Group.....	159
6.2.7.2 Changing the Security Group of an ECS.....	161
6.3 Network ACL.....	162
6.3.1 Network ACL Overview.....	162
6.3.2 Network ACL Configuration Examples.....	172
6.3.3 Managing Network ACLs.....	176
6.3.3.1 Creating a Network ACL.....	176
6.3.3.2 Modifying a Network ACL.....	178
6.3.3.3 Enabling or Disabling a Network ACL.....	178
6.3.3.4 Viewing a Network ACL.....	179
6.3.3.5 Deleting a Network ACL.....	180
6.3.3.6 Managing Network ACL Tags.....	180
6.3.4 Managing Network ACL Rules.....	182
6.3.4.1 Adding a Network ACL Rule (Default Priorities).....	182
6.3.4.2 Adding a Network ACL Rule (Custom Priorities).....	187
6.3.4.3 Modifying a Network ACL Rule.....	188
6.3.4.4 Enabling or Disabling a Network ACL Rule.....	193
6.3.4.5 Exporting and Importing Network ACL Rules.....	194
6.3.4.6 Deleting a Network ACL Rule.....	195
6.3.5 Managing Subnets Associated with a Network ACL.....	196
6.3.5.1 Associating Subnets with a Network ACL.....	196
6.3.5.2 Disassociating Subnets from a Network ACL.....	197
<b>7 IP Address Group.....</b>	<b>199</b>
7.1 IP Address Group.....	199

7.2 Managing an IP Address Group.....	201
7.2.1 Creating an IP Address Group.....	201
7.2.2 Associating an IP Address Group with Resources.....	203
7.2.3 Disassociating an IP Address Group from Resources.....	204
7.2.4 Modifying an IP Address Group.....	205
7.2.5 Exporting IP Address Group Details.....	206
7.2.6 Viewing Details of an IP Address Group.....	207
7.2.7 Deleting an IP Address Group.....	207
7.3 Managing IP Addresses in an IP Address Group.....	208
7.3.1 Adding IP Addresses to an IP Address Group.....	208
7.3.2 Modifying IP Addresses in an IP Address Group.....	210
7.3.3 Importing IP Addresses to an IP Address Group in Batches.....	212
7.3.4 Deleting IP Addresses from an IP Address Group.....	214
<b>8 VPC Peering Connection.....</b>	<b>215</b>
8.1 VPC Peering Connection.....	215
8.2 VPC Peering Connection Usage Examples.....	217
8.3 Creating a VPC Peering Connection with Another VPC in Your Account.....	228
8.4 Creating a VPC Peering Connection with a VPC in Another Account.....	235
8.5 Obtaining the Peer Project ID of a VPC Peering Connection.....	244
8.6 Modifying a VPC Peering Connection.....	245
8.7 Viewing VPC Peering Connections.....	246
8.8 Deleting a VPC Peering Connection.....	246
8.9 Modifying Routes Configured for a VPC Peering Connection.....	247
8.10 Viewing Routes Configured for a VPC Peering Connection.....	248
8.11 Deleting Routes Configured for a VPC Peering Connection.....	250
<b>9 VPC Sharing.....</b>	<b>252</b>
9.1 VPC Sharing.....	252
9.2 Usage Examples for VPC Sharing.....	261
9.3 Sharing a Subnet with Other Accounts.....	263
9.4 Viewing the Details of a Shared Subnet.....	264
9.5 Stopping Sharing a Subnet.....	264
<b>10 Setting Up an IPv6 Network.....</b>	<b>266</b>
<b>11 VPC Flow Log.....</b>	<b>272</b>
11.1 VPC Flow Log.....	272
11.2 Creating a VPC Flow Log.....	273
11.3 Viewing a VPC Flow Log.....	275
11.4 Enabling or Disabling VPC Flow Log.....	278
11.5 Deleting a VPC Flow Log.....	278
<b>12 Traffic Mirroring.....</b>	<b>280</b>
12.1 Traffic Mirroring.....	280

12.2 Mirror Filters.....	290
12.2.1 Creating a Mirror Filter.....	290
12.2.2 Adding an Inbound or Outbound Mirror Filter Rule.....	296
12.2.3 Modifying an Inbound or Outbound Mirror Filter Rule.....	302
12.2.4 Deleting an Inbound or Outbound Mirror Filter Rule.....	307
12.2.5 Modifying Basic Information About a Mirror Filter.....	307
12.2.6 Viewing Details About a Mirror Filter.....	308
12.2.7 Deleting a Mirror Filter.....	309
12.3 Mirror Sessions.....	310
12.3.1 Creating a Mirror Session.....	310
12.3.2 Enabling or Disabling a Mirror Session.....	312
12.3.3 Associating Mirror Sources with a Mirror Session.....	312
12.3.4 Disassociating Mirror Sources from a Mirror Session.....	313
12.3.5 Changing a Mirror Filter for a Mirror Session.....	314
12.3.6 Changing the Mirror Target of a Mirror Session.....	314
12.3.7 Modifying Basic Information About a Mirror Session.....	315
12.3.8 Viewing Details About a Mirror Session.....	317
12.3.9 Deleting a Mirror Session.....	317
<b>13 Elastic IP.....</b>	<b>319</b>
13.1 EIP Overview.....	319
13.2 Assigning an EIP and Binding It to an ECS.....	320
13.3 Unbinding an EIP from an ECS and Releasing the EIP.....	326
13.4 Modifying an EIP Bandwidth.....	327
13.5 Exporting EIP Information.....	330
13.6 Managing EIP Tags.....	330
13.7 IPv6 EIP .....	332
<b>14 Shared Bandwidth.....</b>	<b>339</b>
14.1 Shared Bandwidth Overview.....	339
14.2 Assigning a Shared Bandwidth.....	340
14.3 Adding EIPs to a Shared Bandwidth.....	342
14.4 Removing EIPs from a Shared Bandwidth.....	343
14.5 Modifying a Shared Bandwidth.....	343
14.6 Deleting a Shared Bandwidth.....	345
<b>15 Shared Data Package.....</b>	<b>347</b>
15.1 Shared Data Package Overview.....	347
15.2 Buying a Shared Data Package.....	348
<b>16 Monitoring and Auditing.....</b>	<b>350</b>
16.1 Monitoring.....	350
16.1.1 Supported Metrics.....	350
16.1.2 Viewing Metrics.....	352
16.1.3 Creating an Alarm Rule.....	353

---

16.2 Interconnecting with CTS.....	353
16.2.1 Key Operations Recorded by CTS.....	353
16.2.2 Viewing Traces.....	356
<b>A Change History.....</b>	<b>358</b>



# 1 Permissions Management

---

## 1.1 Creating a User and Granting VPC Permissions

This section describes how to use IAM to implement fine-grained permissions control for your VPC resources. With IAM, you can:

- Create IAM users for personnel based on your enterprise's organizational structure. Each IAM user has their own identity credentials for accessing VPC resources.
- Grant users only the permissions required to perform a given task based on their job responsibilities.
- Entrust a HUAWEI ID or cloud service to perform efficient O&M on your VPC resources.

If your HUAWEI ID meets your permissions requirements, you can skip this section.

[Figure 1-1](#) shows the process flow for granting permissions.

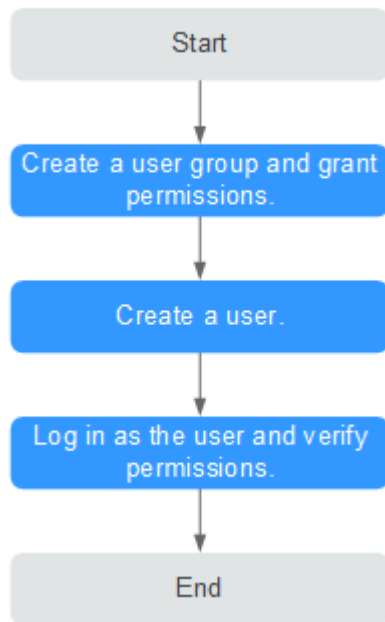
### Prerequisites

Learn about the permissions (see [Permissions](#)) supported by VPC and choose policies or roles according to your requirements.

To grant permissions for other services, learn about all [system-defined permissions](#) supported by IAM.

## Process Flow

**Figure 1-1** Process for granting VPC permissions



1. On the IAM console, **create a user group and grant it permissions**.  
Create a user group on the IAM console and assign the **VPCReadOnlyAccess** permissions to the group.
2. **Create an IAM user and add it to the created user group**.  
Create a user on the IAM console and add the user to the group created in 1.
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.
  - Choose another service from **Service List**. If a message appears indicating that you have insufficient permissions to access the service, the **VPCReadOnlyAccess** policy is in effect.

## 1.2 VPC Custom Policies

Custom policies can be created to supplement the system-defined policies of VPC. For the actions supported for custom policies, see [Permissions Policies and Supported Actions](#).

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.

For operation details, see [Creating a Custom Policy](#). The following section contains examples of common VPC custom policies.

## Example Custom Policies

- Example 1: Allowing users to create and view VPCs

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "vpc:vpcs:create",
        "vpc:vpcs:list"
      ]
    }
  ]
}
```

- 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:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "vpc:vpcs:delete"
      ]
    }
  ]
}
```

- 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:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "vpc:vpcs:create",
        "vpc:vpcs:update"
      ],
      "Effect": "Allow"
    },
    {
      "Action": [
        "ecs:servers:delete"
      ],
      "Effect": "Allow"
    }
  ]
}
```



# 2 VPC and Subnet

---

## 2.1 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.

If you require network connectivity between separate VPCs in the same account or in different accounts, you can use VPC peering connections or Cloud Connect.

- If two VPCs are in the same region, use a [VPC peering connection](#).
- If two VPCs are in different regions, use [Cloud Connect](#).

#### NOTE

By default, you can create a maximum of five VPCs in each region. If this cannot meet your service requirements, request a quota increase. For details, see [How Do I Apply for a Higher Quota?](#)

When creating a VPC, you need to specify an IPv4 CIDR block for it. Consider the following when selecting a CIDR block:

- Number of IP addresses: Reserve sufficient IP addresses for subsequent business growth.
- CIDR blocks: Avoid CIDR block conflicts if you need to connect a VPC to an on-premises data center or connect two VPCs.

When you create a VPC, we recommend that you use the private IPv4 address ranges specified in [RFC 1918](#) as the CIDR block, as described in [Table 2-1](#).

**Table 2-1** VPC CIDR blocks (RFC 1918)

VPC CIDR Block	IP Address Range	Netmask	Example CIDR Block
10.0.0.0/8-24	10.0.0.0– 10.255.255.255	8-24	10.0.0.0/8
172.16.0.0/12-24	172.16.0.0– 172.31.255.255	12-24	172.30.0.0/16
192.168.0.0/16-24	192.168.0.0– 192.168.255.255	16-24	192.168.0.0/24

In addition to the preceding addresses, you can create a VPC with a publicly routable CIDR block that falls outside of the private IPv4 address ranges specified in RFC 1918. However, the reserved system and public CIDR blocks listed in [Table 2-2](#) must be excluded:

**Table 2-2** Reserved system and public CIDR blocks

Reserved system CIDR blocks	Reserved public CIDR blocks
<ul style="list-style-type: none"><li>• 100.64.0.0/10</li><li>• 214.0.0.0/7</li><li>• 198.18.0.0/15</li><li>• 169.254.0.0/16</li></ul>	<ul style="list-style-type: none"><li>• 0.0.0.0/8</li><li>• 127.0.0.0/8</li><li>• 240.0.0.0/4</li><li>• 255.255.255.255/32</li></ul>

When you create a VPC, you specify a primary IPv4 CIDR block for the VPC, which cannot be changed. You can [add a secondary IPv4 CIDR block to the VPC](#) if required.

## 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.

A subnet mask can be between the netmask of its VPC CIDR block and /28 netmask. If a VPC CIDR block is 10.0.0.0/16, its subnet mask can be between 16 and 28.

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 10.0.2.0/24 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 [How Do I Apply for a Higher Quota?](#)

When planning subnets, consider the following:

- Plan subnets as required. You can 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.
- Avoid CIDR block conflicts. When you need to connect subnets in different VPCs or connect a VPC and an on-premises data center, ensure that the CIDR blocks of the subnets at both ends do not conflict.

## 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 control the traffic routing for your subnets in a VPC. 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.

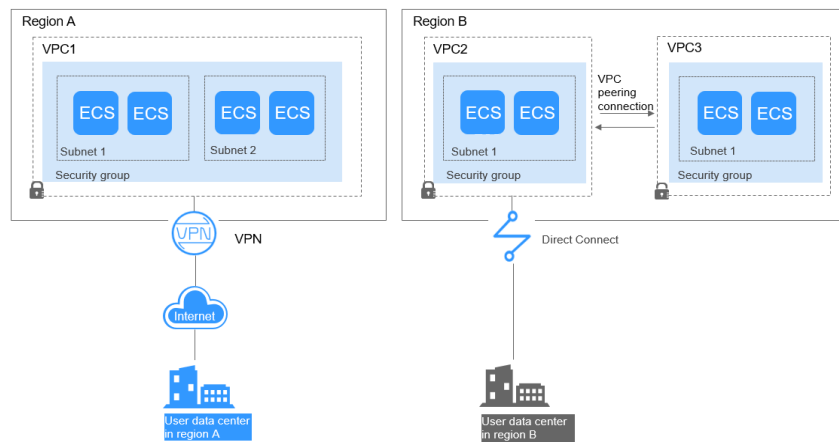
You cannot add two routes with the same destination to a VPC route table even if their next hop types are different, because the destination determines the route priority. According to the longest match routing rule, the destination with a higher matching degree is preferentially selected for packet 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 2-1](#), 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 2-1** 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 on-premises data center in Region A.
- The CIDR block of VPC 2 cannot overlap with the CIDR block of the on-premises 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.

For more information about EIP, see [EIP Overview](#).

### 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 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.

For more information, see [NAT Gateway User Guide](#).

### **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.

For more information, see [Elastic Load Balance User Guide](#).

## **Helpful Links**

- [Application Scenarios](#)
- [Private Network Access](#)
- [Public Network Access](#)

## **2.2 VPC**

### **2.2.1 Creating a VPC**

#### **Scenarios**

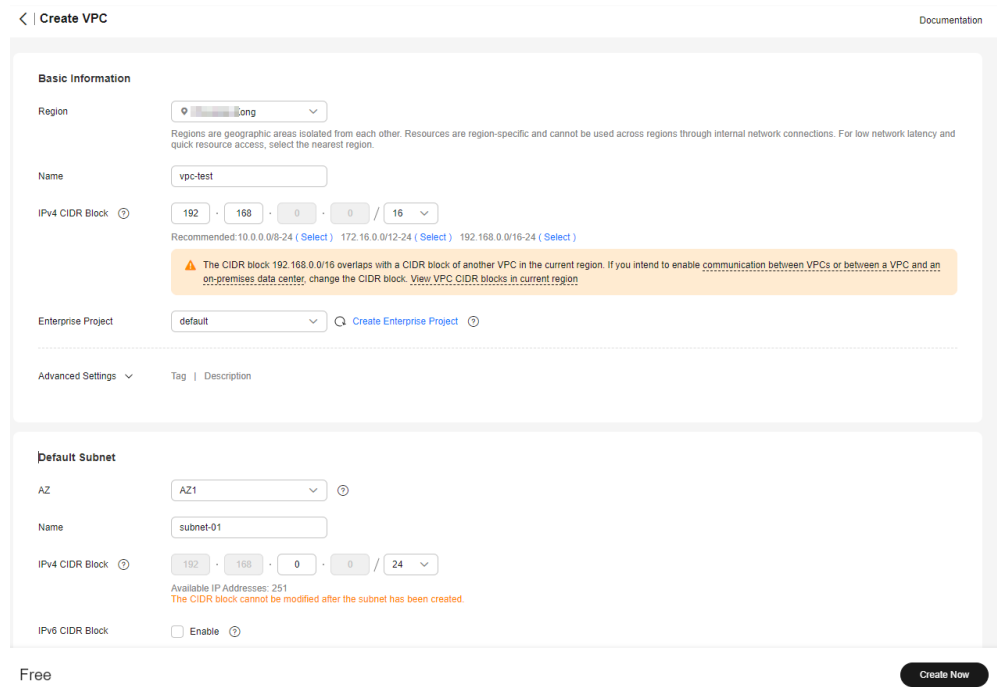
Virtual Private Cloud (VPC) allows you to provision logically isolated virtual private networks for cloud resources, such as cloud servers, containers, and databases.

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

#### **Procedure**

1. Go to the [Create VPC](#) page.
2. On the **Create VPC** page, set parameters for the VPC and subnets as prompted.


**Figure 2-2** Creating a VPC and subnet




**Table 2-3** VPC parameter descriptions

Parameter	Description	Example Value
Region	The region where the VPC belongs. Select the region nearest to you to ensure the lowest latency possible.	CN-Hong Kong
Name	The VPC name. The name: <ul style="list-style-type: none"> <li>Can contain 1 to 64 characters.</li> <li>Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li> </ul>	vpc-test

Parameter	Description	Example Value
IPv4 CIDR Block	<p>The IPv4 CIDR block of the VPC. Consider the following when specifying a CIDR block:</p> <ul style="list-style-type: none"> <li>• Number of IP addresses: Reserve sufficient IP addresses for subsequent business growth.</li> <li>• IP address ranges: Avoid IP address conflicts if you need to connect a VPC to an on-premises data center or connect two VPCs.</li> </ul> <p>When you create a VPC, we recommend that you use the private IPv4 address ranges specified in <a href="#">RFC 1918</a> as the CIDR block:</p> <ul style="list-style-type: none"> <li>• 10.0.0.0/8-24: The IP address ranges from 10.0.0.0 to 10.255.255.255, and the mask ranges from 8 to 24.</li> <li>• 172.16.0.0/12-24: The IP address ranges from 172.16.0.0 to 172.31.255.255, and the mask ranges from 12 to 24.</li> <li>• 192.168.0.0/16-24: The IP address ranges from 192.168.0.0 to 192.168.255.255, and the mask ranges from 16 to 24.</li> </ul> <p>In addition to the preceding addresses, you can create a VPC with a publicly routable CIDR block that falls outside of the private IPv4 address ranges specified in RFC 1918. However, the following system and public reserved addresses must be excluded:</p> <ul style="list-style-type: none"> <li>• Reserved system CIDR blocks <ul style="list-style-type: none"> <li>- 100.64.0.0/10</li> <li>- 214.0.0.0/7</li> <li>- 198.18.0.0/15</li> <li>- 169.254.0.0/16</li> </ul> </li> </ul>	10.0.0.0/8

Parameter	Description	Example Value
	<ul style="list-style-type: none"> <li>• Reserved public CIDR blocks                             <ul style="list-style-type: none"> <li>- 0.0.0.0/8</li> <li>- 127.0.0.0/8</li> <li>- 240.0.0.0/4</li> <li>- 255.255.255.255/32</li> </ul> </li> </ul> <p>For details about VPC planning, see <a href="#">VPC and Subnet Planning Suggestions</a>.</p>	
Enterprise Project	<p>The enterprise project to which the VPC belongs.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is <b>default</b>.</p> <p>For details about creating and managing enterprise projects, see the <a href="#">Enterprise Management User Guide</a>.</p>	default
Advanced Settings > Tag	<p>The VPC tag. Click  to expand the configuration area and set this parameter.</p> <p>Add tags to help you quickly identify, classify, and search for your VPCs.</p> <p>For details, see <a href="#">Managing VPC Tags</a>.</p> <p><b>NOTE</b></p> <p>If your organization has configured tag policies for VPCs, you need to add tags to your VPCs based on the policies. If you add a tag that does not comply with the tag policies, VPCs may fail to be created. Contact your administrator to learn more about tag policies.</p>	<ul style="list-style-type: none"> <li>• Key: vpc_key1</li> <li>• Value: vpc-01</li> </ul>

Parameter	Description	Example Value
Advanced Settings > Description	<p>Supplementary information about the VPC. Click  to expand the configuration area and set this parameter.</p> <p>Enter the description about the VPC in the text box as required.</p> <p>The VPC description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

**Table 2-4** Subnet parameter descriptions



Parameter	Description	Example Value
AZ	<p>An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network.</p> <p>Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.</p> <ul style="list-style-type: none"> <li>• By default, all instances in different subnets of the same VPC can communicate with each other and the subnets can be in different AZs. For example, if you have a VPC with two subnets, A01 in AZ 1 and A02 in AZ 2. Subnet A01 and A02 can communicate with each other by default.</li> <li>• A cloud resource and its subnet can be in different AZs. For example, a cloud server in AZ 1 can use a subnet in AZ 3. If AZ 3 becomes faulty, cloud servers in AZ 1 can still use the subnet in AZ 3, and your services are not interrupted.</li> <li>• Select <b>Central</b> if you want to provision cloud resources on the cloud and run your workloads on the cloud.</li> <li>• Select <b>Edge</b> if you want to provision cloud resources to an edge site and run workloads at the edge site. For details about edge sites, see <a href="#">CloudPond</a>.</li> </ul> <p>For details, see <a href="#">Region and AZ</a>.</p> <p>You can select an AZ for a subnet only in certain regions.</p>	AZ1


Parameter	Description	Example Value
	See the available regions on the management console.	
Name	The subnet name. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	subnet-01
CIDR Block	The CIDR block of the VPC. This parameter is displayed only in regions where IPv4/IPv6 dual stack is not supported. Set the IPv4 CIDR block of the VPC. For details, see section "IPv4 CIDR Block".	10.0.0.0/24


Parameter	Description	Example Value
IPv4 CIDR Block	<p>The IPv4 CIDR block of the VPC. This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported.</p> <p>A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets:</p> <ul style="list-style-type: none"><li>• Plan subnets as required. You can 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.</li><li>• Avoid CIDR block conflicts. When you need to connect subnets in different VPCs or connect a VPC and an on-premises data center, ensure that the CIDR blocks of the subnets at both ends do not conflict.</li></ul> <p>A subnet mask can be between the netmask of its VPC CIDR block and /28 netmask. If a VPC CIDR block is 10.0.0.0/16, its subnet mask can be between 16 to 28.</p> <p>For details about subnet planning, see <a href="#">VPC and Subnet Planning Suggestions</a>.</p>	10.0.0.0/24






Parameter	Description	Example Value
IPv6 CIDR Block	<p>The IPv6 CIDR block of the VPC. This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported.</p> <p>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.</p> <p>For details, see <a href="#">IPv4 and IPv6 Dual-Stack Network</a>.</p>	-
Associated Route Table	<p>The default route table to which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. Each VPC comes with a default route table. Subnets in the VPC are then automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.</p> <p>If the default route table cannot meet your requirements, you can create a custom route table and associate subnets with it. Then, the default route table controls inbound traffic to the subnets, while the custom route table controls outbound traffic from the subnets. For details, see <a href="#">Creating a Custom Route Table</a>.</p>	-

Parameter	Description	Example Value
Advanced Settings > Gateway	<p>The gateway address of the subnet. Click  to expand the configuration area and set this parameter.</p> <p>Retain the default value unless there are special requirements.</p>	10.0.0.1
Advanced Settings > DNS Server Address	<p>The DNS server addresses. Click  to expand the configuration area and set this parameter.</p> <p>Huawei Cloud private DNS server addresses are entered by default. This allows ECSs in a VPC to communicate with each other and also access other cloud services using private domain names without exposing their IP addresses to the Internet.</p> <p>You can change the default DNS server addresses if needed. This may interrupt your access to cloud services.</p> <p>You can also click <b>Reset</b> on the right to restore the DNS server addresses to the default value.</p> <p>A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).</p>	100.125.x.x

Parameter	Description	Example Value
Advanced Settings > Domain Name	<p>The domain name. Click  to expand the configuration area and set this parameter.</p> <p>Enter domain names (), separated with spaces. A maximum of 254 characters are allowed. A domain name can consist of multiple labels (max. 63 characters each).</p> <p>To access a domain name, you only need to enter the domain name prefix. ECSs in the subnet automatically match the configured domain name suffix.</p> <p>If the domain names are changed, ECSs newly added to this subnet will use the new domain names.</p> <p>If an existing ECS in this subnet needs to use the new domain names, restart the ECS or run a command to restart the DHCP Client service or network service.</p> <p><b>NOTE</b></p> <p>The command for updating the DHCP configuration depends on the ECS OS. The following commands are for your reference.</p> <ul style="list-style-type: none"><li>Restart the DHCP Client service: <b>service dhcpd restart</b></li><li>Restart the network service: <b>service network restart</b></li></ul>	test.com

Parameter	Description	Example Value
Advanced Settings > DHCP Lease Time	<p>The period during which a client can use an IP address automatically assigned by the DHCP server. Click  to expand the configuration area and set this parameter.</p> <p>The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client.</p> <ul style="list-style-type: none"><li>• <b>Limited:</b> Set the DHCP lease time. The unit can be day or hour.</li><li>• <b>Unlimited:</b> The DHCP lease time does not expire.</li></ul> <p>After you change the DHCP lease time on the console, the change is applied automatically when the DHCP lease of an instance (such as ECS) is renewed. You can wait for the system to renew the lease or manually renew the lease. Renewing lease will not change the IP address used by the instance. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.</p> <p>For details, see <a href="#">How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?</a></p>	-

Parameter	Description	Example Value
Advanced Settings > NTP Server Address	<p>The IP address of the NTP server. Click  to expand the configuration area and set this parameter.</p> <p>If you want to add NTP server addresses for a subnet, you can specify <b>NTP Server Address</b>. The IP addresses are added in addition to the default NTP server addresses.</p> <ul style="list-style-type: none"> <li>• If you add or change the NTP server addresses of a subnet, you need to renew the DHCP lease for or restart all the ECSs in the subnet to make the change take effect immediately.</li> <li>• If the NTP server addresses have been cleared out, restarting the ECSs will not help. You must renew the DHCP lease for all ECSs to make the change take effect immediately.</li> </ul>	192.168.2.1
Advanced Settings > Tag	<p>The subnet tag. Click  to expand the configuration area and set this parameter.</p> <p>Add tags to help you quickly identify, classify, and search for your subnets.</p> <p>For details, see <a href="#">Managing Subnet Tags</a>.</p> <p><b>NOTE</b></p> <p>If you have configured tag policies for subnets, you need to add tags to your subnets based on the tag policies. If you add a tag that does not comply with the tag policies, subnets may fail to be created. Contact the administrator to learn more about tag policies.</p>	<ul style="list-style-type: none"> <li>• Key: subnet_key1</li> <li>• Value: subnet-01</li> </ul>

Parameter	Description	Example Value
Advanced Settings > Description	<p>Supplementary information about the subnet. Click  to expand the configuration area and set this parameter.</p> <p>Enter the description about the subnet in the text box as required.</p> <p>The subnet description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

3. Click **Create Now**.  
Return to the VPC list and view the new VPC.

## Follow-up Operations

After the VPC and subnets are created, you need to create other cloud resources in the subnets. For details, see [Setting Up an IPv4/IPv6 Dual-Stack Network In a VPC](#) and [Setting Up an IPv4/IPv6 Dual-Stack Network In a VPC](#).

## 2.2.2 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.

#### NOTE

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 [Updating VPC Information](#).

### 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 /3 netmask.
- [Table 2-5](#) lists the secondary CIDR blocks that are not supported.

**Table 2-5** Restricted secondary CIDR blocks

Type	CIDR Block (Not Supported)
Reserved private CIDR blocks	<ul style="list-style-type: none"><li>• 172.31.0.0/16</li><li>• 192.168.0.0/16</li><li>• In-use primary CIDR blocks</li></ul>
Reserved system CIDR blocks	<ul style="list-style-type: none"><li>• 100.64.0.0/10</li><li>• 214.0.0.0/7</li><li>• 198.18.0.0/15</li><li>• 169.254.0.0/16</li></ul>
Reserved public CIDR blocks	<ul style="list-style-type: none"><li>• 0.0.0.0/8</li><li>• 127.0.0.0/8</li><li>• 240.0.0.0/4</li><li>• 255.255.255.255/32</li></ul>

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and choose **Networking > 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**.



## 2.2.3 Obtaining a VPC ID


### Scenarios

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

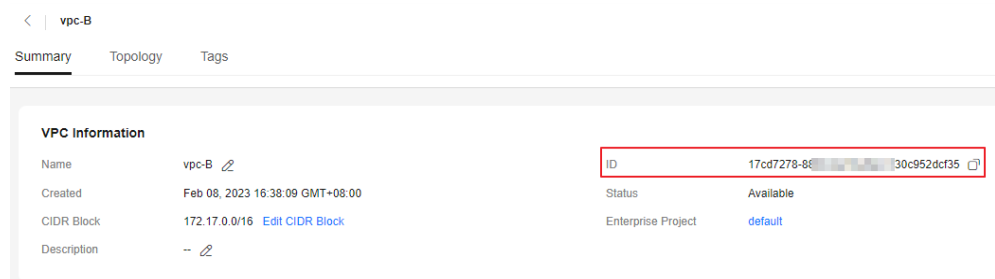
If you create a VPC peering connection between two VPCs in different accounts, you need to obtain the project ID of the region that the peer VPC resides. You can recommend this section to the user of the peer VPC to obtain the project ID.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. On the **Virtual Private Cloud** page, locate the VPC and click its name.  
The VPC details page is displayed.
5. In the **VPC Information** area, view the VPC ID.

Click  next to ID to copy the VPC ID.

**Figure 2-3** VPC ID



## 2.2.4 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 [Updating VPC Information](#).



### Modifying the Name and Description of a VPC

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.



4. Modify the name and description of a VPC using either of the following methods:
  - Method 1:
    - i. In the VPC list, click  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  on the right of the VPC name or description, enter the information, and click .

## Modifying the CIDR Block of a VPC

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. 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.
5. 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.
  - If a VPC has subnets, you only can change its subnet mask.
6. Click **OK**.

## 2.2.5 Managing VPC Tags

### Scenarios

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

You can add tags when creating a VPC or add tags to existing VPCs.

Each cloud resource can have a maximum of 20 tags.



A tag consists of a key and value pair. [Table 2-6](#) lists the tag key and value requirements.

**Table 2-6** VPC tag key and value requirements



Parameter	Requirements	Example Value
Key	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Must be unique for each VPC and can be the same for different VPCs.</li><li>• Can contain a maximum of 36 characters.</li><li>• Can contain letters, digits, underscores (_), and hyphens (-).</li></ul>	vpc_key1
Value	<ul style="list-style-type: none"><li>• Can contain a maximum of 43 characters.</li><li>• Can contain letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	vpc-01

## Procedure

### Search for VPCs by tag key and value on the page showing the VPC list.

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the search box above the VPC list, click anywhere in 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 VPC.

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. On the **Virtual Private Cloud** page, locate the VPC whose tags are to be managed and click the VPC name.  
The page showing details about the particular VPC is displayed.
5. 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 VPC, 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. In the **Edit Tag** dialog box, change the tag 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**.

## 2.2.6 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the VPC list, click the name of the VPC for which the topology is to be viewed.  
The VPC details page is displayed.
5. 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.


## 2.2.7 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 select the desired region and project.

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the upper left corner of the VPC list, click **Export**.
  - Export selected data to an XLSX file: Select one or more VPCs and export information about the selected VPCs.
  - Export all data to an XLSX file: Export information about all the VPCs in the current region.The system will automatically export information about the VPCs as an Excel file to a local directory.


## 2.2.8 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 **Networking > 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**.

## 2.2.9 Deleting a VPC

### Scenarios

If you no longer need a VPC, you can delete it.

---

**NOTICE**


VPCs are free of charge.

---

## 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.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and choose **Networking > 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.  
If your VPC is used by other resources, you need to delete these resources before deleting a VPC.
4. Enter **DELETE** as prompted and click **OK**.

## 2.3 Subnet

### 2.3.1 Creating a Subnet for the VPC

#### Scenarios

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

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



## Notes and Constraints

After a subnet is created, some reserved IP addresses cannot be used. For example, in a subnet with CIDR block 192.168.0.0/24, the following IP addresses are reserved by default:

- 192.168.0.0: Network ID. This address is the beginning of the private IP address range and will not be assigned to any instance.
- 192.168.0.1: The gateway address of the subnet.
- 192.168.0.253: Reserved for the system interface. This IP address is used by the VPC for external communication.
- 192.168.0.254: DHCP service address.
- 192.168.0.255: Network broadcast address.

The preceding default IP addresses are only examples. The system will assign reserved IP addresses based on how you specify your subnet.

## Procedure

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


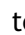
**Table 2-7** Subnet parameter descriptions



Parameter	Description	Example Value
Region	The region where VPC is located.	CN-Hong Kong
VPC	The VPC for which you want to create a subnet.	vpc-test
Subnet Name	The subnet name. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	subnet-01


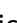
Parameter	Description	Example Value
AZ	<p>An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network.</p> <p>Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.</p> <ul style="list-style-type: none"><li>• By default, all instances in different subnets of the same VPC can communicate with each other and the subnets can be located in different AZs. For example, if you have a VPC with two subnets, A01 in AZ 1 and A02 in AZ 2. Subnet A01 and A02 can communicate with each other by default.</li><li>• A cloud resource can be in a different AZ from its subnet. For example, a cloud server in AZ 1 can be in a subnet in AZ 3. If AZ 3 becomes faulty, cloud servers in AZ 1 can still use the subnet in AZ 3, and your services are not interrupted.</li><li>• Select <b>Central</b> if you want to provision cloud resources on the cloud and run your workloads on the cloud.</li><li>• Select <b>Edge</b> if you want to provision cloud resources to an edge site and run workloads at the edge site. For details about edge sites, see <a href="#">CloudPond</a>.</li></ul> <p>For details, see <a href="#">Region and AZ</a>.</p> <p>You can select an AZ for a subnet only in certain regions. See the available regions on the management console.</p>	AZ1
CIDR Block	<p>The CIDR block of the VPC. This parameter is displayed only in regions where IPv4/IPv6 dual stack is not supported.</p> <p>Set the IPv4 CIDR block of the subnet. For details, see section "IPv4 CIDR Block".</p>	10.0.0.0/24

Parameter	Description	Example Value
IPv4 CIDR Block	<p>The IPv4 CIDR block of the VPC. This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported.</p> <p>A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets:</p> <ul style="list-style-type: none"><li>• Plan subnets as required. You can 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.</li><li>• Avoid CIDR block conflicts. When you need to connect subnets in different VPCs or connect a VPC and an on-premises data center, ensure that the CIDR blocks of the subnets at both ends do not conflict.</li></ul> <p>A subnet mask can be between the netmask of its VPC CIDR block and /28 netmask. If a VPC CIDR block is 10.0.0.0/16, its subnet mask can be between 16 to 28.</p> <p>If the VPC has a secondary CIDR block, you can select the primary or the secondary CIDR block that the subnet will belong to based on service requirements.</p>	10.0.0.0/24
IPv6 CIDR Block	<p>The IPv4 CIDR block of the subnet. This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported.</p> <p>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.</p> <p>For details, see <a href="#">IPv4 and IPv6 Dual-Stack Network</a>.</p>	-



Parameter	Description	Example Value
Associated Route Table	<p>The default route table to which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. Each VPC comes with a default route table. Subnets in the VPC are then automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.</p> <p>If the default route table cannot meet your requirements, you can create a custom route table and associate subnets with it. Then, the default route table controls inbound traffic to the subnets, while the custom route table controls outbound traffic from the subnets. For details, see <a href="#">Creating a Custom Route Table</a>.</p>	-
Advanced Settings > Gateway	<p>The gateway address of the subnet. Click  to expand the configuration area and set this parameter.</p> <p>Retain the default value unless there are special requirements.</p>	10.0.0.1
Advanced Settings > DNS Server Address	<p>The DNS server addresses. Click  to expand the configuration area and set this parameter.</p> <p>Huawei Cloud private DNS server addresses are entered by default. This allows ECSs in a VPC to communicate with each other and also access other cloud services using private domain names without exposing their IP addresses to the Internet.</p> <p>You can change the default DNS server addresses if needed. This may interrupt your access to cloud services.</p> <p>You can also click <b>Reset</b> on the right to restore the DNS server addresses to the default value.</p> <p>A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).</p>	100.125.x.x

Parameter	Description	Example Value
Advanced Settings > Domain Name	<p>The domain name. Click  to expand the configuration area and set this parameter.</p> <p>Enter domain names (), separated with spaces. A maximum of 254 characters are allowed. A domain name can consist of multiple labels (max. 63 characters each).</p> <p>To access a domain name, you only need to enter the domain name prefix. ECSs in the subnet automatically match the configured domain name suffix.</p> <p>If the domain names are changed, ECSs newly added to this subnet will use the new domain names.</p> <p>If an existing ECS in this subnet needs to use the new domain names, restart the ECS or run a command to restart the DHCP Client service or network service.</p> <p><b>NOTE</b></p> <p>The command for updating the DHCP configuration depends on the ECS OS. The following commands are for your reference.</p> <ul style="list-style-type: none"><li>Restart the DHCP Client service: <b>service dhcpd restart</b></li><li>Restart the network service: <b>service network restart</b></li></ul>	test.com
Advanced Settings > NTP Server Address	<p>The IP address of the NTP server. Click  to expand the configuration area and set this parameter.</p> <p>If you want to add NTP server addresses for a subnet, you can specify <b>NTP Server Address</b>. The IP addresses are added in addition to the default NTP server addresses.</p> <ul style="list-style-type: none"><li>If you add or change the NTP server addresses of a subnet, you need to renew the DHCP lease for or restart all the ECSs in the subnet to make the change take effect immediately.</li><li>If the NTP server addresses have been cleared out, restarting the ECSs will not help. You must renew the DHCP lease for all ECSs to make the change take effect immediately.</li></ul>	192.168.2.1

Parameter	Description	Example Value
Advanced Settings > DHCP Lease Time	<p>The period during which a client can use an IP address automatically assigned by the DHCP server. Click  to expand the configuration area and set this parameter.</p> <p>The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client.</p> <ul style="list-style-type: none"><li>• <b>Limited:</b> Set the DHCP lease time. The unit can be day or hour.</li><li>• <b>Unlimited:</b> The DHCP lease time does not expire.</li></ul> <p>If the time period is changed, the new lease time takes effect when the instance (such as an ECS) in the subnet is renewed next time. You can wait for the instance to be renewed automatically or manually modify the lease time. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.</p> <p>For details, see <a href="#">How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?</a></p>	-
Advanced Settings > Description	<p>Supplementary information about the subnet. Click  to expand the configuration area and set this parameter.</p> <p>Enter the description about the subnet in the text box as required.</p> <p>The subnet description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	-

7. Click **Create Now**.

Return to the subnet list and view the new subnet.

## 2.3.2 Modifying a Subnet




### Scenarios

Modify the subnet name, NTP server address, and DNS server address.

## Notes and Constraints

After a subnet is created, its AZ cannot be changed.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. In the subnet list, locate the target subnet and click its name.  
The subnet details page is displayed.
6. On the **Summary** tab, click  on the right of the parameter to be modified and modify the parameter as prompted.

**Table 2-8** Parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	Subnet

Parameter	Description	Example Value
DNS Server Address	<p>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.</p> <p>Huawei Cloud private DNS server addresses are entered by default. This allows ECSs in a VPC to communicate with each other and also access other cloud services using private domain names without exposing their IP addresses to the Internet.</p> <p>You can change the default DNS server addresses if needed. This may interrupt your access to cloud services.</p> <p>You can also click <b>Reset</b> on the right to restore the DNS server addresses to the default value.</p> <p>A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).</p>	100.125.x.x

Parameter	Description	Example Value
Domain Name	<p>Enter domain names (), separated with spaces. A maximum of 254 characters are allowed. A domain name can consist of multiple labels (max. 63 characters each).</p> <p>To access a domain name, you only need to enter the domain name prefix. ECSs in the subnet automatically match the configured domain name suffix.</p> <p>If the domain names are changed, ECSs newly added to this subnet will use the new domain names.</p> <p>If an existing ECS in this subnet needs to use the new domain names, restart the ECS or run a command to restart the DHCP Client service or network service.</p> <p><b>NOTE</b></p> <p>The command for updating the DHCP configuration depends on the ECS OS. The following commands are for your reference.</p> <ul style="list-style-type: none"><li>Restart the DHCP Client service: <b>service dhcpd restart</b></li><li>Restart the network service: <b>service network restart</b></li></ul>	test.com

Parameter	Description	Example Value
DHCP Lease Time	<p>The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client.</p> <ul style="list-style-type: none"><li>● <b>Limited:</b> Set the DHCP lease time. The unit can be day or hour.</li><li>● <b>Unlimited:</b> The DHCP lease time does not expire.</li></ul> <p>If the time period is changed, the new lease time takes effect when the instance (such as an ECS) in the subnet is renewed next time. You can wait for the instance to be renewed automatically or manually modify the lease time. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.</p> <p>For details, see <a href="#">How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?</a></p>	-

Parameter	Description	Example Value
NTP Server Address	<p>The IP address of the NTP server. This parameter is optional.</p> <p>You can configure the NTP server IP addresses to be added to the subnet as required. The IP addresses are added in addition to the default NTP server addresses. If you do not specify this parameter, no additional NTP server IP addresses will be added.</p> <p>A maximum of four unique NTP server IP addresses can be configured. Multiple IP addresses must be separated by a comma (,). If you add or change the NTP server addresses of a subnet, you need to renew the DHCP lease for or restart all the ECSs in the subnet to make the change take effect immediately. If the NTP server addresses have been cleared out, restarting the ECSs will not help. You must renew the DHCP lease for all ECSs to make the change take effect immediately.</p>	192.168.2.1
Description	<p>Supplementary information about the subnet. This parameter is optional.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	-

## 2.3.3 Managing Subnet Tags

### Scenarios

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

You can add tags when creating a subnet or add tags to existing subnets.

Each cloud resource can have a maximum of 20 tags.

A tag consists of a key and value pair. [Table 2-9](#) lists the tag key and value requirements.





**Table 2-9** Subnet tag key and value requirements



Parameter	Requirements	Example Value
Key	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Must be unique for each subnet.</li><li>• Can contain a maximum of 36 characters.</li><li>• Can contain letters, digits, underscores (_), and hyphens (-).</li></ul>	subnet_key1
Value	<ul style="list-style-type: none"><li>• Can contain a maximum of 43 characters.</li><li>• Can contain letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	subnet-01

## 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. In the subnet list, locate the target subnet and click its name.
6. 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. In the **Edit Tag** dialog box, change the tag 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**.

## 2.3.4 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. In the upper left corner of the subnet list, click **Export**.
  - Export selected data to an XLSX file: Select one or more subnets and export information about the selected subnets.
  - Export all data to an XLSX file: Export information about all the subnets in the current region.

The system will automatically export information about the subnets as an Excel file to a local directory.

## 2.3.5 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, network interfaces, load balancers, and NAT gateways.

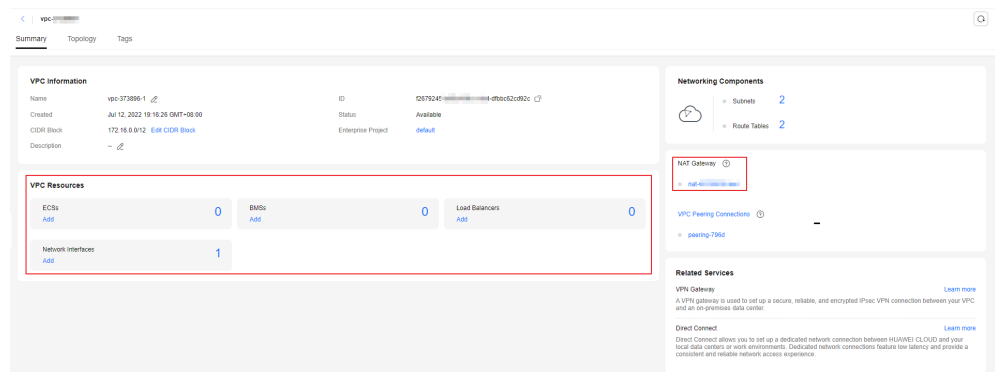
**NOTICE**

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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. Locate the target subnet and click its name.  
The subnet details page is displayed.
6. 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.


**Figure 2-4** Viewing resources in a subnet



7. Delete resources from the subnet.

**Table 2-10** Viewing and deleting resources in a subnet

Resource	Reference
ECS	<p>Currently, you cannot directly switch to ECSs from the subnet details page. You need to search for the target ECS in the ECS list and delete it.</p> <ol style="list-style-type: none"><li>In the ECS list, click the ECS name. The ECS details page is displayed.</li><li>In the <b>NICs</b> area, view the name of the subnet associated with the ECS.</li><li>Confirm the information and <b>delete the ECS</b>.</li></ol>
BMS	<p>Currently, you cannot directly switch to BMSs from the subnet details page. You need to search for the target BMS in the BMS list and delete it.</p> <ol style="list-style-type: none"><li>In the BMS list, click the BMS name. The BMS details page is displayed.</li><li>In the <b>NICs</b> tab, view the subnet associated with the BMS.</li><li>Confirm the information and <b>release the BMS</b>.</li></ol>
Load balancer	<p>You can directly switch to load balancers from the subnet details page.</p> <ol style="list-style-type: none"><li>Click the load balancer quantity. The load balancer list is displayed.</li><li>Locate the row that contains the load balancer and click <b>Delete</b> in the <b>Operation</b> column. For details, see <b>Deleting a Load Balancer</b>.</li></ol>
Network interface	<p>You can directly switch to network interfaces from the subnet details page.</p> <ol style="list-style-type: none"><li>Click the network interface quantity. The <b>Network Interfaces</b> page is displayed.</li><li>Locate the row that contains the network interface and choose <b>More &gt; Delete</b> in the <b>Operation</b> column. For details, see <b>Deleting a Network Interface</b>.</li></ol>

Resource	Reference
NAT gateway	<p>You can directly switch to NAT gateways from the subnet details page.</p> <ol style="list-style-type: none"><li>Click the NAT gateway name in the <b>Networking Components</b> area. The NAT gateway details page is displayed.</li><li>Click  to return to the NAT gateway list.</li><li>Locate the row that contains the NAT gateway and click <b>Delete</b> in the <b>Operation</b> column.<ul style="list-style-type: none"><li><a href="#">Deleting or Unsubscribing from a Public NAT Gateway</a></li><li><a href="#">Deleting a Private NAT Gateway</a></li></ul></li></ol>

## 2.3.6 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, DHCP, and system interface.
  - Used by cloud resources, such as ECSs, load balancers, and RDS instances.

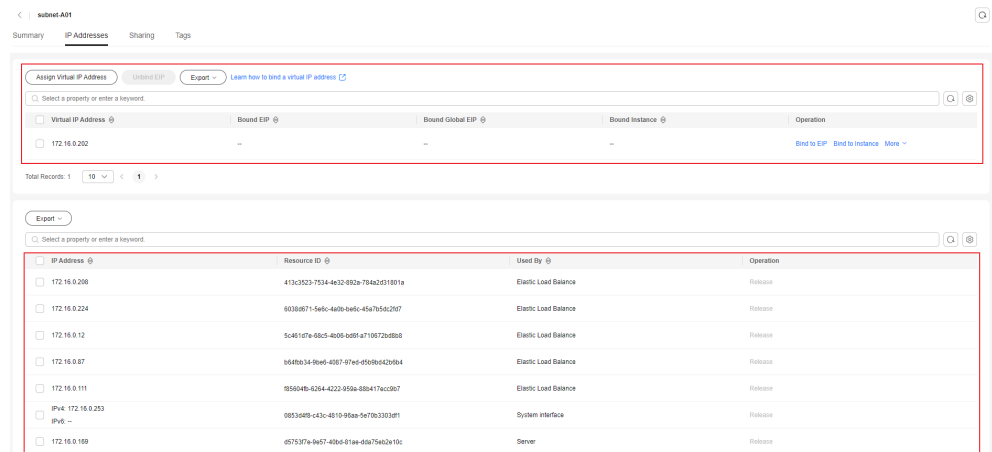
### Notes and Constraints

- 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

- Log in to the management console.
- Click  in the upper left corner and select the desired region and project.
- Click  in the upper left corner and choose **Networking > 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.
- Click the **IP Addresses** tab to view the IP addresses in the subnet.

- In the virtual IP address list, you can view the virtual IP addresses assigned from the subnet.
- In the private IP address list in the lower part of the page, you can view the private IP addresses, the resources that use the IP addresses of the subnet, and the resource ID.

**Figure 2-5** Viewing IP addresses in a 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?](#)

## 2.3.7 Deleting a Subnet

### Scenarios

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



#### NOTICE

Subnets are free of charge.

## Notes and Constraints

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

## Procedure

- Log in to the management console.
- Click  in the upper left corner and select the desired region and project.
- Click  in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. 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.  
If your subnet is used by other resources, you need to delete these resources before deleting a subnet.
6. Enter **DELETE** as prompted and click **OK**.

# 3 Route Tables and Routes


## 3.1 Route Tables and Routes

### What Is a Route Table?

A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. 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.

Both IPv4 and IPv6 routes are supported.

**Figure 3-1** Route tables



- **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 VPN, Cloud Connect, 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

By default, you do not have a quota for creating custom route tables. To create custom route tables, you need to apply for a quota increase first. For details, see [How Do I Apply for a Higher Quota?](#)

## Route

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

- System routes are automatically added by the VPC service and cannot be modified or deleted. After a route table is created, the following system routes will be added to the route table:
  - Routes whose destination is 100.64.0.0/10 (IP address range used to deploy public services, for example, the DNS server). The routes direct instances in a subnet to access these services.
  - Routes whose destination is 198.19.128.0/20 (IP address range used by internal services, such as VPC Endpoint).
  - Routes whose destination is 127.0.0.0/8 (local loopback addresses)
  - Routes whose destination is a subnet CIDR block and that enable instances in a VPC to communicate with each other.

If you enable IPv6 when creating a subnet, the system automatically assigns an IPv6 CIDR block to the subnet. Then, you can view IPv6 routes in its route table. Example destinations of subnet CIDR blocks are as follows:

- IPv4: 192.168.2.0/24
- IPv6: 2407:c080:802:be7::/64
- Custom routes 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 information such as the destination and next hop in the route to determine where network traffic is directed. [Table 3-1](#) 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 3-1** Next hop type

Next Hop Type	Description	Supported Route Table
Server	Traffic intended for the destination is forwarded to an ECS in the VPC.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
Extension NIC	Traffic intended for the destination is forwarded to the extension NIC of an ECS in the VPC.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
BMS user-defined network	Traffic intended for the destination is forwarded to a BMS user-defined network.	<ul style="list-style-type: none"><li>• Custom route table</li></ul>
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
Cloud connection	Traffic intended for the destination is forwarded to a cloud connection.	Custom route table
Supplementary network interface	Traffic intended for the destination is forwarded to the supplementary network interface of an ECS in the VPC.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
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.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>

Next Hop Type	Description	Supported Route Table
VPC endpoint	Traffic intended for the destination is forwarded to a VPC endpoint.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
Cloud container	Traffic intended for the destination is forwarded to a cloud container.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>
Internet gateway	Traffic intended for the destination is forwarded to an internet gateway.	<ul style="list-style-type: none"><li>• Default route table</li><li>• Custom route table</li></ul>

 **NOTE**

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 as the destination of a route. In this case, this route will be delivered as a 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.

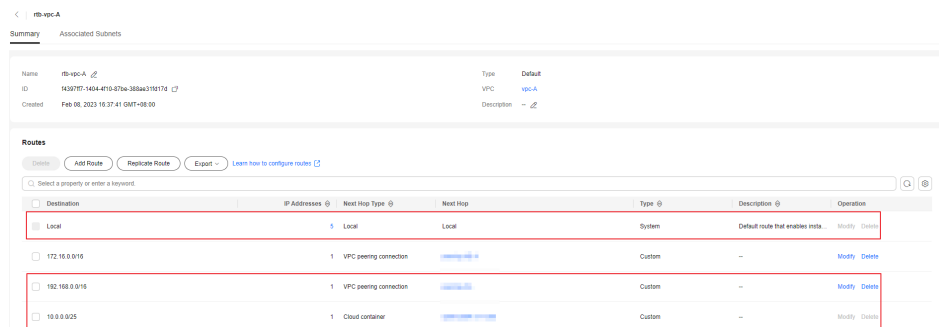
## Notes and Constraints

When you create a VPC, the system automatically generates a default route table for the VPC. You can also create a custom route table.

- A VPC can be associated with a maximum of five route tables, including the default route table and four custom route tables.
- All route tables in a VPC can have a maximum of 1,000 routes, excluding system routes.

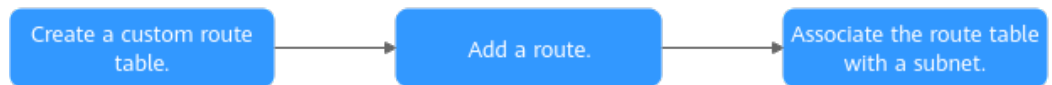
- In a VPC route table, the route priority is as follows:
  - Local route: A route that is automatically added by the system for communication within a VPC. It has a higher priority than a custom route.
  - Custom route: A route added by a user or routes that are delivered during instance creation. It uses the longest prefix match rule to find a destination for packet forwarding.

**Figure 3-2** VPC route table



## Custom Route Table Configuration Process

**Figure 3-3** Process for configuring a route table



**Table 3-2** Process for configuring a route table

N o.	Step	Description	Reference
1	Create a custom route table.	If your default route table cannot meet your service requirements, you can create a custom route table. The custom route table associated with a subnet only controls the outbound traffic. The default route table of a subnet controls the inbound traffic.	<a href="#">Creating a Custom Route Table</a>
2	Add a route.	You can add a custom route and configure information such as the destination and next hop in the route to determine where network traffic is directed.	<a href="#">Adding Routes to a Route Table</a>

No.	Step	Description	Reference
3	Associate the route table with a subnet.	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.	<a href="#">Associating a Route Table with a Subnet</a>

## 3.2 Managing Route Tables

### 3.2.1 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.

#### Notes and Constraints


By default, you do not have a quota for creating custom route tables. To create custom route tables, you need to apply for a quota increase first. For details, see [How Do I Apply for a Higher Quota?](#)

#### Procedure

1. Go to the [route table list page](#).
2. In the upper right corner, click **Create Route Table**. On the displayed page, configure parameters as prompted.

**Table 3-3** Parameter descriptions

Parameter	Description	Example Value
Name	(Mandatory) The name of the route table. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	rtb-001
VPC	(Mandatory) The VPC that the route table belongs to. The route table can be associated with the subnets in this VPC.	vpc-001

Parameter	Description	Example Value
Description	(Optional) Supplementary information about the route table. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Route Settings	(Optional) The route information. You can add a route when creating the route table or after the route table is created. For details, see <a href="#">Adding Routes to a Route Table</a> . You can click  to add more routes.	-

3. 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**.

## 3.2.2 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. In the route table list, locate the row that contains the target route table and click **Associate Subnet** in the **Operation** column.
6. Select the subnet to be associated.
7. Click **OK**.

### 3.2.3 Changing the Route Table Associated with a Subnet

#### Scenarios

You can change the route table for a subnet. If the route table 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.
5. Click the name of the target route table.
6. On the **Associated Subnets** tab page, click **Change Route Table** in the **Operation** column and select a new route table as prompted.
7. Click **OK**.

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



### 3.2.4 Viewing the Route Table Associated with a Subnet

#### Scenarios

You can view the route table associated with a subnet and the routes in the route table.

#### Procedure

1. Log in to the management console.

2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. Locate the target subnet and click its name.  
The subnet details page is displayed.
6. In the right of the subnet details page, view the route table associated with the subnet.
7. Click the name of the route table.  
The route table details page is displayed. You can further view the route information.



## 3.2.5 Viewing Route Table Information

### Scenarios

You can view the following information about a routing table:

- 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. 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.





## 3.2.6 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. In the upper left corner of the route table list, click **Export**.
  - Export selected data to an XLSX file: Select one or more route tables and export information about the selected route tables.
  - Export all data to an XLSX file: Export information about all the route tables in the current region.

The system will automatically export information about the route tables as an Excel file to a local directory.

## 3.2.7 Deleting a Route Table


### Scenarios


If you no longer need a a custom route table, you can delete it.

### Notes and Constraints

- The default route table cannot be deleted.  
However, deleting a VPC will also delete its default route table. Both default and custom route tables are free of charge.
- 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 select the desired region and project.

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**. The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. 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.
6. Confirm the information and click **OK**.




## 3.3 Managing Routes

### 3.3.1 Adding Routes to a Route Table

#### Scenarios

Each route table comes with a default route, which is used to allow instances in a subnet to access public services on the cloud or different subnets in a VPC to communicate with each other. You can also add custom routes as required to control traffic routing.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. Locate the target route table and click its name.  
The route table details page is displayed.
6. Click **Add Route** and set parameters as prompted.  
You can click  to add more routes.

**Table 3-4** Parameter descriptions

Parameter	Description	Example Value
Destination Type	Mandatory The destination type can only be <b>IP address</b> . You can set an IP address or network segment.	IP address

Parameter	Description	Example Value
Destination	<p>Mandatory</p> <p>Enter the destination of the route. You can enter a single IP address or an IP address range in CIDR notation.</p> <p><b>NOTICE</b></p> <ul style="list-style-type: none"><li>The destination of each route in a route table must be unique.</li><li>If an IP address group contains an IP address range in the format of <i>Start IP address-End IP address</i>, the IP address group is not supported. For example, an IP address group cannot contain 192.168.0.1-192.168.0.62. You need to change 192.168.0.1-192.168.0.62 to 192.168.0.0/26.</li></ul>	IPv4: 192.168.0.0/16
Next Hop Type	<p>Mandatory</p> <p>Set the type of the next hop.</p> <p><b>NOTE</b></p> <p>When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to <b>VPN gateway</b>, <b>Direct Connect gateway</b>, or <b>Cloud connection</b>.</p>	VPC peering connection
Next Hop	<p>Mandatory</p> <p>Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.</p>	peer-AB
Description	<p>Optional</p> <p>Enter the description of the route in the text box as required.</p>	-

7. Click **OK**.

You can view the new routes in the route list.

## 3.3.2 Modifying a Route

### Scenarios



You can modify an existing route in a route table.

### Notes and Constraints

- System routes cannot be modified.
- When you create a VPN, Cloud Connect, or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.
- Routes with the next hop type of cloud container cannot be modified or deleted.

- Routes with the next hop type of VPC endpoint cannot be modified or deleted.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. Locate the target route table and click its name.  
The route table details page is displayed.
6. Locate the target route and click **Modify** in the **Operation** column.
7. Modify the route information in the displayed dialog box.

**Table 3-5** Parameter descriptions

Parameter	Description	Example Value
Destination Type	Mandatory The destination type can only be <b>IP address</b> . You can set an IP address or network segment.	IP address
Destination	Mandatory Enter the destination of the route. You can enter a single IP address or an IP address range in CIDR notation. <b>NOTICE</b> <ul style="list-style-type: none"><li>• The destination of each route in a route table must be unique.</li><li>• If an IP address group contains an IP address range in the format of <i>Start IP address-End IP address</i>, the IP address group is not supported. For example, an IP address group cannot contain 192.168.0.1-192.168.0.62. You need to change 192.168.0.1-192.168.0.62 to 192.168.0.0/26.</li></ul>	IPv4: 192.168.0.0/16
Next Hop Type	Mandatory Set the type of the next hop. <b>NOTE</b> When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to <b>VPN gateway</b> , <b>Direct Connect gateway</b> , or <b>Cloud connection</b> .	VPC peering connection

Parameter	Description	Example Value
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.	-

8. Click **OK**.

### 3.3.3 Replicating a Route

#### Scenarios

You can replicate a route from a custom route table to one another within a VPC. You can also replicate a route from the default route table to a custom route table, or the other way around.

#### Notes and Constraints

**Table 3-6** shows whether routes of different types can be replicated to default or custom route tables.

If the next hop type of a route is a server, this route can be replicated to both default and 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 3-6** 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
BMS user-defined network	No	Yes
VPN gateway	No	Yes
Direct Connect gateway	No	Yes
Cloud connection	No	Yes

Next Hop Type	Can Be Replicated to Default Route Table	Can Be Replicated to Custom Route Table
Supplementary network interface	Yes	Yes
NAT gateway	Yes	Yes
VPC peering connection	Yes	Yes
Virtual IP address	Yes	Yes
VPC endpoint	No	No
Cloud container	No	No
Enterprise router	Yes	Yes
Cloud firewall	Yes	Yes

 **NOTE**

- If the Direct Connect service is enabled by call or email, the routes delivered to the default route table cannot be replicated to a custom route table.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. Locate the target route table and click its name.  
The route table details page is displayed.
6. Click **Replicate Route** above the route list and select the target route table and route.
7. Click **OK**.

### 3.3.4 Deleting a Route

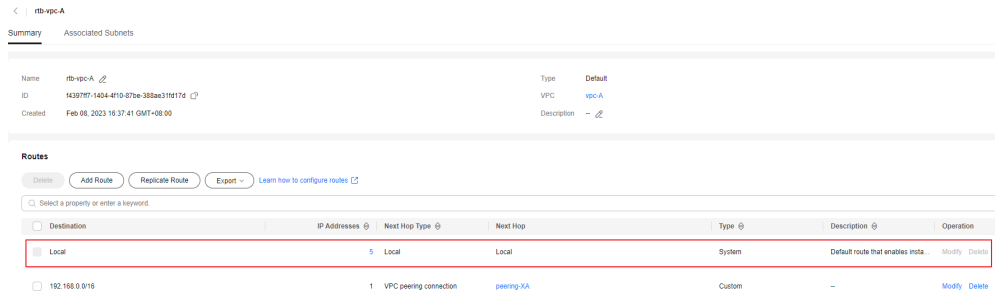
#### Scenarios

You can delete a custom route from a route table.

## Notes and Constraints



- System routes cannot be deleted.

**Figure 3-4** System route

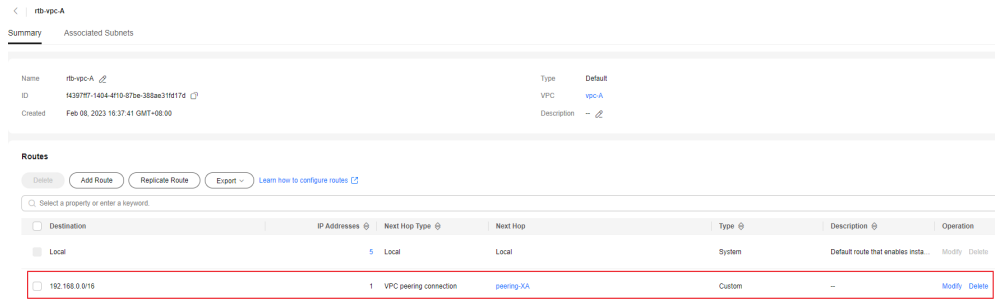


- The routes automatically delivered by VPN, Cloud Connect, or Direct Connect to the default route table cannot be deleted. The next hop types of such routes are:
  - VPN gateway
  - Direct Connect gateway
  - Cloud connectionTo delete these routes, you need to delete the associated network instances first.
- Routes with the next hop type of cloud container cannot be modified or deleted.
- Routes with the next hop type of VPC endpoint cannot be modified or deleted.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Route Tables**.  
The route table list is displayed.
5. Locate the target route table and click its name.  
The route table details page is displayed.

**Figure 3-5** Deleting a custom route



6. 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.
7. Confirm the information and click **OK**.

## 3.4 Route Configuration Examples

### 3.4.1 Configuring an SNAT Server to Enable ECSs to Share an EIP to Access the Internet

#### 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).

#### Differences Between SNAT ECSs and NAT Gateways



The NAT Gateway service provides network address translation (NAT) for servers, such as ECSs, BMSs and Workspace desktops, in a VPC or servers from an on-premises data center that connects to a VPC through Direct Connect or VPN. A NAT gateway allows these servers to share an EIP to access the Internet or provide services accessible from the Internet.

The NAT Gateway service is easier to configure and use than SNAT. This service can be flexibly deployed across subnets and AZs and has different NAT gateway specifications. You can click **NAT Gateway** under **Networking** on the management console to try this service.

For details, see the [NAT Gateway User Guide](#).



## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click . In the service list, choose **Compute > Elastic Cloud Server**.
4. 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.
5. On the displayed ECS details page, click the **NICs** tab.
6. 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.
7. 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](#).
8. On the ECS console, use the remote login function to log in to the ECS where you plan to configure SNAT.
9. Run the following command and enter the password of user **root** to switch to user **root**:

```
su - root
```

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

### NOTE

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

### **ping support.huawei.com**

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

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

11. 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 [14](#).
- If IP forwarding in Linux is disabled, go to [12](#) 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.

12. 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.

13. Run the following command to make the change take effect:

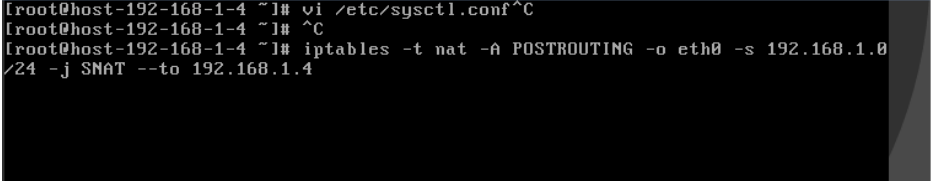
```
sysctl -p /etc/sysctl.conf
```

14. 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 3-6** Configuring SNAT



```
[root@host-192-168-1-4 ~]# vi /etc/sysctl.conf^C
[root@host-192-168-1-4 ~]# ^C
[root@host-192-168-1-4 ~]# iptables -t nat -A POSTROUTING -o eth0 -s 192.168.1.0/24 -j SNAT --to 192.168.1.4
```

#### NOTE

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 [14](#) to configure SNAT.

3. Save the configuration and exit:

```
:wq
```

4. Add the execution permissions for the `rc.local` file:

```
# chmod +x /etc/rc.local
```

15. Check whether the configuration is successful. If information similar to [Figure 3-7](#) (for example, 192.168.1.0/24) is displayed, the configuration was successful.

```
iptables -t nat --list
```

**Figure 3-7** Verifying configuration

```
[root@host-192-168-1-4 ~]# iptables -t nat --list
Chain PREROUTING (policy ACCEPT)
target     prot opt source                destination

Chain POSTROUTING (policy ACCEPT)
target     prot opt source                destination
SNAT       all  --  192.168.1.0/24        anywhere             to:192.168.1.4
SNAT       all  --  192.168.1.0/24        anywhere             to:192.168.1.4

Chain OUTPUT (policy ACCEPT)
target     prot opt source                destination
[root@host-192-168-1-4 ~]# _
```

16. Add a route. For details, see section [Adding Routes to a Route Table](#).

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.

# 4 Virtual IP Address

---

## 4.1 Virtual IP Address Overview

### What Is a Virtual IP Address?

A virtual IP address can be shared among multiple ECSs. An ECS can have a private and a virtual IP address, which allows your users to access the ECS through either IP address.

You can use either IP address to enable layer 2 and layer 3 communications in a VPC, access a different VPC using peering connections, and access cloud servers through EIPs, Direct Connect connections, and VPN connections.

You can bind a virtual IP address to ECSs deployed in the active/standby pair, 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**


To improve service availability and eliminate single points of failure, you can deploy ECSs in the active/standby pair or deploy one active ECS and multiple standby ECSs. And then, you can bind the same virtual IP address to these ECSs. If the active ECS becomes faulty, a standby ECS takes over services from the active ECS and services continue uninterrupted.

**Figure 4-1** Networking diagram of the HA mode



- As shown in the above figure, bind a virtual IP address to two ECSs in the same subnet.
- Configure Keepalived for the two ECSs to work in the active/standby pair. 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 4-2** 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 pair. 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.
- Check whether the source/destination check is disabled on the active and standby LVS ECSs. For details, see [Disabling Source/Destination Check for an ECS NIC](#).

If you bind an ECS to a virtual IP address on the management console, the source/destination check is automatically disabled. If you bind an ECS to a virtual IP address by calling APIs, you need to manually disable the source/destination check.

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. A VPC peering connection is needed so that two 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. Using the virtual IP addresses may cause route conflicts on the ECS, which would lead to communication failures.
- A virtual IP address from a subnet can only be bound to cloud servers from the same subnet.
- 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](#).
- Virtual IP addresses and extension network interfaces cannot be used to directly access Huawei Cloud services, such as DNS. You can use VPCEP to access these services. For details, see [Buying a VPC Endpoint](#).

## 4.2 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.
5. In the subnet list, click the name of the subnet where a virtual IP address is to be assigned.
6. Click the **IP Addresses** tab and click **Assign Virtual IP Address**.
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.

## 4.3 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.

---

#### NOTICE

If you need capabilities such as active/standby switchover or load balancing, you must configure Keepalived to work together with virtual IP addresses. For details, see [Building Highly Available Web Server Clusters with Keepalived](#).

---

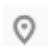

### Notes and Constraints

- 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.

#### NOTE

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.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.

- The **Subnets** page is displayed.
- Click the name with a hyperlink of the subnet that the virtual IP address belongs to.  
The subnet details page is displayed.
  - On the **IP Addresses** tab, bind an EIP to the virtual IP address:
    - 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.
    - Select an EIP and click **OK**.  
In the virtual IP address list, you can view that the virtual IP address has an EIP bound.
  - On the **IP Addresses** tab, bind an instance to the virtual IP address:
    - Locate the row that contains the virtual IP address and click **Bind to Instance** in the **Operation** column.  
The **Bind to Instance** dialog box is displayed.
    - Select an instance and click **OK**.  
In the virtual IP address list, you can view that the virtual IP address has an instance 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.)

- 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:

```
[root@192.168.254 ~]# nmcli connection
NAME                UUID                                  TYPE      DEVICE
Wired connection 1  5e72ec5a-6165-3bd6-a34b-ce43981acb27  ethernet  eth0
docker0             cd351a91-c5eb-4b69-83eb-df092a2ccf6b  bridge    docker0
```



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 NIC" +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:

```
[root@ecs-pod6-gaea-dpdk-ipv6 ~]# nmcli connection up "Wired connection 1"
Connection successfully activated (D-Bus active path: /org/freedesktop/NetworkManager/ActiveConnection/6)
```

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.

```
[root@ecs-pod6-gaea-dpdk-ipv6 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    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 group default qlen 1000
    link/ether fa:16:3e:e5:d5:ed brd ff:ff:ff:ff:ff:ff
    inet 172.16.0.247/24 brd 172.16.0.255 scope global noprefixroute dynamic eth0
        valid_lft 86398sec preferred_lft 86398sec
    inet 172.16.0.125/32 brd 172.16.0.125 scope global noprefixroute eth0
        valid_lft forever preferred_lft forever
    inet6 2001:db8:a5b3:62c:7ad3:a19a:4031:d6fb/128 scope global tentative noprefixroute dynamic
        valid_lft 86400sec preferred_lft 86400sec
    inet6 fe80::5371:9bf9:b652:e35b/64 scope link noprefixroute
        valid_lft forever preferred_lft forever
```

 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-ubuntu:~# 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-ubuntu:/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: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**

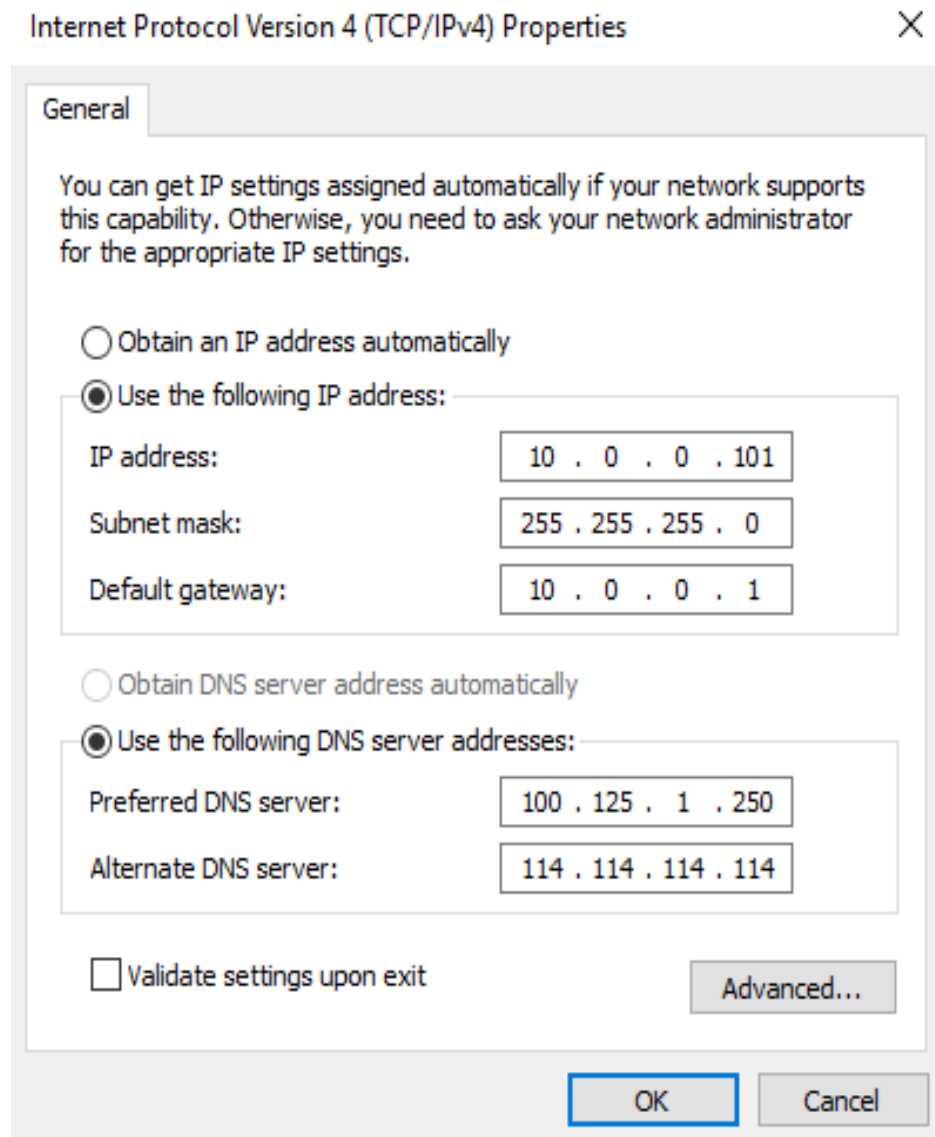
To delete an added virtual IP address, perform the following steps:

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**.

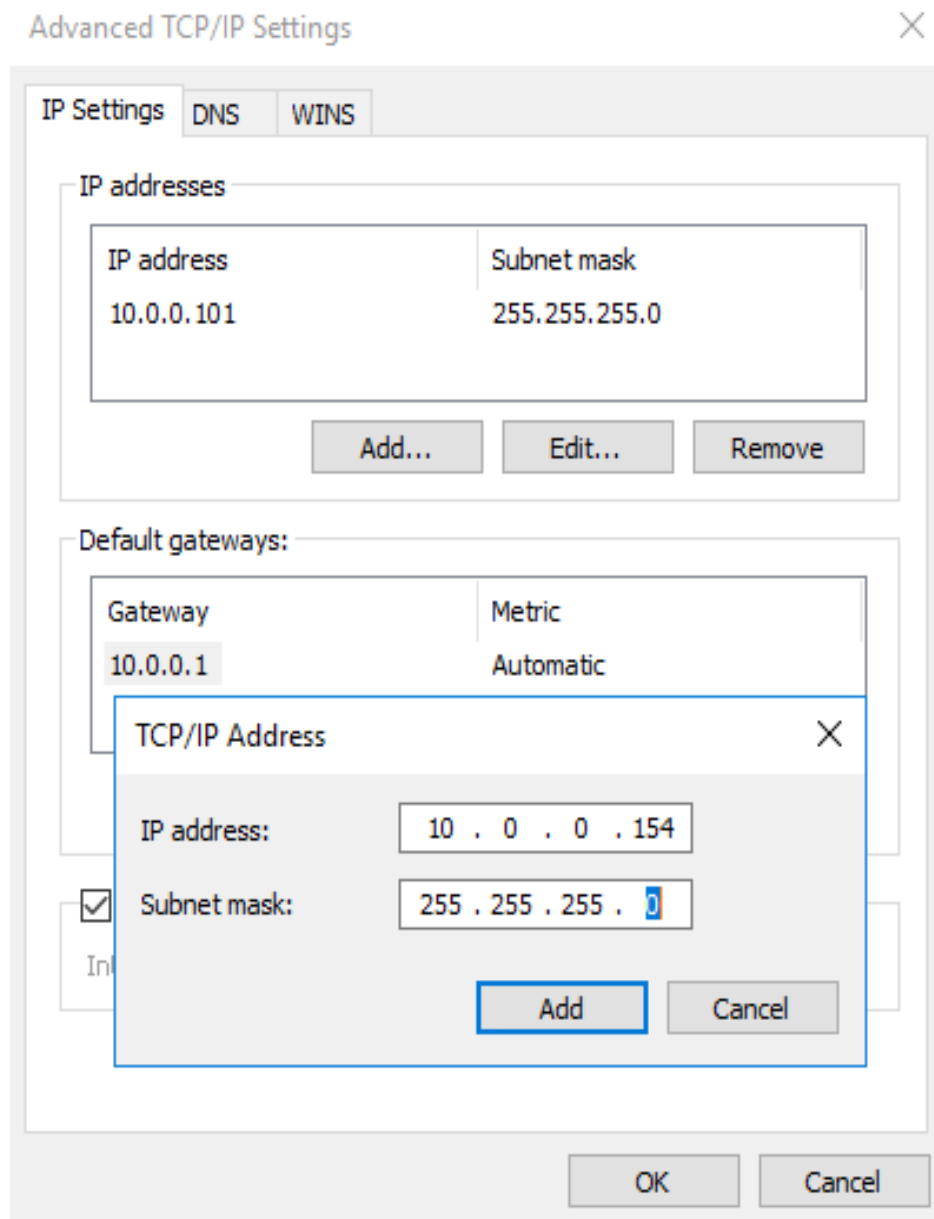
#### **Windows OS (Windows Server is used as an example here.)**

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 4-3** Configuring private IP address



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 4-4** 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.

## Helpful Links

- [Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?](#)

- [What Are the Differences Between EIP, Private IP Address, and Virtual IP Address?](#)
- [Unbinding a Virtual IP Address from an EIP](#)

## 4.4 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Elastic IP**.  
The EIP list page is displayed.
4. Locate the row that contains the EIP to be bound to the virtual IP address, and click **Bind** in the **Operation** column.
5. In the **Bind EIP** dialog box, set **Instance Type** to **Virtual IP address**.
6. In the virtual IP address list, select the virtual IP address to be bound and click **OK**.

## 4.5 Unbinding a Virtual IP Address from an Instance

### Scenarios

This section describes how to unbind a virtual IP address from an instance, such as an ECS or a Layer 2 connection.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud** > **Subnets**.  
The **Subnets** page is displayed.



5. Click the name of the subnet that the virtual IP address belongs to.  
The **Summary** page is displayed.
6. Click the **IP Addresses** tab.  
The virtual IP address list is displayed.
7. Locate the row that contains the virtual IP address, click **More** in the **Operation** column, and select **Unbind from Instance**.  
The **Bound Instance** dialog box is displayed.
8. 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**.

## 4.6 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.  
The **Subnets** page is displayed.
5. Click the name of the subnet that the virtual IP address belongs to.  
The **Summary** page is displayed.
6. Click the **IP Addresses** tab.  
The virtual IP address list is displayed.
7. 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.
8. Confirm the information and click **Yes**.

## 4.7 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 4-1](#).


**Table 4-1** Releasing a virtual IP address that is being used by a resource

Prompts	Cause Analysis and Solution
<p>Scenario 1: Virtual IP address cannot be released:</p> <p>This virtual IP address is bound to a resource. Unbind the virtual IP address first.</p>	<p>This virtual IP address is being used by cloud resources such as an EIP or an ECS.</p> <p>Unbind the virtual IP address from the cloud resources first.</p> <ul style="list-style-type: none"> <li>EIP: <a href="#">Unbinding a Virtual IP Address from an EIP</a></li> <li>ECS or Layer 2 connection: <a href="#">Unbinding a Virtual IP Address from an Instance</a></li> </ul> <p>Release the virtual IP address.</p>
<p>Scenario 2: Virtual IP address cannot be released:</p> <p>This operation cannot be performed because the IP address is being used by a system component.</p>	<p>The virtual IP address is being used by an instance. Delete the instance, which will also release the virtual IP address.</p> <p>Search for the instance based on the instance information displayed on the virtual IP address console and delete the instance.</p> <ul style="list-style-type: none"> <li>RDS DB instance: <a href="#">RDS Documentation</a></li> <li>CCE instance: <a href="#">CCE Documentation</a></li> <li>API gateway: <a href="#">API Gateway Documentation</a></li> </ul>



**Figure 4-5** Scenario 1: Virtual IP address cannot be released





**Figure 4-6** Scenario 2: Virtual IP address cannot be released

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Subnets**.
5. Click the name of the subnet that the virtual IP address belongs to.
6. 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.
7. Confirm the information and click **Yes**.

## 4.8 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.

## 4.9 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. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click . In the service list, choose **Compute > Elastic Cloud Server**.
4. In the ECS list, click the ECS name.
5. On the displayed ECS details page, click the **NICs** tab.
6. Check that **Source/Destination Check** is disabled.

# 5 Elastic Network Interface and Supplementary Network Interface

---

## 5.1 Elastic Network Interface

### 5.1.1 Elastic Network Interface Overview

An elastic network interface (referred to as a network interface in this documentation) is a virtual network card. You can create and configure network interfaces and attach them to your instances (ECSs and BMSs) to obtain flexible and highly available network configurations.

#### Network Interface Types

- A primary network interface is created together with an instance by default, and cannot be detached from the instance.
- An extended network interface is created on the **Network Interfaces** console, and can be attached to or detached from an instance.

#### Application Scenarios

- Flexible migration  
You can detach a network interface from an instance and then attach it to another instance. The network interface retains its private IP address, EIP, and security group rules. In this way, service traffic on the faulty instance can be quickly migrated to the standby instance, implementing quick service recovery.
- Traffic management  
You can attach multiple network interfaces that belong to different subnets in a VPC to the same instance, and configure the network interfaces to carry the private network traffic, public network traffic, and management network traffic of the instance. You can configure access control policies and routing policies for each subnet, and configure security group rules for each network interface to isolate networks and service traffic.

## Notes and Constraints

- The number of extended network interfaces that can be attached to an ECS is determined by the ECS specifications. For details, see [ECS Specifications](#).
- Extended network interfaces cannot be used to directly access Huawei Cloud services, such as DNS. You can use VPCEP to access these services. For details, see [Buying a VPC Endpoint](#).

## 5.1.2 Creating a Network Interface

### Scenarios

A primary network interface is created together with an instance by default. You can perform the following operations to create extended network interfaces on the **Network Interfaces** console.

### Notes and Constraints

An extended network interface created on the console can only be attached to an instance from the same VPC, but they can be associated with different security groups.

 **NOTE**

If you create an extended network interface using an API, the interface can be attached to an instance from a different VPC.

### Procedure

1. Go to the [network interface list page](#).
2. Click **Create Network Interface**.
3. Configure parameters for the network interface, as shown in [Table 5-1](#).

**Table 5-1** Parameter descriptions

Parameter	Parameter Description	Example Value
Name	Enter the name of the network interface. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	networkInterface-891e
VPC	Select the VPC to which the network interface belongs.	vpc-001
Subnet	Select the subnet that the network interface belongs to.	subnet-001

Parameter	Parameter Description	Example Value
Private IP Address	Select whether to automatically assign a private IP address.	-
Security Group	Select the security group that the network interface belongs to.	sg-001



4. Click **OK**.

## 5.1.3 Viewing Basic Information About a Network Interface

### Scenarios

You can view basic information about your network interface on the management console, including the name, ID, type, VPC, attached instance, and associated security groups.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click the name of the target network interface.

### Other Operations

On the network interface details page, you can also modify the following information:



- You can edit the network interface name, change IP addresses, and attach the network interface to or detach it from the instance.
- Instance-dependent Deletion
  - **Instance-dependent Deletion** is disabled by default. The network interface will not be deleted if it is detached from the instance or if the instance is deleted. You can attach the network interface to another instance.
  - If **Instance-dependent Deletion** has been enabled, the network interface will be deleted after it is detached from the instance.

## 5.1.4 Attaching a Network Interface to a Cloud Server

### Scenarios

You can attach a network interface to an ECS or a BMS to achieve flexible and high-availability network configurations.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. In the network interface list, locate the row that contains the target network interface, click **Attach Instance** in the **Operation** column, and select the instance to be attached.
6. Click **OK**.

### Related Operations

After a network interface is attached to an instance, it is recommended to enable NIC multi-queue to improve network performance. For details, see [Enabling NIC Multi-Queue](#).

## 5.1.5 Binding an EIP to a Network Interface


### Scenarios


You can bind an EIP to a network interface to achieve more flexible and scalable networks.

Each network interface has a private IP address. After the network interface is bound to an EIP, the network interface has both a private IP address and a public IP address. The binding between a network interface and an EIP will not change even after the network interface is detached from an instance. After a network interface is migrated from one instance to another, its private IP address and EIP will be migrated together at the same time.

An instance can have multiple network interfaces attached. If each network interface has an EIP bound, the instance will have multiple EIPs and can provide flexible access services.

### Procedure

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

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. In the network interface list, locate the row that contains the target network interface, click **Bind EIP** in the **Operation** column, and select the EIP to be bound.
6. Click **OK**.

## 5.1.6 Binding a Network Interface to a Virtual IP Address



### Scenarios

You can bind a network interface to a virtual IP address so that you can access the instance attached to the network interface using the virtual IP address.

Only a network interface with an instance attached can be bound to a virtual IP address.

For more information about virtual IP addresses, see [Virtual IP Address Overview](#).

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. In the network interface list, locate the row that contains the target network interface, and choose **More > Bind Virtual IP Address** in the **Operation** column.  
The **IP Addresses** page will be displayed.
6. Locate the row that contains the target virtual IP address and click **Bind to Server** in the **Operation** column.
7. Select the server and NIC, and click **OK**.

## 5.1.7 Detaching a Network Interface from an Instance or Unbinding an EIP from a Network Interface

### Scenarios

This section describes how to detach a network interface from an instance or unbind a network interface from an EIP.



## Notes and Constraints

- If **Instance-dependent Deletion** is enabled for a network interface, the network interface will be deleted if it is detached from its instance.
  - Deleting a network interface will also delete any supplementary network interfaces and VLAN sub-interfaces attached to it.
  - Deleting a network interface will also unbind its EIP.
- If **Instance-dependent Deletion** is disabled for a network interface, the network interface will not be deleted if it is detached from its instance.  
If a network interface has an EIP bound, detaching the network interface from its instance will also unbind the EIP from the network interface.

### NOTE

After an EIP is unbound from a network interface, if you do not release the EIP, the EIP will continue to be billed, but you can still bind it to other cloud resources.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. In the network interface list, locate the row that contains the target network interface, and click **Detach Instance** or **Unbind EIP** in the **Operation** column.
6. Click **Yes**.  
If you no longer need an EIP, you can release the EIP after unbinding it.



## 5.1.8 Changing Security Groups That Are Associated with a Network Interface

### Scenarios

You can change the security groups that are associated with a network interface on either the network interface list page or the network interface details page.

### Procedure

#### Changing the security group associated with a network interface on the network interface list page



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.



The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. In the network interface list, locate the row that contains the target network interface, and choose **More > Change Security Group** in the **Operation** column.
6. On the **Change Security Group** page, select the security groups to be associated and click **OK**.

#### Changing the security group associated with a network interface on the network interface details page

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. Click the name of the target network interface.
6. Click the **Associated Security Groups** tab. Then, click **Change Security Group**.
7. On the **Change Security Group** page, select the security groups to be associated and click **OK**.

## Other Operations

On the network interface details page, click the **Associated Security Groups** tab, and then click **Manage Rule**. For details about how to configure security group rules, see [Adding a Security Group Rule](#).

## 5.1.9 Deleting a Network Interface

### Scenarios

This section describes how to delete a network interface.

### Notes and Constraints



- A primary network interface is created together with an instance by default, and cannot be detached from the instance. If you want to delete a primary network interface, you need to delete its instance first, which will also delete the primary network interface.
- If you want to delete an extended network interface that is attached to an instance, **detach the interface from the instance** first.
- Deleting a network interface will also delete its supplementary network interfaces.
- Deleting a network interface will also unbind its EIP. You can choose to release the EIP if needed.

If you do not release the EIP, the EIP will continue to be billed, but you can still bind it to other cloud resources.

- If a network interface has resources associated, deleting the interface will also delete any rules for associated resources that reference it.

For example, if the next hop of a custom route in a VPC route table is a network interface, deleting the network interface will also delete the route.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. Locate the row that contains the network interface, click **More** in the **Operation** column, and click **Delete**.  
A confirmation dialog box is displayed.
6. Confirm the information and click **OK**.

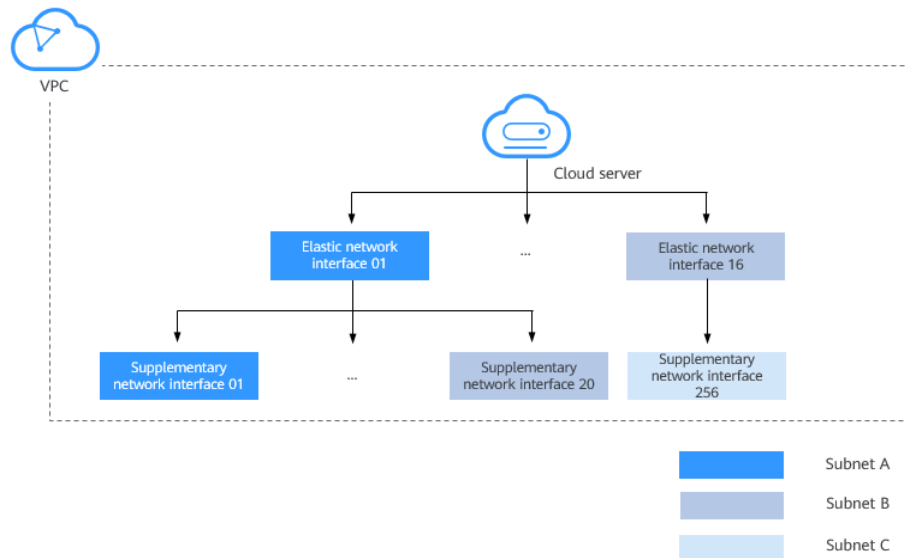
## 5.2 Supplementary Network Interfaces

### 5.2.1 Supplementary Network Interface Overview

Supplementary network interfaces are a supplement to elastic network interfaces. If the number of elastic network interfaces that can be attached to your ECS cannot meet your requirements, you can use supplementary network interfaces, which can be attached to VLAN subinterfaces of elastic network interfaces.

### Application Scenarios

Supplementary network interfaces are attached to VLAN subinterfaces of elastic network interfaces. [Figure 5-1](#) shows the networking diagram.

**Figure 5-1** Supplementary network interface networking diagram

The number of elastic network interfaces that can be attached to each ECS is limited. If this limit cannot meet your requirements, you can attach supplementary network interfaces to elastic network interfaces.

- You can attach supplementary network interfaces that belong to different subnets in the same VPC to an ECS. Each supplementary network interface has its private IP address and EIP for private or Internet communication.
- You can security group rules for supplementary network interfaces for network isolation.

## Notes and Constraints

- A maximum of 256 supplementary network interfaces can be attached to an ECS of certain flavors. The number of supplementary network interfaces that can be attached to an ECS varies by ECS flavor. ECS specifications that support supplementary network interfaces are as follows:  
ECS: C7, S7, and M7 series. For details, see [ECS Specifications](#).  
Cloud container: c6ne
- An ECS cannot use Cloud-Init through the private IP addresses of its supplementary network interfaces.
- A supplementary network interface cannot have a virtual IP address bound.
- The flow logs of supplementary network interfaces cannot be collected separately. Their flow logs are generated together with their network interfaces.

## 5.2.2 Creating a Supplementary Network Interface

### Scenarios

The number of elastic network interfaces that can be attached to each ECS is limited. If this limit cannot meet your requirements, you can use supplementary network interfaces.

## Notes and Constraints

- Supplementary network interfaces and its elastic network interface must be in the same VPC but can belong to different subnets and security groups.
- Before using a supplementary network interface, you need to create a VLAN sub-interface on its ECS NIC and configure routes. For details, see [Configuring a Supplementary Network Interface](#).

## Creating a Supplementary Network Interface

1. In the upper right corner of the page, click **Create Supplementary Network Interface**.
2. Configure the parameters based on [Table 5-2](#).

**Table 5-2** Parameter descriptions

Parameter	Description	Example Value
Network Interface	Elastic network interface that the supplementary network interface to be attached to. Select an elastic network interface from the drop-down list.	--(172.16.0.145)
VPC	VPC that the supplementary network interface belongs to. You do not need to set this parameter.	vpc-A
Subnet	Select the subnet for the supplementary network interface.	subnet-A01
Description	(Optional) Enter the description of the supplementary network interface in the text box as required. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Quantity	Number of supplementary network interfaces to be created. The value ranges from 1 to 20.	1
Private IP Address	Whether to assign a private IPv4 address to the supplementary network interface. This parameter cannot be deselected in the current version.	-

Parameter	Description	Example Value
IPv4 Address	Select a virtual IP address assignment mode. <ul style="list-style-type: none"><li>• <b>Automatically assign IP address:</b> The system assigns an IP address automatically.</li><li>• <b>Manually specify IP address:</b> The system assigns an IP address that you specify. If you select <b>Manually specify IP address</b>, enter a private IPv4 address.</li></ul>	Automatically assign IP address
Security Group	Select the security group that the supplementary network interface belongs to.	sg-001

3. Click **OK**.

---

**NOTICE**

After a supplementary network interface is created, you need to create a VLAN sub-interface on the ECS NIC and configure routes. For details, see [Configuring a Supplementary Network Interface](#).

---

## Configuring a Supplementary Network Interface

After a supplementary network interface is created, you need to create a VLAN sub-interface and configure a private IP address and default routes for the interface.

You need to obtain the information about the supplementary network interface, as shown in [Table 5-3](#).

**Table 5-3** Supplementary network interface information

Information	How to Obtain	Description
VLAN	Management console	Obtain the value from the supplementary network interface list.
MAC address		For details, see <a href="#">Viewing Basic Information About a Supplementary Network Interface</a> .
Private IP address		
Gateway		Obtain the value from the details page of the subnet that the supplementary network interface belongs to.

The following describes how to create a VLAN sub-interface on eth0 of an ECS (CentOS 8.2 is used as an example. For details about other OSs, see the OS documentation).

In this example:

- VLAN: 2110
- Private IP address: 192.168.0.2/24
- Gateway: 192.168.0.1
- MAC address: fa:16:3e:a1:b2:\*\*

### Procedure

1. Log in to the ECS.  
For details, see [Logging In to a Linux ECS](#).
2. Create a VLAN sub-interface for eth0.  
**ip link add link eth0 name eth0.2110 type vlan id 2110**
3. Create a namespace **ns2110**.  
**ip netns add ns2110**
4. Add the VLAN sub-interface **eth0.2110** to the namespace **ns2110**.  
**ip link set eth0.2110 netns ns2110**
5. Change the MAC address of the VLAN sub-interface to **fa:16:3e:a1:b2:\*\***.  
**ip netns exec ns2110 ifconfig eth0.2110 hw ether fa:16:3e:a1:b2:\*\***
6. Enable the VLAN sub-interface.  
**ip netns exec ns2110 ifconfig eth0.2110 up**
7. Configure the private IP address **192.168.0.2/24** for the VLAN sub-interface.  
**ip netns exec ns2110 ip addr add 192.168.0.2/24 dev eth0.2110**
8. Configure the default route for the VLAN sub-interface. 192.168.0.1 is the gateway of the subnet that the supplementary network interface works.  
**ip netns exec ns2110 ip route add default via 192.168.0.1**

### Verification

1. Access other private IP addresses in the same VPC from the namespace to check whether the configuration on the supplementary network interface takes effect.

```
ip netns exec ns2110 ping a.b.c.d
```

Figure 5-2 Success example

```
PING 192.168.0.1 (192.168.0.1) 56(84) bytes of data:
64 bytes from 192.168.0.1: icmp_seq=1 ttl=63 time=0.275 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=63 time=0.351 ms
```

Figure 5-3 Failure example



```
--- 192.168.0.1 ping statistics ---
11 packets transmitted, 0 received, 100% packet loss, time 10004ms
```

## 5.2.3 Viewing Basic Information About a Supplementary Network Interface

### Scenarios

You can view basic information about your supplementary network interface on the management console, including its ID, network interface, VLAN ID, VPC, subnet, private IP address, EIP, MAC address, and security groups.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
6. Click the private IP address of the supplementary network interface whose details you want to view.
  - On the **Summary** tab, you can view its ID, network interface, VLAN ID, VPC, subnet, private IP address, EIP, and MAC address.
  - On the **Associated Security Groups** tab, you can view its associated security groups and their rules.

### Other Operations

On the supplementary network interface details page, you can also modify the following information:

- On the **Summary** tab, you can modify the description of the interface and change its bound EIP.
- On the **Associated Security Groups** tab, you can change the associated security groups of the interface. For details, see [Changing Security Groups That Are Associated with a Supplementary Network Interface](#).

## 5.2.4 Binding or Unbinding an EIP to or from a Supplementary Network Interface

### Scenarios

You can bind a supplementary network interface to an EIP.



A supplementary network interface has a private IP address. You can also bind an EIP to the interface. The binding between a supplementary network interface and an EIP does not change when the network interface of the supplementary

network interface is detached from an ECS and then attached to another ECS. The supplementary network interface still has its private IP address and EIP.



A network interface can have multiple supplementary network interfaces attached. If each supplementary network interface has an EIP, the ECS with the network interface attached can have multiple EIPs for flexible Internet access.

If you do not need an EIP or want to delete a supplementary network interface, you can unbind the EIP from the interface.

## Binding a Supplementary Network Interface to an EIP

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
6. Locate the row that contains the supplementary network interface, click **Bind EIP** in the **Operation** column, and select the EIP to be bound.
7. Click **OK**.

## Unbinding a Supplementary Network Interface from an EIP

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
6. Locate the row that contains the supplementary network interface and click **Unbind EIP** in the **Operation** column.
7. Click **Yes**.

## 5.2.5 Changing Security Groups That Are Associated with a Supplementary Network Interface

### Scenarios

After a supplementary network interface is created, you can change its security group.





You can change the security group of a supplementary network interface:



- On the page of the supplementary network interface list.
- On the details page of a supplementary network interface.

## Procedure

### Changing the security group associated with a supplementary network interface on the supplementary network interface list page

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click the **Supplementary Network Interfaces** tab.
6. Locate the row that contains the supplementary network interface and click **Change Security Group** in the **Operation** column.
7. On the **Change Security Group** page, select the security group to be associated.
8. Click **OK**.

### Changing the security group associated with a supplementary network interface on the supplementary network interface details page

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click the **Supplementary Network Interfaces** tab.
6. Click the private IP address of the supplementary network interface whose security group is to be changed.
7. Click the **Associated Security Groups** tab. Then, click **Change Security Group**.
8. On the **Change Security Group** page, select the security group to be associated.
9. Click **OK**.

## 5.2.6 Deleting a Supplementary Network Interface



### Scenarios

If you want to delete a supplementary network interface with an EIP bound, you have to first unbind the EIP from the interface.

### Notes and Constraints

- Deleting a supplementary network interface will also detach it from its network interface.
- Deleting a supplementary network interface will also unbind its EIP. You can choose to release the EIP if needed.  
If you do not release the EIP, the EIP will continue to be billed, but you can still bind it to other cloud resources.
- If a supplementary network interface has resources associated, deleting the interface will also delete any rules for associated resources that reference it. For example, if the next hop of a custom route in a VPC route table is a supplementary network interface, deleting the interface will also delete the route.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
5. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
6. Locate the row that contains the supplementary network interface and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
7. Confirm the information and click **OK**.  
Deleting a supplementary network interface will also delete the VLAN sub-interfaces configured on the ECS.

# 6 Access Control

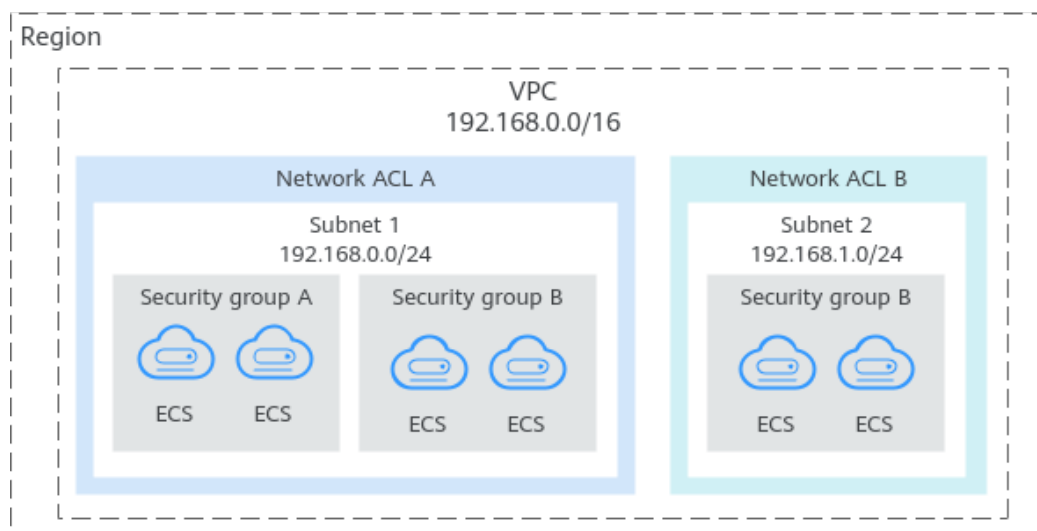
## 6.1 What Is Access Control?

A VPC is your private network on the cloud. You can configure security groups and network ACL rules to ensure the security of instances, such as ECSs, databases, and containers, running in a VPC.

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

**Figure 6-1** shows how security groups and network ACLs are used. 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.

**Figure 6-1** Security groups and network ACLs



### Differences Between Security Groups and Network ACLs

**Table 6-1** describes detailed differences between security groups and network ACLs.

**Table 6-1** Differences between security groups and network ACLs

Item	Security Group	Network ACL
Protection Scope	Protects instances in a security group, such as ECSs, databases, and containers.	Protects subnets and all the instances in the subnets.
Mandatory	Yes. Instance must be added to at least one security group.	No. You can determine whether to associate a subnet with a network ACL based on service requirements.
Stateful	Yes. The response traffic of inbound and outbound requests is allowed to flow to and leave an instance.	Yes. The response traffic of inbound and outbound requests is allowed to flow to and leave a subnet.
Rules	Supports both <b>Allow</b> and <b>Deny</b> rules. <ul style="list-style-type: none"> <li>• <b>Allow</b>: allows the matched traffic to flow in or out of the instances.</li> <li>• <b>Deny</b>: denies the matched traffic to flow in or out of the instances.</li> </ul>	Supports both <b>Allow</b> and <b>Deny</b> rules. <ul style="list-style-type: none"> <li>• <b>Allow</b>: allows the matched traffic to flow in or out of the subnet.</li> <li>• <b>Deny</b>: denies the matched traffic to flow in or out of the subnet.</li> </ul>
Rule 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.
Matching Order	If an instance is associated with multiple security groups that have multiple rules: <ol style="list-style-type: none"> <li>1. Rules are first matched based on the sequence each security group is associated with an instance. Security groups with lower sequence numbers have higher priorities.</li> <li>2. Rules are then matched by priority in that security group. Rules with lower values have higher priorities than those with higher values.</li> <li>3. Deny rules take precedence over allow rules if the rules have the same priority.</li> </ol>	A subnet can have only one network ACL associated. If there are multiple rules, traffic is matched based on the rule priority. A smaller value indicates a higher priority.

Item	Security Group	Network ACL
Usage	<ul style="list-style-type: none"><li>• When creating an instance, for example, an ECS, you must select a security group. If no security group is selected, the ECS will be associated with the default security group.</li><li>• After creating an instance, you can:<ul style="list-style-type: none"><li>– Add or remove the instance to or from the security group on the security group console.</li><li>– Associate or disassociate a security group with or from the instance on the instance console.</li></ul></li></ul>	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.

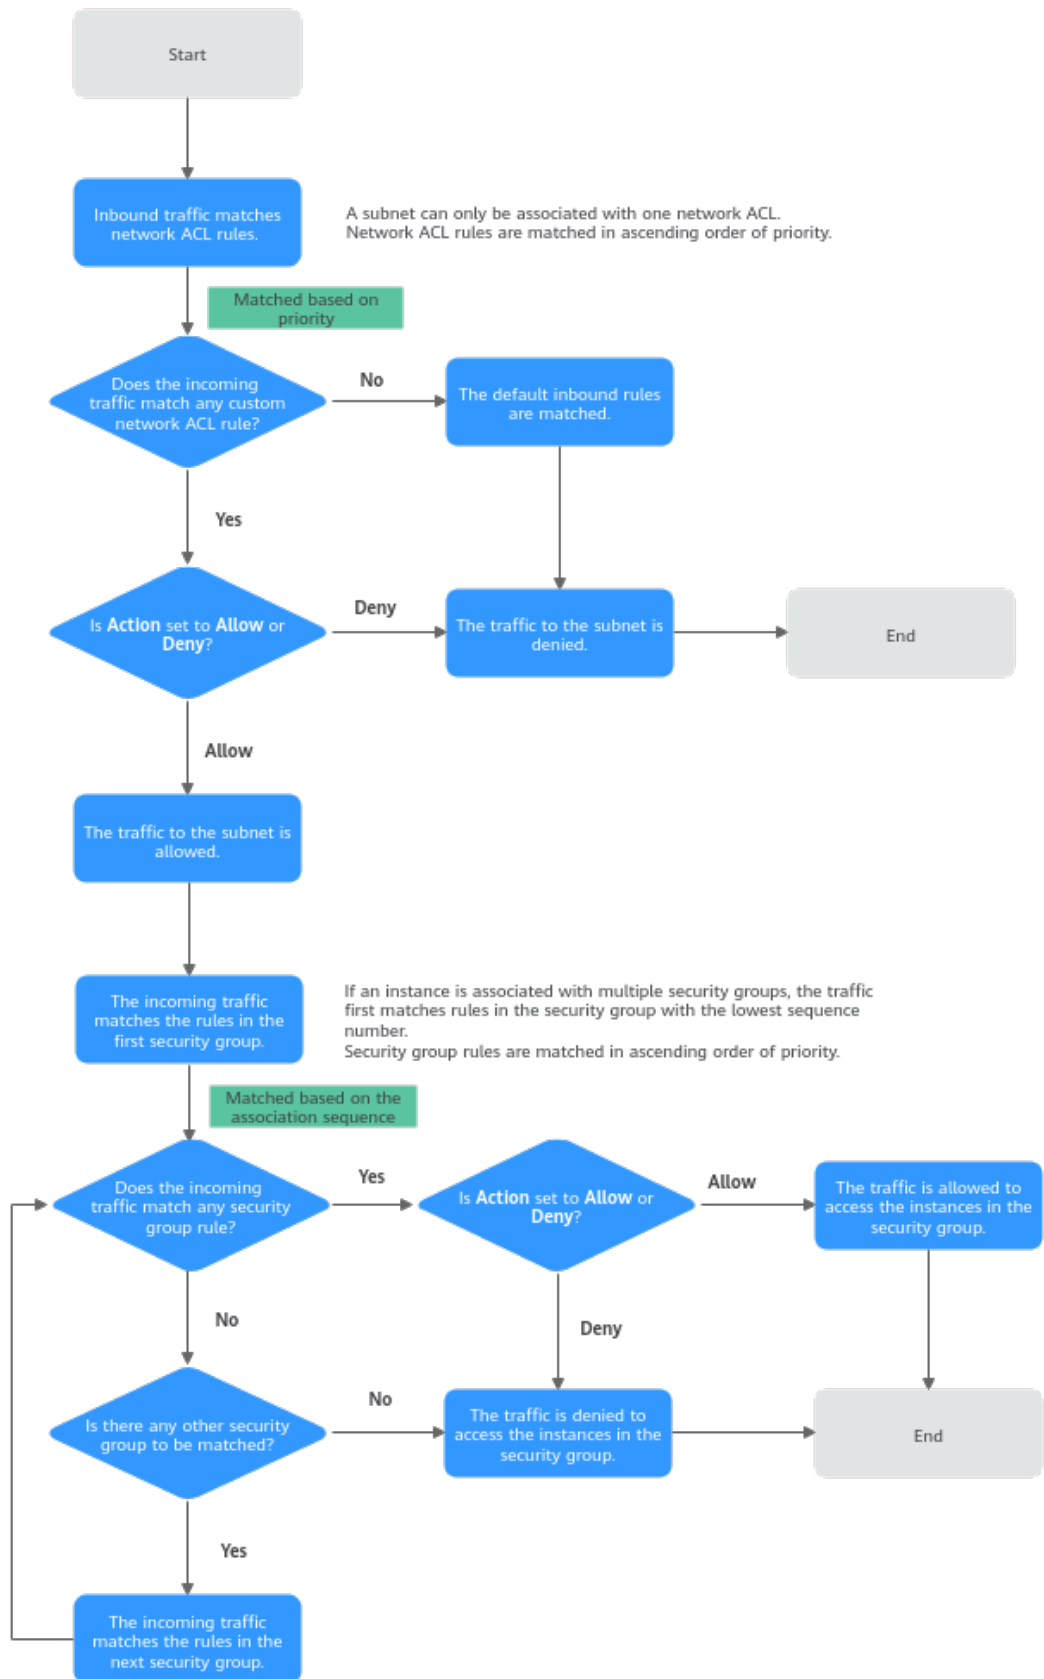
## How Traffic Matches Security Group and Network ACL Rules

If both security group and network ACL rules are configured, traffic matches network ACL rules first and then security group rules. [Figure 6-2](#) describes how inbound traffic matches security group and network ACL rules.

1. Traffic first matches network ACL rules.
  - If the traffic does not match any rule, the default rule is applied, and traffic to the subnet is denied.
  - If the traffic matches a rule, the rule is applied, which determines where the traffic will go.
    - If **Action** is set to **Deny**, the traffic to the subnet is denied.
    - If **Action** is set to **Allow**, the traffic to the subnet is allowed.
2. The traffic continues to match the security group rules.
  - a. If an instance is associated with multiple security groups, the traffic first matches rules in the security group with the lowest sequence number.
    - i. If the traffic does not match any rule, it is denied to access the instance.
    - ii. If the traffic matches a rule, the rule determines where the traffic will go.
      - If **Action** is set to **Deny**, the traffic is denied to access the instance.
      - If **Action** is set to **Allow**, the traffic is allowed to access the instance.
  - b. If the traffic fails to match the rules in the first security group, it continues to match the rules in the second security group.

- c. If the traffic does not match the rules of all security groups, the traffic is denied.

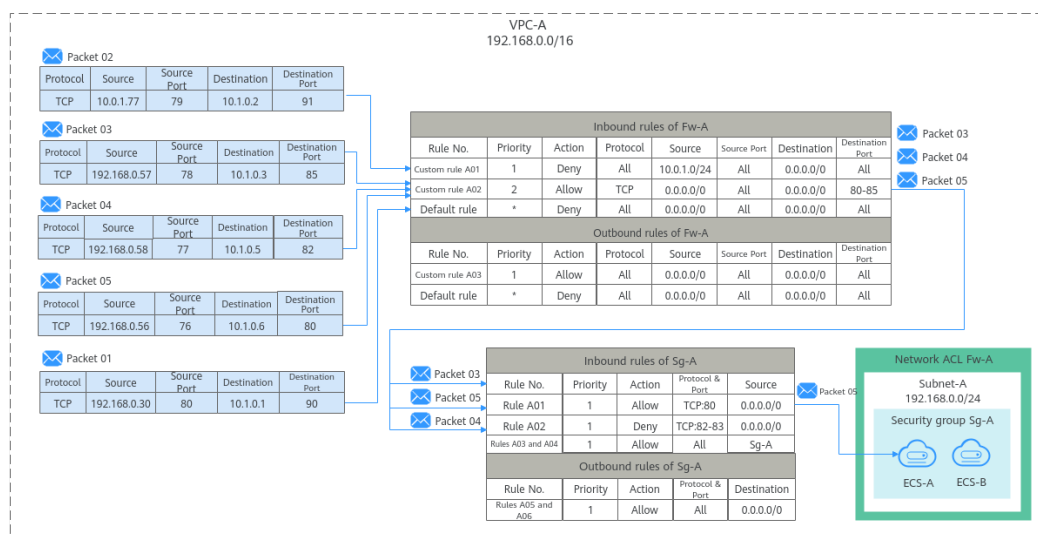
**Figure 6-2** How inbound traffic matches security group and network ACL rules



As shown in **Figure 6-3**, there is a subnet (**Subnet-A**) in **VPC-A**, and two ECSs (**ECS-A** and **ECS-B**) are running in this subnet. To protect your resources in **VPC-A**, you:

- Associate a network ACL (**Fw-A**) with **Subnet-A**. The default rules in **Fw-A** cannot be deleted. Traffic preferentially matches the rules you have configured. **Table 6-2** shows some example rules.
- Create a security group **Sg-A** to protect the ECSs. When creating security group **Sg-A**, you can select an existing template. The template comes with some default rules. You can modify or delete default rules, or add rules. For details about security group rules, see **Table 6-3**.

**Figure 6-3** How inbound traffic matches security group and network ACL rules



**Table 6-2** Rules configured for Fw-A

Direction	Priority	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Inbound	1	IPv4	Deny	All	10.0.1.0/24	All	0.0.0.0/0	All	Custom rule A01: denies traffic from 10.0.1.0/24 to the subnet.
Inbound	2	IPv4	Allow	TCP	0.0.0.0/0	All	0.0.0.0/0	80-85	Custom rule A02: allows all TCP traffic to the ECS in the subnet over ports 80 to 85.
Inbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule: denies all inbound traffic.



Direction	Priority	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Outbound	1	IPv4	Allow	All	0.0.0.0/0	All	0.0.0.0/0	All	Custom rule A03: allows all outbound traffic.
Outbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule: denies all outbound traffic.

**Table 6-3** Rules configured for Sg-A

Direction	Priority	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	1	Allow	IPv4	TCP: 80	Source: 0.0.0.0/0	Rule A01: allows all IPv4 traffic to the ECS over port 80.
Inbound	1	Deny	IPv4	TCP: 82-83	Source: 0.0.0.0/0	Rule A02: denies all IPv4 traffic to the ECS over ports 82 and 83.
Inbound	1	Allow	IPv4	All	Source: current security group ( <b>Sg-A</b> )	Rule A03: allows all IPv4 traffic in the security group so the ECS can communicate with other instances in the security group.
Inbound	1	Allow	IPv6	All	Source: current security group ( <b>Sg-A</b> )	Rule A04: allows all IPv6 traffic in the security group so the ECS can communicate with other instances in the security group.
Outbound	1	Allow	IPv4	All	Destination: 0.0.0.0/0	Rule A05: allows all traffic from the ECS in the security group to any IPv4 address.
Outbound	1	Allow	IPv6	All	Destination: ::/0	Rule A06: allows all traffic from the ECS in the security group to any IPv6 address.

Based on the preceding scenarios, different inbound packets match rules as follows:

- **Packet 01:** If no custom rules in **Fw-A** are matched, the default rule is applied, denying packet 01 to the subnet.
- **Packet 02:** If custom rule A01 in **Fw-A** is matched, this rule is applied, denying packet 02 to the subnet.
- **Packet 03:** If custom rule A02 in **Fw-A** is matched, this rule is applied, allowing packet 03 to the subnet. Packet 03 continues to match the security group rules. If it does not match any inbound rule in **Sg-A**, packet 03 is denied.
- **Packet 04:** If custom rule A02 in **Fw-A** is matched, this rule is applied, allowing packet 04 to the subnet. Packet 04 continues to match the security group rules. If it matches rule A02 in **Sg-A**, packet 04 is denied.
- **Packet 05:** If custom rule A02 in **Fw-A** is matched, this rule is applied, allowing packet 05 to the subnet. Packet 05 continues to match the security group rules. If it matches rule A01 in **Sg-A**, packet 05 is allowed.

## 6.2 Security Group

### 6.2.1 Security Groups and Security Group Rules

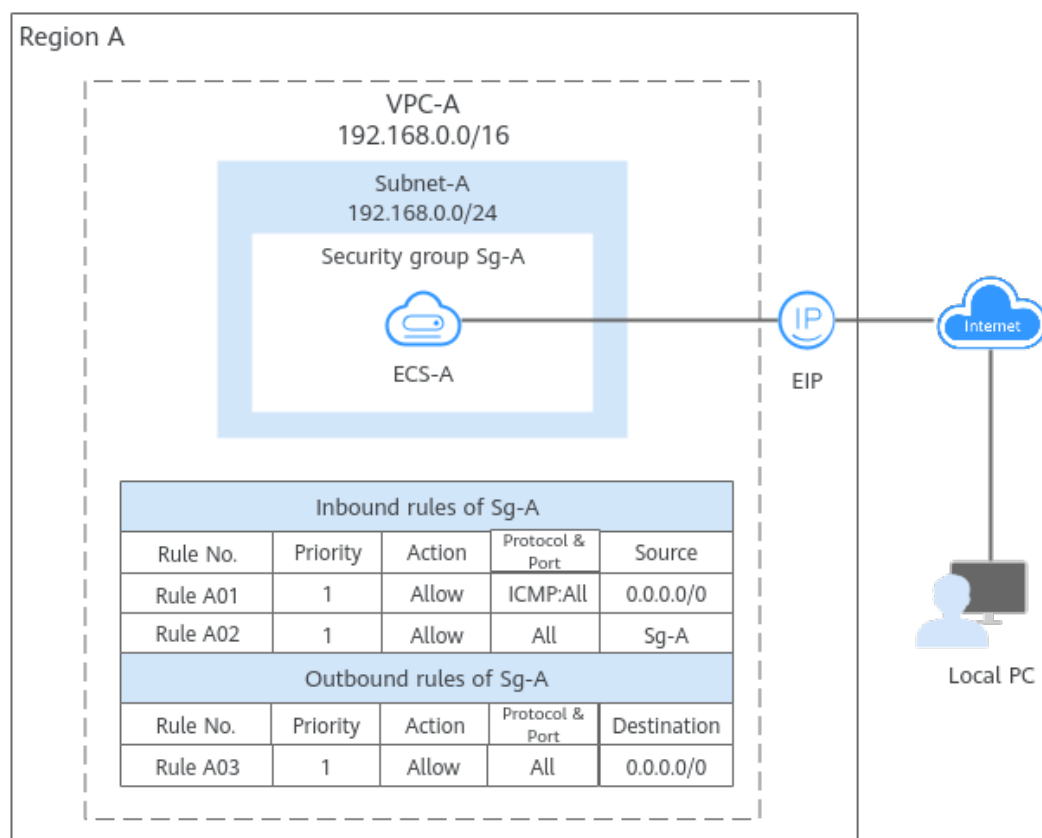
#### What Is a Security Group?

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.

When creating an instance (for example, an ECS), you must associate it with a security group. If there are no security groups yet, a **default security group** will be automatically created and associated with the instance. You can also create a security group based on service requirements and associate it with the instance. A cloud resource can be associated with multiple security groups, and traffic to and from the cloud resource is matched by priority in a descending order.

Each security group can have both inbound and outbound rules. You need to specify the source, port, and protocol for each inbound rule and specify the destination, port, and protocol for each outbound rule to control the inbound and outbound traffic to and from the instances in the security group. As shown in **Figure 6-4**, you have a VPC (**VPC-A**) with a subnet (**Subnet-A**) in region A. An ECS (**ECS-A**) is running in **Subnet-A** and associated with security group **Sg-A**.

- Security group **Sg-A** has a custom inbound rule to allow ICMP traffic to **ECS-A** from your PC over all ports. However, the security group does not contain rules that allow SSH traffic to **ECS-A** so you cannot remotely log in to **ECS-A** from your PC.
- If **ECS-A** needs to access the Internet through an EIP, the outbound rule of **Sg-A** must allow all traffic from **ECS-A** to the Internet.

**Figure 6-4** Security group architecture**NOTE**

You can use security groups free of charge.

## What Are 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.
  - Inbound rules: control traffic to the instances in a security group.
  - Outbound rules: control traffic from the instances in a security group for accessing external networks.
- You can specify protocol, port, source or destination for a security group rule. The following describes key information about a security group.
  - **Action: Allow or Deny.** If the protocol, port, source or destination of the traffic matches a security group rule, traffic will be allowed or denied.
  - **Priority:** The value ranges from 1 to 100. A smaller value indicates a higher priority. Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see [How Traffic Matches Security Group Rules](#).
  - **Type:** IPv4 or IPv6.
  - **Protocol & Port:** network protocol type and port range.
    - Network protocol: The protocol can be TCP, UDP, ICMP, or GRE.

- Port range: The value ranges from 1 to 65535.
- **Source or Destination:** source address of traffic in the inbound direction or destination address of traffic in the outbound direction.  
The source or destination can be an IP address, security group, or IP address group.
  - IP address: a fixed IP address or CIDR block. Both IPv4 and IPv6 addresses are supported, for example, 192.168.10.10/32 (IPv4 address), 192.168.1.0/24 (IPv4 CIDR), or 2407:c080:802:469::/64 (IPv6 CIDR).
  - Security group: If the selected security group and the current security group are in the same region, the traffic is allowed or denied to the private IP addresses of all instances in the selected security group. For example, if there is instance A in security group A and instance B in security group B, and the inbound rule of security group A allows traffic from security group B, traffic is allowed from instance B to instance A.
  - IP address group: If you have multiple IP addresses with the same security requirements, you can add them to an **IP address group** and select this IP address group when you configure a rule, to help you manage them in a more simple way.

## How Security Groups Work

- 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 be applied 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 are like a whitelist. If there are no rules that allow or deny some traffic, the security group denies all traffic to or from the instances in the security group.
  - Inbound rules: If the source of a request matches the source specified in a rule with **Action** set to **Allow**, the request is allowed. For this reason, you do not need to configure a deny rule in the inbound direction.  
The rules in [Table 6-4](#) ensure that instances in a security group can communicate with each other. Do not delete or modify these rules.
  - Outbound rules: The rules in [Table 6-4](#) allow all traffic to leave the instances in the security group so that the instances can access any external IP address. If you delete these rules, the instances in the security group cannot access external networks.

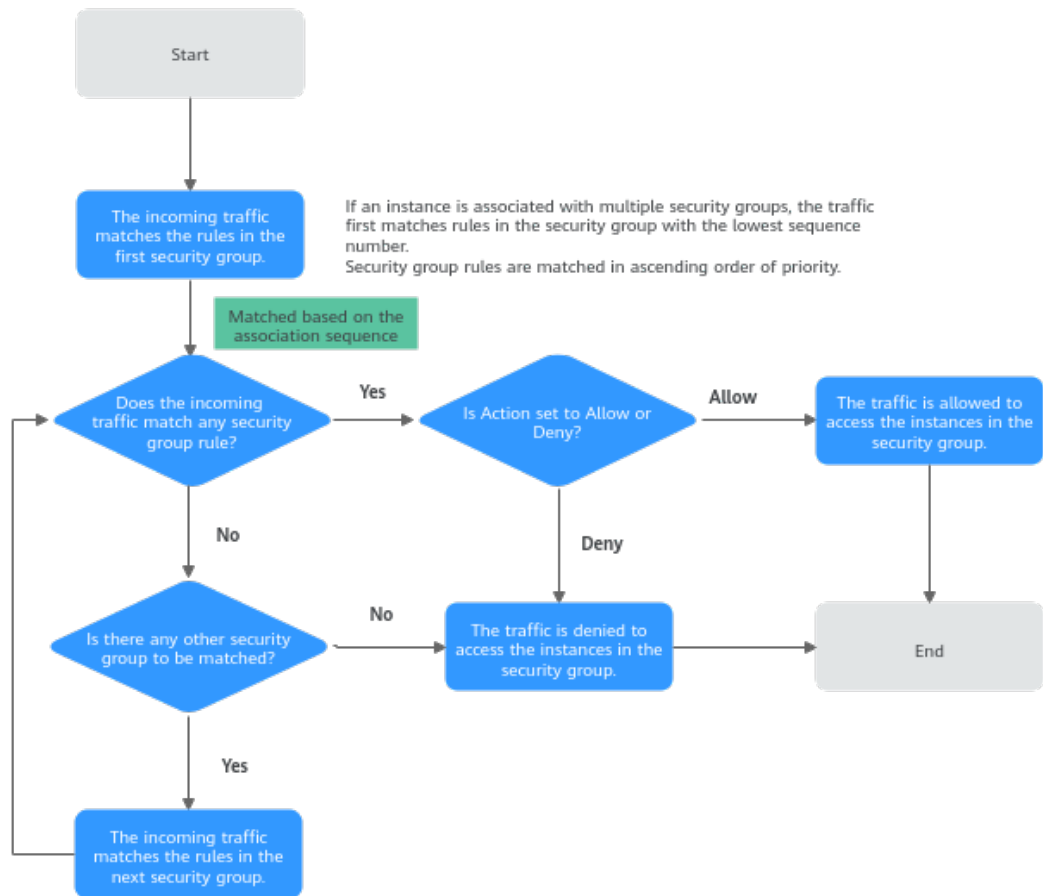
**Table 6-4** Security group rules

Direction	Action	Type	Protocol & Port	Source/Destination
Inbound	Allow	IPv4	All	Source: current security group
Inbound	Allow	IPv6	All	Source: current security group
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0
Outbound	Allow	IPv6	All	Destination: ::/0

## How Traffic Matches Security Group Rules

An instance can have multiple security groups associated, and a security group can contain multiple security group rules. Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. The following takes inbound traffic as an example to match security group rules:

1. First, traffic is matched based on the sequence number of security groups. You can adjust the security group sequence. A smaller security group sequence number indicates a higher priority.  
If the sequence number of security group A is 1 and that of security group B is 2, the priority of security group A is higher than that of security group B. Traffic preferentially matches the inbound rules of security group A.
2. Second, traffic is matched based on the priorities and actions of security group rules.
  - a. Security group rules are matched by priority first. A smaller value indicates a higher priority.  
If the priority of security group rule A is 1 and that of security group rule B is 2, the priority of security group rule A is higher than that of security group rule B. Therefore, traffic preferentially matches security group rule A.
  - b. Deny rules take precedence over allow rules of the same priority.
3. Traffic matches all inbound rules of a security group based on the protocol, ports and source.
  - If the traffic matches a rule:
    - With **Action of Allow**, the traffic is allowed to access the instances in the security group.
    - With **Action of Deny**, the traffic is denied to access the instances in the security group.
  - If the traffic does not match any rule, the traffic is denied to access the instances in the security group.

**Figure 6-5** Security group matching sequence

## How Security Groups Are Used

You can allow given IP addresses to access instances in a security group, or allow access from another security group to enable instances in different security groups to communicate with each other. You can add security group rules to flexibly control the traffic in and out of a network to ensure network security. The following provides some examples on how security groups can be used.

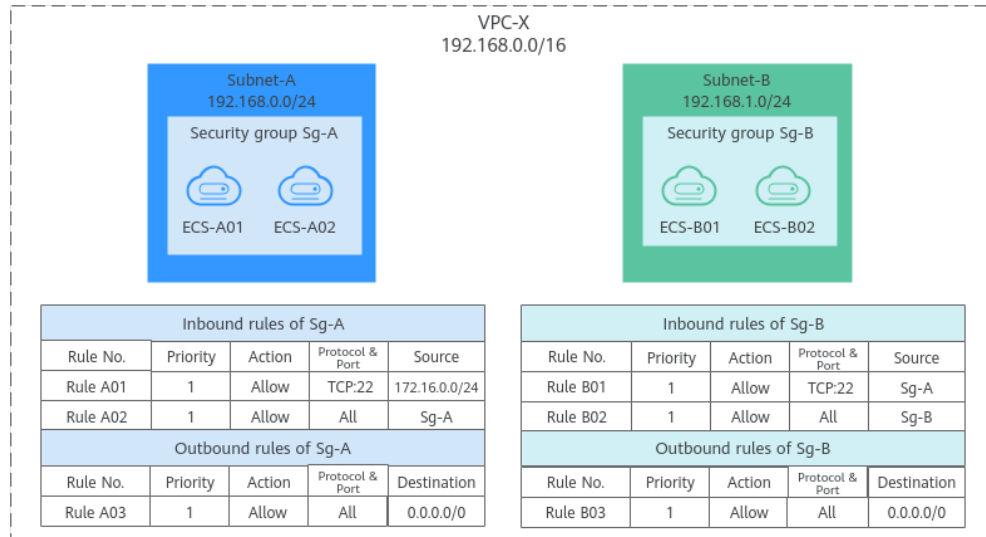
### Allowing Traffic from Given IP Addresses or a Security Group

As shown in [Figure 6-6](#), there are two subnets (**Subnet-A** and **Subnet-B**) in **VPC-X**. ECSs in **Subnet-A** are associated with **Sg-A** because these ECSs are used to run the same services and have the same network communication requirements. Similarly, ECSs in **Subnet-B** are associated with security group **Sg-B**.

- The inbound rule A01 of **Sg-A** allows traffic from IP addresses in **172.16.0.0/24** to access SSH port 22 of the ECSs in **Sg-A** for remotely logging in to these ECSs.
- The inbound rule A02 of **Sg-A** allows the ECSs in this security group to communicate with each other using any protocol and port.
- The inbound rule B01 of **Sg-B** allows the ECSs in **Sg-A** to access SSH port 22 of the ECSs in **Sg-B** for remotely logging in to the ECSs in **Subnet-B**.
- The inbound rule B02 of **Sg-B** allows the ECSs in this security group to communicate with each other using any protocol and port.

- The outbound rules of both security groups allow all traffic from the ECSs in the security groups.

**Figure 6-6** Allowing traffic from given IP addresses and security groups



**NOTE**

[Security Group Examples](#) lists more security group rule configuration examples.

## Allowing Traffic from a Virtual IP Address

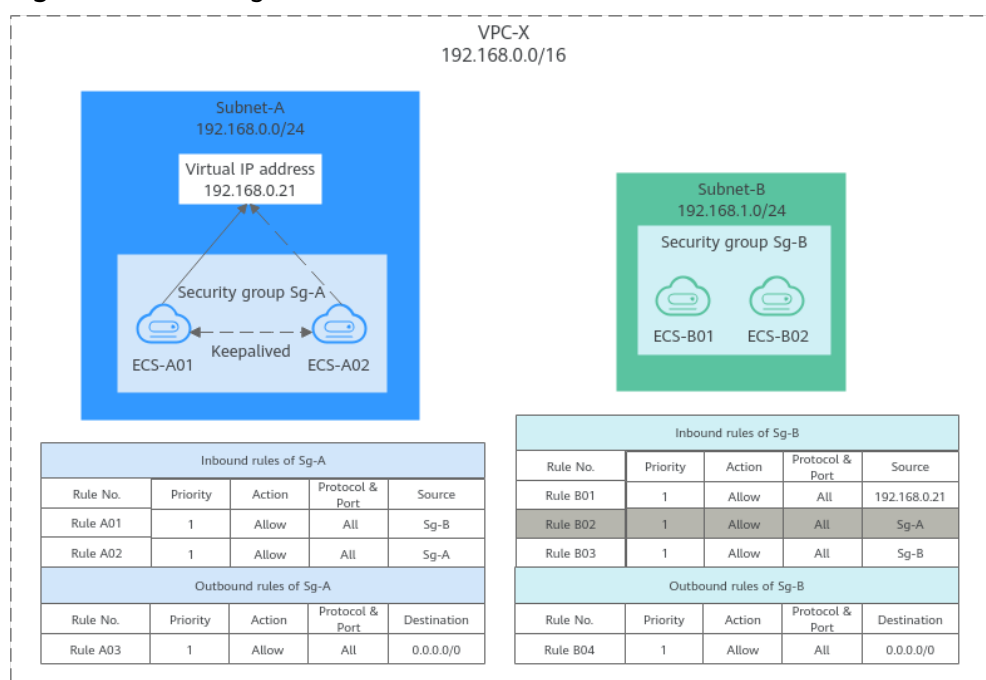
As shown in [Figure 6-7](#), you use virtual IP address **192.168.0.21** to connect the ECSs in **Subnet-A** and **Subnet-B**. If you set the source of an inbound rule to the security group associated with the ECSs, the ECSs in the two security groups cannot communicate with each other, because they are connected by a virtual IP address.

In [Figure 6-7](#), **VPC-X** has two subnets: **Subnet-A** and **Subnet-B**. ECSs in **Subnet-A** are associated with security group **Sg-A**, and ECSs in **Subnet-B** are associated with security group **Sg-B**. **ECS-A01** and **ECS-A02** work in active/standby mode, forming a Keepalived HA cluster. The ECSs use virtual IP address **192.168.0.21** to communicate with external networks.

- Inbound rule A01 of **Sg-A** allows ECSs in **Sg-B** to access ECSs in **Sg-A** using any protocol over any port.
- **Sg-B** has the following inbound rules:
  - Rule B02: Allows ECSs in **Sg-A** to use private IP addresses to access ECSs in **Sg-B**. However, in this networking, ECSs in **Sg-A** are supposed to communicate with ECSs in **Sg-B** through virtual IP address **192.168.0.21**. However, rule B02 does not allow traffic from this virtual IP address.
  - Rule B01: Allows traffic from virtual IP address **192.168.0.21** to ECSs in **Sg-B** using any protocol over port. In this networking, you can also set the source to **192.168.0.0/24**, the CIDR block of **Subnet-A**.



**Figure 6-7** Allowing traffic from a virtual IP address



**NOTE**


[Security Group Examples](#) lists more security group rule configuration examples.

## Allowing Communications Between Instances in Two VPCs Connected by a VPC Peering Connection

In [Figure 6-8](#), VPC-A and VPC-B in region A are connected by VPC peering connection **peering-AB**. After routes are configured for the VPC peering connection, **Subnet-A** and **Subnet-B** can communicate with each other. However, the ECSs in the two subnets are associated with different security groups. To allow ECSs in **Sg-A** and **Sg-B** to communicate with each other, you can add the following rules:

- Rule A01 with **Source** to **Sg-B** to allow ECSs in **Sg-B** to access ECSs in **Sg-A**.
- Rule B01 with **Source** to **Sg-A** to allow ECSs in **Sg-A** to access ECSs in **Sg-B**.

**Figure 6-8** Allowing communications between ECSs in two VPCs connected by a VPC peering connection

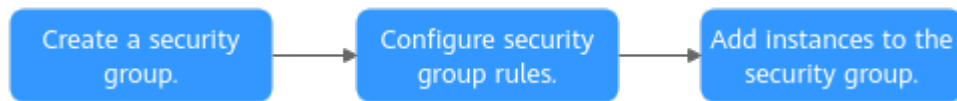


**NOTE**

[Security Group Examples](#) lists more security group rule configuration examples.

## Security Group Configuration Process

**Figure 6-9** Process of using a security group



**Table 6-5** Security group configuration process description

No.	Step	Description	Reference
1	Create a security group.	When creating a security group, you can select a template, such as <b>General-purpose web server</b> or <b>All ports open</b> . A template contains preset security group rules. For details, see <a href="#">Table 6-21</a> .	<a href="#">Creating a Security Group</a>
2	Configure security group rules.	After a security group is created, if its rules cannot meet your service requirements, you can add new rules to the security group or modify original rules.	<a href="#">Adding a Security Group Rule</a> <a href="#">Fast-Adding Security Group Rules</a>

No.	Step	Description	Reference
3	Add instances to the security group.	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.	<a href="#">Adding an Instance to or Removing an Instance from a Security Group</a>

## Notes and Constraints

- For better network performance, you are advised to associate an instance with no more than five security groups.
- A security group can have no more than 6,000 instances associated, or its performance will deteriorate.
- For inbound security group rules, the sum of the rules with **Source** set to **Security group**, the rules with **Source** set to **IP address group**, and the rules with inconsecutive ports, cannot exceed 128. Outbound rules also have this restriction.
  - When **Source** is set to **Security group**, you can select the current security group or a different security group.
  - An example of inconsecutive ports is 22,25,27.
- If you specify an IP address group or inconsecutive ports for a security group rule, the rule is only applied for certain ECSs. For details, see [Table 6-6](#).

**Table 6-6** Scenarios that security group rules do not take effect

Rule Configuration	ECS Type
<b>Source</b> or <b>Destination</b> is set to <b>IP address group</b> .	The following x86 ECS types are not supported: <ul style="list-style-type: none"> <li>• General computing (S1, C1, and C2 ECSs)</li> <li>• Memory-optimized (M1 ECSs)</li> <li>• High-performance computing (H1 ECSs)</li> <li>• Disk-intensive (D1 ECSs)</li> <li>• GPU-accelerated (G1 and G2 ECSs)</li> <li>• Large-memory (E1, E2, and ET2 ECSs)</li> </ul>
<b>Port</b> is set to non-consecutive ports.	The following x86 ECS types are not supported: <ul style="list-style-type: none"> <li>• General computing (S1, C1, and C2 ECSs)</li> <li>• Memory-optimized (M1 ECSs)</li> <li>• High-performance computing (H1 ECSs)</li> <li>• Disk-intensive (D1 ECSs)</li> <li>• GPU-accelerated (G1 and G2 ECSs)</li> <li>• Large-memory (E1, E2, and ET2 ECSs)</li> </ul>

Rule Configuration	ECS Type
	<p>All Kunpeng ECS flavors do not support inconsecutive ports.</p> <p>If you use inconsecutive port numbers in a security group rule of a Kunpeng ECS, this rule and rules configured after this one do not take effect.</p> <p>If you configure security group rule A with inconsecutive ports <b>22,24</b> and then configure security group rule B with port 9096, both rule A and rule B do not take effect.</p>

 NOTE

- For details about x86 ECSs, see [ECS Specifications \(x86\)](#).
- For details about Kunpeng ECSs, see [ECS Specifications \(Kunpeng\)](#).
- Traffic from load balancers is not restricted by network ACL and security group rules if:

**Transfer Client IP Address** is enabled for the listener of a load balancer.

The load balancer can still forward traffic to backend servers, even if there is a rule that denies traffic from the load balancer to the backend servers.

## Recommendations

- Instances in a security group deny all external access requests by default, but you can add rules to allow specific requests.
- When adding a security group rule, grant the minimum permissions possible. For example, if remote login to an ECS over port 22 is allowed, only allow specific IP addresses to log in to the ECS. Do not use 0.0.0.0/0 (all IP addresses).
- Keep your configurations simple. There should be a different reason for each security group. If you use the same security group for all your different instances, the rules in the security group will likely be redundant and complex. It will make it much harder to maintain and manage.
- You can add instances to different security groups based on their functions. For example, if you want to provide website services accessible from the Internet, you can add the web servers to a security group configured for that specific purpose and only allow external access over specific ports, such as 80 and 443. By default, other external access requests are denied. Do not run internal services, such as MySQL or Redis, on web servers that provide services accessible from the Internet. Deploy internal services on servers that do not need to connect to the Internet and associate these servers with security groups specifically configured for that purpose.
- If you have multiple IP addresses with the same security requirements, you can add them to an IP address group and select this IP address group when you configure a rule, to help you manage them in a more simple way. When an IP address changes, you only need to change the IP address in the IP address group. Then, the rules in the IP address group change accordingly. You

do not need to modify the rules in the security group one by one. This simplifies security group management and improves efficiency. For details, see [Using IP Address Groups to Reduce the Number of Security Group Rules](#).

- Do not directly modify security group rules for active services. 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. For details, see [Cloning a Security Group](#).
- After you add instances to or modify rules of a security group, the security group rules are applied automatically. There is no need to restart the instances.

If a security group rule does not take effect after being configured, see [Why Are My Security Group Rules Not Applied?](#)

## 6.2.2 Default Security Groups

If no security groups have been created yet, a default security group is automatically created for you, and the instance will be associated with it when you are creating the instance. Note the following when using the default security group:

- The name of the default security group is **default**. It is recommended that you do not change the name of the default security group in order to distinguish it from any security groups that you may create.
- You cannot delete the default security group, but you can modify its rules or add rules to it.
- The default security group denies all external requests. To allow access to an instance associated with this security group, you can add rules to allow access over given ports by referring to [Remotely Logging In to an ECS from a Local Server](#).
- If your service has different security requirements on instances for different purposes, you can create security groups and associate these instances with different security groups based on their purposes.

### NOTE

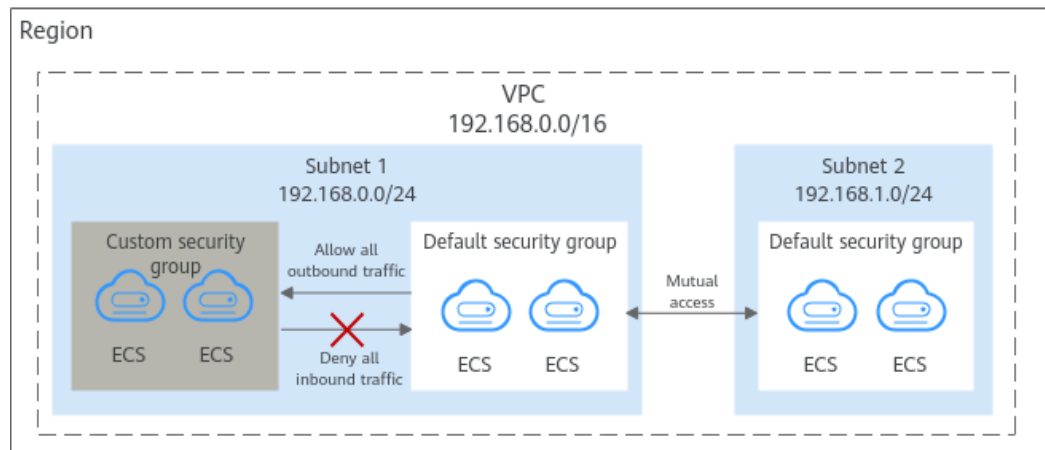
Security groups are free of charge.

## Default Security Group Rules

Note the following when using default security group rules:

- Inbound rules control incoming traffic to instances in the default security group. The instances can only communicate with each other but cannot be accessed from external networks.
- Outbound rules allow all traffic from the instances in the default security group to external networks.

**Figure 6-10** Default security group



**Table 6-7** describes the default rules for the default security group.

**Table 6-7** Default security group rules

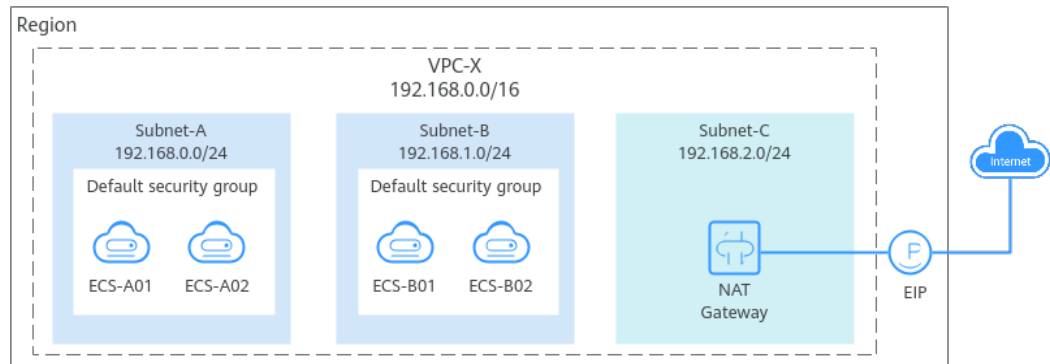
Direction	Action	Type	Protocol & Port	Source/ Destination	Description
Inbound	Allow	IPv4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inbound	Allow	IPv6	All	Source: default security group (default)	Allows IPv6 instances in the security group to communicate with each other using any protocol over any port.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows all traffic from the instances in the security group to any IPv4 address over any port.
Outbound	Allow	IPv6	All	Destination: :/0	Allows all traffic from the instances in the security group to any IPv6 address over any port.

## A Default Security Group Example

As shown in **Figure 6-11**, VPC-X has three subnets: **Subnet-A**, **Subnet-B**, and **Subnet-C**. ECSs in **Subnet-A** and **Subnet-B** have been associated with the default security group. The default security group allows the instances in the security group to communicate with each other and denies all external requests. So, the four ECSs (**ECS-A01**, **ECS-A02**, **ECS-B01**, and **ECS-B02**) can communicate with each other, but they cannot receive traffic from the NAT gateway.

To allow traffic from the NAT gateway, you need to add rules to the default security group or create a security group and associate it with the instances.

**Figure 6-11** Use cases



## 6.2.3 Security Group Examples

When you create instances, such as cloud servers, containers, and databases, in a VPC subnet, you can use the default security group or create a security group. You can add inbound and outbound rules to the default or your security group to control traffic from and to the instances in the security group. Here are some common security group configuration examples:

- [Remotely Logging In to an ECS from a Local Server](#)
- [Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files](#)
- [Setting Up a Website on an ECS to Provide Services Externally](#)
- [Using ping Command to Verify Network Connectivity](#)
- [Enabling Communications Between Instances in Different Security Groups](#)
- [Allowing External Instances to Access the Database Deployed on an ECS](#)
- [Allowing ECSs to Access Specific External Websites](#)

### NOTICE

If your security group rules are not applied, [submit a service ticket](#).

## Precautions

Note the following before configuring security group rules:

- Instances associated with different security groups are isolated from each other by default.
- Generally, a security group denies all external requests by default. You need to add inbound rules to allow specific traffic to the instances in the security group.
- By default, outbound security group rules allow all requests from the instances in the security group to access external resources.

If outbound rules are deleted, the instances in the security group cannot communicate with external resources. To allow outbound traffic, you need to add outbound rules by referring to [Table 6-8](#).

**Table 6-8** Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 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 in a security group from a local server, add an inbound rule based on the OS running on the ECS.

- To remotely log in to a Linux ECS using SSH, enable port 22. For details, see [Table 6-9](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 6-10](#).

**Table 6-9** Remotely logging in to a Linux ECS using SSH

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 22	IP address: 0.0.0.0/0

**Table 6-10** Remotely logging in to a Windows ECS using RDP

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3389	IP address: 0.0.0.0/0



**NOTICE**

If the source is set to 0.0.0.0/0, any IP address can be used to remotely log in to the ECS. To ensure security, set the source to a specific IP address based on service requirements. For details about the configuration example, see [Table 6-11](#).

**Table 6-11** Remotely logging in to an ECS using a specified IP address

ECS Type	Direction	Priority	Action	Type	Protocol & Port	Source
Linux ECS	Inbound	1	Allow	IPv4	TCP: 22	IP address: 192.168.0.0/24
Windows ECS	Inbound	1	Allow	IPv4	TCP: 3389	IP address: 10.10.0.0/24

## Remotely Connecting to an ECS from a Local Server to Upload or Download FTP 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 6-12** Remotely connecting to an ECS from a local server to upload or download files

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	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 6-13** Setting up a website on an ECS to provide services externally

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 80	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv4	TCP: 443	IP address: 0.0.0.0/0

## Using ping Command to Verify Network Connectivity

Ping works by sending an Internet Control Message Protocol (ICMP) Echo Request. To ping an ECS from your PC to verify the network connectivity, you need to add an inbound rule to the security group of the ECS to allow ICMP traffic.

**Table 6-14** Using ping command to verify network connectivity

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	ICMP: All	IP address: 0.0.0.0/0
Inbound	1	Allow	IPv6	ICMP: All	IP address: ::/0

## Enabling Communications Between Instances in Different Security Groups

Instances in the same VPC but associated with different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** 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**.

**Table 6-15** Enabling communications between instances in different security groups

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A

### NOTICE

As shown in [How Security Groups Are Used](#), if you want to use virtual IP address **192.168.0.21** to connect the ECSs in **Subnet-A** and **Subnet-B**, you need to set the source of an inbound rule to virtual IP address **192.168.0.21**.

## Allowing External Instances to Access the Database Deployed on an ECS

A security group denies all external requests by default. If you have deployed a database on an ECS and want the database to be accessed from external instances

on a private network, you need to add an inbound rule to the security group of the ECS to allow access over corresponding ports. Here are some common ports for databases:

- MySQL: port 3306
- Oracle: port 1521
- MS SQL: port 1433
- PostgreSQL: port 5432
- Redis: port 6379

**Table 6-16** Allowing external instances to access the database deployed on an ECS

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A	Allows the ECSs in security group <b>sg-A</b> to access the MySQL database service.
Inbound	1	Allow	IPv4	TCP: 1521	Security group: sg-B	Allows the ECSs in security group <b>sg-B</b> to access the Oracle database service.
Inbound	1	Allow	IPv4	TCP: 1433	IP address: 172.16.3.21/32	Allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database service.
Inbound	1	Allow	IPv4	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.
Inbound	1	Allow	IPv4	TCP: 6379	IP address group: ipGroup-A	Allows ECSs whose private IP addresses are in IP address group <b>ipGroup-A</b> to access the Redis database service.

#### NOTICE

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

## Allowing ECSs to Access Specific External Websites

By default, a security group allows all outbound traffic. [Table 6-18](#) lists the default rules. If you want to allow ECSs to access specific websites, configure the security group as follows:

1. Add outbound rules to allow traffic over specific ports to specific IP addresses.

**Table 6-17** Allowing ECSs to access specific external websites

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	TCP: 80	IP address: 132.15.XX.XX	Allows ECSs in the security group to access the external website at http://132.15.XX.XX:80.
Outbound	1	Allow	IPv4	TCP: 443	IP address: 145.117.XX.XX	Allows ECSs in the security group to access the external website at https://145.117.XX.XX:443.

2. Delete the original outbound rules that allow all traffic.

**Table 6-18** Default outbound rules in a security group

Direction	Priority	Action	Type	Protocol & Port	Destination	Description
Outbound	1	Allow	IPv4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outbound	1	Allow	IPv6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

### 6.2.4 Common Ports Used by ECSs

When adding a security group rule, you must specify a port or port range for communications. Traffic is then allowed or denied if traffic matches this rule. Suppose a client requests to remotely log in to an ECS using SSH. When the request reaches the security group, the IP address and port of the client will be checked. If the IP address and the port match the allow rules in the security group, the request is allowed.

[Table 6-19](#) lists some high-risk ports that are blocked by default. Even if you have added a security group rule to allow access over these ports, traffic over these

ports in restricted regions is still denied. In this case, do not use these high-risk ports for your services.

**Table 6-19** High-risk ports

Protocol	Port
TCP	42, 135, 137, 138, 139, 444, 445, 593, 1025, 1068, 1433, 1434, 3127, 3128, 3129, 3130, 4444, 4789, 5554, 5800, 5900, 8998, 9995, and 9996
UDP	135~139 1026 1027 1028 1068 1433 1434 4789 5554 9995 9996

## Common Ports

**Table 6-20** lists the common ports used by ECSs. You can configure security group rules to allow traffic to and from specified ECS ports. For details, see [Adding a Security Group Rule](#). For more information about requirements for Windows, see [Service overview and network port requirements for Windows](#).

**Table 6-20** Common ports used by ECSs

Port	Protocol	Description
21	FTP	Used by FTP services for uploading and downloading files. For configuration examples, see <a href="#">Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files</a> .
22	SSH	Used to remotely connect to Linux ECSs. For configuration examples, see <a href="#">Remotely Logging In to an ECS from a Local Server</a> . For details about how to log in to a Linux ECS, see <a href="#">Linux ECS Login Overview</a> .
23	Telnet	Used to remotely log in to ECSs.
25	SMTP	Used to send emails. For security purposes, TCP port 25 is disabled in the outbound direction by default. For details about how to open the port, see <a href="#">Why Is Outbound Access Through TCP Port 25 Restricted?</a>
80	HTTP	Used to access websites over HTTP. For configuration examples, see <a href="#">Setting Up a Website on an ECS to Provide Services Externally</a> .
110	POP3	Used to receive emails using Post Office Protocol version 3 (POP3).
143	IMAP	Used to receive emails using Internet Message Access Protocol (IMAP).

Port	Protocol	Description
443	HTTPS	Used to access websites over HTTPS. For configuration examples, see <a href="#">Setting Up a Website on an ECS to Provide Services Externally</a> .
1433	SQL Server	A TCP port of the SQL Server for providing services. For configuration examples, see <a href="#">Allowing External Instances to Access the Database Deployed on an ECS</a> .
1434	SQL Server	A UDP port of the SQL Server for returning the TCP/IP port number used by the SQL Server. For configuration examples, see <a href="#">Allowing External Instances to Access the Database Deployed on an ECS</a> .
1521	Oracle	Used for Oracle database communications. This port must be enabled on the ECSs where Oracle SQL Server is deployed. For configuration examples, see <a href="#">Allowing External Instances to Access the Database Deployed on an ECS</a> .
3306	MySQL	Used by MySQL databases to provide services. For configuration examples, see <a href="#">Allowing External Instances to Access the Database Deployed on an ECS</a> .
3389	Windows Server Remote Desktop Services	Used to connect to Windows ECSs. For configuration examples, see <a href="#">Remotely Logging In to an ECS from a Local Server</a> . For details about how to log in to a Windows ECS, see <a href="#">Windows ECS Login Overview</a> .
8080	Proxy	Used by the WWW proxy service for web browsing, like port 80. If you use port 8080, you need to add <b>:8080</b> after the IP address when you visit a website or use a proxy server. If Apache Tomcat is installed, its default service port is 8080.
137, 138, and 139	NetBIOS	Used for Windows files, printer sharing, and Samba. <ul style="list-style-type: none"><li>• Ports 137 and 138: UDP ports that are used when files are transferred using Network Neighborhood (My Network Places).</li><li>• Port 139: Connections from this port try to access the NetBIOS/SMB service.</li></ul>

## 6.2.5 Managing a Security Group

## 6.2.5.1 Creating a Security Group

### Scenarios

A security group consists of inbound and outbound rules. You can add security group rules to allow or deny the traffic to reach and leave the instances (such as ECSs) in the security group.

When creating an instance (for example, an ECS), you must associate it with a security group. If no security group has been created yet, a **default security group** will be created and associated with the instance. You can also create a security group and add inbound and outbound rules to allow specific traffic. For more information about security groups and rules, see [Security Groups and Security Group Rules](#).

### Security Group Templates

Several security group templates are provided for you to create a security group. A security group template has preconfigured inbound and outbound rules. You can select a template based on your service requirements. [Table 6-21](#) describes the security group templates.

**Table 6-21** Security group rules

Template	Direction	Protocol/Port/Type	Source/Destination	Description	Scenario
General - purpose web server	Inbound	TCP: 22 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in the security group over port 22 (SSH) for remotely logging in to Linux instances.	<ul style="list-style-type: none"> <li>Remotely log in to an instance (such as an ECS) in a security group from an external network.</li> <li>Enable external servers to ping the instances in a security group to verify network connectivity.</li> <li>Use instances in a security group as web servers to provide website services accessible from the Internet.</li> </ul>
		TCP: 3389 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 3389 (RDP) for remotely logging in to Windows instances.	
		TCP: 80 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 80 (HTTP) for visiting websites.	
		TCP: 443 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 443 (HTTPS) for visiting websites.	
		ICMP: All (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over any port for using the ping command to test connectivity.	
		All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network using any protocol.	



Template	Direction	Protocol/Port/Type	Source/Destination	Description	Scenario
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources using any protocol.	
All ports open	Inbound	All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network using any protocol.	Allow any traffic to enter and leave a security group over any port may be risky.
		All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows any IP address to access the instances in a security group over any port.	
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources using any protocol.	
Fast-add rule	Inbound	All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network.	You can select protocols and ports that the inbound rule will apply to.  If you do not select any protocols and ports, no protocols and ports will be opened. After the security group is created, add required rules by referring to <a href="#">Adding a Security Group Rule</a> .
		Custom port and protocol	0.0.0.0/0	Allows all IP addresses to access ECSs in a security group over specified ports (TCP or ICMP) for different purposes.	
	Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources using any protocol.	

## Procedure

1. Go to the [security group list page](#).
2. In the upper right corner, click **Create Security Group**.  
The **Create Security Group** page is displayed.
3. Configure the parameters as prompted.

**Table 6-22** Parameter description

Parameter	Description	Example Value
Name	<p>Mandatory</p> <p>The name of the security group. The name:</p> <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul> <p><b>NOTE</b> 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.</p>	sg-AB
Enterprise Project	<p>Mandatory</p> <p>When creating a security group, you can add the security group to an enabled enterprise project.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is <b>default</b>.</p> <p>For details about creating and managing enterprise projects, see the <a href="#">Enterprise Management User Guide</a>.</p>	default
Tag	<p>Optional</p> <p>You can add tags to the security group. Tags help you to identify, classify, and search for your security groups.</p> <p>For details, see <a href="#">Managing Security Group Tags</a>.</p>	<p><b>Tag key:</b> test</p> <p><b>Tag value:</b> 01</p>
Template	<p>Mandatory</p> <p>The system provides several security group templates for you to create a security group. A security group template has preconfigured inbound and outbound rules. You can select a template based on your service requirements.</p> <p><a href="#">Table 6-21</a> describes the security group templates.</p>	General-purpose web server

Parameter	Description	Example Value
Description	Optional Supplementary information about the security group. The security group description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

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

## Related Operations

- After a security group is created, if its rules cannot meet your service requirements, you can add new rules to the security group or modify original rules. For details, see [Adding a Security Group Rule](#).
- Each ECS must be associated with at least one security group. You can add an ECS to multiple security groups based on service requirements. For details, see [Adding an Instance to or Removing an Instance from a Security Group](#).

### 6.2.5.2 Cloning a Security Group

#### Scenarios

You can clone a security group from the same or a different region to another to quickly apply the security group rules to ECSs in that 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.
- Only security group rules are cloned, but not the instances associated with the security group.

- Cloned security groups can only be used in the same account. To quickly create a security group across accounts, you can import or export security group rules by referring to [Importing and Exporting Security Group Rules](#).

## Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. Locate the row that contains the security group, click **More** in the **Operation** column, and click **Clone**.
6. Select the region and name of the new security group as prompted.
7. Click **OK**.  
You can then switch to the required region to view the cloned security group in the security group list.

### 6.2.5.3 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. 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.
6. Modify the name and description of the security group as required.
7. Click **OK** to save the modification.



## 6.2.5.4 Viewing the Details of a Security Group

### Scenarios

You can view the details of a security group, such as the security group name, security group rules, and the instances associated with this security group.

You can also search for a given security group by key information, such as the security group name, ID, and description.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The **Security Groups** page is displayed.
5. In the search box above the security group list, select filters to quickly search for the target security group.
6. Locate the target security group and click its name.  
The security group details page is displayed.
7. On the security group details page, click different tabs to view the following information:
  - **Summary**: security group name, ID, enterprise project, and description.
  - **Inbound Rules**: the priority, action, source, and modification time of an inbound rule.
  - **Outbound Rules**: the priority, action, destination, and modification time of an outbound rule.
  - **Associated Instances**: the information about instances associated with the security group. The instances include servers, extension NICs, supplementary network interfaces, and others.

## 6.2.5.5 Deleting a Security Group

### Scenarios

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

#### NOTE

Both default and custom security groups are free.

### Notes and Constraints

- The default security group is named **default** and cannot be deleted.

**Figure 6-12** Default security group

Name	ID	Instances	Rules	Created Time	Region	Operation
Sys-WebServer	13e7c18f-959c-454e-875c-e0385aa466da	13	1	Jul 09, 2022 15:29:49 GMT+08:00	default	Manage Rules Manage Instances More
Default	3b4d184a-8754-4686-923f-75d91e0233ab	0	0	Jul 09, 2022 15:29:48 GMT+08:00	default	Manage Rules Manage Instances Clone

- 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](#).



If you want to know the instances associated with a security group, refer to [How Do I Know the Instances Associated with 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. Locate the row that contains the target security group, click **More** in the **Operation** column, and click **Delete**.  
A confirmation dialog box is displayed.
6. Confirm the information and click **Yes**.

### 6.2.5.6 Managing Security Group Tags

#### Scenarios

Tags help you identify, classify, and search for security groups. You can refer to the following sections to manage the tags of a security group.

- [Adding a Tag to a Security Group](#)
- [Modifying a Tag of a Security Group](#)
- [Deleting a Tag from a Security Group](#)

[Table 6-23](#) lists the details about a security group tag.



**Table 6-23** Security group tag naming requirements

Parameter	Requirements	Example Value
Tag key	<ul style="list-style-type: none"><li>For each security group, each tag key must be unique, and each tag key can only have one tag value.</li><li>Cannot be left blank.</li><li>Can contain a maximum of 128 characters.</li><li>Can consist of letters, digits, underscores (_), and hyphens (-).</li></ul>	test
Tag value	<ul style="list-style-type: none"><li>Can be left blank.</li><li>Can contain a maximum of 256 characters.</li><li>Can consist of letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	01


## Notes and Constraints


- Each tag consists of a **tag key** and a **tag value**. Only the **tag value** can be edited.  
If you want to change the **tag key**, delete it and add one again.
- Each cloud resource can have a maximum of 20 tags.

## Adding a Tag to a Security Group



- Log in to the management console.
- Click  in the upper left corner and select the desired region and project.
- Click  in the upper left corner and choose **Networking > 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.
- Locate the target security group and click its name.  
The security group details page is displayed.
- On the **Tags** tab, click **Add Tag** in the upper left corner of the tag list.  
The **Add Tag** dialog box is displayed.
- Enter the tag key and value as required, and click **OK**.  
You can view the tag you have added in the tag list.

## Modifying a Tag of a Security Group

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

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. Locate the target security group and click its name.  
The security group details page is displayed.
6. Click the **Tags** tab, locate the target tag and click **Edit** in the **Operation** column.  
The **Edit Tag** dialog box is displayed.
7. Enter the tag value as needed and click **OK**.  
You can view the edited tag in the tag list.

## Deleting a Tag from a Security Group

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. Locate the target security group and click its name.  
The security group details page is displayed.
6. Click the **Tags** tab, locate the target tag and click **Delete** in the **Operation** column.  
The **Delete Tag** page is displayed.
7. Click **Yes** in the displayed dialog box.  
Return to the tag list, and confirm that the tag has been deleted.

## 6.2.6 Managing Security Group Rules

### 6.2.6.1 Adding a Security Group Rule

#### Scenarios

A security group consists of inbound and outbound rules. You can add security group rules to allow or deny the traffic to reach and leave the instances (such as ECSs) in the security group.

Security group rules allow or deny network traffic from specific sources over specific protocols or specific ports.



## Precautions

- Before configuring security group rules, you need to plan access policies for instances in the security group. For details about common security group rules, see [Security Group Examples](#).
- Add as fewer rules as possible. [Notes and Constraints](#) lists the constraints on the number of rules in a security group.
- After allowing traffic over a port in a security group rule, ensure that the port used by the instance is opened. For details, see [Verifying Security Group Rules](#).
- By default, instances in the same security group can communicate with each other. If instances in the same security group cannot communicate with each other, possible causes are as follows:
  - The inbound rules for communications between these instances are deleted. [Table 6-24](#) shows the inbound rules.




**Table 6-24** Inbound rules for communication between instances

Direction	Priority	Action	Type	Protocol & Port	Source/Destination
Inbound	1	Allow	IPv4	All	Source: current security group (Sg-A)
Inbound	1	Allow	IPv6	All	Source: current security group (Sg-A)

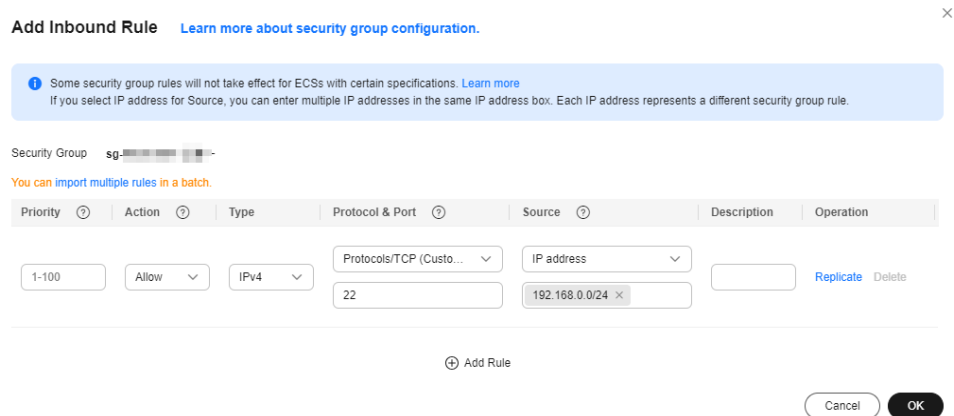
- Different VPCs cannot communicate with each other. The instances belong to the same security group but different VPCs.

You can use [VPC peering connections](#) to connect VPCs in different regions.

## Adding Rules to a Security Group

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. 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.
6. On the **Inbound Rules** tab, click **Add Rule**.  
The **Add Inbound Rule** dialog box is displayed.
7. Configure required parameters.  
You can click  to add more inbound rules.

**Figure 6-13** Add Inbound Rule



**Table 6-25** Inbound rule parameter description

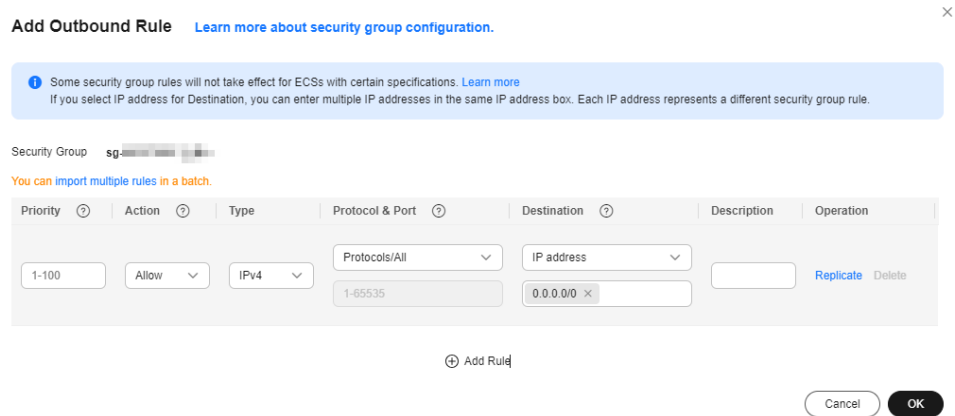
Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	Allow or Deny <ul style="list-style-type: none"> <li>If the <b>Action</b> is set to <b>Allow</b>, access from the source is allowed to ECSs in the security group over specified ports.</li> <li>If the <b>Action</b> is set to <b>Deny</b>, access from the source is denied to ECSs in the security group over specified ports.</li> </ul> Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see <a href="#">How Traffic Matches Security Group Rules</a> .	Allow
Type	Source IP address version. You can select: <ul style="list-style-type: none"> <li><b>IPv4</b></li> <li><b>IPv6</b></li> </ul>	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>GRE</b> , and <b>ICMP</b> .	TCP

Parameter	Description	Example Value
	<p>Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.</p> <p>Inbound rules control incoming traffic over specific ports to instances in the security group.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"><li>• Individual port: Enter a port, such as 22.</li><li>• Consecutive ports: Enter a port range, such as 22-30.</li><li>• Non-consecutive ports: Enter ports and port ranges, such as <b>22,23-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li><li>• All ports: Leave it empty or enter 1-65535.</li></ul>	22, or 22-30

Parameter	Description	Example Value
Source	<p>Used to match the IP address or address range of an external request. The source can be:</p> <ul style="list-style-type: none"> <li>• <b>IP address:</b> If you select <b>IP address</b> for <b>Source</b>, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. <ul style="list-style-type: none"> <li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul> </li> <li>• <b>Security group:</b> The source is from another security group. You can select a security group in the same region from the drop-down list. If there is instance A in security group A and instance B in security group B, and the inbound rule of security group A allows traffic from security group B, traffic is allowed from instance B to instance A.</li> <li>• <b>IP address group:</b> An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.</li> </ul>	IP address: 0.0.0.0/0
Description	<p>Supplementary information about the security group rule. This parameter is optional.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

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

**Figure 6-14** Add Outbound Rule



**Table 6-26** Outbound rule parameter description

Parameter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	Allow or Deny <ul style="list-style-type: none"> <li>If the <b>Action</b> is set to <b>Allow</b>, access from ECSs in the security group is allowed to the destination over specified ports.</li> <li>If the <b>Action</b> is set to <b>Deny</b>, access from ECSs in the security group is denied to the destination over specified ports.</li> </ul> Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see <a href="#">How Traffic Matches Security Group Rules</a> .	Allow
Type	Destination IP address version. You can select: <ul style="list-style-type: none"> <li><b>IPv4</b></li> <li><b>IPv6</b></li> </ul>	IPv4
Protocol & Port	The network protocol used to match traffic in a security group rule. The value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>GRE</b> , and <b>ICMP</b> .	TCP

Parameter	Description	Example Value
	<p>Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.</p> <p>Outbound rules control outgoing traffic over specific ports from instances in the security group.</p> <p>Specify one of the following:</p> <ul style="list-style-type: none"><li>• Individual port: Enter a port, such as 22.</li><li>• Consecutive ports: Enter a port range, such as 22-30.</li><li>• Non-consecutive ports: Enter ports and port ranges, such as <b>22,23-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li><li>• All ports: Leave it empty or enter 1-65535.</li></ul>	22, or 22-30

Parameter	Description	Example Value
Destination	<p>Used to match the destination address of an internal request. The destination can be:</p> <ul style="list-style-type: none"><li>• <b>IP address:</b> If you select <b>IP address</b> for <b>Destination</b>, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.<ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>Security group:</b> The destination is from another security group. You can select a security group in the same region under the current account from the drop-down list. If there is instance A in security group A and instance B in security group B, and the outbound rule of security group A allows traffic to security group B, traffic is allowed from instance A to instance B.</li><li>• <b>IP address group:</b> An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.</li></ul>	IP address: 0.0.0.0/0
Description	<p>Supplementary information about the security group rule. This parameter is optional.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

11. Click **OK**.

The outbound rule list is displayed.

## Verifying Security Group Rules

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

For example, if you have deployed a website on an ECS and want users to access your website through HTTP (80), you need to add an inbound rule to the ECS security group to allow access over the port. [Table 6-27](#) shows the rule.

**Table 6-27** Security group rule

Direction	Priority	Action	Type	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 80	IP address: 0.0.0.0/0

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

1. Log in to the ECS and check whether the ECS port is opened.
  - **Checking the port of a Linux server**  
Run the following command to check whether TCP port 80 is being listened on:

```
netstat -an | grep 80
```

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

**Figure 6-15** Command output for the Linux ECS

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

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

```
netstat -an | findstr 80
```

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

**Figure 6-16** Command output for the Windows ECS

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

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

## 6.2.6.2 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.

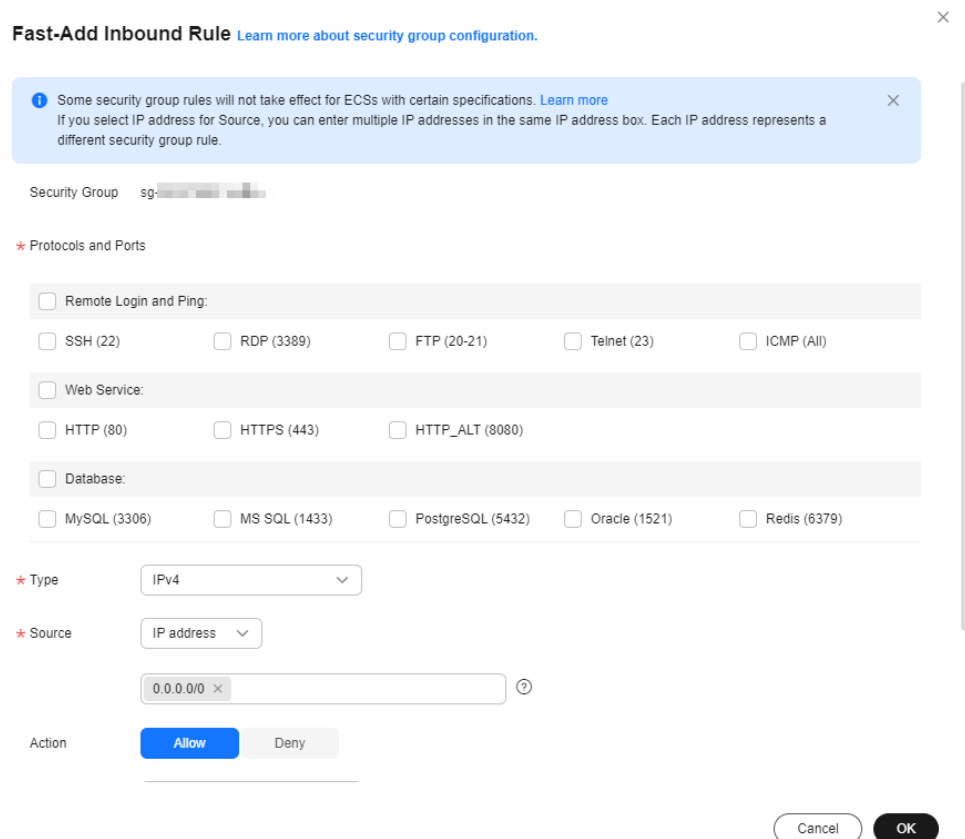
For details about common ports used by cloud servers, see [Common Ports Used by ECSs](#).



## Procedure


1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. 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.
6. On the **Inbound Rules** tab, click **Fast-Add Rule**.  
The **Fast-Add Inbound Rule** dialog box is displayed.
7. Configure required parameters.

**Figure 6-17** Fast-Add Inbound Rule



**Fast-Add Inbound Rule** [Learn more about security group configuration.](#) ×

**ⓘ** Some security group rules will not take effect for ECSs with certain specifications. [Learn more](#)  
If you select IP address for Source, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. ×

Security Group sg-

**\* Protocols and Ports**

Remote Login and Ping:

SSH (22)  RDP (3389)  FTP (20-21)  Telnet (23)  ICMP (All)

Web Service:

HTTP (80)  HTTPS (443)  HTTP\_ALT (8080)

Database:

MySQL (3306)  MS SQL (1433)  PostgreSQL (5432)  Oracle (1521)  Redis (6379)

**\* Type** IPv4 ▼

**\* Source** IP address ▼

0.0.0.0/0 × ?

Action

Cancel OK

**Table 6-28** Inbound rule parameter description

Parameter	Description	Example Value
Protocols and Ports	Common protocols and ports are provided for: <ul style="list-style-type: none"><li>• Remote login and ping</li><li>• Web services</li><li>• Databases</li></ul>	SSH (22)
Type	Source IP address version. You can select: <ul style="list-style-type: none"><li>• <b>IPv4</b></li><li>• <b>IPv6</b></li></ul>	IPv4
Source	Used to match the IP address or address range of an external request. The source can be: <ul style="list-style-type: none"><li>• <b>IP address:</b> If you select <b>IP address</b> for <b>Source</b>, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.<ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>Security group:</b> The source is from another security group. You can select a security group in the same region from the drop-down list. If there is instance A in security group A and instance B in security group B, and the inbound rule of security group A allows traffic from security group B, traffic is allowed from instance B to instance A.</li><li>• <b>IP address group:</b> An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.</li></ul>	Security group

Parameter	Description	Example Value
Action	<p>Allow or Deny</p> <ul style="list-style-type: none"><li>• If the <b>Action</b> is set to <b>Allow</b>, access from the source is allowed to ECSs in the security group over specified ports.</li><li>• If the <b>Action</b> is set to <b>Deny</b>, access from the source is denied to ECSs in the security group over specified ports.</li></ul> <p>Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see <a href="#">How Traffic Matches Security Group Rules</a>.</p>	Allow
Priority	<p>Security group rule priority.</p> <p>The priority value is from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.</p>	1
Description	<p>(Optional) Supplementary information about the security group rule.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	-

8. Click **OK**.  
The inbound rule list is displayed and you can view your added rule.
9. On the **Outbound Rules** tab, click **Fast-Add Rule**.  
The **Fast-Add Outbound Rule** dialog box is displayed.
10. Configure required parameters.

**Figure 6-18** Fast-Add Outbound Rule

**Fast-Add Outbound Rule** [Learn more about security group configuration.](#) ✕

**i** Some security group rules will not take effect for ECSs with certain specifications. [Learn more](#)  
If you select IP address for Destination, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.

Security Group: sg-██████████

**\* Protocols and Ports**

Remote Login and Ping:
  SSH (22)     RDP (3389)     FTP (20-21)     Telnet (23)     ICMP (All)

Web Service:
  HTTP (80)     HTTPS (443)     HTTP\_ALT (8080)

Database:
  MySQL (3306)     MS SQL (1433)     PostgreSQL (5432)     Oracle (1521)     Redis (6379)

**\* Type** IPv4

**\* Destination** IP address

0.0.0.0/0 ✕ ?

**Action** Allow Deny

Cancel
OK

**Table 6-29** Outbound rule parameter description

Parameter	Description	Example Value
Protocols and Ports	Common protocols and ports are provided for: <ul style="list-style-type: none"> <li>• Remote login and ping</li> <li>• Web services</li> <li>• Databases</li> </ul>	SSH (22)
Type	Source IP address version. You can select: <ul style="list-style-type: none"> <li>• <b>IPv4</b></li> <li>• <b>IPv6</b></li> </ul>	IPv4

Parameter	Description	Example Value
Destination	<p>Used to match the destination address of an internal request. The destination can be:</p> <ul style="list-style-type: none"> <li>• <b>IP address:</b> If you select <b>IP address</b> for <b>Destination</b>, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule. <ul style="list-style-type: none"> <li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul> </li> <li>• <b>Security group:</b> The destination is from another security group. You can select a security group in the same region under the current account from the drop-down list. If there is instance A in security group A and instance B in security group B, and the outbound rule of security group A allows traffic to security group B, traffic is allowed from instance A to instance B.</li> <li>• <b>IP address group:</b> An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way.</li> </ul>	Security group
Priority	<p>Security group rule priority.</p> <p>The priority value is from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.</p>	1

Parameter	Description	Example Value
Action	<p>Allow or Deny</p> <ul style="list-style-type: none"> <li>If the <b>Action</b> is set to <b>Allow</b>, access from ECSs in the security group is allowed to the destination over specified ports.</li> <li>If the <b>Action</b> is set to <b>Deny</b>, access from ECSs in the security group is denied to the destination over specified ports.</li> </ul> <p>Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see <a href="#">How Traffic Matches Security Group Rules</a>.</p>	Allow
Description	<p>(Optional) Supplementary information about the security group rule.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	-

11. Click **OK**.

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

### 6.2.6.3 Allowing Common Ports with A Few Clicks

#### Scenarios

You can configure a security group to allow common ports with a few clicks. This function is suitable for the following scenarios:

- Remotely log in to ECSs.
- Use the ping command to test ECS connectivity.
- ECSs functioning as web servers provide website access services.



[Table 6-30](#) describes the common ports that can be opened with a few clicks.

**Table 6-30** Common ports

Direction	Protocol & Port & Type	Source/ Destination	Description
Inbound	TCP: 22 (IPv4)	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.

Direction	Protocol & Port & Type	Source/ Destination	Description
	TCP: 3389 (IPv4)	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.
	TCP: 80 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over port 80 (HTTP) for visiting websites.
	TCP: 443 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over port 443 (HTTPS) for visiting websites.
	TCP: 20-21 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over ports 20 and 21 (FTP) for uploading or downloading files.
	ICMP: All (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access ECSs in the security group over any port for using the ping command to test ECS connectivity.
Outbound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows access from ECSs in the security group to any IP address over any port.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. In the security group list, click the name of the security group.  
The security group details page is displayed.
6. Click the **Inbound Rules** or **Outbound Rules** tab, and then click **Allow Common Ports**.  
The **Allow Common Ports** page is displayed.

7. Click **OK**.

After the operation is complete, you can view the added rules in the security group rule list.

## 6.2.6.4 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.

### Notes and Constraints

Note that modifying a security group rule may interrupt your services or cause network security risks.



Security group rules are like a whitelist. If there are no rules that allow or deny some traffic, the security group denies all traffic to or from the instances in the security group.

- The inbound rules in [Table 6-31](#) ensure that instances in the security group can communicate with each other. Do not modify these rules.
- The outbound rules in [Table 6-31](#) allow instances in the security group to access external networks. If you modify these rules, the instances in the security group cannot access external networks.

**Table 6-31** Security group rules

Direction	Action	Type	Protocol & Port	Source/Destination
Inbound	Allow	IPv4	All	Source: current security group
Inbound	Allow	IPv6	All	Source: current security group
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0
Outbound	Allow	IPv6	All	Destination: ::/0

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.





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

### 6.2.6.5 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Virtual Private Cloud**.  
The **Virtual Private Cloud** page 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 **Replicate** in the **Operation** column.  
The **Replicate Inbound Rule** dialog box is displayed.
7. Modify the security group rule information as prompted and click **OK**.

### 6.2.6.6 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 can import and export security group rules in the following scenarios:

- If you want to back up security group rules locally, you can export the rules to an Excel file.
- If you want to quickly create or restore security group rules, you can import your security group rule file to the security group.
- If you want to quickly apply the rules of one security group to another, you can export and import existing rules.

- If you want to modify multiple rules of the current security group at a time, you can export and import 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 you import a security group rule with **Source/Destination** set to a security group or IP address group, ensure that the group ID is correct. Otherwise, the import will fail.
- If the security group rules to be imported are the same as existing ones, the system automatically deletes them and continues to execute the import.
- Do not import two security group rules with the same **Direction**, **Type**, **Protocol & Port**, and **Source/Destination**, but different **Action** configurations. [Table 6-32](#) shows an example.
  - If a rule to be imported conflicts with an existing rule in the security group, the import will fail. In this case, rectify the fault as prompted.
  - If rules to be imported conflicts with each other, the import will fail. In this case, rectify the fault as prompted.

**Table 6-32** Rules with different actions

Rule	Direction	Priority	Action	Type	Protocol & Port	Destination
Rule A	Inbound	1	Allow	IPv4	TCP: 22	0.0.0.0/0
Rule B	Inbound	5	Deny	IPv4	TCP: 22	0.0.0.0/0

- If you want to import rules of the security group in one region to another under one account, only rules with both **Source** and **Destination** set to **IP address** can be applied.
- If you want to import rules of the security group in one account to the security group in another account, only rules with both **Source** and **Destination** set to **IP address** can be applied.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.

5. On the security group list, click the name of the target security group.  
The security group details page is displayed.
6. Export and import security group rules.
  - Click **Export Rule** to export all rules of the current security group to an Excel file.
  - Click **Import Rule** to import security group rules from an Excel file into the current security group.

**Table 6-33** describes the parameters in the template for importing rules.

**Table 6-33** Template parameters

Parameter	Description	Example Value
Direction	The direction in which the security group rule takes effect. <ul style="list-style-type: none"><li>• <b>Inbound:</b> Inbound rules control incoming traffic to instances in the security group.</li><li>• <b>Outbound:</b> Outbound rules control outgoing traffic from instances in the security group.</li></ul>	Inbound
Priority	The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	Allow or Deny <ul style="list-style-type: none"><li>• If the <b>Action</b> is set to <b>Allow</b>, access from the source is allowed to ECSs in the security group over specified ports.</li><li>• If the <b>Action</b> is set to <b>Deny</b>, access from the source is denied to ECSs in the security group over specified ports.</li></ul> Security group rules are matched by priority and then by action. Deny rules take precedence over allow rules. For more information, see <a href="#">How Traffic Matches Security Group Rules</a> .	Allow
Protocol & Port	The network protocol used to match traffic in a security group rule. The value can be <b>All</b> , <b>TCP</b> , <b>UDP</b> , <b>GRE</b> , and <b>ICMP</b> .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.  Inbound rules control incoming traffic over specific ports to instances in the security group.  Outbound rules control outgoing traffic over specific ports from instances in the security group.	22, or 22-30

Parameter	Description	Example Value
Type	Source IP address version. You can select: <ul style="list-style-type: none"> <li>● IPv4</li> <li>● IPv6</li> </ul>	IPv4
Source	The source in an inbound rule is used to match the IP address or address range of an external request. The source can be: <ul style="list-style-type: none"> <li>● <b>IP address:</b> <ul style="list-style-type: none"> <li>– Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>– IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>– All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul> </li> <li>● <b>Security group:</b> The source is from another security group. You can select a security group in the same region under the current account. Instance A is in security group A and instance B is in security group B. If security group A has an inbound rule with <b>Action</b> set to <b>Allow</b> and <b>Source</b> set to security group B, access from instance B is allowed to instance A. A security group is in the format of <i>Security group name(Security group ID)</i>. An example is sg-test(96a8a93f-XXX-d7872990c314).</li> <li>● <b>IP address group:</b> An IP address group is a collection of one or more IP addresses. You can select an available IP address group. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. A security group is in the format of <i>IP address group name(IP address group ID)</i>. An example is ipGroup-test(96a8a93f-XXX-d7872990c314).</li> </ul>	sg-test[96a8a93f-XXX-d7872990c314]

Parameter	Description	Example Value
Destination	<p>The destination in an outbound rule is used to match the IP address or address range of an internal request. The destination can be:</p> <ul style="list-style-type: none"><li>● <b>IP address</b><ul style="list-style-type: none"><li>– Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>– IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>– All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul></li><li>● <b>Security group</b>: The destination is from another security group. Instance A is in security group A and instance B is in security group B. If security group A has an outbound rule with <b>Action</b> set to <b>Allow</b> and <b>Destination</b> set to security group B, access from instance A is allowed to instance B. A security group is in the format of <i>Security group name(Security group ID)</i>. An example is sg-test(96a8a93f-XXX-d7872990c314).</li><li>● <b>IP address group</b>: An IP address group is a collection of one or more IP addresses. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. A security group is in the format of <i>IP address group name(IP address group ID)</i>. An example is ipGroup-test(96a8a93f-XXX-d7872990c314).</li></ul>	sg-test[96a8a93f-XXX-d7872990c314]
Description	<p>Supplementary information about the security group rule. This parameter is optional.</p> <p>The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	-
Last Modified	The time when the security group was modified.	-

## 6.2.6.7 Deleting a Security Group Rule

### Scenarios

If you no longer need a security group rule to control the traffic to and from the instances in a security group, you can delete it.

### Notes and Constraints

Note that deleting a security group rule may interrupt your services or cause network security risks.



Security group rules are like a whitelist. If there are no rules that allow or deny some traffic, the security group denies all traffic to or from the instances in the security group.

- The inbound rules in [Table 6-34](#) ensure that instances in the security group can communicate with each other. Do not delete these rules.
- The outbound rules in [Table 6-34](#) allow all traffic from the instances in the security groups to external networks. If you delete these rules, the instances in the security group cannot access external networks.

**Table 6-34** Security group rules

Direction	Action	Type	Protocol & Port	Source/Destination
Inbound	Allow	IPv4	All	Source: current security group
Inbound	Allow	IPv6	All	Source: current security group
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0
Outbound	Allow	IPv6	All	Destination: ::/0

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. In the security group list, click the name of the security group.  
The security group details page is displayed.
6. Click the **Inbound Rules** or **Outbound Rules** tab as required.  
The security group rule list is displayed.

7. 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.
8. Click **OK**.

## 6.2.7 Managing Instances Associated with a Security Group

### 6.2.7.1 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.

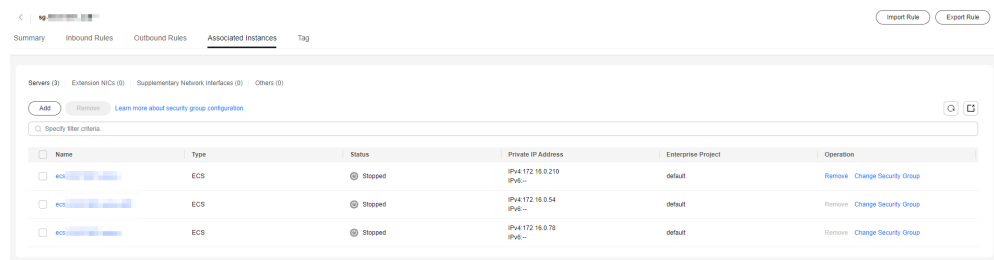
You can add servers, extension NICs, and supplementary network interfaces to a security group by referring to the following operations:

- [Adding an Instance to a Security Group](#)
- [Removing an Instance from a Security Group](#)

#### Adding an Instance to a Security Group

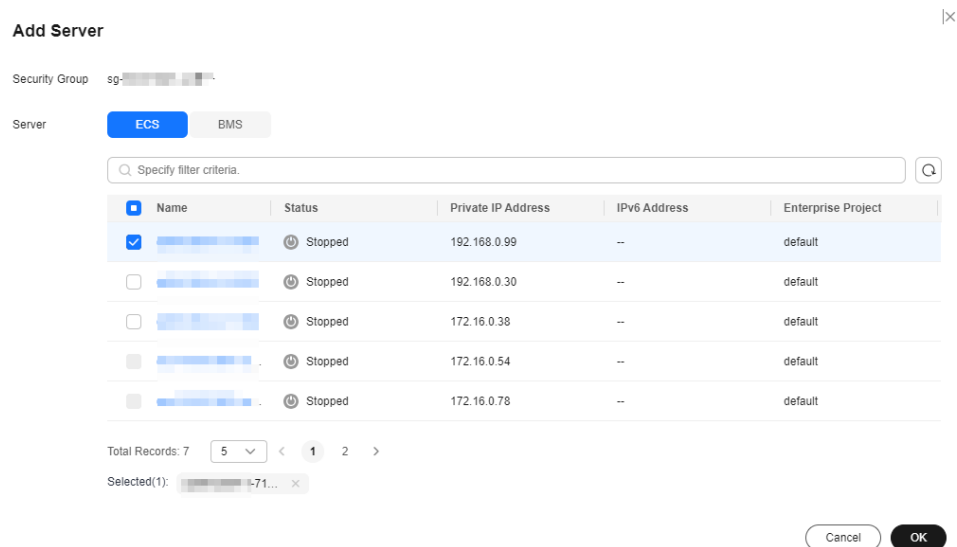
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. In the 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.
6. Click the required instance type tab.  
The following operations use **Servers** as an example.

**Figure 6-19** Associated Instances (Servers)



7. Click the **Servers** tab and click **Add**.  
The **Add Server** dialog box is displayed.



**Figure 6-20** Adding cloud servers



8. 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

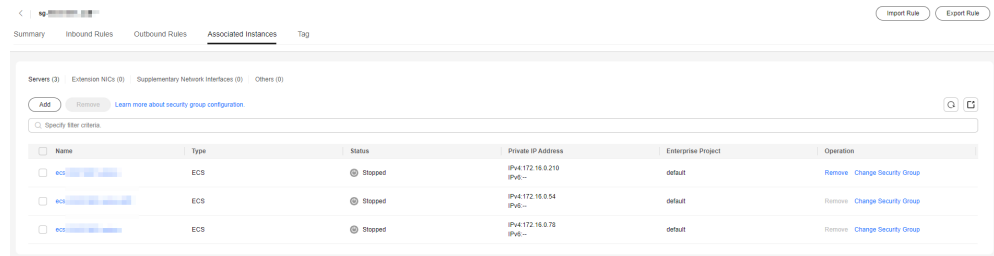
An instance must be added to at least one security group. If you want to remove an instance from a security group, the instance must be associated with at least two security groups now.

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Security Groups**.  
The security group list is displayed.
5. In the 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.



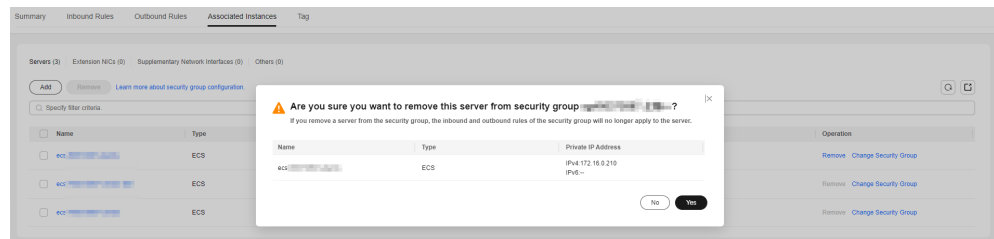
- Click the required instance type tab.  
The following operations use **Servers** as an example.

**Figure 6-21** Associated Instances (Servers)



- 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.

**Figure 6-22** Removing cloud servers



- Confirm the information and click **Yes**.


## 6.2.7.2 Changing the Security Group of an ECS

### Scenarios

When creating an ECS, you must associate it with a security group. If no security group has been created yet, a **default security group** will be created and associated with the ECS. If the default security group cannot meet your service requirements, you can change the security group associated with the ECS.

You can also associate an ECS with a custom security group. If the custom security group cannot meet your requirements, you can change the custom security group.

### Procedure

- Log in to the management console.
- Click . Under **Compute**, 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 6-23** Change Security Group

Change Security Group ×

ECS Name ecs-e498

NIC (primary)

Security Group    [Create Security Group](#)

Security Group Name	Description
<input checked="" type="checkbox"/> default	Default security group
<input type="checkbox"/> sg-0228	

Selected security groups: default

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 apply to 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**.

## 6.3 Network ACL

### 6.3.1 Network ACL Overview

#### Network ACL

A network ACL is an optional layer of security for your subnets. After you add inbound and outbound rules to a network ACL and associate subnets with it, you can control traffic in and out of the subnets.

A network ACL is different from a security group. A security group protects the instances in it, such as ECSs, databases, and containers, while a network ACL protects subnets and all the instances in the subnets. Security groups are a mandatory layer of protection but network ACLs are optional. Network ACLs and security groups can be used together for fine-grained access control.

You need to specify the protocol, source port and address, and destination port and address for each inbound and outbound rule of the network ACL. Suppose you have two subnets in region A, as shown in [Figure 6-24](#). **Subnet-X01** is associated with network ACL **Fw-A**, and ECSs deployed in this subnet provide web

services accessible from the Internet. **Subnet-X02** is associated with network ACL **Fw-B**. **Subnet-X02** and **Subnet-Y01** are connected through a VPC peering connection. Now, you need to configure inbound and outbound rules to allow **ECS-C01** in **Subnet-Y01** to remotely log in to ECSs in **Subnet-X02**.

- Inbound and outbound rules on **Fw-A**:

The custom inbound rule **A01** allows any IP address to access the ECSs in **Subnet-X01** through port 80 over TCP or HTTP. If the traffic does not match custom rule **A01**, the default rule is applied and the traffic is denied to flow into the subnet.

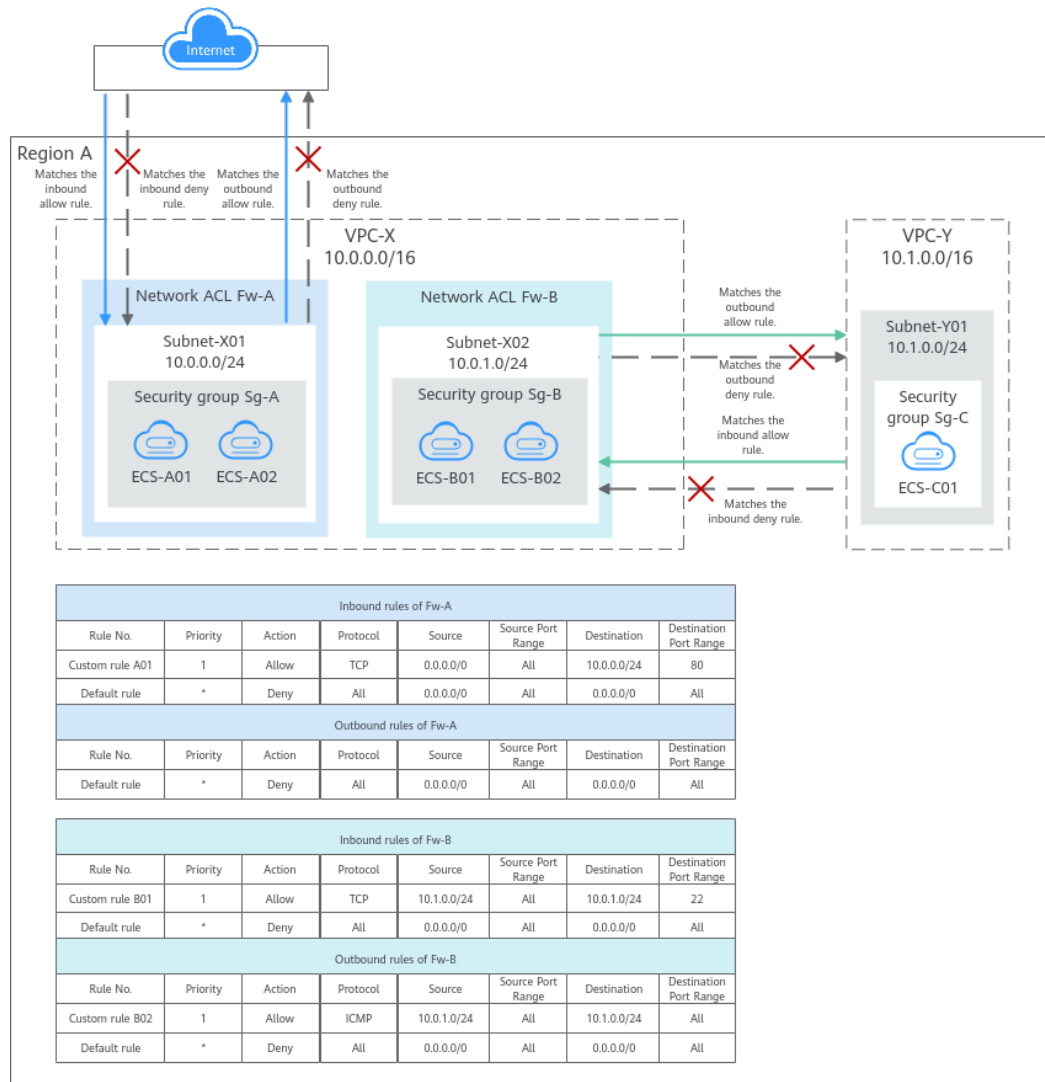
Stateful network ACLs allow responses to inbound requests to leave the subnet without being controlled by rules. The responses from ECSs in **Subnet-X01** can go out of the subnet. Other outbound traffic is not allowed to leave **Subnet-X01**, because the default rule is applied.

- Inbound and outbound rules on **Fw-B**:

The custom inbound rule **B01** allows **ECS-C01** in **Subnet-Y01** to use access the ECSs in **Subnet-X02** through port 22 over TCP or SSH.

The custom outbound rule **B02** allows all ICMP traffic over any port. The ping traffic from ECSs in **Subnet-X02** to **ECS-C01** in **Subnet-Y01** can be routed successfully to test the network connectivity.

Figure 6-24 Network ACL rules



**NOTE**

Figure 6-24 shows how network ACLs control traffic in and out of subnets. In actual use cases, security groups may also be used to control traffic to and from the instances. For details about security groups and network ACLs, see [What Is Access Control?](#)

### Network ACL Rules

- Network ACL uses inbound and outbound rules to control traffic in and out of subnets.
  - Inbound rules: control traffic sent to the instances in a subnet.
  - Outbound rules: control traffic from the instances in a subnet to external networks.
- You need to define the protocol, source and destination ports, source and destination IP addresses, and other information for network ACL rules.
  - **Priority:** Indicates the priority of a rule. Rules are given sequence numbers and traffic is matched against the rules based on their priority. A

smaller number indicates a higher priority. A rule with a higher priority is preferentially applied over a rule with a lower priority.

The priority of the default rule on network ACL is \*. The default rule has the lowest priority.

- **Status: Enabled or Disabled.** Enabled rules are applied, while disabled rules are not.
- **Type: IPv4 or IPv6.**
- **Action: Allow or Deny.** If a request matches a network ACL rule, the action defined in the rule is taken to allow or deny the request.
- **Protocol:** The protocol to match traffic. The value can be **TCP, UDP, or ICMP.**
- **Source/Destination:** The source or destination of the traffic.  
The source or destination can be an IP address, CIDR block, or IP address group.
  - **IP address:** an IPv4/IPv6 address, an IPv4/IPv6 CIDR block, for example, 192.168.10.10/32 (IPv4 address), 192.168.1.0/24 (IPv4 CIDR block), or 2407:c080:802:469::/64 (IPv6 CIDR block).
  - **IP address group:** You can add multiple IP addresses with the same security requirements to an **IP address group** and select this IP address group when you configure a rule.
- **Source Port Range/Destination Port Range:** The source or destination port or port range, which ranges from 1 to 65535.

## How Network ACL Rules Work

- After a network ACL is created, you can associate it with one or more subnets to control traffic in and out of the subnets. You can associate a network ACL with multiple subnets. However, a subnet can only be associated with one network ACL.
- Network ACLs are stateful. If the network ACL rule allows outbound traffic from your instance, the response to the outbound traffic is allowed to flow in, regardless of the inbound rule settings. Similarly, if inbound traffic is allowed, responses to such inbound traffic are allowed to flow out, regardless of the outbound rule settings.
- 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 be applied 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.

- Each network ACL has the default inbound and outbound rules, as shown in [Table 6-35](#). If no custom rules are available, the default rules are applied to deny all inbound and outbound traffic. You can use the default rules only when there is no need for traffic to go in and out of the subnet. If the traffic needs to go in and out of the subnet, you need to add custom rules to control traffic as required.

**Table 6-35** Default network ACL rules

Direction	Priority	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range
Inbound	*	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All
Outbound	*	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All

- The default and custom rules of a network ACL does not block the traffic described in [Table 6-36](#).

**Table 6-36** Traffic not blocked by network ACL rules

Direction	Description
Inbound	Traffic between the source and destination in the same subnet
	Broadcast traffic to 255.255.255.255/32
	Multicast traffic to 224.0.0.0/24
Outbound	Traffic between the source and destination in the same subnet

Direction	Description
	Broadcast traffic to 255.255.255.255/32
	Multicast traffic to 224.0.0.0/24
	TCP metadata traffic to 169.254.169.254/32 over port 80
	Traffic to 100.125.0.0/16 that is reserved for public services on the cloud, such as the DNS server address and NTP server address.


## How Traffic Matches Network ACL Rules

A subnet can be bound to only one network ACL. When there are multiple rules on the network ACL, rules are applied based on their priority. A smaller number indicates a higher priority. The value of the default rule priority is \*, which has the lowest priority.

The following takes inbound traffic as an example to describe how the rules are applied.

- If a custom rule is matched:
  - When **Action** is set to **Deny**, the traffic is denied to access the subnet.
  - When **Action** is set to **Allow**, the traffic is allowed to access the subnet.
- If no custom rule is matched, the default rule is applied and the traffic is not allowed to access the subnet.

**Figure 6-25** Network ACL matching



## How Network ACLs Are Used

A network ACL controls traffic in and out of a subnet. If both security group and network ACL rules are configured, traffic is matched against network ACL rules first and then security group rules. You can add security group rules as required and use network ACLs as an additional layer of protection for your subnets. The following provides some examples on how network ACLs can be used.

## Controlling External Access to Instances in a Subnet

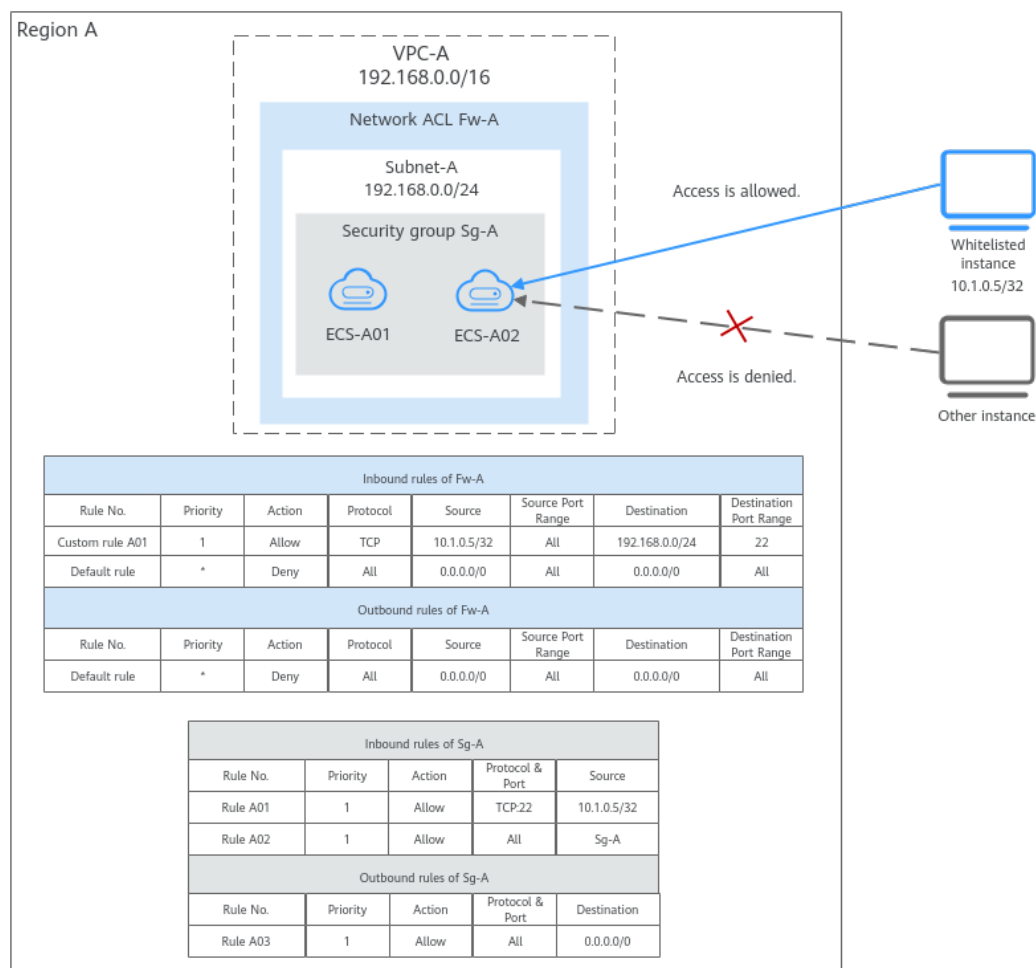
As shown in [Figure 6-26](#), ECS-A01 and ECS-A02 in Subnet-A need to communicate with each other, and the instance with the IP address 10.1.0.5/32 needs to be whitelisted to allow it to remotely log in to ECS-A01 and ECS-A02 to perform O&M operations. The whitelisted instance can be a local PC, an instance in a different subnet of VPC-A, or an instance in another VPC. You need to configure network ACL and security group rules to allow the whitelisted instance to access ECSs in VPC-A and deny any other traffic.

- Network ACL rules:
  - Inbound rule: Custom rule **A01** allows the whitelisted instance to remotely log in to the instances in **Subnet-A** over SSH. The default rule denies any other traffic to the subnet.



- Outbound rule: Network ACLs are stateful. The responses to inbound requests are allowed to leave the subnet. This means you do not need to additionally add outbound rules to allow such response traffic. The default rule denies any other outbound traffic.
- Security group rules:
  - Inbound rule: Rule **A01** allows the whitelisted instance to remotely log in to instances in **Subnet-A** over SSH. Rule **A02** allows instances in the security group to communicate with each other. Other traffic is denied to access the instances in security group **Sg-A**.
  - Outbound rule: Rule **A03** allows instances in **Sg-A** to access external resources.

**Figure 6-26** Controlling external access to instances in a subnet



If you set loose security group rules, network ACL rules can add an additional layer of protection. As described in [Table 6-37](#), the security group rule allows any IP address to remotely log in to instances in the security group. The inbound rule of **Fw-A** associated with **Subnet-A** allows only the specified IP address (10.1.0.5/32) to access instances in **Subnet-A**. The default rule denies other traffic to the subnet, eliminating possible security risks.

**Table 6-37** Security group rules

Direction	Priority	Action	Type	Protocol & Port	Source	Description
Inbound	1	Allow	IPv4	TCP:22	IP address: 0.0.0.0/0	Allows any IP address to remotely log in to instances in the security group using SSH

 **NOTE**

For more network ACL examples, see [Network ACL Configuration Examples](#).

## Controlling Communications Between Instances in Different Subnets

As shown in [Figure 6-27](#), VPC-X has two subnets: **Subnet-X01** and **Subnet-X02**. **ECS-01** and **ECS-02** work in **Subnet-X01**, and **ECS-03** works in **Subnet-X02**.

Suppose you want to:

- Allow **ECS-02** and **ECS-03** to communicate with each other, but
- Deny **ECS-01** and **ECS-03** from communicating with each other.

To achieve this purpose, you need to configure security group and network ACL rules as follows:

1. Add inbound and outbound rules to **Sg-A** to ensure that the ECSs in this security group can communicate with each other.

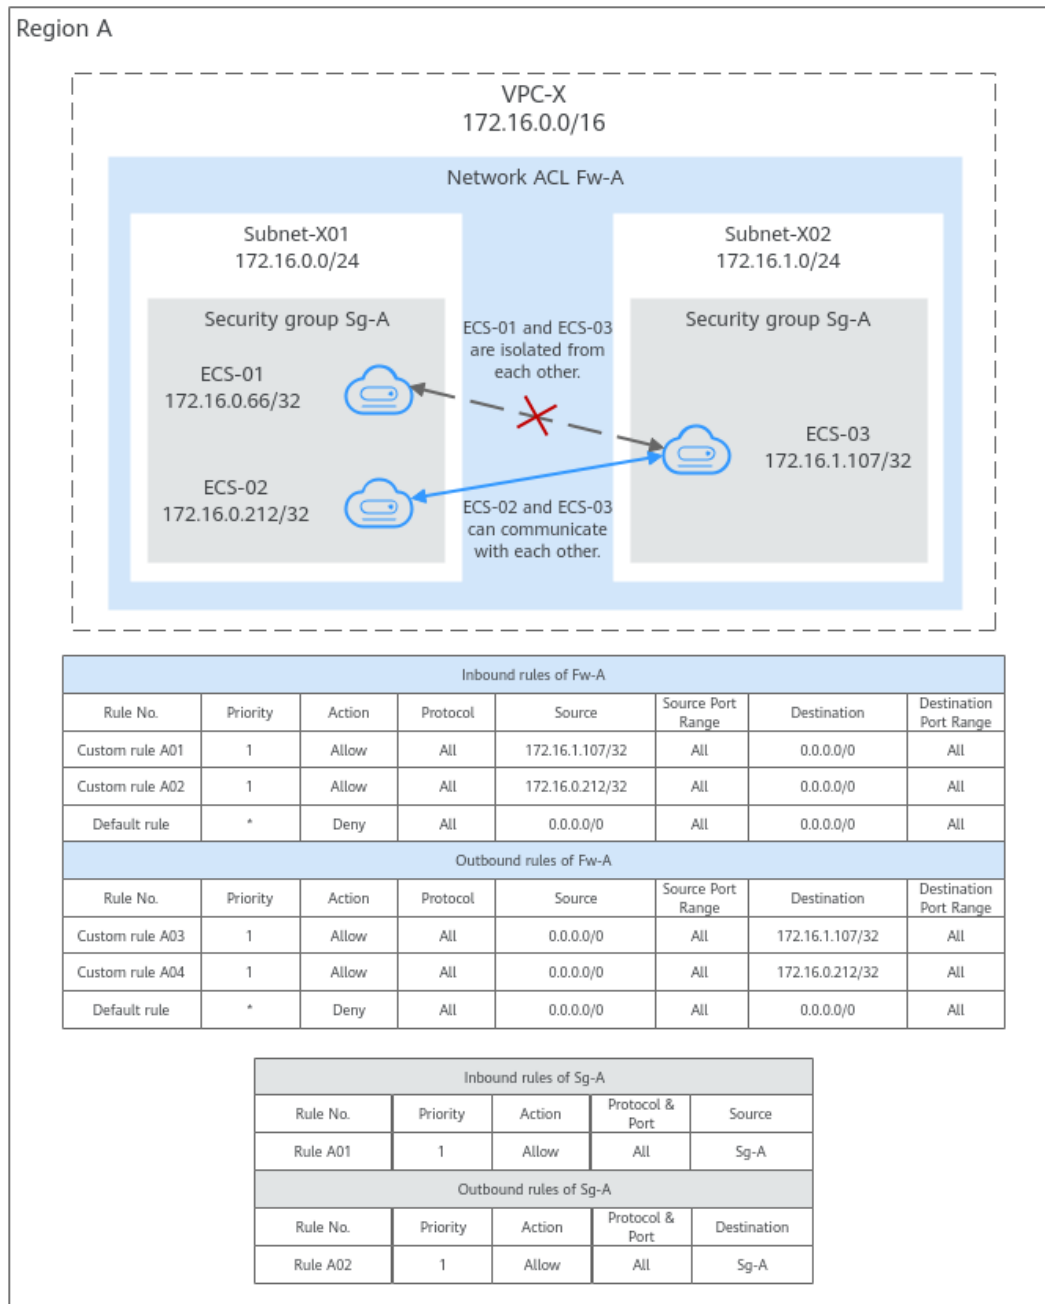
The subnet has not been associated with a network ACL, so after the security group rules are added, both **ECS-01** and **ECS-02** can communicate with **ECS-03**.

2. Associate **Subnet-X01** and **Subnet-X02** with **Fw-A**.

If there is only the default rule in **Fw-A**, instances in the same subnet can communicate with each other, while instances in different subnets are isolated from each other. In this case, **ECS-01** and **ECS-02** can communicate with each other, while **ECS-01** and **ECS-03** as well as **ECS-02** and **ECS-03** are isolated from each other.

3. Add custom rules to **Fw-A** to allow **ECS-02** to communicate with **ECS-03**.
  - Add custom rule A01 to allow **ECS-03** to access **Subnet-X01**.
  - Add custom rule A02 to allow **ECS-02** to access **Subnet-X02**.
  - Add custom rule A03 to allow traffic destined for **ECS-03** to leave **Subnet-X01**.
  - Add custom rule A04 to allow traffic destined for **ECS-02** to leave **Subnet-X02**.

**Figure 6-27** Controlling communications between instances in different subnets




**NOTE**

For more network ACL examples, see [Network ACL Configuration Examples](#).

## Network ACL Configuration Procedure

**Figure 6-28** Procedure for configuring a network ACL



**Table 6-38** Procedure for configuring a network ACL

No.	Step	Description	Reference
1	Create a network ACL.	A network ACL comes with default inbound and outbound rules to deny traffic in and out of a subnet. The default rules cannot be deleted or modified.	<a href="#">Creating a Network ACL</a>
2	Add inbound and outbound rules.	You can add custom rules to control traffic in and out of a subnet. Traffic will be preferentially matched against the custom rules.	<a href="#">Adding a Network ACL Rule (Default Priorities)</a> <a href="#">Adding a Network ACL Rule (Custom Priorities)</a>
3	Associate the network ACL with one or more subnets.	You can associate the network ACL with one or more subnets. If it is enabled, it controls traffic in and out of the subnets.  A subnet can be associated with only one network ACL.	<a href="#">Associating Subnets with a Network ACL</a>

## Notes and Constraints

- By default, each account can have up to 200 network ACLs in a region.
- A network ACL can have no more than 100 rules in one direction, or performance will deteriorate.
- For each network ACL rule, up to 124 rules can have IP address groups associated in either inbound or outbound direction.
- Traffic from load balancers is not restricted by network ACL and security group rules if:

**Transfer Client IP Address** is enabled for the listener of a load balancer.

The load balancer can still forward traffic to backend servers, even if there is a rule that denies traffic from the load balancer to the backend servers.

## 6.3.2 Network ACL Configuration Examples

You can use network ACLs to control the traffic in and out of a subnet. When both security groups and network ACLs are configured, traffic matches network ACL rules first and then security group rules. You can add security group rules as required and use network ACLs to protect instances in the associated subnets. The following provides some examples on how network ACLs can be used.

- [Denying External Access to a Specific Port in a Subnet](#)
- [Denying Access from a Specific IP Address](#)
- [Allowing External Access to Specific Ports on an Instance in a Subnet](#)

**NOTICE**

If your network ACL rules do not work, [submit a service ticket](#).

## Precautions

Note the following before configuring network ACL rules:

- Each network ACL has default rules, as shown in [Table 6-39](#). If a network ACL has no custom rules, the default rule is applied, denying all traffic in and out of a subnet.

**Table 6-39** Default network ACL rules

Direction	Priority	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range
Inbound	*	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All
Outbound	*	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All

- You do not need to add a rule to allow response traffic to inbound requests. This is because the network ACLs are stateful and allow the responses to leave the subnet without being controlled by rules.

For more information about how network ACL rules work, see [How Network ACL Rules Work](#).

## Denying External Access to a Specific Port in a Subnet

If you want to block TCP port 445 to protect instances against WannaCry ransomware attacks, you can add inbound rules described in [Table 6-40](#) to protect the instances in 10.0.0.0/24.

- The default rule denies any traffic to the subnet. You need to add custom rule 02 to allow inbound traffic.
- Add custom rule 01 to deny all inbound traffic to TCP port 445. Place the deny rule above the allow rule to let the deny rule be applied first. For details, see [Adding a Network ACL Rule \(Custom Priorities\)](#).

**Table 6-40** Inbound rules for denying external access to a specific port in a subnet

Direction	Priority	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Inbound	1	IPv4	Deny	TCP	0.0.0.0/0	All	10.0.0.0/24	445	Custom rule 01
Inbound	2	IPv4	Allow	All	0.0.0.0/0	All	10.0.0.0/24	All	Custom rule 02
Inbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule

## Denying Access from a Specific IP Address

You can add inbound rules as described in [Table 6-41](#) to deny the access from abnormal IP addresses, for example, 10.1.1.12/32, to protect the instances in 10.5.0.0/24.

- The default rule denies any traffic to the subnet. You need to add custom rule 02 to allow inbound traffic.
- Add custom rule 01 to deny traffic from 10.1.1.12/32 to 10.5.0.0/24. Place the deny rule above the allow rule to let the deny rule be applied first. For details, see [Adding a Network ACL Rule \(Custom Priorities\)](#).

**Table 6-41** Inbound rules for denying access from a specific IP address

Direction	Priority	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Inbound	1	IPv4	Deny	TCP	10.1.1.12/32	All	10.5.0.0/24	All	Custom rule 01
Inbound	2	IPv4	Allow	All	0.0.0.0/0	All	10.5.0.0/24	All	Custom rule 02
Inbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule

## Allowing External Access to Specific Ports on an Instance in a Subnet

If you deploy a web server in a subnet and want this server to be accessible from the Internet, you need to add network ACL and security group rule to allow HTTP traffic over port 80 and HTTPS traffic over port 443.

1. Add network ACL rules listed in [Table 6-42](#).
  - Add custom rule A01 to allow any HTTP traffic to the instance in the subnet (10.8.0.0/24) over port 80.
  - Add custom rule A02 to allow any HTTPS traffic to the instance in the subnet (10.8.0.0/24) over port 443.

**Table 6-42** Network ACL rules for allowing access to specific ports on an instance in a subnet

Direction	Priority	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description
Inbound	1	IPv4	Allow	TCP	0.0.0.0/0	All	10.8.0.0/24	80	Custom rule 01
Inbound	2	IPv4	Allow	TCP	0.0.0.0/0	All	10.8.0.0/24	443	Custom rule 02
Inbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule
Outbound	*	--	Deny	All	0.0.0.0/0	All	0.0.0.0/0	All	Default rule

2. Add security group rules listed in [Table 6-43](#).
  - Add inbound rule 01 to allow any HTTP traffic to the instance over port 80.
  - Add inbound rule 02 to allow any HTTPS traffic to the instance over port 443.
  - Add outbound rule 03 to allow any traffic to leave the security group.  
You do not need to worry about the loose control of the security group outbound rules. Network ACL rules only allow response traffic to inbound requests to leave the subnet.

**Table 6-43** Security group rules for allowing access to specific ports

Direction	Priority	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	1	Allow	IPv4	TCP:80	IP address: 0.0.0.0/0	Rule 01
Inbound	1	Allow	IPv4	TCP:443	IP address: 0.0.0.0/0	Rule 02
Outbound	1	Allow	IPv4	All	IP address: 0.0.0.0/0	Rule 03

## 6.3.3 Managing Network ACLs

### 6.3.3.1 Creating a Network ACL

#### Scenarios

A security group protects the instances in it, such as ECSs, databases, and containers, while a network ACL protects associated subnets and all the instances in the subnets. Security groups are mandatory, while network ACLs are optional. If you want to add an additional layer of protection, you can create a network ACL and associate it with one or more subnets. Network ACLs and security groups can be used together for fine-grained and comprehensive access control.

#### Procedure

1. Go to the [network ACL list page](#).
2. In the upper right corner of the network ACL list, click **Create Network ACL**.
3. On the displayed page, configure the parameters as prompted.

**Table 6-44** Parameter descriptions

Parameter	Description	Example Value
Name	Mandatory The network ACL 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.	fw-A



Parameter	Description	Example Value
Enterprise Project	<p>Mandatory</p> <p>Enterprise project that the network ACL belongs to.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is <b>default</b>.</p> <p>For details about creating and managing enterprise projects, see the <a href="#">Enterprise Management User Guide</a>.</p>	default
Tag	<p>Optional</p> <p>When creating a network ACL, you can add tags to it to help you identify and search for given network ACLs.</p> <p>Each cloud resource can have a maximum of 20 tags.</p> <p>For details, see <a href="#">Table 6-45</a>.</p>	<b>Tag key:</b> test <b>Tag value:</b> 01
Description	<p>Supplementary information about the network ACL. This parameter is optional.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

**Table 6-45** Network ACL naming requirements

Parameter	Requirements	Example Value
Tag key	<ul style="list-style-type: none"><li>For each resource, each tag key must be unique, and each tag key can have only one tag value.</li><li>Cannot be left blank.</li><li>Can contain a maximum of 128 characters.</li><li>Can consist of letters, digits, underscores (_), and hyphens (-).</li></ul>	test
Tag value	<ul style="list-style-type: none"><li>Can be left blank.</li><li>Can contain a maximum of 256 characters.</li><li>Can consist of letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	01

- Click **OK**.

5. Click **OK**.

## Follow-up Operations



1. A new network ACL comes with default inbound and outbound rules that deny all traffic in and out of associated subnets. You can add custom rules to allow traffic by referring to [Adding a Network ACL Rule \(Default Priorities\)](#) or [Adding a Network ACL Rule \(Custom Priorities\)](#). Traffic will preferentially match the custom rules.
2. You need to associate the enabled network ACL with the subnets by referring to [Associating Subnets with a Network ACL](#).

### 6.3.3.2 Modifying a Network ACL

#### Scenarios

You can 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. On the **Summary** tab, modify the name and description as needed.

### 6.3.3.3 Enabling or Disabling a Network ACL



#### Scenarios

After a network ACL is created, it is enabled by default. You can disable it as required.

- If a network ACL is disabled, custom rules will become invalid but default rules are still applied. As a result, all traffic to and from the associated subnets are denied. If a network ACL has a subnet associated, disabling it will interrupt the network traffic to and from the subnet.
- If a network ACL is enabled, both custom and default rules are applied. If a network ACL has a subnet associated and has only default rules, enabling it will interrupt the network traffic to and from the subnet.

#### Procedure

1. Log in to the management console.

2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, enable or disable the target network ACL.
  - Enabling a network ACL
    - i. Locate the target network ACL and choose **More > Enable** in the **Operation** column.  
A confirmation dialog box is displayed.
    - ii. Confirm the information and click **Yes**.
  - Disabling a network ACL
    - i. Locate the target network ACL and choose **More > Disable** in the **Operation** column.  
A confirmation dialog box is displayed.
    - ii. Confirm the information and click **Yes**.



### 6.3.3.4 Viewing a Network ACL

#### Scenarios

You can check the details of a network ACL, such as the name, rules, and associated subnets.

You can search for a network ACL by name, ID, and description.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. On the **Summary** tab, you can view the following information:
  - Basic information: name, ID, status, and description.
  - Inbound and outbound rules: rule priority, status, protocol, source, source port, destination, and destination port.
  - Associated subnets: the subnets associated with the network ACL. A network ACL can be associated with multiple subnets.



### 6.3.3.5 Deleting a Network ACL

#### Scenarios

You can delete a network ACL when it is no longer required.

Deleting a network ACL will also disassociate it from its associated subnets. Be careful with this operation as it may interrupt services.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and choose **More > Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
6. Confirm the information and click **OK**.

### 6.3.3.6 Managing Network ACL Tags

#### Scenarios

Tags help you identify, classify, and search for network ACLs. You can perform the following operations to manage tags of network ACLs:

- [Adding Network ACL Tags](#)
- [Modifying Network ACL Tags](#)
- [Deleting Network ACL Tags](#)

For details about tag rules, see [Table 6-46](#).

**Table 6-46** Network ACL naming requirements



Parameter	Requirements	Example Value
Tag key	<ul style="list-style-type: none"><li>• For each resource, each tag key must be unique, and each tag key can have only one tag value.</li><li>• Cannot be left blank.</li><li>• Can contain a maximum of 128 characters.</li><li>• Can consist of letters, digits, underscores (_), and hyphens (-).</li></ul>	test

Parameter	Requirements	Example Value
Tag value	<ul style="list-style-type: none"><li>• Can be left blank.</li><li>• Can contain a maximum of 256 characters.</li><li>• Can consist of letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	01



## Notes and Constraints

- Each tag consists of a **tag key** and a **tag value**. Only the **tag value** can be edited.  
If you want to change the **tag key**, delete it and add one again.
- Each cloud resource can have a maximum of 20 tags.

## Adding Network ACL Tags



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list page is displayed.
5. In the network ACL list, click the hyperlink of the network ACL name.  
The network ACL summary page is displayed.
6. Click the **Tags** tab and then click **Add Tag**.  
The **Add Tag** dialog box is displayed.
7. Configure the tag key and tag value as prompted, and click **OK**.  
You can view the added tag in the tag list.

## Modifying Network ACL Tags

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list page is displayed.
5. In the network ACL list, click the hyperlink of the network ACL name.  
The network ACL summary page is displayed.

6. Click the **Tags** tab, locate the row that contains the tag you want to edit and click **Edit** in the **Operation** column.  
The **Edit Tag** dialog box is displayed.
7. Edit the tag value and click **OK**.  
You can view the edited tag in the tag list.

## Deleting Network ACL Tags

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control** > **Network ACLs**.  
The network ACL list page is displayed.
5. In the network ACL list, click the hyperlink of the network ACL name.  
The network ACL summary page is displayed.
6. Click the **Tags** tab, locate the row that contains the tag you want to delete and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
7. Confirm the information and click **Yes**.  
A deleted tag cannot be recovered.

## 6.3.4 Managing Network ACL Rules

### 6.3.4.1 Adding a Network ACL Rule (Default Priorities)

#### Scenarios

You can add inbound and outbound rules to a network ACL to control the traffic in and out of a subnet.

When you perform the following operations to add a rule, the system generates a priority based on the sequence when the rule is added. You cannot specify a priority.

As shown in [Table 6-47](#), there are two custom inbound rules (rule A and rule B) and one default rule. The priority of rule A is 1 and that of rule B is 2. The default rule has the lowest priority. If rule C is added, the system sets its priority to 3, which has lower priority than rules A and B and higher priority than the default rule.

**Table 6-47** Default priorities

Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
Custom rule A	1	Custom rule A	1




Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
--	--	Custom rule B	2
Custom rule B	2	<b>Custom rule C</b>	<b>3</b>
Default rule	*	Default rule	*

If the default priorities do not meet your requirements, you can customize the priorities by referring to [Adding a Network ACL Rule \(Custom Priorities\)](#).

## Notes and Constraints

A network ACL can contain up to 100 rules in one direction, or performance will deteriorate.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. On the **Inbound Rules** or **Outbound Rules** tab, click **Add Rule**.  
The **Add Inbound Rule** or **Add Outbound Rule** dialog box is displayed.
7. Configure required parameters.
  - 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 6-48** Parameter descriptions

Parameter	Description	Example Value
Type	Network ACL type. There are two options: <ul style="list-style-type: none"><li>• IPv4</li><li>• IPv6</li></ul>	IPv4

Parameter	Description	Example Value
Action	The action in the network ACL. There are two options: <ul style="list-style-type: none"><li>• <b>Allow</b>: allows matched traffic in and out of a subnet.</li><li>• <b>Deny</b>: denies matched traffic in and out of a subnet.</li></ul>	Allow
Protocol	The protocol supported by the network ACL to match traffic. The value can be <b>TCP, UDP, or ICMP</b> .	TCP
Source	The source from which the traffic is allowed or denied. The source can be: <ul style="list-style-type: none"><li>• <b>IP address</b><ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>IP address group</b>: The source is an IP address group. An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. Either the source or the destination of a network ACL rule can use the IP address group. For example, if the source uses an IP address group, the destination address cannot use an IP address group.</li></ul>	192.168.0.0/24



Parameter	Description	Example Value
Source Port Range	<p>The source port or port range used to match traffic. The value ranges from 1 to 65535.</p> <p>Enter ports in the following format:</p> <ul style="list-style-type: none"><li>• Individual port: Enter a port, such as <b>22</b>.</li><li>• Consecutive ports: Enter a port range, such as <b>22-30</b>.</li><li>• Non-consecutive ports: Enter ports and port ranges, such as <b>22,24-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li><li>• All ports: Leave it empty or enter <b>1-65535</b>.</li></ul>	22-30

Parameter	Description	Example Value
Destination	<p>The destination to which the traffic is allowed or denied. The destination can be:</p> <ul style="list-style-type: none"><li>• <b>IP address</b><ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32  Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24  Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses.  ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>IP address group</b>: The destination is an IP address group. An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. Either the source or the destination of a network ACL rule can use the IP address group. For example, if the source uses an IP address group, the destination address cannot use an IP address group.</li></ul>	0.0.0.0/0

Parameter	Description	Example Value
Destination Port Range	<p>The destination port or port range used to match traffic. The value ranges from 1 to 65535.</p> <p>Enter ports in the following format:</p> <ul style="list-style-type: none"> <li>Individual port: Enter a port, such as <b>22</b>.</li> <li>Consecutive ports: Enter a port range, such as <b>22-30</b>.</li> <li>Non-consecutive ports: Enter ports and port ranges, such as <b>22,23-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li> <li>All ports: Leave it empty or enter <b>1-65535</b>.</li> </ul>	22-30
Description	<p>Supplementary information about the network ACL rule. This parameter is optional.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

8. Click **OK**.

Return to the rule list to check the new rule.

- The system generates priorities based on the sequence when rules are added. The rule that is added earlier is preferentially matched.
- If the status of the new rule is **Enabled**, the rule is applied.

### 6.3.4.2 Adding a Network ACL Rule (Custom Priorities)

#### Scenarios

If you want a new rule to have a higher or lower priority than a specific rule, you can insert the new rule above or below the specific rule.



As shown in [Table 6-49](#), there are two custom inbound rules (rule A and rule B) and one default rule. The priority of rule A is 1 and that of rule B is 2. The default rule has the lowest priority. If you want rule C to be applied earlier than rule B, you can insert rule C above rule B. After rule C is added, the priority of rule C is 2, and that of rule B is 3.

**Table 6-49** Custom priorities

Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
Custom rule A	1	Custom rule A	1

Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
--	--	Custom rule C	2
Custom rule B	2	Custom rule B	3
Default rule	*	Default rule	*

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. Click the **Inbound Rules** or **Outbound Rules** tab and insert a rule.
  - Locate the target rule and choose **More > Insert Rule Above** in the **Operation** column. The new rule has higher priority than the current rule.
  - Locate the target rule and choose **More > Insert Rule Below** in the **Operation** column. The new rule has lower priority than the current rule.

### 6.3.4.3 Modifying a Network ACL Rule

#### Scenarios



If a network ACL rule no longer meets your requirements, you can modify the port, protocol, and source/destination it.

Modifying rules may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

#### Notes and Constraints

Default network ACL rules cannot be modified or deleted.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

- The **Virtual Private Cloud** page is displayed.
- In the navigation pane on the left, choose **Access Control** > **Network ACLs**. The network ACL list is displayed.
  - In the network ACL list, locate the target network ACL and click its name. The network ACL summary page is displayed.
  - Click the **Inbound Rules** or **Outbound Rules** tab, locate the target rule, click **Modify** in the **Operation** column, and modify parameters based on [Table 6-50](#).

**Table 6-50** Parameter descriptions

Parameter	Description	Example Value
Type	Network ACL type. There are two options: <ul style="list-style-type: none"><li>• <b>IPv4</b></li><li>• <b>IPv6</b></li></ul>	IPv4
Action	The action in the network ACL. There are two options: <ul style="list-style-type: none"><li>• <b>Allow</b>: allows matched traffic in and out of a subnet.</li><li>• <b>Deny</b>: denies matched traffic in and out of a subnet.</li></ul>	Allow
Protocol	The protocol supported by the network ACL to match traffic. The value can be <b>TCP</b> , <b>UDP</b> , or <b>ICMP</b> .	TCP

Parameter	Description	Example Value
Source	<p>The source from which the traffic is allowed or denied. The source can be:</p> <ul style="list-style-type: none"><li>• <b>IP address</b><ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32  Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24  Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses.  ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>IP address group</b>: The source is an IP address group. An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. Either the source or the destination of a network ACL rule can use the IP address group. For example, if the source uses an IP address group, the destination address cannot use an IP address group.</li></ul>	192.168.0.0/24

Parameter	Description	Example Value
Source Port Range	<p>The source port or port range used to match traffic. The value ranges from 1 to 65535.</p> <p>Enter ports in the following format:</p> <ul style="list-style-type: none"><li>• Individual port: Enter a port, such as <b>22</b>.</li><li>• Consecutive ports: Enter a port range, such as <b>22-30</b>.</li><li>• Non-consecutive ports: Enter ports and port ranges, such as <b>22,24-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li><li>• All ports: Leave it empty or enter <b>1-65535</b>.</li></ul>	22-30

Parameter	Description	Example Value
Destination	<p>The destination to which the traffic is allowed or denied. The destination can be:</p> <ul style="list-style-type: none"><li>• <b>IP address</b><ul style="list-style-type: none"><li>- Single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32  Example IPv6 address: 2002:50::44/128</li><li>- IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24  Example IPv6 address range: 2407:c080:802:469::/64</li><li>- All IP addresses 0.0.0.0/0 represents all IPv4 addresses.  ::/0 represents all IPv6 addresses.</li></ul></li><li>• <b>IP address group</b>: The destination is an IP address group. An IP address group is a collection of one or more IP addresses. You can select an available IP address group from the drop-down list. An IP address group can help you manage IP address ranges and IP addresses with same security requirements in a more simple way. Either the source or the destination of a network ACL rule can use the IP address group. For example, if the source uses an IP address group, the destination address cannot use an IP address group.</li></ul>	0.0.0.0/0



Parameter	Description	Example Value
Destination Port Range	<p>The destination port or port range used to match traffic. The value ranges from 1 to 65535.</p> <p>Enter ports in the following format:</p> <ul style="list-style-type: none"><li>• Individual port: Enter a port, such as <b>22</b>.</li><li>• Consecutive ports: Enter a port range, such as <b>22-30</b>.</li><li>• Non-consecutive ports: Enter ports and port ranges, such as <b>22,23-30</b>. You can enter a maximum of 20 ports and port ranges. Each port range must be unique.</li><li>• All ports: Leave it empty or enter <b>1-65535</b>.</li></ul>	22-30
Description	<p>Supplementary information about the network ACL rule. This parameter is optional.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	N/A

7. Click **OK**.

### 6.3.4.4 Enabling or Disabling a Network ACL Rule

#### Scenarios


After a rule is added, it is in **Enabled** status. You can disable it if you need.


- If custom rules are disabled, they will become invalid but default rules are still applied. As a result, all traffic to and from the associated subnets are denied. Disabling all custom rules may interrupt network traffic. Be careful with this operation as it may interrupt services.
- If a custom rule is enabled, it is applied. Enabling custom rules may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

#### Notes and Constraints

Default network ACL rules cannot be modified or deleted.

#### Procedure

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

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. Click the **Inbound Rules** or **Outbound Rules** tab as required.  
The network ACL rule list is displayed.
7. In the rule list, perform the following operations to enable or disable a rule:
  - Enabling a network ACL rule
    - i. Locate the target network ACL rule and choose **More > Enable** in the **Operation** column.  
A confirmation dialog box is displayed.
    - ii. Confirm the information and click **Yes**.
  - Disabling a network ACL rule
    - i. Locate the target network ACL rule and choose **More > Disable** in the **Operation** column.  
A confirmation dialog box is displayed.
    - ii. Confirm the information and click **Yes**.

### 6.3.4.5 Exporting and Importing Network ACL Rules

#### Scenarios

You can specify rule parameters in an Excel file and import it into an existing network ACL. You can also export rules of a network ACL to an Excel file.



You can import or export network ACL rules if you want to:

- Back up these rules to a local directory as an Excel file.
- Quickly add and restore rules by modifying and importing the Excel file you have exported.
- Quickly add rules to other network ACLs.
- Modify rules in batches. You can export rules as an Excel file, modify these rules in the Excel file, and import the file to the network ACL.

#### Notes and Constraints

- For optimal performance, you can import or export up to 40 network ACL inbound and outbound at a time.
- Importing rules will not delete existing rules.
- Importing duplicate rules will fail.
- Default rules cannot be exported.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. Export or import network ACL rules.
  - Click **Export Rule** to export the network ACL rules to an Excel file.
  - Click **Import Rule** to import the network ACL rules from an Excel file into the current network ACL.

### 6.3.4.6 Deleting a Network ACL Rule

#### Scenarios



You can delete a network ACL rule if you no longer need it.

Deleting rules may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

#### Notes and Constraints

Default network ACL rules cannot be modified or deleted.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
5. In the network ACL list, locate the target network ACL and click its name.  
The network ACL summary page is displayed.
6. Click the **Inbound Rules** or **Outbound Rules** tab as required.  
The network ACL rule list is displayed.
7. In the rule list, perform the following operations to delete a rule:
  - To delete a single rule, locate the target rule and click **Delete** in the **Operation** column.

- To delete multiple rules, select the rules and click **Delete** in the upper left corner.
- 8. In the displayed dialog box, confirm the information and click **OK**.

## 6.3.5 Managing Subnets Associated with a Network ACL

### 6.3.5.1 Associating Subnets with a Network ACL

#### Scenarios




You can associate a subnet with a network ACL. If it is enabled, it controls traffic in and out of the subnet.

Associating subnets with a network ACL may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

#### 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 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 \(Default Priorities\)](#).

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. Associate a subnet with a network ACL using either of the following methods:
  - Method 1
    - i. In the navigation pane on the left, click **Subnets**.  
The **Subnets** page is displayed.
    - ii. In the subnet list, locate the row that contains the subnet and click **Associate** under the **Network ACL** column.  
The **Associate Network ACL** page is displayed.
    - iii. Select a network ACL from the drop-down list.  
If there is no network ACL, click  in the drop-down list to create one.
    - iv. Click **OK**.  
The subnet list is displayed. You can view the associated network ACL of the subnet.
  - Method 2

- i. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
- ii. In the subnet list, locate the row that contains the network ACL and click **Associate Subnet** in the **Operation** column.  
The **Associated Subnets** tab is displayed.
- iii. On the **Associated Subnets** tab, click **Associate**.  
The **Associate Subnet** dialog box is displayed.
- iv. In the **Associate Subnet** dialog box, select the subnet from the subnet list and click **OK**.  
In the associated subnet list, you can view all subnets associated with the network ACL.

 **NOTE**



A subnet with a network ACL associated will not be displayed in the subnet list of the **Associate Subnet** dialog box 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.

### 6.3.5.2 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. Disassociate a subnet with a Networking using either of the following methods:
  - Method 1
    - i. In the navigation pane on the left, click **Subnets**.  
The **Subnets** page is displayed.
    - ii. In the subnet list, click the subnet name with a hyperlink.  
The subnet details page is displayed.
    - iii. In the upper right corner of the subnet details page, click **Disassociate** next to the network ACL.  
A confirmation dialog box is displayed.
    - iv. Confirm the information and click **OK**.  
On the subnet details page, you can see that no network ACL is associated with the subnet.

- Method 2
  - i. In the navigation pane on the left, click **Subnets**.  
The **Subnets** page is displayed.
  - ii. In the subnet list, locate the row that contains the subnet and click hyperlink under the **Network ACL** column.  
The network ACL details page is displayed.
  - iii. Click the **Associated Subnets** tab, select one or more subnets, and click **Disassociate**.  
A confirmation dialog box is displayed.
  - iv. Click **Yes** in the displayed dialog box.  
On the **Associated Subnets** tab, you can see that the disassociated subnets are not displayed in the subnet list.
- Method 3
  - i. In the navigation pane on the left, choose **Access Control > Network ACLs**.  
The network ACL list is displayed.
  - ii. Locate the row that contains the network ACL and click **Associate Subnet** in the **Operation** column.  
The **Associated Subnets** tab is displayed.
  - iii. Select one or more subnets and click **Disassociate**.  
A confirmation dialog box is displayed.
  - iv. Click **Yes** in the displayed dialog box.  
On the **Associated Subnets** tab, you can see that the disassociated subnets are not displayed in the subnet list.

# 7 IP Address Group

## 7.1 IP Address Group

### What Is an IP Address Group?

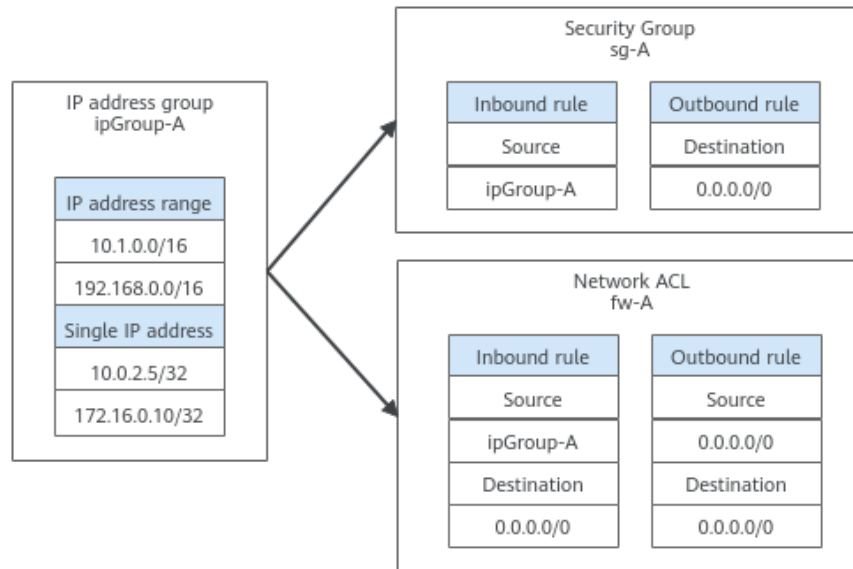
An IP address group is a collection of IP addresses. It can be associated with security groups and network ACLs to simplify IP address configuration and management.

You can add IP address ranges and IP addresses that need to be managed in a unified manner to an IP address group. An IP address group can work together with different cloud resources. [Table 7-1](#) lists the resources that can be associated with an IP address group.

**Table 7-1** Resources that can be associated with an IP address group

Resource	Description	Example
Security group	The <b>Source</b> or <b>Destination</b> of a security group rule can be set to <b>IP address group</b> .	As shown in <a href="#">Figure 7-1</a> , the inbound rule of security group <b>sg-A</b> uses IP address group <b>ipGroup-A</b> as the source.
Network ACL	The <b>Source</b> or <b>Destination</b> of a network ACL is set to <b>IP address group</b> .	As shown in <a href="#">Figure 7-1</a> , the inbound rule of network ACL <b>fw-A</b> uses IP address group <b>ipGroup-A</b> as the source.

**Figure 7-1** Using IP address group



## Notes

If you have multiple IP addresses with the same security requirements, you can add them to an IP address group and select this IP address group when you configure a rule, to help you manage them in a more simple way. When an IP address changes, you only need to change the IP address in the IP address group. Then, the rules in the IP address group change accordingly. You do not need to modify the rules in the security group one by one. This simplifies security group management and improves efficiency. For details, see [Using IP Address Groups to Reduce the Number of Security Group Rules](#).

## Constraints

- Security group rules that are associated with an IP address group do not take effect for certain ECSs.
  - General computing (S1, C1, and C2 ECSs)
  - Memory-optimized (M1 ECSs)
  - High-performance computing (H1 ECSs)
  - Disk-intensive (D1 ECSs)
  - GPU-accelerated (G1 and G2 ECSs)
  - Large-memory (E1, E2, and ET2 ECSs)
- If a network ACL rule uses an IP address group:
  - Either the source or the destination of an inbound rule can use the IP address group.
  - Either the source or the destination of an outbound rule can use the IP address group.

For example, if the source of an inbound rule network ACL is set to an IP address group, the rule destination can only be an IP address.



## 7.2 Managing an IP Address Group

### 7.2.1 Creating an IP Address Group

#### Scenarios

This section describes how to create an IP address group. An IP address group is a collection of IP addresses that can be associated with security groups and network ACLs to simplify IP address configuration and management.

#### Procedure

1. Go to the [Create IP Address Group](#) page.
2. Configure the parameters as prompted.

For details, see [Table 7-2](#).

**Table 7-2** Parameters for creating an IP address group

Parameter	Description	Example Value
Region	Mandatory The region where the IP address group belongs. Select the region nearest to you to ensure the lowest latency possible. An IP address group can be associated only with resources in the same region.	Region A
Name	Mandatory Enter the name of the IP address group. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul> You can customize the name of an IP address group that is uniquely identified by its ID.	ipGroup-A
IP Address Version	Mandatory Select the type of IP addresses that can be added to an IP address group. <ul style="list-style-type: none"><li>• <b>IPv4</b></li><li>• <b>IPv6</b></li></ul>	IPv4

Parameter	Description	Example Value
IP Addresses	<p>Optional</p> <p>Enter an IP address or IP address range on each line, and press <b>Enter</b>. The format is "IP address Description". Description is optional, can contain 0 to 255 characters, and cannot contain angle brackets (&lt; or &gt;). You can enter:</p> <ul style="list-style-type: none"><li>• An IPv4 address range, for example, 192.168.0.0/16 or 192.168.0.0/16   ECS01</li><li>• A single IPv4 address, for example, 192.168.10.10 or 192.168.10.10   ECS01</li><li>• An IPv6 address range, for example, 2001:db8:a583:6e::/64 or 2001:db8:a583:6e::/64   ECS01</li><li>• A single IPv6 address, for example, 2001:db8:a583:6e::5c or 2001:db8:a583:6e::5c   ECS01</li></ul>	<ul style="list-style-type: none"><li>• Without description: 192.168.0.0/16</li><li>• With description: 192.168.0.0/16   ECS01</li></ul>
Enterprise Project	<p>Mandatory</p> <p>When creating an IP address group, you can add the group to an enabled enterprise project.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is <b>default</b>.</p> <p>For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i>.</p>	default
Description	<p>Optional</p> <p>Enter the description of the IP address group in the text box as required.</p>	-

3. Click **Create Now**.

The IP address group list is displayed. The status of the created IP address group is **Normal**.

**NOTICE**

An IP address group takes effect only after it is associated with corresponding resources. For details, see [Associating an IP Address Group with Resources](#).

## 7.2.2 Associating an IP Address Group with Resources

### Scenarios

This section describes how to associate an IP address group with a resource.

An IP address group can be associated with security groups and network ACLs.

### Prerequisites

- You have created an IP address group has been created. For details, see [Creating an IP Address Group](#).
- You have added IP addresses to the IP address group. For details, see [Adding IP Addresses to an IP Address Group](#).

### Procedure

You need to associate an IP address group with resources. For details, see [Table 7-3](#).

**Table 7-3** Associating an IP address group with resources

Resource	Description	Reference
Security group	The <b>Source</b> or <b>Destination</b> of a security group rule can be set to <b>IP address group</b> .	<a href="#">Adding a Security Group Rule</a> <ul style="list-style-type: none"><li>Inbound rule: Set <b>Source</b> to an IP address group.</li><li>Outbound rule: Set <b>Destination</b> to an IP address group.</li></ul>
Network ACL	The <b>Source</b> or <b>Destination</b> of a network ACL is set to <b>IP address group</b> .	<a href="#">Adding a Network ACL Rule (Default Priorities)</a> <ul style="list-style-type: none"><li>Inbound rule: Set <b>Source</b> or <b>Destination</b> to an IP address group. Either the source or the destination can use the IP address group.</li><li>Outbound rule: Set <b>Source</b> or <b>Destination</b> to an IP address group. Either the source or the destination can use the IP address group.</li></ul>

## 7.2.3 Disassociating an IP Address Group from Resources

### Scenarios



This section describes how to disassociate an IP address group from a resource.

An IP address group can be associated with security groups and network ACLs.

### Notes and Constraints

Disassociating an IP address group from resources will make the rules of the resources invalid, and this action cannot be undone.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, locate the row that contains the IP address group and click the resource hyperlink in the **Associated Resources** column.  
The **Associated Resources** page is displayed.
6. In the **Associated Resources** list, click the hyperlink of the corresponding resource name.  
The resource summary page is displayed. You can refer to [Table 7-4](#) to disassociate the IP address group from resources.

**Table 7-4** Disassociating an IP address group from resources

Resource	Description	Reference
Security group	Modify or delete inbound or outbound rules associated with the IP address group.	<ul style="list-style-type: none"><li>• <a href="#">Modifying a Security Group Rule</a><ul style="list-style-type: none"><li>- Inbound rule: Change the value of <b>Source</b>.</li><li>- Outbound rule: Change the value of <b>Destination</b>.</li></ul></li><li>• <a href="#">Deleting a Security Group Rule</a></li></ul>

Resource	Description	Reference
network ACL	Modify or delete inbound or outbound rules associated with the IP address group.	<ul style="list-style-type: none"><li>• <a href="#">Modifying a Network ACL Rule</a><ul style="list-style-type: none"><li>– Inbound rule: Change the value of <b>Source</b> or <b>Destination</b>.</li><li>– Outbound rule: Change the value of <b>Source</b> or <b>Destination</b>.</li></ul></li><li>• <a href="#">Deleting a Network ACL Rule</a></li></ul>




## 7.2.4 Modifying an IP Address Group

### Scenarios

This section describes how to modify basic information about an IP address group, including:

- Name
- Description

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, click the hyperlink of the IP address group name.  
The basic information page of the IP address group is displayed.
6. On the **Basic Information** tab page of the IP address group, click  on the right of the target parameter and modify the parameter as prompted.  
For details, see [Table 7-5](#).

**Table 7-5** IP address group parameters

Parameter	Description	Example Value
Name	Mandatory Enter the name of the IP address group. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul> You can customize the name of an IP address group that is uniquely identified by its ID.	ipGroup-A
Description	Optional Enter the description of the IP address group in the text box as required.	-

7. Click .



## 7.2.5 Exporting IP Address Group Details

### Scenarios

This section describes how to export details about IP address groups, including:

- Name, ID, and creation time
- Added IP addresses
- Associated resources

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, select one or more IP address groups and click **Export** above the list.

Details about the IP address groups are exported to an Excel file.



## 7.2.6 Viewing Details of an IP Address Group

### Scenarios

This section describes how to view information about an IP address group, including:

- Name, ID, and creation time
- Added IP addresses
- Associated resources

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, click the hyperlink of the IP address group name.  
The basic information page of the IP address group is displayed.
6. Click different tabs to view the required information.
  - a. On the **Basic Information** tab page, view the basic information and IP addresses added to the IP address group.
  - b. On the **Associated Resources** tab page, view the resources associated with the IP address group.

## 7.2.7 Deleting an IP Address Group



### Scenarios

This section describes how to delete an IP address group.

### Notes and Constraints

If an IP address group has been associated with a resource, deleting the IP address group will delete the rules that use the IP address group for the associated resource. This interrupts network connectivity.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

- The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
  5. In the IP address group list, delete IP address groups.
    - Delete a single IP address group:
      - i. In the IP address list, locate the row that contains the IP address group and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
      - ii. Confirm the information and click **OK**.
    - Delete IP address groups in batches.
      - i. In the IP address list, select the IP address groups to be deleted.
      - ii. Click the **Delete** button located above the IP address group list.  
A confirmation dialog box is displayed.
      - iii. Confirm the information and click **OK**.  
If a message indicating that IP address groups with associated resources cannot be deleted is displayed, go to the resource details page and [disassociate the IP address group from the resources first](#).

## 7.3 Managing IP Addresses in an IP Address Group

### 7.3.1 Adding IP Addresses to an IP Address Group

#### Scenarios



This section describes how to add IP addresses to an IP address group.

#### Notes and Constraints

If an IP address group has resources associated, adding IP addresses to the group may affect your network communications.

If an IP address group is associated with security groups and network ACLs, the rules associated with the IP address groups will change.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.



5. In the IP address group list, click the hyperlink of the IP address group name. The basic information page of the IP address group is displayed.
6. In the upper left corner of the IP address list, click **Add**. The **Add IP Address** dialog box is displayed.
7. Add IP addresses to the IP address group as prompted.
  - Method 1
    - i. Enter IP addresses in the **IP Addresses** box. For details, see [Table 7-6](#).

**Table 7-6** Parameters for adding IP addresses

Parameter	Description	Example Value
Name	The name of the IP address group.	ipGroup-A
IP Address Version	IP address version supported by an IP address group. You can select an IP address version when you create an IP address group. IP address version cannot be modified after the IP address group is created. The supported versions are as follows: <ul style="list-style-type: none"><li>• IPv4</li><li>• <b>IPv6</b></li></ul>	IPv4

Parameter	Description	Example Value
IP Addresses	<p>Mandatory</p> <p>Enter an IP address or IP address range on each line, and press <b>Enter</b>.</p> <p>The format is "IP address Description". Description is optional, can contain 0 to 255 characters, and cannot contain angle brackets (&lt; or &gt;). You can enter:</p> <ul style="list-style-type: none"><li>• An IPv4 address range, for example, 192.168.0.0/16 or 192.168.0.0/16   ECS01</li><li>• A single IPv4 address, for example, 192.168.10.10 or 192.168.10.10   ECS01</li><li>• An IPv6 address range, for example, 2001:db8:a583:6e::/64 or 2001:db8:a583:6e::/64   ECS01</li><li>• A single IPv6 address, for example, 2001:db8:a583:6e::5c or 2001:db8:a583:6e::5c   ECS01</li></ul>	<ul style="list-style-type: none"><li>• Without description: 192.168.0.0/16</li><li>• With description: 192.168.0.0/16   ECS01</li></ul>

ii. Click **OK**.

In the IP address list, you can view the newly added IP addresses.

– Method 2

Click **Batch Import** in the lower part of the **IP Addresses** box. On the displayed **Batch Import IP Addresses** dialog box, import IP addresses by referring to [Importing IP Addresses to an IP Address Group in Batches](#).

## 7.3.2 Modifying IP Addresses in an IP Address Group

### Scenarios



This section describes how to modify IP addresses, IP address ranges, and their descriptions in an IP address group.

### Notes and Constraints

If an IP address group has resources associated, modifying IP addresses in an IP address group may affect your network communications.

If an IP address group is associated with security groups and network ACLs, the rules associated with the IP address groups will change.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, click the hyperlink of the IP address group name.  
The basic information page of the IP address group is displayed.
6. In the upper left corner of the IP address list, click **Modify**.  
The **Modify IP Address** dialog box is displayed.
7. Modify the information as prompted.  
For details, see [Table 7-7](#).

**Table 7-7** Parameters for modifying IP addresses

Parameter	Description	Example Value
Name	The name of the IP address group.	ipGroup-A
IP Address Version	IP address version supported by an IP address group. You can select an IP address version when you create an IP address group. IP address version cannot be modified after the IP address group is created. The supported versions are as follows: <ul style="list-style-type: none"><li>• IPv4</li><li>• IPv6</li></ul>	IPv4

Parameter	Description	Example Value
IP Addresses	<p>You can modify existing IP addresses, IP address ranges, and their descriptions in an IP address group.</p> <p>The format is "IP address Description". Description is optional, can contain 0 to 255 characters, and cannot contain angle brackets (&lt; or &gt;). You can enter:</p> <ul style="list-style-type: none"><li>• An IPv4 address range, for example, 192.168.0.0/16 or 192.168.0.0/16   ECS01</li><li>• A single IPv4 address, for example, 192.168.10.10 or 192.168.10.10   ECS01</li><li>• An IPv6 address range, for example, 2001:db8:a583:6e::/64 or 2001:db8:a583:6e::/64   ECS01</li><li>• A single IPv6 address, for example, 2001:db8:a583:6e::5c or 2001:db8:a583:6e::5c   ECS01</li></ul>	<ul style="list-style-type: none"><li>• Without description: 192.168.0.0/16</li><li>• With description: 192.168.0.0/16   ECS01</li></ul>

8. Click **OK**.

The IP address list is displayed and you can view that the IP address was modified.

### 7.3.3 Importing IP Addresses to an IP Address Group in Batches

#### Scenarios



You can enter IP addresses, IP address ranges, and their descriptions in an Excel file and import the file to an IP address group. This allows you to quickly add multiple IP addresses.

#### Notes and Constraints

- The number of IP addresses that can be imported is limited. You can check the limited quota on the console.
- Duplicate IP addresses will not be imported, for example:
  - Both the IP address ranges and their descriptions are the same.
  - The IP address ranges are the same but their descriptions are different.

#### Procedure

1. Log in to the management console.

2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, click the hyperlink of the IP address group name.  
The basic information page of the IP address group is displayed.
6. In the upper left corner of the IP address list, click **Import**.  
The **Batch Import IP Addresses** dialog box is displayed.
7. Click **Download Template** to download the Excel template.
8. In the Excel file, enter IP addresses, IP address ranges, and their descriptions, and save the file.

For details about parameter in the Excel file, see [Table 7-8](#).

**Table 7-8** Parameters for importing IP addresses

Parameter	Description	Example Value
IP Addresses	<p>Mandatory</p> <p>In the <b>IP Addresses</b> column, enter an IP address range or a single IP address on a separate line. You can enter:</p> <ul style="list-style-type: none"><li>• An IPv4 address range, for example, 192.168.0.0/16</li><li>• A single IPv4 address, for example, 192.168.10.10</li><li>• An IPv6 address range, for example, 2001:db8:a583:6e::/64</li><li>• A single IPv6 address, for example, 2001:db8:a583:6e::5c</li></ul>	192.168.0.0/16
Description	<p>Optional</p> <p>In the <b>Description</b> column, enter the description of the IP address or IP address range. Description can contain 0 to 255 characters, and cannot contain angle brackets (&lt;&gt;).</p>	ECS01

9. In the **Batch Import IP Addresses** dialog box, click **Select File**, select the Excel file, and click **Import**.  
After the import is complete, you can view the newly imported IP addresses, IP address ranges, and their descriptions.

## 7.3.4 Deleting IP Addresses from an IP Address Group

### Scenarios



This section describes how to delete IP addresses from an IP address group.

### Notes and Constraints

If an IP address group has resources associated, deleting IP addresses from the group may affect your network communications.

If an IP address group is associated with security groups and network ACLs, the rules associated with the IP address groups will change.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **IP Address Groups**.  
The IP address group list is displayed.
5. In the IP address group list, click the hyperlink of the IP address group name.  
The basic information page of the IP address group is displayed.
6. Delete IP addresses:
  - Delete a single IP address.
    - i. In the IP address list, locate the row that contains the IP address and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
    - ii. Confirm the information and click **OK**.
  - Delete IP addresses in batches.
    - i. In the IP address list, select the IP addresses to be deleted.
    - ii. Click the **Delete** button located above the IP address list.  
A confirmation dialog box is displayed.
    - iii. Confirm the information and click **OK**.

# 8 VPC Peering Connection

## 8.1 VPC Peering Connection

### What Is a VPC Peering Connection?

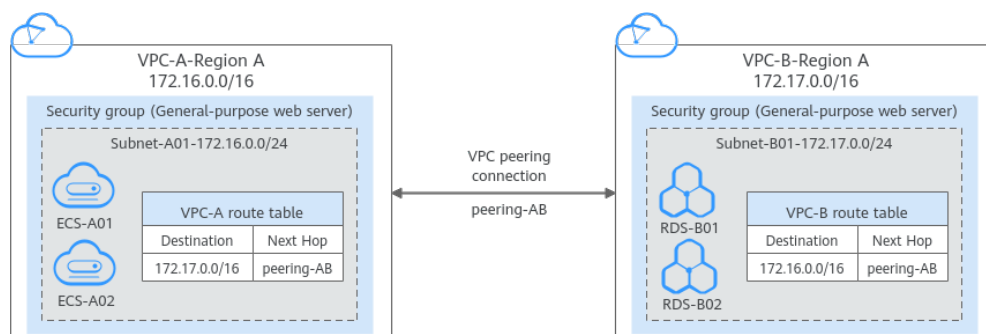
A VPC peering connection is a networking connection that connects two VPCs for them to communicate using private IP addresses. The VPCs to be peered can be in the same account or different accounts, but must be in the same region.

- If you want to connect VPCs in different regions, use [Cloud Connect](#).
- You can use VPC peering connections to build different networks. 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



**NOTICE**

Currently, VPC peering connections are free of charge.

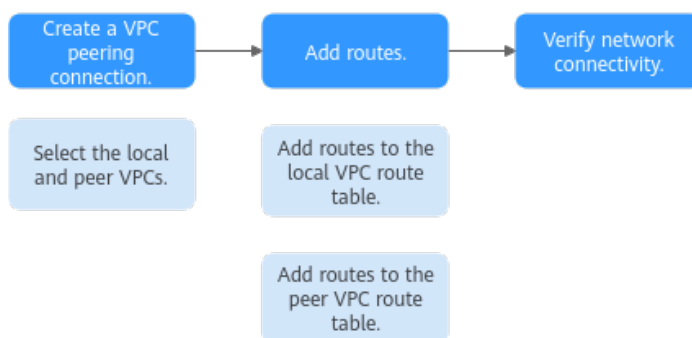
### VPC Peering Connection Creation Process

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

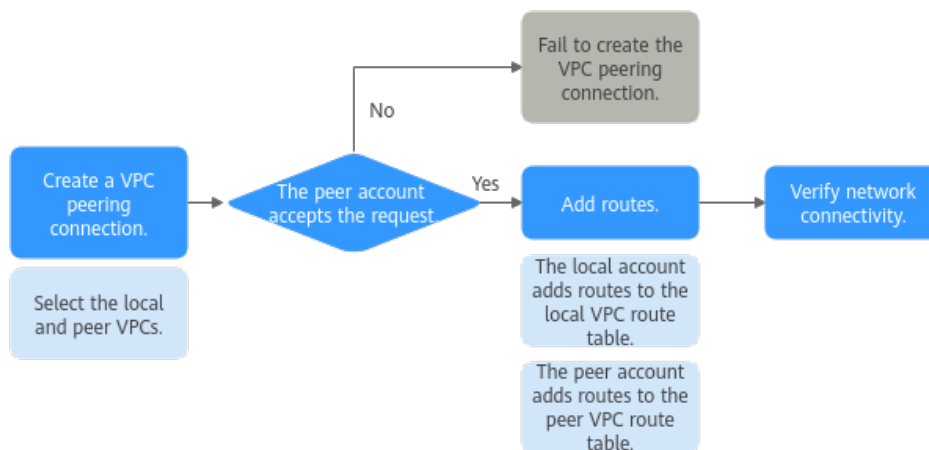


- If two VPCs are in different accounts, the process of creating a VPC peering connection is shown in [Figure 8-3](#).

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.
  - A VPC peering connection can enable a VPC created on the Huawei Cloud Chinese Mainland website and the other created on the Huawei Cloud International website to communicate, but the VPCs must be in the same region. For example, one VPC on the Chinese Mainland website is in CN-Hong Kong region, and the other VPC on the International website is also in CN-Hong Kong region.
  - If you want to connect VPCs in different regions, you can use [Cloud Connect](#).
  - If you only need few ECSs in different regions to communicate with each other, you can [assign and bind EIPs to the ECSs](#).
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect.

In this case, you can refer to [networking configuration examples](#).
- By default, if VPC A is peered with VPC B that has EIPs, VPC A cannot use EIPs in VPC B to access the Internet. To enable this, you can use the NAT Gateway service or configure an SNAT server. For details, see [Enabling Internet Connectivity for an ECS Without an EIP](#).

## 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

Location	CIDR Block	Description	Usage Example
VPCs in the same region	<ul style="list-style-type: none"><li>• VPC CIDR blocks do not overlap.</li><li>• Subnet CIDR blocks of VPCs do not overlap.</li></ul>	You can create VPC peering connections to connect entire CIDR blocks of VPCs. Then, all resources in the VPCs can communicate with each other.	<ul style="list-style-type: none"><li>• <a href="#">Peering Two or More VPCs</a></li><li>• <a href="#">Peering One Central VPC with Multiple VPCs</a></li><li>• For more usage scenarios, see <a href="#">Connecting Entire CIDR Blocks of VPCs</a>.</li></ul>

Location	CIDR Block	Description	Usage Example
VPCs in the same region	<ul style="list-style-type: none"> <li>VPC CIDR blocks overlap.</li> <li>Some subnet CIDR blocks overlap.</li> </ul>	<p>You can create VPC peering connections to connect specific subnets or ECSs from different VPCs.</p> <ul style="list-style-type: none"> <li>To connect specific subnets from two VPCs, the subnet CIDR blocks cannot overlap.</li> <li>To connect specific ECSs from two VPCs, each ECS must have a unique private IP address.</li> </ul>	<ul style="list-style-type: none"> <li><a href="#">Peering Two VPCs with Overlapping CIDR Blocks</a></li> <li>For more usage scenarios, see <a href="#">Connecting Specific Subnets from Different VPCs</a>.</li> </ul>
			<ul style="list-style-type: none"> <li><a href="#">Peering ECSs in a Central VPC with ECSs in Two Other VPCs</a></li> <li>For more usage scenarios, see <a href="#">Connecting Specific ECSs from Different VPCs</a>.</li> </ul>
VPCs in the same region	<ul style="list-style-type: none"> <li>VPC CIDR blocks overlap.</li> <li>All subnet CIDR blocks overlap.</li> </ul>	VPC peering connections are not usable.	<ul style="list-style-type: none"> <li><a href="#">Invalid VPC Peering Connections</a></li> <li>For more examples, see <a href="#">Unsupported VPC Peering Configurations</a>.</li> </ul>

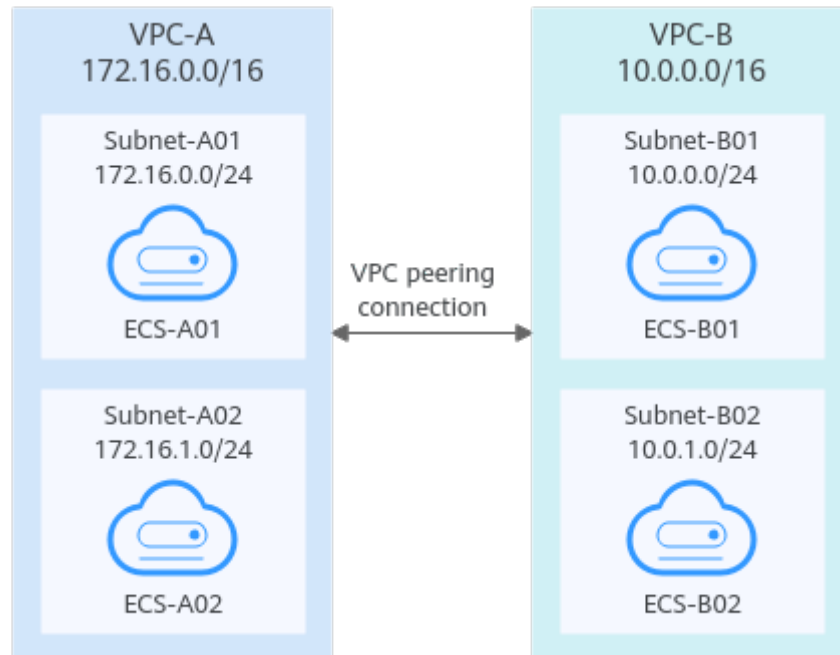
**NOTICE**

A VPC peering connection can only connect VPCs in the same region. If your VPCs are in different regions, use [Cloud Connect](#).

**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.

**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)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.0.0/16	Peering-AB	Custom	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	Custom	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.

**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)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.0.0/16	Peering-AB	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.

Route Table	Destination	Next Hop	Route Type	Description
	192.168.0.0/16	Peering-AC	Custom	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	Custom	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	Custom	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	Custom	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	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop.


 **NOTE**

If a large number of VPCs, for example, 10 VPCs, need to communicate with each other, the networking for establishing VPC peering connections among them is complex. In this case, an enterprise router is recommended. You can attach all the VPCs to an enterprise router to allow them to communicate. For details, see [Using an Enterprise Router to Enable Communications Between VPCs in the Same Region](#).

## 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)

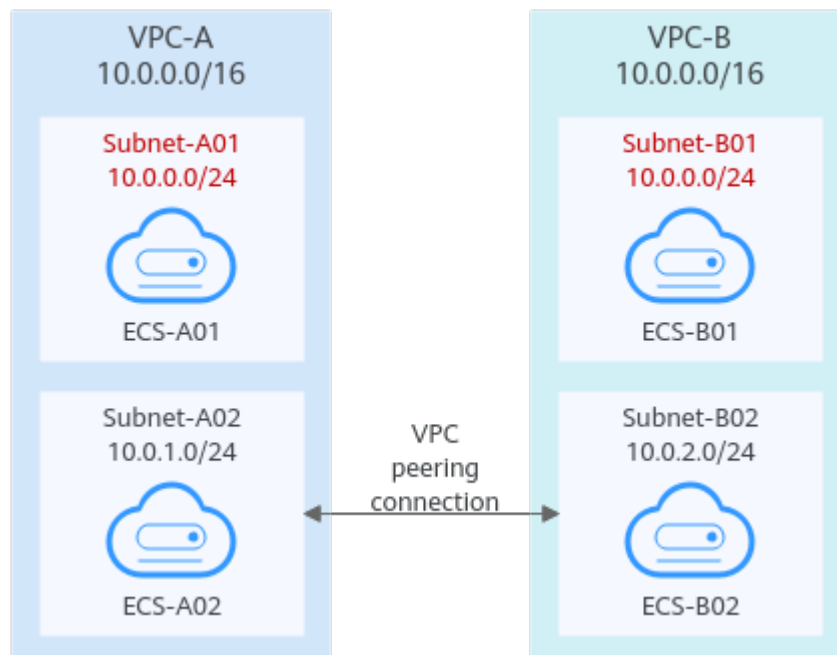
Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.0.0/16	Peering-AB	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
	192.168.0.0/16	Peering-AC	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.
	10.2.0.0/16	Peering-AD	Custom	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	Custom	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	Custom	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	Custom	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	Custom	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	Custom	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	Custom	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	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-F	172.16.0.0/16	Peering-AF	Custom	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	Custom	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.

**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



**Table 8-9** VPC route table details (IPv4)

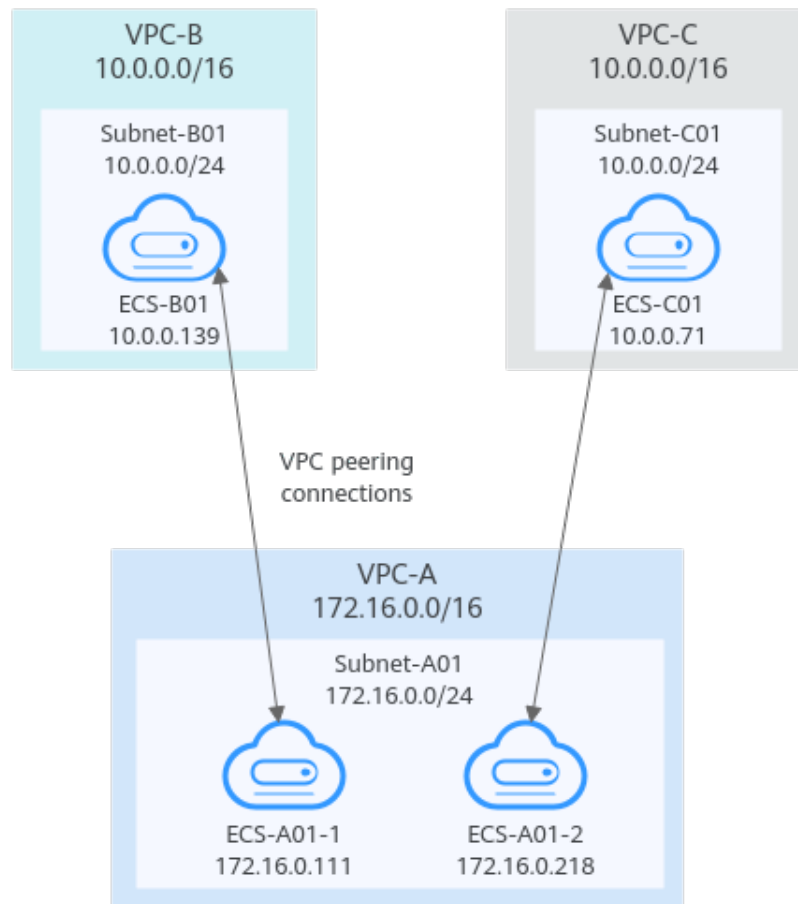
Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.2.0/24	Peering-AB	Custom	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
rtb-VPC-B	10.0.1.0/24	Peering-AB	Custom	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AB as the next hop.

### 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)


Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.0.13/32	Peering-AB	Custom	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	Custom	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	Custom	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	Custom	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

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	10.0.0.0/16 (VPC-B)	Peering-AB	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	10.0.0.0/16 (VPC-A)	Peering-AB	Custom	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.

The 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.

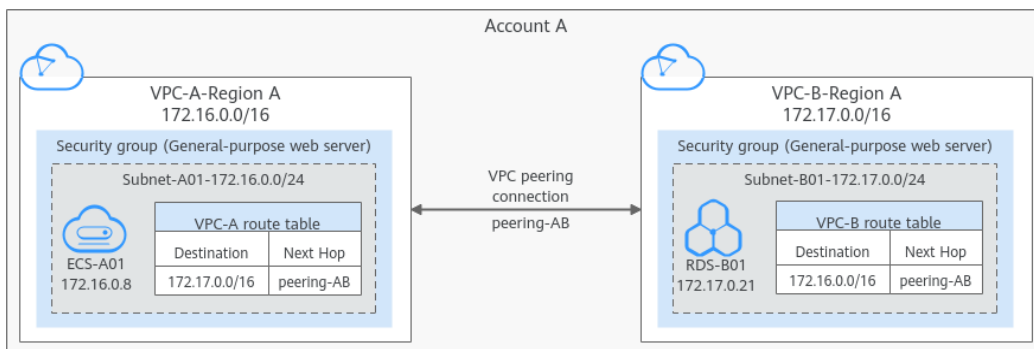
Procedure:

**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



### NOTICE

Currently, VPC peering connections are free of charge.

### 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.
  - A VPC peering connection can enable a VPC created on the Huawei Cloud Chinese Mainland website and the other created on the Huawei Cloud International website to communicate, but the VPCs must be in the same region. For example, one VPC on the Chinese Mainland website is in CN-Hong Kong region, and the other VPC on the International website is also in CN-Hong Kong region.

- If you want to connect VPCs in different regions, you can use [Cloud Connect](#).
- If you only need few ECSs in different regions to communicate with each other, you can [assign and bind EIPs to the ECSs](#).
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect.  
In this case, you can refer to [networking configuration examples](#).

## 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. Go to the [VPC peering connection list page](#).
2. In the upper right corner of the page, click **Create VPC Peering Connection**.  
The **Create VPC Peering Connection** page is displayed.
3. Configure the parameters as prompted.  
For details, see [Table 8-13](#).

**Figure 8-11** Create VPC Peering Connection

|X

### Create VPC Peering Connection

**i** A VPC peering connection can connect VPCs from the same account or from different accounts as long as they are in the same region.

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

If you want to connect VPCs in different regions, use [Cloud Connect](#).

\* VPC Peering Connection Name

**Local VPC Settings**

\* Local VPC  Q

Local VPC CIDR Block

**Peer VPC Settings**

\* Account  My account  Another account ?

\* Peer Project  v

If you select **My account**, the project is filled in by default.

\* Peer VPC  v

Cancel
OK

**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

Parameter	Description	Example Value
Local VPC CIDR Block	CIDR block of the selected local VPC	172.16.0.0/16
Account	Mandatory <ul style="list-style-type: none"><li>Options: <b>My account</b> and <b>Another account</b></li><li>Select <b>My account</b>.</li></ul>	My account
Peer Project	The system fills in the corresponding project by default because <b>My account</b> is set to <b>Account</b> . 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.	ab-cdef-1
Peer VPC	This parameter is mandatory if <b>Account</b> is set to <b>My account</b> . VPC at the other end of the VPC peering connection. You can select one from the drop-down list.	VPC-B
Peer VPC CIDR Block	CIDR block of the selected peer VPC If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. For details, see <a href="#">VPC Peering Connection Usage Examples</a> .	172.17.0.0/16
Description	Optional Enter the description of the VPC peering connection in the text box as required.	peering-AB connects VPC-A and VPC-B.

- Click **OK**.  
A dialog box for adding routes is displayed.
- In the displayed dialog box, click **Add Now**. On the displayed page about the VPC peering connection details, go to [Step 2: Add Routes for the VPC Peering Connection](#) to add a route.

## Step 2: Add Routes for the VPC Peering Connection

- In the lower part of the VPC peering connection details page, click **Add Route**.

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

**Figure 8-12** Add Route

### Add Route

✕

\* VPC vpc-A

\* Route Table rtb-vpc-A(Default) [View Route Table](#)

\* Destination 172.17.0.0/16

\* Next Hop peering-AB(5d057078-b955-49dc-9bb2-9d32cd3)

Description 0/255

Add a route for the other VPC

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. [Learn more](#)

\* VPC vpc-B

\* Route Table rtb-vpc-B(Default) [View Route Table](#)

\* Destination 172.16.0.0/16

\* Next Hop peering-AB(5d057078-b955-49dc-9bb2-9d32cd3)

Description 0/255

Cancel OK

2. Add routes to the route tables as prompted.

**Table 8-14** describes the parameters.

**Table 8-14** Parameter description

Parameter	Description	Example Value
VPC	Select a VPC that is connected by the VPC peering connection.	VPC-A



Parameter	Description	Example Value
Route Table	<p>Select the route table of the VPC. The route will be added to this route table.</p> <p>Each VPC comes with a default route table to control the outbound traffic from the subnets in the VPC. In addition to the default route table, you can also create a custom route table and associate it with the subnets in the VPC. Then, the custom route table controls outbound traffic of the subnets.</p> <ul style="list-style-type: none"><li>• If there is only the default route table in the drop-down list, select the default route table.</li><li>• If there are both default and custom route tables in drop-down list, select the route table associated with the subnet connected by the VPC peering connection.</li></ul>	rtb-VPC-A (Default route table)
Destination	<p>An IP address or address range in the other VPC connected by the VPC peering connection. The value can be a VPC CIDR block, subnet CIDR block, or ECS IP address. For details about the route configuration example, see <a href="#">VPC Peering Connection Usage Examples</a>.</p>	VPC-B CIDR block: 172.17.0.0/16
Next Hop	<p>The default value is the current VPC peering connection. You do not need to specify this parameter.</p>	peering-AB
Description	<p>Supplementary information about the route. This parameter is optional.</p> <p>The description can contain a maximum of 255 characters and cannot contain angle brackets (&lt; or &gt;).</p>	Route from VPC-A to VPC-B
Add a route for the other VPC	<p>If you select this option, you can also add a route for the other VPC connected by the VPC peering connection.</p> <p>To enable communications between VPCs connected by a VPC peering connection, you need to add both forward and return routes to the route tables of the VPCs. For details, see <a href="#">VPC Peering Connection Usage Examples</a>.</p>	Selected

Parameter	Description	Example Value
VPC	By default, the system selects the other VPC connected by the VPC peering connection. You do not need to specify this parameter.	VPC-B
Route Table	Select the route table of the VPC. The route will be added to this route table. Each VPC comes with a default route table to control the outbound traffic from the subnets in the VPC. In addition to the default route table, you can also create a custom route table and associate it with the subnets in the VPC. Then, the custom route table controls outbound traffic of the subnets. <ul style="list-style-type: none"><li>• If there is only the default route table in the drop-down list, select the default route table.</li><li>• If there are both default and custom route tables in drop-down list, select the route table associated with the subnet connected by the VPC peering connection.</li></ul>	rtb-VPC-B (Default route table)
Destination	An IP address or address range in the other VPC connected by the VPC peering connection. The value can be a VPC CIDR block, subnet CIDR block, or ECS IP address. For details about the route configuration example, see <a href="#">VPC Peering Connection Usage Examples</a> .	VPC-A CIDR block: 172.16.0.0/16
Next Hop	The default value is the current VPC peering connection. You do not need to specify this parameter.	peering-AB
Description	Supplementary information about the route. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	Route from VPC-B to VPC-A.

3. Click **OK**.

You can view the routes 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.  
Multiple methods are available for logging in to an ECS. For details, see [Logging In to an ECS](#).
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 Communications Between Instances in Different Security Groups](#).
- 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:

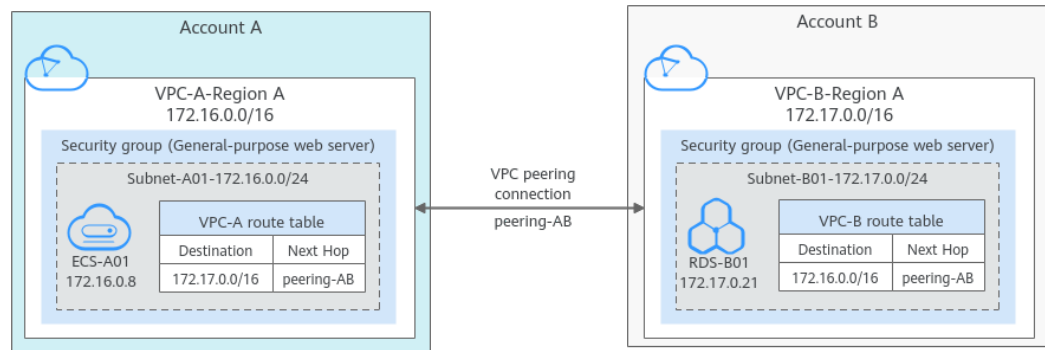
**Step 1: Create a VPC Peering Connection**

**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 8-13** Networking diagram of a VPC peering connection between VPCs in different accounts



### NOTICE

Currently, VPC peering connections are free of charge.

## 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.
  - A VPC peering connection can enable a VPC created on the Huawei Cloud Chinese Mainland website and the other created on the Huawei Cloud International website to communicate, but the VPCs must be in the same region. For example, one VPC on the Chinese Mainland website is in CN-Hong Kong region, and the other VPC on the International website is also in CN-Hong Kong region.
  - If you want to connect VPCs in different regions, you can use [Cloud Connect](#).
  - If you only need few ECSs in different regions to communicate with each other, you can [assign and bind EIPs to the ECSs](#).
- If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. In this case, you can refer to [networking configuration examples](#).
- 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. Go to the [VPC peering connection list page](#).
2. In the upper right corner of the page, click **Create VPC Peering Connection**. The **Create VPC Peering Connection** page is displayed.
3. Configure the parameters as prompted.  
For details, see [Table 8-15](#).

Figure 8-14 Create VPC Peering Connection

**Create VPC Peering Connection** ×

as long as they are in the same region.

- [Creating a VPC Peering Connection with Another VPC in Your Account](#)
- [Creating a VPC Peering Connection with a VPC in Another Account](#)

If you want to connect VPCs in different regions, use [Cloud Connect](#).

\* VPC Peering Connection Name

**Local VPC Settings**

\* Local VPC  🔍

Local VPC CIDR Block

**Peer VPC Settings**

\* Account  My account  Another account ?

The VPC peering connection will be activated only after the peer account accepts the connection request.

\* Peer Project ID

If you select Another account, enter the project ID of the region that the VPC of the peer account is in. [Learn more](#)

\* Peer VPC ID

**Table 8-15** 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 <ul style="list-style-type: none"><li>Options: <b>My account</b> and <b>Another account</b></li><li>Select <b>Another account</b>.</li></ul>	Another account
Peer Project ID	This parameter is mandatory because <b>Account</b> is set to <b>Another account</b> . The project ID of the region that the peer VPC resides. For details about how to obtain the project ID, see <a href="#">Obtaining the Peer Project ID of a VPC Peering Connection</a> .	Project ID of VPC-B in region A: 067cf8aecf3XXX08322f13b
Peer VPC ID	This parameter is mandatory because <b>Account</b> is set to <b>Another account</b> . ID of the VPC at the other end of the VPC peering connection. For details about how to obtain the ID, see <a href="#">Obtaining a VPC ID</a> .	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.


4. 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](#).

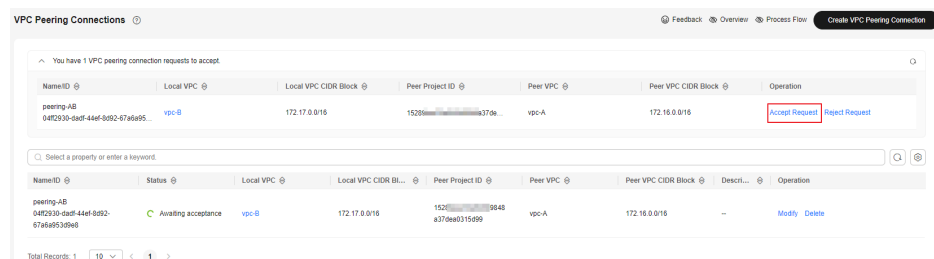
**Figure 8-15** Awaiting acceptance

Name/ID	Status	Local VPC	Local VPC CIDR Block	Peer Project ID	Peer VPC	Peer VPC CIDR Block	Descr...	Operation
peering-AB 04f2930-dadf-44ef-8092-67a5a953a9e8	Awaiting acceptance	vpc-A	172.16.0.0/16	15296... 0786a9f71	vpc-B	172.17.0.0/16	-	Modify Delete

## 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 **Networking > 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 part of the VPC peering connection list, locate the VPC peering connection request to be accepted.

**Figure 8-16** Accept Request

The screenshot shows the 'VPC Peering Connections' page with a table of connections. A red box highlights the 'Accept Request' button in the 'Operation' column of the first row.

Name/ID	Local VPC	Local VPC CIDR Block	Peer Project ID	Peer VPC	Peer VPC CIDR Block	Operation
peering-AB 04f2930-dadf-44ef-8092-67a5a953a9e8	vpc-B	172.17.0.0/16	15296... a370e...0648 a379ea0315f99	vpc-A	172.16.0.0/16	Accept Request Reject Request

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 page showing the VPC peering connection details is displayed.
  - b. In the lower part of the VPC peering connection details page, click **Add Route**.  
The **Add Route** dialog box is displayed.

**Figure 8-17** Add Route

- c. Add routes to the route tables as prompted.  
[Table 8-16](#) describes the parameters.

**Table 8-16** Parameter description

Parameter	Description	Example Value
VPC	The default value is the VPC connected by the VPC peering connection in the current account. You do not need to select a VPC.	VPC-A



Parameter	Description	Example Value
Route Table	Select the route table of the VPC. The route will be added to this route table. Each VPC comes with a default route table to control the outbound traffic from the subnets in the VPC. In addition to the default route table, you can also create a custom route table and associate it with the subnets in the VPC. Then, the custom route table controls outbound traffic of the subnets. <ul style="list-style-type: none"><li>• If there is only the default route table in the drop-down list, select the default route table.</li><li>• If there are both default and custom route tables in drop-down list, select the route table associated with the subnet connected by the VPC peering connection.</li></ul>	rtb-VPC-A (Default route table)
Destination	An IP address or address range in the other VPC connected by the VPC peering connection. The value can be a VPC CIDR block, subnet CIDR block, or ECS IP address. For details about the route configuration example, see <a href="#">VPC Peering Connection Usage Examples</a> .	VPC-B CIDR block: 172.17.0.0/16
Next Hop	The default value is the current VPC peering connection. You do not need to specify this parameter.	peering-AB
Description	Supplementary information about the route. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	Route from VPC-A to VPC-B

- d. Click **OK**.  
You can view the routes 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 page showing the VPC peering connection details is displayed.

- b. In the lower part of the VPC peering connection details page, click **Add Route**.

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

**Figure 8-18** Add Route

- c. Add routes to the route table as prompted.

**Table 8-17** describes the parameters.

**Table 8-17** Parameter description

Parameter	Description	Example Value
VPC	The default value is the VPC connected by the VPC peering connection in the current account. You do not need to select a VPC.	VPC-B

Parameter	Description	Example Value
Route Table	Select the route table of the VPC. The route will be added to this route table. Each VPC comes with a default route table to control the outbound traffic from the subnets in the VPC. In addition to the default route table, you can also create a custom route table and associate it with the subnets in the VPC. Then, the custom route table controls outbound traffic of the subnets. <ul style="list-style-type: none"><li>• If there is only the default route table in the drop-down list, select the default route table.</li><li>• If there are both default and custom route tables in drop-down list, select the route table associated with the subnet connected by the VPC peering connection.</li></ul>	rtb-VPC-B (Default route table)
Destination	An IP address or address range in the other VPC connected by the VPC peering connection. The value can be a VPC CIDR block, subnet CIDR block, or ECS IP address. For details about the route configuration example, see <a href="#">VPC Peering Connection Usage Examples</a> .	VPC-A CIDR block: 172.16.0.0/16
Next Hop	The default value is the current VPC peering connection. You do not need to specify this parameter.	peering-AB
Description	Supplementary information about the route. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	Route from VPC-B to VPC-A.

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.

Multiple methods are available for logging in to an ECS. For details, see [Logging In to an ECS](#).

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 Communications Between Instances in Different Security Groups](#).
- 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

1. 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.  
The **My Credentials** page is displayed.

**Figure 8-19 My Credentials**



3. In the project list, obtain the project ID.  
Locate the region of the peer VPC and obtain the project ID corresponding to the region.

**Figure 8-20 Project ID**

Projects

Project ID	Project Name	Region
067	4	
92F	9	
152	3	
857	-1	
59F	-4	



## 8.6 Modifying a VPC Peering Connection

### Scenarios

This section describes how to modify the basic information about a VPC peering connection, including its name and description.

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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
5. 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.
6. 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
5. 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

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
5. 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.
6. 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.

5. 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.
6. In the route list, click the route table hyperlink of the route.  
The route table details page is displayed.
7. In the route list, locate the route and click **Modify** in the **Operation** column.
8. 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 select the desired region and project.
  - b. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
  - c. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
  - d. 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.
  - e. In the route list, click the route table hyperlink of the route.  
The route table details page is displayed.
  - f. In the route list, locate the route and click **Modify** in the **Operation** column.
  - g. 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
5. 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.
6. In the route list, view the route information.  
You can view the route destination, VPC, next hop, route table, and more.

## 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 select the desired region and project.
  - b. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
  - c. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
  - d. 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.
  - e. In the route list, view the route information.  
You can view the route destination, VPC, next hop, route table, and more.
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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
5. 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.
6. In the route list, locate the route and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
7. Confirm the information and click **OK**.

### 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 select the desired region and project.
  - b. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.

- c. In the navigation pane on the left, choose **Virtual Private Cloud > VPC Peering Connections**.  
The VPC peering connection list is displayed.
  - d. 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.
  - e. In the route list, locate the route and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
  - f. Confirm the information and click **OK**.
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 Sharing

---

## 9.1 VPC Sharing

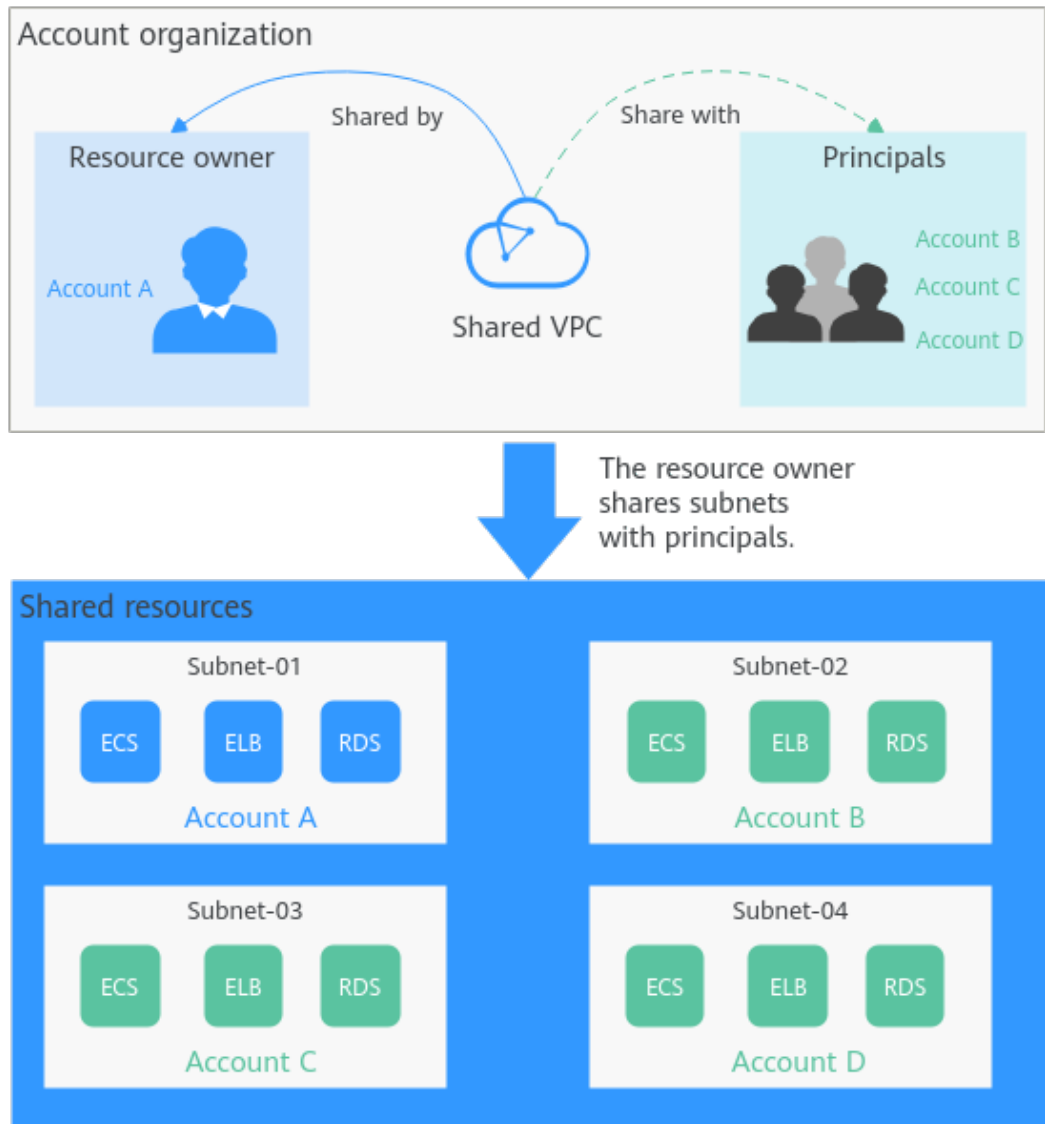
### What Is VPC Sharing?

VPC sharing allows multiple accounts to create and manage cloud resources, such as ECSs, load balancers, and RDS instances, in one VPC. With Resource Access Manager (RAM), you can share subnets in a VPC with one or more accounts so you can centrally manage resources in multiple accounts, which improves resource management efficiency and reduces O&M costs.

The following describes how you can share subnets among several accounts, as shown in [Figure 9-1](#).

- Account A: IT management account of the enterprise and the owner of the VPC and subnets.  
Account A creates a VPC and four subnets and shares these subnets with other accounts. Account A creates resources in Subnet-01.
- Account B: service account of the enterprise and the principal of the shared subnet. Account B creates resources in Subnet-02.
- Account C: service account of the enterprise and the principal of the shared subnet. Account C creates resources in Subnet-03.
- Account D: service account of the enterprise and the principal of the shared subnet. Account D creates resources in Subnet-04.

Figure 9-1 Application scenario



**NOTICE**

The subnets of the owner and those of the principals are in the same VPC, so resources in these subnets can communicate with each other by default. However, if the resources in the subnets are associated with different security groups, the resources are isolated from each other. If you want the resources to communicate with each other, you need to add security group rules by referring to [Adding a Security Group Rule](#).

For example, to allow ECSs in accounts A and B to communicate with each other, you need to add inbound rules to their security groups and set the source to the security group in the other account.

## Advantages

For basic IT systems of financial enterprises and large enterprises, resources are managed by multiple accounts based on permissions. The following problems may arise from time to time:

- There are multiple accounts, such as network accounts, security accounts, and service accounts. This makes cross-account resource O&M hard and time-consuming.
- The cross-account network configurations result in a complex networking structure, hard user operations, and low efficiency.

To deal with these problems, you can share subnets with multiple accounts. You can organize accounts in an orderly and centralized manner based on organization structure or business model.

- You can create subnets in a VPC under an account and share the subnets with principals. In this way, principals do not need to create VPCs and subnets. Fewer resources and simplified network architecture improves management efficiency and reduces costs.

If there are VPCs in different accounts, VPC peering connections are required for mutual communications among VPCs. With VPC sharing, different accounts can create resources in one VPC. This eliminates the need for configuring VPC peering connections and simplifies the network structure.

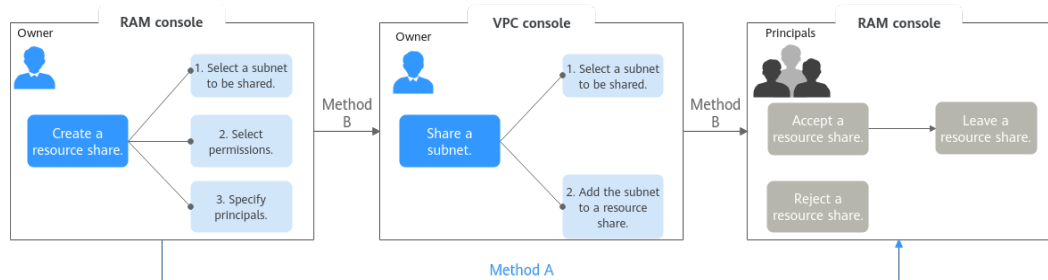
- Resources can be centrally managed in one account, which helps enterprises configure service security policies in a centralized manner and better monitor and audit resource usage for higher security.

## Process for Sharing a Subnet

Before sharing a subnet, you need to enable the RAM service in your account. For details, see [Resource Access Manager User Guide](#).

As the owner of VPC subnets, you can share the subnets with other accounts. Principals need to accept the sharing requests before they use the subnets. [Figure 9-2](#) shows the process of sharing a subnet.

**Figure 9-2** Process for sharing a subnet



You can share a subnet on the RAM or VPC console. For details, see [Table 9-1](#).

**Table 9-1** The process for sharing a subnet

Method	Description	Reference
Method A	<ol style="list-style-type: none"> <li>1. On the RAM console, the owner creates a resource share.                             <ol style="list-style-type: none"> <li>a. Select a subnet to be shared.</li> <li>b. Select permissions to grant to principals on the shared subnet.</li> <li>c. Specify principals that can use the shared subnet.</li> </ol> </li> <li>2. On the RAM console, principals accept or reject the resource share.                             <ul style="list-style-type: none"> <li>• If principals accept the resource share, they can use the shared subnet. If principals do not want to use the shared subnet, they can leave the resource share.</li> <li>• If principals reject the resource share, they cannot use the subnet.</li> </ul> </li> </ol>	<ol style="list-style-type: none"> <li>1. <a href="#">Creating a Resource Share</a></li> <li>2. <a href="#">Responding to a Resource Sharing Invitation</a> <a href="#">Leaving a Resource Share</a></li> </ol>
Method B	<ol style="list-style-type: none"> <li>1. On the RAM console, the owner creates a resource share.                             <ol style="list-style-type: none"> <li>a. Select a subnet to be shared.</li> <li>b. Select permissions to grant to principals on the shared subnet.</li> <li>c. Specify principals that can use the shared subnet.</li> </ol> </li> <li>2. On the VPC console, the owner shares a subnet and adds it to the resource share created in <b>1</b>.</li> <li>3. On the RAM console, principals accept or reject the resource share.                             <ul style="list-style-type: none"> <li>• If principals accept the resource share, they can use the shared subnet. If principals do not want to use the shared subnet, they can leave the resource share.</li> <li>• If principals reject the resource share, they cannot use the subnet.</li> </ul> </li> </ol>	<ol style="list-style-type: none"> <li>1. <a href="#">Creating a Resource Share</a></li> <li>2. <a href="#">Sharing a Subnet with Other Accounts</a></li> <li>3. <a href="#">Responding to a Resource Sharing Invitation</a> <a href="#">Leaving a Resource Share</a></li> </ol>

## Operation Permissions on a Shared Subnet

The owner and principals of a shared subnet have different operation permissions on the subnet and associated resources. For details, see [Table 9-2](#).

**Table 9-2** Operation permissions on a shared subnet and associated resources

Role	When a Share Is Accepted	When a Share Is Stopped	When the Principals Leave a Share
Owner	<ul style="list-style-type: none"> <li>Has operation permissions listed in <a href="#">Table 9-3</a>.</li> <li>Cannot modify or delete resources created by principals, such as ECSs, load balancers, and RDS instances.</li> <li>Views the information such as the IP address and ID of the resource created by principals on the <b>IP Addresses</b> tab of the shared subnet.</li> </ul>	<ul style="list-style-type: none"> <li>Uses, deletes, and manages all resources in the VPC.</li> <li>If principals have resources in the subnet, the owner cannot delete the shared subnet or the VPC where the shared subnet belongs after the share is stopped.</li> </ul>	<ul style="list-style-type: none"> <li>Uses, deletes, and manages all resources in the VPC.</li> <li>If principals have resources in the subnet, the owner cannot delete the shared subnet or the VPC where the shared subnet belongs after the principals leave the share.</li> </ul>
Principal	<ul style="list-style-type: none"> <li>Has operation permissions listed in <a href="#">Table 9-3</a>.</li> <li>Create resources, such as ECSs, load balancers, and RDS instances, in the shared subnets.</li> <li>Views the information such as the IP address and ID of the resource created by themselves on the <b>IP Addresses</b> tab of the shared subnet.</li> </ul>	Uses the existing resources created by themselves, but cannot create resources in the shared subnet.	Uses the existing resources created by themselves, but cannot create resources in the shared subnet.

The owner and principals of a shared subnet have different operation permissions on the subnet and associated resources. For details, see [Table 9-3](#).

**Table 9-3** Operation permissions on a shared subnet and associated resources (sharing)

Resource	Owner	Principal
VPC	Has all operation permissions on the VPC of a shared subnet.	Only can view the VPC that the shared subnet belongs to, but cannot perform any operations on the VPC.



Resource	Owner	Principal
Subnet	Has all operation permissions on the shared subnet and can view the virtual IP addresses and network interfaces in the shared subnet.	Only can view the shared subnet, but cannot: <ul style="list-style-type: none"><li>• Modify the subnet.</li><li>• Delete the subnet.</li><li>• Add, modifying, and delete subnet tags.</li></ul> Can assign virtual IP addresses and network interfaces in the subnet.
Route table	Has all operation permissions on the route table.	<ul style="list-style-type: none"><li>• Cannot create a route table in the VPC that the shared subnet belongs to.</li><li>• Can view the route table associated with the shared subnet and the routes in the route table, but cannot perform any operations on the route table or the routes.</li></ul>
Network ACL	Has all operation permissions on the network ACL.	<ul style="list-style-type: none"><li>• Can view the network ACL associated with the shared subnet, but cannot perform any operation on the network ACL.</li><li>• Cannot associate the owner's network ACL with their own subnets.</li></ul>

Resource	Owner	Principal
Security group	<ul style="list-style-type: none"> <li>• Can create their own security groups.</li> <li>• Only has the operation permissions on their own security groups and cannot perform any operations on the security groups of the principals.</li> <li>• For security groups associated with resources in a shared subnet, the owner can add rules to their own security groups and can set <b>Source</b> of the rules to the security groups of the principals. For example, in the shared Subnet-X:               <ul style="list-style-type: none"> <li>– The owner has created ECS-X with security group Sys-X associated.</li> <li>– Principal A has created ECS-A with security group Sys-A associated.</li> <li>– Principal B has created database RDS-B with security group Sys-B associated.</li> </ul> <p>The owner can add rules with <b>Source</b> set to <b>Sys-A</b> or <b>Sys-B</b> to security group <b>Sys-X</b>.</p> </li> </ul>	<ul style="list-style-type: none"> <li>• Can create their own security groups.</li> <li>• Only has the operation permissions on their own security groups and cannot perform any operations on the security groups of the owner or other principals.</li> <li>• For security groups associated with resources in a shared subnet, a principal can add rules to their own security groups and can set <b>Source</b> of the rules to the security groups of the owner or other principals. For example, in the shared Subnet-X:               <ul style="list-style-type: none"> <li>– The owner has created ECS-X with security group Sys-X associated.</li> <li>– Principal A has created ECS-A with security group Sys-A associated.</li> <li>– Principal B has created database RDS-B with security group Sys-B associated.</li> </ul> <p>Principal A can add rules with <b>Source</b> set to <b>Sys-X</b> or <b>Sys-B</b> to security group <b>Sys-A</b>.</p> </li> </ul>
IP address group	<p>IP address groups are independent of each other. Owners can create an IP address group and associate it with their own security groups.</p>	<p>IP address groups are independent of each other. Principals can create an IP address group and associate it with their own security groups.</p>

Resource	Owner	Principal
VPC flow log	<ul style="list-style-type: none"> <li>Can create a flow log with <b>Resource Type</b> set to <b>VPC</b> or <b>Subnet</b>. Traffic on all network interfaces of the principal in the shared subnet will be recorded in this flow log.</li> <li>Can create a flow log with <b>Resource Type</b> set to <b>NIC</b>. Traffic on all network interfaces of the owner will be recorded in this flow log.</li> </ul>	Can create a flow log with <b>Resource Type</b> set to <b>NIC</b> . Traffic on all network interfaces of the principal will be recorded in this flow log.
VPC peering connection	Selects the VPC with subnets shared with other accounts to create a VPC peering connection.	Cannot select the VPC with subnets shared with other accounts to create a VPC peering connection.
NAT gateway	Creates and manages NAT gateways in the shared subnet.	Cannot create NAT gateways in the shared subnet.
VPN gateway	Creates and manages VPN gateways in the shared subnet.	Cannot create VPN gateways in the shared subnet.
Enterprise router	Attaches the VPC with subnets shared with other accounts to an enterprise router.	Cannot attach the VPC with subnets shared with other accounts to an enterprise router.
Enterprise switch	Creates and manages enterprise switches in the shared subnet.	Cannot create enterprise switches in the shared subnet.
Direct Connect connection	Creates and manages Direct Connect connections in the shared subnet.	Cannot create Direct Connect connections in the shared subnet.
Cloud connection	Loads the VPC with subnets shared with other accounts to a cloud connection.	Cannot load the VPC with subnets shared with other accounts to a cloud connection.
VPC endpoint	Creates and manages VPC endpoints in the shared subnet.	Cannot create VPC endpoints in the shared subnet.
Tag	Adds and manages tags in the shared subnet.	Cannot add tags in the shared subnet.

## Billing

You only need to pay for the resources (such as ECSs, load balancers, and RDS instances) you create in the shared subnets. For details, see the billing description of each cloud resource.

## Quotas

**Table 9-4** lists the quotas of shared subnets. The quotas cannot be increased.

**Table 9-4** Quotas

Item	Default Quota
Maximum number of subnet shares that a principal can receive	100
Maximum number of principal that a subnet can be shared with	100

## Notes and Constraints

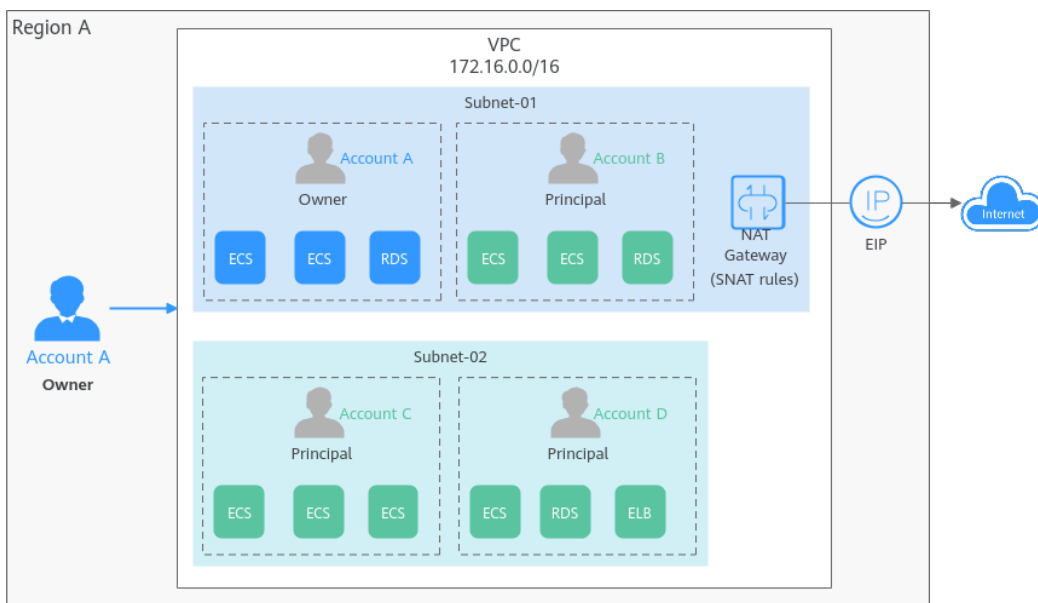
- A principal can receive a maximum of 100 subnet shares.
- A subnet can be shared with a maximum of 100 principals.
- The following cloud resources can be created in a shared subnet:
  - [ECSs](#)
  - [BMSs](#)
  - [Dedicated load balancers](#)
  - [CCE turbo clusters](#)
  - [API gateways](#)
  - [Kafka instances](#)
  - [ServiceStage environments](#)
  - [ServiceComb engines](#)
  - [FunctionGraph functions](#)
  - [GaussDB instances](#)
  - [GaussDB\(for MySQL\) instances](#)
  - [GeminiDB Influx instances](#)
  - [GeminiDB Redis instances](#)
  - [GeminiDB Cassandra instances](#)
  - [RDS for MySQL instances](#)
  - [RDS for PostgreSQL instances](#)
  - [DDS cluster instances](#)
  - [Dedicated HSM instances](#)

- Database audit instances
- CBH instances
- GaussDB(DWS) instances
- DataArts Studio instances
- CSS clusters
- Network connections between DLI and resources
- CDM clusters

## 9.2 Usage Examples for VPC Sharing

Suppose you have two types of workloads running on the cloud. One type of workloads needs to access the Internet and the other type does not. To make resource management easier, you can use account A to manage basic, public IT resources, such as VPCs, subnets, and route tables. And you can share subnets in a VPC in account A with accounts B, C, and D, so the principals can create resources, such as ECSs, RDS instances, and load balancers, in the shared subnets. You can plan your VPC sharing by referring to **Figure 9-3**, and plan accounts and resources as described in **Table 9-5**.

**Figure 9-3** Planning on VPC sharing



**Table 9-5** Planning on VPC sharing

Account	Role	Resource Permissions
Account A	Owner	<ul style="list-style-type: none"> <li>• Creates a VPC and subnets and shares the subnets with other accounts.</li> <li>• Creates a NAT gateway with an EIP bound and configures SNAT rules to enable Subnet-01 to connect to the Internet.</li> </ul>

Account	Role	Resource Permissions
Account B	Principal	Creates ECSs and RDS instances in Subnet-01 to deploy applications that can be accessed over the Internet.
Accounts C and D	Principal	Create ECSs, RDS instances, and load balancers in Subnet-02. These resources do not need to connect to the Internet.

Subnets in the same VPC can communicate with each other by default. However, if instances are associated with different security groups, they are isolated from each other. If you want the resources to communicate with each other, you need to add security group rules to allow the communications.

- Resources in account A are protected by security group Sg-A.
- Resources in account B are protected by security group Sg-B.
- Resources in account C are protected by security group Sg-C.
- Resources in account D are protected by security group Sg-D.

To enable resources in accounts C and D to communicate with each other, add inbound rules to security groups Sg-C and Sg-D.

**Table 9-6** Inbound rules

Security Group	Direction	Priority	Action	Type	Protocol & Port	Source
Sg-C	Inbound	1	Allow	IPv4	Specify the protocol and port as needed. Example: <b>Protocol/All</b>	Security group Sg-D
Sg-D	Inbound	1	Allow	IPv4	Specify the protocol and port as needed. Example: <b>Protocol/All</b>	Security group Sg-C

## 9.3 Sharing a Subnet with Other Accounts

### Scenarios

The owner of a VPC can share subnets in the VPC with principals. You need to create a resource share, select the subnets to share, associate permissions, and specify principals. Once a subnet is shared, principals can create instances in this subnet.



### Prerequisites

A resource share is available. If there are no resource shares, create one by referring to [Creating a Resource Share](#).

### Notes and Constraints

- A principal can receive a maximum of 100 subnet shares.
- A subnet can be shared with a maximum of 100 principals.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. Locate the subnet to share and click its name.  
The **Summary** page is displayed.
5. Click the **Sharing** tab and click **Share Subnet**.  
The **Share Subnet** dialog box is displayed.
6. Select an available resource share.  
If there is no resource share available, create one.
  - a. Click **Cancel** to close the **Share Subnet** dialog box.  
The **Sharing** tab is displayed.
  - b. Click **Create Resource Share**.  
Create a resource share on the RAM console by referring to [Creating a Resource Share](#).
  - c. Repeat **5** to **6** to add subnets to the existing resource share.
7. Click **OK**.  
Return to the **Sharing** tab. You can view that the resource share is in **Sharing** status.

## Follow-up Operations



After an owner shares a subnet with a principal, the principal needs to accept the sharing within a specified period to use the subnet. For details, see [Responding to a Resource Sharing Invitation](#).

# 9.4 Viewing the Details of a Shared Subnet

## Scenarios

The owner and principals of a shared subnet can view the details about a shared subnet, such as the name and status of the shared subnet.

## Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. Locate the target subnet and click its name.  
The **Summary** page is displayed.
5. Click the **Sharing** tab and view the name and status of the resource share that the subnet belongs to.
  - If you are the owner of a shared subnet, you can view shared resources, permissions, and principals on the RAM management console. For details, see [Viewing a Resource Share](#).
  - If you are a principal of a shared subnet, you can view shared resources, permissions, and resource owner on the RAM management console. For details, see [Viewing Resources Shared with You](#).

# 9.5 Stopping Sharing a Subnet

## Scenarios

The owner of a shared subnet can stop sharing a subnet. After the share is stopped, principals cannot create resources in the subnet, but existing resources can be used normally.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.



The **Virtual Private Cloud** page is displayed.

4. Locate the subnet to share and click its name.

The **Summary** page is displayed.

5. Click the **Sharing** tab, locate the row that contains the resource share, and click **Stop Sharing** in the **Operation** column.

A confirmation dialog box is displayed.

6. Confirm the information and click **OK**.

Return to the **Sharing** tab page. You can view that the resource share is in the **Sharing stopped** status.


# 10 Setting Up an IPv6 Network

## What Is an IPv4 and IPv6 Dual-Stack Network?

An IPv4 and IPv6 dual-stack network allows your resources, such as ECSs, to use both IPv4 and IPv6 addresses for private and public network communications.

**Figure 10-1** shows how an IPv4 and IPv6 dual-stack network works.

**Figure 10-1** An IPv4 and IPv6 dual-stack network



**Table 10-1** Steps for deploying a dual-stack network

Step	Description
1	If you enable IPv6 when adding a VPC subnet, an IPv6 CIDR block is automatically assigned to the subnet. You cannot customize the IPv6 CIDR block.

Step	Description
2	<p>Subnets in the same VPC can communicate with each other by default. Network ACLs protect subnets, and security groups protect the instances in it.</p> <ol style="list-style-type: none"><li>Subnets in different network ACLs are isolated from each other. To connect these subnets, you need to add inbound and outbound rules to allow traffic in and out of the subnets.</li><li>Security groups are isolated from each other. If two instances are associated with different security groups, you need to add inbound and outbound rules to allow the instances to communicate with each other.</li></ol> <p>As shown in <a href="#">Figure 10-1</a>, if allow rules are configured for network ACLs <b>Fw-A</b> and <b>Fw-B</b> and security groups <b>Sg-A</b> and <b>Sg-B</b>, <b>ECS-A</b> and <b>ECS-B</b> can communicate with each other:</p> <ul style="list-style-type: none"><li>Using private IPv4 addresses (<b>192.168.0.10</b> and <b>192.168.1.20</b>).</li><li>Using IPv6 addresses (<b>2407:c080:1200:2075::a</b> and <b>2407:c080:1200:1668::b</b>).</li></ul>
3	<p>To enable instances to communicate with the Internet using IPv4 addresses, you need to buy an EIP and bind it to the instance. An EIP can be bound to only one instance.</p> <p>As shown in <a href="#">Figure 10-1</a>, you can bind <b>EIP-A</b> to <b>ECS-A</b> and <b>EIP-B</b> to <b>ECS-B</b> so that <b>ECS-A</b> and <b>ECS-B</b> can communicate with the Internet.</p>
4	<p>To enable instances to communicate with the Internet using an IPv6 address, you need to add the IPv6 address to a shared bandwidth. You can add multiple IPv6 addresses to a shared bandwidth.</p> <p>As shown in <a href="#">Figure 10-1</a>, you can add the IPv6 addresses of <b>ECS-A</b> and <b>ECS-B</b> to a shared bandwidth so that <b>ECS-A</b> and <b>ECS-B</b> can communicate with the Internet.</p>

## Notes and Constraints

- The IPv4/IPv6 dual-stack function is free for now, 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 **Buy 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.

- On the ECS Specifications page.

For example, if you want to check whether general computing-plus ECSs support IPv6:

- Open the [ECS Specifications](#) page.


- ii. Under **General Computing-Plus**, click the link for detailed information, as shown in **Figure 10-2**.

**Figure 10-2** Link for detailed information



- iii. On the **General Computing-plus ECSs** page, check whether IPv6 is supported in **Network** column in the table of ECS features, as shown in **Figure 10-3**.

**Figure 10-3** General computing-plus ECSs



## IPv4 and IPv6 Dual-Stack Application Scenarios

If your ECS supports IPv6, you can build an IPv4 and IPv6 dual-stack network. **Table 10-2** shows where IPv4 and IPv6 dual-stack networks can be used.

**Table 10-2** Application scenarios of IPv4 and IPv6 dual-stack networks

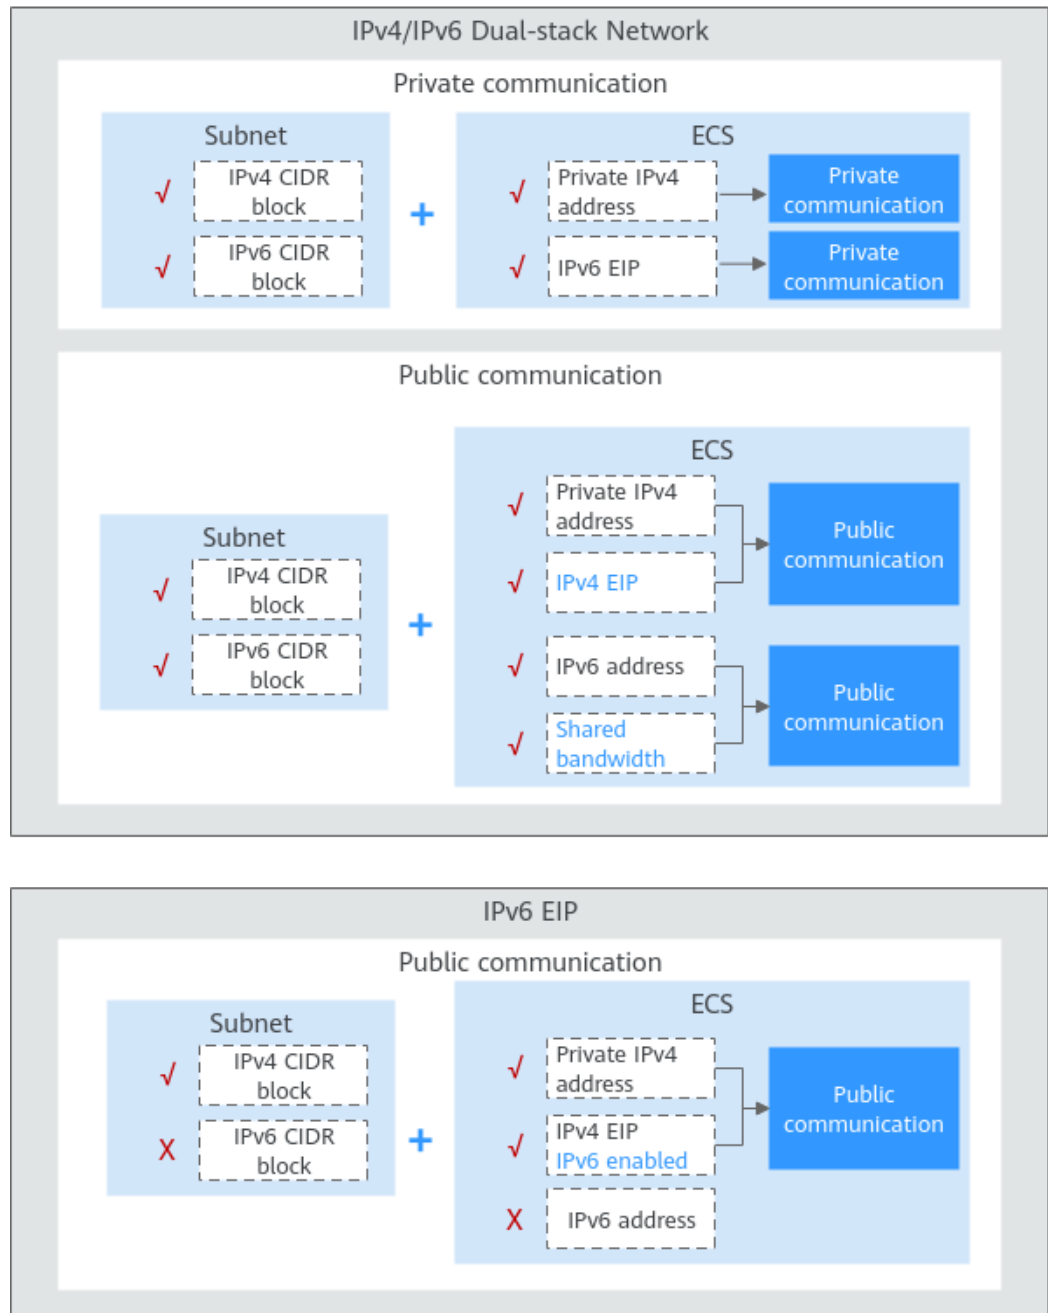
Applica tion Scenari o	Scenario	Subnet	ECS
Private commu nicatio n using IPv6 address es	Your applications deployed on ECSs need to communicate with other systems (such as databases) through private networks using IPv6 addresses.	<ul style="list-style-type: none"> <li>• IPv4 CIDR block</li> <li>• IPv6 CIDR block</li> </ul>	<ul style="list-style-type: none"> <li>• Private IPv4 address: used for private communication</li> <li>• IPv6 address: used for private communication.</li> </ul>
Public commu nicatio n using IPv6 address es	<p>Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.</p> <p>Your applications deployed on ECSs need to both provide services accessible from the Internet and analyze the access request data using IPv6 addresses.</p>	<ul style="list-style-type: none"> <li>• IPv4 CIDR block</li> <li>• IPv6 CIDR block</li> </ul>	<ul style="list-style-type: none"> <li>• Private IPv4 address + IPv4 EIP: used for public network communication</li> <li>• IPv6 address + shared bandwidth: used for public network communication</li> </ul>

If your ECS flavor does not support IPv6 addresses, you can enable the IPv6 EIP function to allow communications using IPv6 addresses. For details, see [Table 10-3](#).

**Table 10-3** Application scenarios of IPv6 EIPs

Application Scenario	Description	Subnet	ECS
Public communication using IPv6 addresses	Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.	IPv4 CIDR block	<ul style="list-style-type: none"><li>• Private IPv4 address</li><li>• IPv4 EIP (with IPv6 function enabled): used for public communication using IPv4 and IPv6 EIPs</li></ul>

**Figure 10-4** Application scenarios of IPv6 networks



## Operation Guide on IPv6 Networks

Operations on an IPv6 network are similar to those on an IPv4 network. Only some functions are configured in a different way. [Table 10-4](#) describes how you can build and use an IPv6 network.

**Table 10-4** Operation guide on IPv6 networks

Scenario	Description	Reference
Creating an IPv6 subnet	Select <b>Enable</b> for <b>IPv6 CIDR Block</b> when creating a subnet. An IPv6 CIDR block will be automatically assigned to the subnet. <ul style="list-style-type: none"><li>You cannot customize an IPv6 CIDR block.</li><li>IPv6 cannot be disabled after the subnet is created.</li><li>You can enable IPv6 for existing subnets.</li></ul>	<a href="#">Creating a Subnet for the VPC</a>
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 <b>IP Addresses</b> tab.	<a href="#">Viewing IP Addresses in a Subnet</a>
Adding a security group rule (IPv6)	Add a security group rule with <b>Type</b> set to <b>IPv6</b> and <b>Source</b> or <b>Destination</b> set to an IPv6 address or IPv6 CIDR block.	<a href="#">Adding a Security Group Rule</a>
Adding a network ACL rule (IPv6)	Add a network ACL rule with <b>Type</b> set to <b>IPv6</b> and <b>Source</b> or <b>Destination</b> set to an IPv6 address or IPv6 CIDR block.	<a href="#">Adding a Network ACL Rule (Default Priorities)</a>
Adding an IPv6 route to the VPC route table	Add a route with <b>Destination</b> and <b>Next Hop</b> set to an IPv4 or IPv6 CIDR block. <ul style="list-style-type: none"><li>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.</li><li>If the destination is an IPv6 CIDR block, the next hop type can only be ECS, extension NIC, or virtual IP address. They must also have IPv6 addresses.</li></ul>	<a href="#">Adding Routes to a Route Table</a>
Assigning a virtual IPv6 address	If IPv6 is enabled for a VPC subnet, you can set <b>IP Address Type</b> to <b>IPv6</b> when assigning for a virtual IP address.	<a href="#">Assigning a Virtual IP Address</a>

# 11 VPC Flow Log

---

## 11.1 VPC Flow Log

### 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.

Currently, the VPC flow log function is supported in certain regions. You can go to [Function Overview](#) and click **VPC Flow Log** to check.

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 stream in LTS. [Figure 11-1](#) shows the process for configuring VPC flow logs.



**Figure 11-1** Configuring VPC flow logs



The VPC flow log function itself is free of charge, but you may be charged for other resources used. For example, the storage of VPC flow log records will be charged. For details, see *Log Tank Service User Guide*.

### Notes and Constraints

- Currently, S2, M2, Hc2, D2, Pi1, S3, C3, M3, H3, Ir3, I3, S6, E3, C3ne, M3ne, G5, P2v, C6, M6, Pi1, and H3 ECSs support VPC flow logs.  
For details about ECS types, see [ECS Types](#).
- Each account can have up to 10 VPC flow logs in a region.

## 11.2 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 stream.

For more information about the LTS service, see the *Log Tank Service User Guide*.

## Procedure

1. Go to the [VPC flow log list page](#).
2. In the upper right corner, click **Create VPC Flow Log**. On the displayed page, configure parameters as prompted.

**Table 11-1** Parameter descriptions

Parameter	Description	Example Value
Name	The VPC flow log name. The name: <ul style="list-style-type: none"><li>• Can contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	flowlog-495d
Resource Type	The type of resources whose traffic is to be logged. You can select <b>NIC</b> , <b>Subnet</b> , or <b>VPC</b> .	NIC
Resource	The specific NIC whose traffic is to be logged. <b>NOTE</b> 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	<ul style="list-style-type: none"><li>• <b>All traffic</b>: specifies that both accepted and rejected traffic of the specified resource will be logged.</li><li>• <b>Accepted traffic</b>: 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.</li><li>• <b>Rejected traffic</b>: specifies that only rejected traffic of the specified resource will be logged. Rejected traffic refers to the traffic denied by the network ACL.</li></ul>	All
Log Group	The log group created in LTS.	lts-group-abc
Log Stream	The log stream created in LTS.	lts-topic-abc
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

 NOTE

Only two flow logs, each with a different filter, can be created for a single resource under the same log group and log stream. Each VPC flow log must be unique.

3. After setting the parameters, click **OK**.  
Return to the VPC flow log list. You can check the new VPC flow log.

## 11.3 Viewing a VPC Flow Log

### Scenarios

This section describes how you can view the VPC flow log details.



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 flow logs cannot be collected after the VPC flow log is enabled, the possible causes are as follows:

- If an ECS is in the stopped state, its flow log records will not be displayed.
- The VPC flow log quota is insufficient. If you want to continue collecting flow logs, [configure the quota](#).

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **VPC Flow Logs**.  
The **VPC Flow Logs** page is displayed.
5. 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:

```
<version> <project-id> <interface-id> <srcaddr> <dstaddr> <srcport> <dstport> <protocol> <packets>  
<bytes> <start> <end> <action> <log-status>
```

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 11-2** describes the fields of a flow log record.

**Table 11-2** Log field description

Field	Description	Example Value
version	The VPC flow log version.	1
project-id	The project ID.	5f67944957444bd6bb4fe3b367de8f3d
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 <a href="#">Assigned Internet Protocol Numbers</a> .	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: <ul style="list-style-type: none"><li>● <b>ACCEPT</b>: The recorded traffic was allowed by the security groups or network ACLs.</li><li>● <b>REJECT</b>: The recorded traffic was denied by the security groups or network ACLs.</li></ul>	ACCEPT
log-status	The logging status of the VPC flow log: <ul style="list-style-type: none"><li>● <b>OK</b>: Data is logging normally to the chosen destinations.</li><li>● <b>NODATA</b>: There was no traffic of the <b>Filter</b> setting to or from the NIC during the capture window.</li><li>● <b>SKIPDATA</b>: Some flow log records were skipped during the capture window. This may be caused by an internal capacity constraint or an internal error.</li></ul> Example: When <b>Filter</b> is set to <b>Accepted traffic</b> , if there is accepted traffic, the value of <b>log-status</b> is <b>OK</b> . If there is no accepted traffic, the value of <b>log-status</b> is <b>NODATA</b> regardless of whether there is rejected traffic. If some accepted traffic is abnormally skipped, the value of <b>log-status</b> is <b>SKIPDATA</b> .	OK

You can enter a keyword on the log stream details page on the LTS console to search for flow log records.

## 11.4 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **VPC Flow Logs**.  
The **VPC Flow Logs** page is displayed.
5. Locate the target flow log and click **Enable** or **Disable** in the **Operation** column.  
A confirmation dialog box is displayed.
6. Confirm the information and click **OK**.

## 11.5 Deleting a VPC Flow Log


### Scenarios


You can delete a VPC flow log if you no longer need it. 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 select the desired region and project.

3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **VPC Flow Logs**.  
The **VPC Flow Logs** page is displayed.
5. Locate the target flow log and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
6. Confirm the information and click **OK**.

# 12 Traffic Mirroring

---

## 12.1 Traffic Mirroring

### What Is Traffic Mirroring?

Traffic Mirroring can be used to mirror traffic that meets a mirror filter from an elastic network interface. You can configure inbound and outbound rules for a mirror filter to determine which traffic from an elastic network interface will be mirrored to a network interface or load balancer. You can then send the traffic for inspection, audit analysis, and troubleshooting.

---

#### NOTICE

Currently, the Traffic Mirroring function is free. You will be notified in advance if the billing starts.

Currently, Traffic Mirroring is available only in certain regions. For details, visit [Function Overview](#) and click **Traffic Mirroring**.

---

### Concepts

The following are the concepts for Traffic Mirroring:

- A mirror filter is a set of inbound rules and outbound rules to determine the traffic that is mirrored. You can specify matching criteria, such as priority and action for each rule.
  - Inbound rules match the traffic received by a mirror source.
  - Outbound rules match the traffic sent by a mirror source.
- A mirror source is an elastic network interface and traffic of an elastic network interface needs to be mirrored.
- A mirror target is an ECS network interface or a load balancer, which is used to receive mirrored traffic.
- A mirror session can be associated with a mirror filter, multiple mirror sources, and a mirror target. A mirror session mirrors traffic from a mirror source to a mirror target that meets the mirror filter.



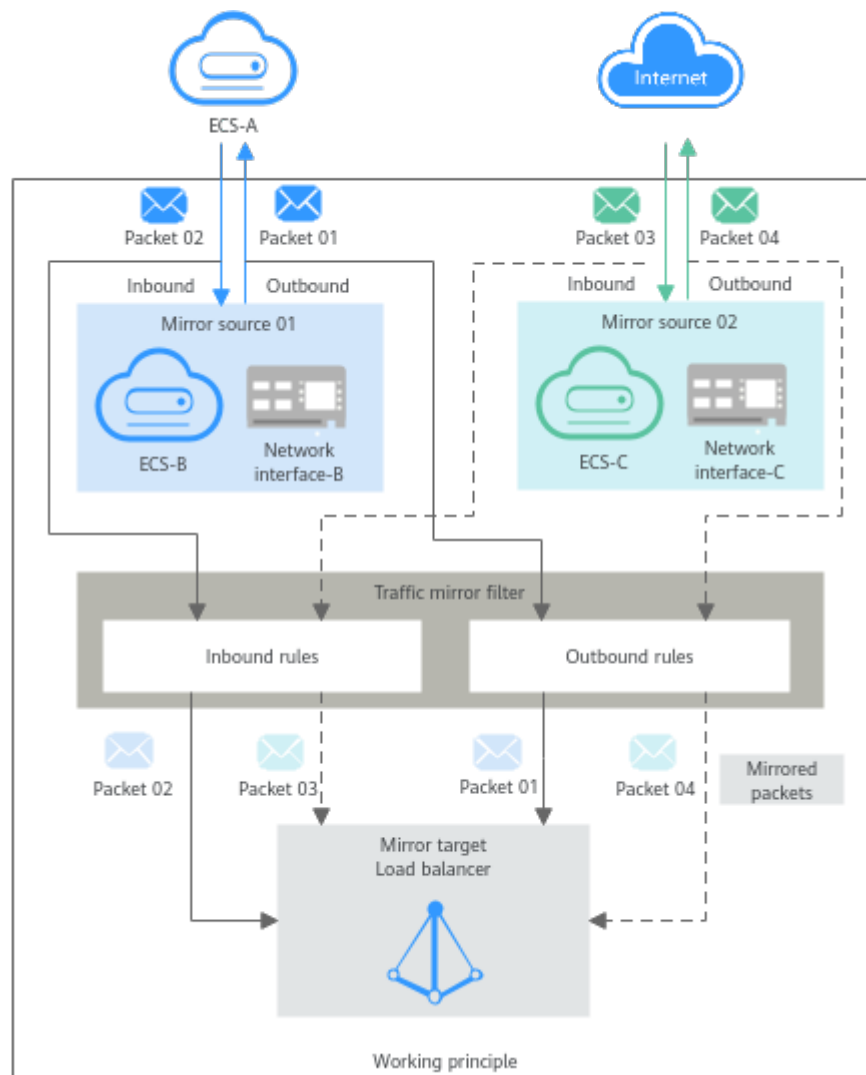
## Working Principles

The following describes the working principles of traffic mirroring. As shown in [Figure 12-1](#), a mirror session is associated with two mirror sources, one mirror filter, and one mirror target.

- Mirror source 01 is network interface-B that is attached to ECS-B. To access ECS-A from ECS-B, the outbound and inbound traffic of network interface-B is mirrored.
- Mirror source 02 is network interface-C that is attached to ECS-C. To access ECS-C from the Internet, the outbound and inbound traffic of network interface-C is mirrored.
- The mirror filter contains both inbound and outbound rules.
- The mirror target is a load balancer that receives mirrored traffic.

In [Table 12-1](#), mirror sources network interface-B and network interface-C are used as examples to describe the traffic mirroring principle.

**Figure 12-1** Traffic mirroring architecture



**Table 12-1** Mirror path of packets

Mirror Source	Access Path	Packet	Direction	Description
Network interface-B	From ECS-B to ECS-A	Request packet 01	Outbound	Request packet 01 from ECS-B is an outbound packet for network interface-B. If packet 01 matches the outbound rules of the mirror filter, packet 01 is mirrored to the load balancer.
		Response packet 02	Inbound	Response packet 02 from ECS-A is an inbound packet for network interface-B. If packet 02 matches the inbound rules of the mirror filter, packet 02 is mirrored to the load balancer.
Network interface-C	From the Internet to ECS-C	Request packet 03	Inbound	Request packet 03 from the Internet is an inbound packet for network interface-C. If packet 03 matches the inbound rules of the mirror filter, packet 03 is mirrored to the load balancer.
		Response packet 04	Outbound	Response packet 04 from ECS-C is an outbound packet for network interface-C. If packet 04 matches the outbound rules of the mirror filter, packet 04 is mirrored to the load balancer.

**Table 12-2** shows rules in a mirror filter and describes how the mirror session filters traffic using the mirror filter.

**Table 12-2** Traffic filtering description

Direction	Priority	Protocol	Action	Type	Source	Source Port Range	Destination	Destination Port Range	Filtering Description
Inbound	1	TCP	Accept	IPv4	172.16.0.0/24	10000-10001	10.0.0.3/32	80-80	If traffic enters a network interface of the mirror source, the mirror session will mirror packets that meet the following rule:  TCP (IPv4) packets from source 172.16.0.0/24 over port 10000 or 10001 to destination 10.0.0.3/32 over port 80
Outbound	1	All	Reject	IPv4	192.168.0.0/24	All	10.2.0.0/24	All	If traffic leaves a network interface of the mirror source, the mirror session will not mirror packets that meet the following rule:  IPv4 packets from source 192.168.0.0/24 over any port to destination 10.2.0.0/24 over any port.

## Application Scenarios

- Traffic inspection  
If there are network intrusions, you can use traffic mirroring to mirror required traffic to security software for comprehensive analysis and check. This helps to quickly locate security vulnerabilities and ensure network security.
- Traffic auditing  
You can use traffic mirroring to mirror traffic to a specific platform for auditing and analysis. This applies to scenarios that have high security requirements, such as finance.
- Fault locating  
O&M engineers can directly view mirrored traffic instead of capturing packets on service servers to locate faults. This prevents services from being affected during O&M.

## Matching Rules

If a packet from the same mirror source meets multiple mirror filter rules, the packet is matched only once. The matching rules are described as follows:


**Table 12-3** Matching rules

Matching Rule	Description
Sequence match	<p>Matching is performed in descending order of priority. A smaller value indicates a higher priority. For example, the priority of 1 is higher than that of 2.</p> <ul style="list-style-type: none"><li>• Mirror session priority: A mirror source can be associated with multiple mirror sessions at the same time. The mirror sessions are matched in descending order of priority. For details, see <a href="#">the matching process of mirror sessions</a>.</li><li>• Mirror filter rule priority: A mirror session can be associated with only one mirror filter that contains multiple rules. The rules are matched in descending order of priority. An inbound or outbound mirror filter rule determines the traffic that is mirrored. You can specify matching criteria, such as priority and action for each rule. For details, see <a href="#">the matching process of mirror filter rules</a>.</li></ul>
Unique match	<p>If a packet matches a mirror filter rule, the packet does not attempt to match any other rules.</p>

- [Figure 12-2](#) describes the matching process of mirror sessions. If a mirror source is associated with multiple mirror sessions, packets are matched in descending order of mirror session priorities. Inbound packets are used as an example here.
  - If a packet matches an inbound mirror filter rule of a mirror session:
    - The packet will be mirrored if the rule action is **Accept**.
    - The packet will not be mirrored if the rule action is **Reject**.
  - If a packet does not match any inbound mirror filter rule in a mirror session, the packet will not be mirrored.

For example, a mirror source is associated with both mirror sessions A and B. The priority of mirror session A is 1 and that of mirror session B is 2. If a packet in the inbound direction of the mirror source meets the mirror filter rules of both mirror session A and mirror session B, the packet preferentially matches the mirror filter rules of mirror session A according to the priority, and will not match that of mirror session B.


**Figure 12-2** Mirror session matching process



- **Figure 12-3** describes the matching process of mirror filter rules. If a mirror source is associated with only one mirror session, packets are matched in descending order of priorities of inbound mirror filter rules. Inbound packets are used as an example here.
  - If a packet matches an inbound mirror filter rule:
    - The packet will be mirrored if the rule action is **Accept**.
    - The packet will not be mirrored if the rule action is **Reject**.
  - If a packet does not match any inbound mirror filter rule, the packet will not be mirrored.

For example, a mirror source is associated with mirror session A. The mirror filter of mirror session A has inbound rules A and B, which have the same traffic matching conditions but different priorities and actions. The priority of rule A is 1, and the action is **Reject**. The priority of rule B is 2, and the action is **Accept**. If a packet in the inbound direction of the mirror source meets the traffic matching conditions of both rule A and rule B, the packet matches rule A first according to the rule priority. The packet will be rejected and will not be mirrored and match rule B.

**Figure 12-3** Mirror filter rule matching process



## Traffic Mirroring Quotas

**Table 12-4** lists the quotas about Traffic Mirroring resources. Some default quotas can be increased.

**Table 12-4** Quotas


Item	Default Quota	Adjustable
Maximum number of mirror sources that can be associated with a mirror session	10	Yes
Maximum number of mirror sessions that can be associated with a mirror source	3	No
Maximum number of mirror targets that can be associated with a mirror session	1	No
Maximum number of mirror sessions that can be associated with a mirror target	<ul style="list-style-type: none"> <li>10 (if the mirror target is an ECS network interface)</li> <li>200 (if the mirror target is a load balancer)</li> </ul>	No

Item	Default Quota	Adjustable
Maximum number of mirror filters that can be associated with a mirror session	1	No
Maximum number of mirror sessions that can be associated with a mirror filter	1,000	No
Maximum number of rules that can be added to a mirror filter	<ul style="list-style-type: none"> <li>• 10 inbound rules</li> <li>• 10 outbound rules</li> </ul>	No
Maximum number of mirror sessions that can be created in a region	20,000	No

## Notes and Constraints

- As shown in [Figure 12-4](#), mirrored traffic is encapsulated in the standard VXLAN packet format. For more information about the VXLAN protocol, see [RFC 7348](#). If the total length of mirrored packets and VXLAN packets is greater than the MTU of the mirror source, the system truncates the packets. To prevent packets from being truncated, you are advised to set the MTU of the elastic network interface to be at least 64 bytes smaller than the MTU supported by the link in IPv4 scenarios.

**Figure 12-4** Traffic Mirroring packet format



- Currently, only elastic network interfaces of c7t ECSs can be used as mirror sources.
- An elastic network interface cannot be used as both a mirror source and a mirror target at the same time.
- Traffic Mirroring occupies the bandwidth of instances attached to elastic network interfaces and does not have bandwidth limits.
- If a mirror target needs to receive mirrored traffic from multiple mirror sources, ensure that the mirror target has proper specifications based on service requirements.

- If a packet from a mirror source meets multiple mirror filter rules, the packet will be matched only once and will be accepted or rejected to a mirror target according to the rule action.
- If a packet from a mirror source is discarded by a security group or network ACL, the packet will not be mirrored.
- If a packet from a mirror source meets a mirror filter, the packet will be mirrored and will not be restricted by outbound rules of a security group or network ACL of the mirror source. That is, you do not need to configure the security group or network ACL for the mirror source. However, if you want to mirror the packet to the mirror target, you need to configure the following rules for the security group and network ACL of the mirror target:
  - Add a security group rule to allow inbound UDP packets from the IP address of the mirror source (elastic network interface) over port 4789. [Table 12-5](#) shows a rule example if the IP address of the mirror source is 192.168.0.27. To learn about how to add a rule, see [Adding a Security Group Rule](#).

**Table 12-5** Security group rule example

Direction	Action	Type	Protocol & Port	Source
Inbound	Allow	IPv4	UDP: 4789	IP address: 192.168.0.27/32  Set the IP address based on the actual requirements.

- Add a network ACL rule to allow inbound UDP packets from the IP address of the mirror source (elastic network interface) over any port. [Table 12-6](#) shows a rule example if the IP address of the mirror source is 192.168.0.27. To learn about how to add a rule, see [Adding a Network ACL Rule \(Default Priorities\)](#).



**Table 12-6** Network ACL rule example

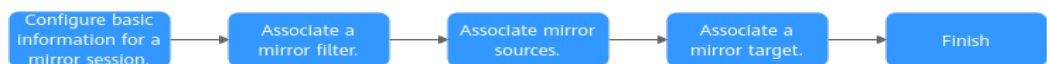
Direction	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range
Inbound	IPv4	Allow	UDP	IP address: 192.168.0.27/32 Set the IP address based on the actual requirements.	If not specified, all ports are used.	IP address: 10.10.0.0/24 Set the IP address based on the actual requirements.	4789 Port 4789 must be opened. Open other ports based on the actual requirements.

- Resources from different VPCs cannot communicate with each other. If a mirror source and a mirror target are not in the same VPC, you need to use a VPC peering connection or an enterprise router to connect their VPCs first.
  - To use a VPC peering connection, see [VPC Peering Connection Overview](#).
  - To use an enterprise router, see [Using an Enterprise Router to Enable Communications Between VPCs in the Same Region](#).

## Usage Process

To use the traffic mirroring function, you need to create a mirror session and associate a mirror filter, mirror sources, and a mirror target with the mirror session. [Figure 12-5](#) shows the process.

**Figure 12-5** Process of using Traffic Mirroring



**Table 12-7** Description of the Traffic Mirroring process

Step	Description	Reference
Configure basic information about a mirror session.	Set parameters such as the name and priority of the mirror session.	<a href="#">Creating a Mirror Session</a>

Step	Description	Reference
Associate a mirror filter.	Select a mirror filter and associate it with the mirror session. Each mirror session can have one mirror filter associated. If there is no mirror filter required, you can create one by referring to <a href="#">Creating a Mirror Filter</a> .	
Associate mirror sources.	Select an elastic network interface as the mirror source and associate it with the mirror session. <ul style="list-style-type: none"><li>• Each mirror session can be associated with multiple mirror sources.</li><li>• Currently, only elastic network interfaces of c7t ECSs can be used as mirror sources.</li></ul>	
Associate a mirror target.	Select an ECS network interface or load balancer as the mirror target and associate it with the mirror session.	
Finish	If the mirror session is enabled, the traffic that meets the mirror filter from the mirror source will be mirrored to the mirror target. If you do not enable the mirror session when creating it, the traffic of the mirror source will not be mirrored. You can enable the mirror session by referring to <a href="#">Enabling or Disabling a Mirror Session</a> .	

## 12.2 Mirror Filters

### 12.2.1 Creating a Mirror Filter

#### Scenarios

A mirror filter is a set of inbound rules and outbound rules to determine the traffic that is mirrored. You can specify matching criteria, such as priority and action for each rule.

- Inbound rules match the traffic received by a mirror source.
- Outbound rules match the traffic sent by a mirror source.

A mirror filter takes effect only after it is associated with mirror sessions.



#### Mirror Filter Rule Examples

[Table 12-8](#) shows rules in a mirror filter and describes how the mirror session filters traffic using the mirror filter.

**Table 12-8** Traffic filtering description

Direction	Priority	Protocol	Action	Type	Source	Source Port Range	Destination	Destination Port Range	Filtering Description
Inbound	1	TCP	Accept	IPv4	172.16.0.0/24	10000-10001	10.0.0.3/32	80-80	If traffic enters a network interface of the mirror source, the mirror session will mirror packets that meet the following rule: TCP (IPv4) packets from source 172.16.0.0/24 over port 10000 or 10001 to destination 10.0.0.3/32 over port 80
Outbound	1	All	Reject	IPv4	192.168.0.0/24	All	10.2.0.0/24	All	If traffic leaves a network interface of the mirror source, the mirror session will not mirror packets that meet the following rule: IPv4 packets from source 192.168.0.0/24 over any port to destination 10.2.0.0/24 over any port.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. In the upper right corner of the mirror filter list, click **Create Mirror Filter**.  
The **Create Mirror Filter** page is displayed.
6. Set basic information about the mirror filter as prompted.

**Table 12-9** Parameters for configuring basic information

Parameter	Description	Example Value
Name	Mandatory Enter the name of the mirror filter. The name: <ul style="list-style-type: none"> <li>• Must contain 1 to 64 characters.</li> <li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li> </ul>	mirror-filter-01
Description	Optional Enter the description of the mirror filter in the text box as required.	-

7. Click **Add Rule** in the **Inbound Rules** area to add inbound rules.  
You can click + to add more inbound rules.

**Table 12-10** Inbound rule parameter description

Parameter	Description	Example Value
Priority	Priority of a mirror filter rule. <ul style="list-style-type: none"> <li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li> <li>• Priorities of inbound rules must be unique for each mirror filter.</li> </ul> A mirror filter can contain multiple rules and the rules are matched in ascending order of priority. For details, see <a href="#">the matching process of mirror filter rules</a> .	1
Protocol	Select a network protocol. <ul style="list-style-type: none"> <li>• If you select TCP, you can customize the source and destination port ranges.</li> <li>• If you select UDP, you can customize the source and destination port ranges.</li> <li>• If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li> <li>• If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li> <li>• If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li> </ul>	TCP

Parameter	Description	Example Value
Action	<p>Whether to accept or reject inbound traffic of a mirror source.</p> <ul style="list-style-type: none"> <li>If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li> <li>If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li> </ul>	Accept
Type	<p>IP address version of inbound traffic. You can specify:</p> <ul style="list-style-type: none"> <li><b>IPv4</b></li> <li><b>IPv6</b></li> </ul>	IPv4
Source	<p>Source of inbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	10.0.0.0/24
Source Port Range	<p>Source port range of inbound traffic.</p> <ul style="list-style-type: none"> <li>Port range: 1 to 65535</li> <li>Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	22-23
Destination	<p>Destination of inbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	0.0.0.0/0

Parameter	Description	Example Value
Destination Port Range	Destination port range of inbound traffic. <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	1-65535
Description	Enter the description of the mirror filter rule in the text box as required.	-

- Click **OK**.
- Click **Add Rule** in the **Outbound Rules** area to add outbound rules.  
You can click + to add more outbound rules.

**Table 12-11** Outbound rule parameter description

Parameter	Description	Example Value
Priority	Priority of a mirror filter rule. <ul style="list-style-type: none"> <li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li> <li>• Priorities of inbound rules must be unique for each mirror filter.</li> </ul> A mirror filter can contain multiple rules and the rules are matched in ascending order of priority. For details, see <a href="#">the matching process of mirror filter rules</a> .	1
Protocol	Select a network protocol. <ul style="list-style-type: none"> <li>• If you select TCP, you can customize the source and destination port ranges.</li> <li>• If you select UDP, you can customize the source and destination port ranges.</li> <li>• If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li> <li>• If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li> <li>• If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li> </ul>	All

Parameter	Description	Example Value
Action	Whether to accept or reject outbound traffic of a mirror source. <ul style="list-style-type: none"><li>If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li><li>If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li></ul>	Reject
Type	IP address version of outbound traffic. You can specify: <ul style="list-style-type: none"><li><b>IPv4</b></li><li><b>IPv6</b></li></ul>	IPv4
Source	Source of outbound traffic. You can enter: <ul style="list-style-type: none"><li>A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul>	192.168.0.0/24
Source Port Range	Source port range of outbound traffic. <ul style="list-style-type: none"><li>Port range: 1 to 65535</li><li>Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li><li>If not specified or <b>1-65535</b> is specified, all ports are used.</li></ul>	All
Destination	Destination of outbound traffic. You can enter: <ul style="list-style-type: none"><li>A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul>	10.2.0.0/24

Parameter	Description	Example Value
Destination Port Range	Destination port range of outbound traffic. <ul style="list-style-type: none"><li>• Port range: 1 to 65535</li><li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li><li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li></ul>	All
Description	Enter the description of the mirror filter rule in the text box as required.	-

10. Click **OK**.
11. After setting the parameters, click **Create Now**.  
The mirror filter list page is displayed.

## Follow-up Operations

A mirror filter takes effect only after it is associated with mirror sessions. Each mirror session only can have one mirror filter associated.

- If you have no mirror session, refer to [Creating a Mirror Session](#).
- If you have a mirror session and want to change the mirror filter of the mirror session, refer to [Changing a Mirror Filter for a Mirror Session](#).

## 12.2.2 Adding an Inbound or Outbound Mirror Filter Rule

### Scenarios

You can add inbound and outbound rules to a mirror filter.

- Inbound rules match the traffic received by a mirror source.
- Outbound rules match the traffic sent by a mirror source.

### Mirror Filter Rule Examples



[Table 12-12](#) shows rules in a mirror filter and describes how the mirror session filters traffic using the mirror filter.



**Table 12-12** Traffic filtering description

Direction	Priority	Protocol	Action	Type	Source	Source Port Range	Destination	Destination Port Range	Filtering Description
Inbound	1	TCP	Accept	IPv4	172.16.0.0/24	10000-10001	10.0.0.3/32	80-80	If traffic enters a network interface of the mirror source, the mirror session will mirror packets that meet the following rule: TCP (IPv4) packets from source 172.16.0.0/24 over port 10000 or 10001 to destination 10.0.0.3/32 over port 80
Outbound	1	All	Reject	IPv4	192.168.0.0/24	All	10.2.0.0/24	All	If traffic leaves a network interface of the mirror source, the mirror session will not mirror packets that meet the following rule: IPv4 packets from source 192.168.0.0/24 over any port to destination 10.2.0.0/24 over any port.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. Locate the row that contains the mirror filter and click the hyperlink in the **Inbound and Outbound Rules** column.  
The **Inbound Rules** tab page is displayed.
6. On the **Inbound Rules** tab, click **Add Rule**. In the displayed dialog box, set required parameters to add an inbound rule.  
You can click **+** to add more inbound rules.

**Table 12-13** Inbound rule parameter description

Parameter	Description	Example Value
Priority	<p>Priority of a mirror filter rule.</p> <ul style="list-style-type: none"><li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li><li>• Priorities of inbound rules must be unique for each mirror filter.</li></ul> <p>A mirror filter can contain multiple rules and the rules are matched in ascending order of priority.</p> <p>For details, see <a href="#">the matching process of mirror filter rules</a>.</p>	1
Protocol	<p>Select a network protocol.</p> <ul style="list-style-type: none"><li>• If you select TCP, you can customize the source and destination port ranges.</li><li>• If you select UDP, you can customize the source and destination port ranges.</li><li>• If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li><li>• If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li><li>• If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li></ul>	TCP
Action	<p>Whether to accept or reject inbound traffic of a mirror source.</p> <ul style="list-style-type: none"><li>• If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li><li>• If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li></ul>	Accept
Type	<p>IP address version of inbound traffic. You can specify:</p> <ul style="list-style-type: none"><li>• <b>IPv4</b></li><li>• <b>IPv6</b></li></ul>	IPv4

Parameter	Description	Example Value
Source	<p>Source of inbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	10.0.0.0/24
Source Port Range	<p>Source port range of inbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	22-23
Destination	<p>Destination of inbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	0.0.0.0/0
Destination Port Range	<p>Destination port range of inbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	1-65535
Description	Enter the description of the mirror filter rule in the text box as required.	-

7. Click **OK**.  
You can view the added inbound rule in the list.
8. Click the **Outbound Rules** tab. In the upper left corner of the outbound rule list, click **Add Rule** to add an outbound rule.  
You can click **+** to add more outbound rules.

**Table 12-14** Outbound rule parameter description

Parameter	Description	Example Value
Priority	<p>Priority of a mirror filter rule.</p> <ul style="list-style-type: none"> <li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li> <li>• Priorities of inbound rules must be unique for each mirror filter.</li> </ul> <p>A mirror filter can contain multiple rules and the rules are matched in ascending order of priority. For details, see <a href="#">the matching process of mirror filter rules</a>.</p>	1
Protocol	<p>Select a network protocol.</p> <ul style="list-style-type: none"> <li>• If you select TCP, you can customize the source and destination port ranges.</li> <li>• If you select UDP, you can customize the source and destination port ranges.</li> <li>• If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li> <li>• If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li> <li>• If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li> </ul>	All
Action	<p>Whether to accept or reject outbound traffic of a mirror source.</p> <ul style="list-style-type: none"> <li>• If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li> <li>• If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li> </ul>	Reject
Type	<p>IP address version of outbound traffic. You can specify:</p> <ul style="list-style-type: none"> <li>• <b>IPv4</b></li> <li>• <b>IPv6</b></li> </ul>	IPv4

Parameter	Description	Example Value
Source	<p>Source of outbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	192.168.0.0/24
Source Port Range	<p>Source port range of outbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	All
Destination	<p>Destination of outbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	10.2.0.0/24
Destination Port Range	<p>Destination port range of outbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	All
Description	Enter the description of the mirror filter rule in the text box as required.	-

9. Click **OK**.  
You can view the added outbound rule in the list.



## 12.2.3 Modifying an Inbound or Outbound Mirror Filter Rule

### Scenarios

You can modify inbound and outbound rules of a mirror filter.

- Inbound rules match the traffic received by a mirror source.
- Outbound rules match the traffic sent by a mirror source.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. Locate the row that contains the mirror filter and click the hyperlink in the **Inbound and Outbound Rules** column.  
The **Inbound Rules** tab page is displayed.
6. In the inbound rule list, locate the row that contains the rule and click **Modify** in the **Operation** column.

**Table 12-15** Inbound rule parameter description

Parameter	Description	Example Value
Priority	<p>Priority of a mirror filter rule.</p> <ul style="list-style-type: none"><li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li><li>• Priorities of inbound rules must be unique for each mirror filter.</li></ul> <p>A mirror filter can contain multiple rules and the rules are matched in ascending order of priority. For details, see <a href="#">the matching process of mirror filter rules</a>.</p>	1

Parameter	Description	Example Value
Protocol	<p>Select a network protocol.</p> <ul style="list-style-type: none"> <li>If you select TCP, you can customize the source and destination port ranges.</li> <li>If you select UDP, you can customize the source and destination port ranges.</li> <li>If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li> <li>If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li> <li>If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li> </ul>	TCP
Action	<p>Whether to accept or reject inbound traffic of a mirror source.</p> <ul style="list-style-type: none"> <li>If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li> <li>If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li> </ul>	Accept
Type	<p>IP address version of inbound traffic. You can specify:</p> <ul style="list-style-type: none"> <li><b>IPv4</b></li> <li><b>IPv6</b></li> </ul>	IPv4
Source	<p>Source of inbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	10.0.0.0/24

Parameter	Description	Example Value
Source Port Range	Source port range of inbound traffic. <ul style="list-style-type: none"><li>• Port range: 1 to 65535</li><li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li><li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li></ul>	22-23
Destination	Destination of inbound traffic. You can enter: <ul style="list-style-type: none"><li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li><li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li><li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li></ul>	0.0.0.0/0
Destination Port Range	Destination port range of inbound traffic. <ul style="list-style-type: none"><li>• Port range: 1 to 65535</li><li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li><li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li></ul>	1-65535
Description	Enter the description of the mirror filter rule in the text box as required.	-

7. Click **OK**.  
You can view the modified inbound rule in the list.
8. On the **Outbound Rules** tab page, locate the row that contains the rule in the outbound rule list and click **Modify** in the **Operation** column.



**Table 12-16** Outbound rule parameter description

Parameter	Description	Example Value
Priority	<p>Priority of a mirror filter rule.</p> <ul style="list-style-type: none"> <li>• A priority value can be from 1 to 65535. A smaller value indicates a higher priority.</li> <li>• Priorities of inbound rules must be unique for each mirror filter.</li> </ul> <p>A mirror filter can contain multiple rules and the rules are matched in ascending order of priority. For details, see <a href="#">the matching process of mirror filter rules</a>.</p>	1
Protocol	<p>Select a network protocol.</p> <ul style="list-style-type: none"> <li>• If you select TCP, you can customize the source and destination port ranges.</li> <li>• If you select UDP, you can customize the source and destination port ranges.</li> <li>• If you set <b>Type</b> to <b>IPv4</b> and select ICMP, all ports are specified for source and destination port ranges by default.</li> <li>• If you set <b>Type</b> to <b>IPv6</b> and select ICMPv6, all ports are specified for source and destination port ranges by default.</li> <li>• If you select <b>All</b>, all network protocols are supported and all ports are specified for source and destination port ranges by default.</li> </ul>	All
Action	<p>Whether to accept or reject outbound traffic of a mirror source.</p> <ul style="list-style-type: none"> <li>• If you set <b>Action</b> to <b>Accept</b>, the traffic will be mirrored to the mirror target.</li> <li>• If you set <b>Action</b> to <b>Reject</b>, the traffic will not be mirrored to the mirror target.</li> </ul>	Reject
Type	<p>IP address version of outbound traffic. You can specify:</p> <ul style="list-style-type: none"> <li>• <b>IPv4</b></li> <li>• <b>IPv6</b></li> </ul>	IPv4

Parameter	Description	Example Value
Source	<p>Source of outbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	192.168.0.0/24
Source Port Range	<p>Source port range of outbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	All
Destination	<p>Destination of outbound traffic. You can enter:</p> <ul style="list-style-type: none"> <li>• A single IP address: IP address/mask Example IPv4 address: 192.168.10.10/32 Example IPv6 address: 2002:50::44/128</li> <li>• An IP address range in CIDR notation: IP address/mask Example IPv4 address range: 192.168.52.0/24 Example IPv6 address range: 2407:c080:802:469::/64</li> <li>• All IP addresses 0.0.0.0/0 represents all IPv4 addresses. ::/0 represents all IPv6 addresses.</li> </ul>	10.2.0.0/24
Destination Port Range	<p>Destination port range of outbound traffic.</p> <ul style="list-style-type: none"> <li>• Port range: 1 to 65535</li> <li>• Use a hyphen (-) to connect the start port and the end port, for example, 22-23. The end port cannot be smaller than the start port.</li> <li>• If not specified or <b>1-65535</b> is specified, all ports are used.</li> </ul>	All
Description	Enter the description of the mirror filter rule in the text box as required.	-



9. Click **OK**.  
You can view the modified outbound rule in the list.

## 12.2.4 Deleting an Inbound or Outbound Mirror Filter Rule

### Scenarios

You can delete inbound and outbound rules of a mirror filter.

### Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. Locate the row that contains the mirror filter and click the hyperlink in the **Inbound and Outbound Rules** column.  
The **Inbound Rules** tab page is displayed.
6. Locate the row that contains the inbound rule and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
7. Click **Yes**.  
Deleted inbound rules cannot be recovered.
8. On the **Outbound Rules** tab page, locate the row that contains the rule in the outbound rule list and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
9. Click **Yes**.  
Deleted outbound rules cannot be recovered.



## 12.2.5 Modifying Basic Information About a Mirror Filter

### Scenarios

You can modify basic information about a mirror filter, including its name and description.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

- The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
  5. Locate the row that contains the mirror filter and click its name with a hyperlink.  
The **Inbound Rules** tab page is displayed.
  6. Click the **Basic Information** tab and modify parameters as prompted.
    - a. Click  next to the parameter to be modified and enter information in the text box.
    - b. Click  to save the modification.

**Table 12-17** Parameters for configuring basic information

Parameter	Description	Example Value
Name	Mandatory Enter the name of the mirror filter. The name: <ul style="list-style-type: none"> <li>• Must contain 1 to 64 characters.</li> <li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li> </ul>	mirror-filter-01
Description	Optional Enter the description of the mirror filter in the text box as required.	-



## 12.2.6 Viewing Details About a Mirror Filter

### Scenarios

You can view the following information about a mirror filter:

- Basic information, such as name, ID, and creation time
- Inbound and outbound rules, such as their priority, protocol, and action
- Associated mirror sessions, such as their name, mirror target, and status

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. Locate the row that contains the mirror filter and click its name with a hyperlink.  
The **Inbound Rules** tab page is displayed.
6. View information about the mirror filter on different tab pages.
  - On the **Basic Information** tab page, view mirror filter information, such as name, ID, and creation time.
  - On the **Inbound Rules** tab page, view rule details, such as their priority, protocol, and action.
  - On the **Outbound Rules** tab page, view rule details, such as their priority, protocol, and action.
  - On the **Associated Mirror Sessions** tab page, view information about mirror sessions, such as their name, mirror target, and status.

## 12.2.7 Deleting a Mirror Filter

### Scenarios



If a mirror filter is no longer required, you can delete it.

### Notes and Constraints

If a mirror filter has mirror sessions associated, disassociate the mirror sessions first and then delete the mirror filter.

- Each mirror session must have a mirror filter associated. You can change the mirror filter for the mirror sessions. For details, see [Changing a Mirror Filter for a Mirror Session](#).
- If your mirror sessions are no longer required, you can also delete them. For details, see [Deleting a Mirror Session](#).

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Filters**.  
The mirror filter list page is displayed.
5. Locate the row that contains the mirror filter and click **Delete** in the **Operation** column.  
A confirmation dialog box is displayed.
6. Confirm the information and click **Yes**.  
A deleted mirror filter cannot be recovered.

## 12.3 Mirror Sessions



### 12.3.1 Creating a Mirror Session

#### Scenarios

To use Traffic Mirroring, you need to create a mirror session, associate a mirror filter, multiple mirror sources, and a mirror target with the mirror session. A mirror session mirrors traffic from a mirror source to a mirror target that meets the mirror filter.

For details about mirror sessions, see [Traffic Mirroring](#).

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the upper right corner of the mirror session list, click **Create Mirror Session**.  
The **Create Mirror Session** page is displayed.
6. Set basic information about the mirror session as prompted.

**Table 12-18** Parameters for configuring basic information about a mirror session

Parameter	Description	Example Value
Name	Mandatory Enter the name of the mirror session. The name: <ul style="list-style-type: none"><li>• Must contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	mirror-session-01

Parameter	Description	Example Value
Priority	<p>Mandatory</p> <p>Priority of the mirror session.</p> <ul style="list-style-type: none"> <li>• A priority value can be from 1 to 32766. A smaller value indicates a higher priority.</li> <li>• Priorities must be unique for each mirror session of the same account in a region.</li> </ul> <p>A mirror source can be associated with multiple mirror sessions at the same time. The mirror sessions are matched from the lowest value to the highest value.</p> <p>For details, see <a href="#">the matching process of mirror sessions</a>.</p>	1
VNI	<p>Optional</p> <p>VXLAN Network Identifiers (VNIs) are used to distinguish different mirror sessions for a mirror target. A VNI can be from 0 to 16777215.</p> <p>If not specified, the default value is 1.</p>	1
Packet Length	<p>Optional</p> <p>The number of bytes that meet the mirror filter and will be mirrored. The value can be from 1 to 1460.</p> <p>If not specified, the default value is 96.</p>	96
Mirror Session	<p>Optional</p> <ul style="list-style-type: none"> <li>• If the mirror session is disabled, traffic of mirror sources cannot be monitored.</li> <li>• If the mirror session is enabled, traffic of mirror sources can be monitored.</li> </ul>	Enable
Description	<p>Optional</p> <p>Enter the description of the mirror session in the text box as required.</p>	-

7. Click **Next**.  
The **Associate Mirror Filter** page is displayed.
8. In the mirror filter list, select a mirror filter.  
Each mirror session can be associated with only one mirror filter.  
If there is no mirror filter you want, create one by referring to [Creating a Mirror Filter](#).
9. Click **Next**.  
The **Associate Mirror Sources** page is displayed.

10. In the mirror source list, select mirror sources.
  - A mirror source is an elastic network interface and traffic of this network interface needs to be mirrored.
  - Each mirror session can be associated with multiple mirror sources.
  - An elastic network interface needs to be attached to an ECS. Currently, only network interfaces of c7t ECSs can be used as mirror sources.
11. Click **Next**.

The **Associate Mirror Target** page is displayed.
12. In the mirror target list, select a mirror target.
  - A mirror target is a network interface of an ECS or a load balancer, which is used to receive mirrored traffic.
  - Each mirror session can be associated with only one mirror target.
13. Click **Next**.

The **Confirm** page is displayed.
14. After confirming that the configuration is correct, click **Create Now**.

You can view the created mirror session in the mirror session list.



## 12.3.2 Enabling or Disabling a Mirror Session

### Scenarios

You can enable or disable a mirror session.

- If the mirror session is disabled, traffic of mirror sources cannot be monitored.
- If the mirror session is enabled, traffic of mirror sources can be monitored.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.

The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and click **Enable** or **Disable** in the **Operation** column.

A confirmation dialog box is displayed.
6. Click **Yes**.

## 12.3.3 Associating Mirror Sources with a Mirror Session



### Scenarios

You can associate mirror sources with a mirror session.



- A mirror source is an elastic network interface and traffic of this network interface needs to be mirrored.
- Each mirror session can be associated with multiple mirror sources.
- An elastic network interface needs to be attached to an ECS. Currently, only network interfaces of c7t ECSs can be used as mirror sources.

## Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and click its name with a hyperlink.  
The **Basic Information** page is displayed.
6. Click the **Mirror Sources** tab. In the upper left corner of the mirror source list, click **Associate**.  
The **Associate Mirror Sources** dialog box is displayed.
7. In the mirror source list, select mirror sources and click **OK**.  
In the mirror source list, you can view the associated mirror sources.

## 12.3.4 Disassociating Mirror Sources from a Mirror Session

### Scenarios

You can disassociate mirror sources from a mirror session.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and click its name with a hyperlink.  
The **Basic Information** page is displayed.



6. Click the **Mirror Sources** tab. In the mirror source list, locate the row that contains the mirror source and click **Disassociate** in the **Operation** column.  
A confirmation dialog box is displayed.
7. Click **Yes**.  
After the disassociation is successful, the mirror source list is displayed.

## 12.3.5 Changing a Mirror Filter for a Mirror Session

### Scenarios

A mirror session can be associated with only one mirror filter. If the current mirror filter cannot meet your requirements, you can change one.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and choose **More > Change Mirror Filter** in the **Operation** column.  
The **Change Mirror Filter** dialog box is displayed.
6. In the mirror filter list, select a mirror filter and click **OK**.  
After the change, you can see that the new mirror filter in the **Mirror Filter** column of the mirror session list.  
If there is no mirror filter you want, create one by referring to [Creating a Mirror Filter](#).

## 12.3.6 Changing the Mirror Target of a Mirror Session



### Scenarios

A mirror session can be associated with only one mirror target. You can change the mirror target of a mirror session.

- A mirror target is a network interface of an ECS or a load balancer, which is used to receive mirrored traffic.
- Each mirror session can be associated with only one mirror target.

### Procedure



1. Log in to the management console.

2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and choose **More > Change Mirror Target** in the **Operation** column.  
The **Change Mirror Target** dialog box is displayed.
6. In the mirror target list, select a mirror target, and click **OK**.  
After the change, you can see that the new mirror target in the **Mirror Target** column of the mirror session list.

### 12.3.7 Modifying Basic Information About a Mirror Session

You can modify basic information about a mirror session, including its name, priority, and description.

#### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, locate the row that contains the mirror session and click **Modify** in the **Operation** column.  
The **Modify Mirror Session** dialog box is displayed.
6. Modify the parameters as prompted.

**Table 12-19** Parameters for configuring basic information about a mirror session

Parameter	Description	Example Value
Name	Mandatory Enter the name of the mirror session. The name: <ul style="list-style-type: none"><li>• Must contain 1 to 64 characters.</li><li>• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).</li></ul>	mirror-session-01
Priority	Mandatory Priority of the mirror session. <ul style="list-style-type: none"><li>• A priority value can be from 1 to 32766. A smaller value indicates a higher priority.</li><li>• Priorities must be unique for each mirror session of the same account in a region.</li></ul> A mirror source can be associated with multiple mirror sessions at the same time. The mirror sessions are matched from the lowest value to the highest value. For details, see <a href="#">the matching process of mirror sessions</a> .	1
VNI	Optional VXLAN Network Identifiers (VNIs) are used to distinguish different mirror sessions for a mirror target. A VNI can be from 0 to 16777215. If not specified, the default value is 1.	1
Packet Length	Optional The number of bytes that meet the mirror filter and will be mirrored. The value can be from 1 to 1460. If not specified, the default value is 96.	96
Mirror Session	Optional <ul style="list-style-type: none"><li>• If the mirror session is disabled, traffic of mirror sources cannot be monitored.</li><li>• If the mirror session is enabled, traffic of mirror sources can be monitored.</li></ul>	Enable
Description	Optional Enter the description of the mirror session in the text box as required.	-

7. Click **OK** to save the modification.



## 12.3.8 Viewing Details About a Mirror Session

### Scenarios

You can view the following information about a mirror session:

- Basic information, such as name, priority, and description
- Mirror filter
- Mirror sources, such as the private IP addresses, attached instances, and security groups of elastic network interfaces
- Mirror target, such as an ECS network interface or a load balancer

### Procedure



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.  
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.  
The **Mirror Sessions** page is displayed.
5. In the mirror session list, you can view the mirror session name, mirror filter, and mirror target.
6. In the mirror session list, locate the row that contains the mirror session and click its name with a hyperlink.  
The **Basic Information** page is displayed.
7. View information about the mirror session on different tab pages.
  - On the **Basic Information** tab page, view mirror session information, such as name, priority, and description.
  - On the **Mirror Sources** tab page, view the private IP addresses, attached instances, and security groups of elastic network interfaces

## 12.3.9 Deleting a Mirror Session

### Scenarios

If a mirror session is no longer required, you can delete it.

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.

The **Virtual Private Cloud** page is displayed.

4. In the navigation pane on the left, choose **Traffic Mirroring > Mirror Sessions**.

The **Mirror Sessions** page is displayed.

5. In the mirror session list, locate the row that contains the mirror session and choose **More > Delete** in the **Operation** column.

A confirmation dialog box is displayed.

6. Confirm the information and click **Yes**.

A deleted mirror session cannot be recovered.

# 13 Elastic IP


## 13.1 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 13-1** 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?](#)

If you want to increase your quota, see [How Do I Apply for a Higher Quota?](#)

- Your request for a larger quota will only be approved if your account has valid orders and you are continuously using cloud resources. If you have released resources immediately after subscribing to them multiple times, your request for quota increase will be declined.
- If you have increased the EIP quota but you have not used the quota for a long time, Huawei Cloud will reduce the quota to the default value.

## EIP 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.
- Flexible billing  
EIPs are available on a pay-per-use (bandwidth usage or amount of traffic is billed) basis. The yearly/monthly billing mode is more preferential.
- Shared bandwidth  
EIPs can use shared bandwidth to lower bandwidth costs.
- Immediate use  
EIP binding, unbinding, and bandwidth adjustments take effect immediately.

## Notes and Constraints

- An EIP and its bound cloud resource can use different billing modes. For example, a yearly/monthly EIP can be bound to a pay-per-use ECS.
- If the used EIP bandwidth exceeds the purchased size or is attacked (usually by a DDoS attack), the EIP will be blocked but can still be bound or unbound.
- EIPs cannot be transferred across accounts. That is, an EIP of account A cannot be transferred to account B.

# 13.2 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.

## Notes and Constraints

- Each EIP can only be bound to one cloud resource and they must be in the same region.
- If an EIP is frozen due to account arrears or security reasons, it cannot be bound or unbound.



## Assigning an EIP

1. Go to the [Buy EIP](#) page.
2. Set the parameters as prompted.

**Figure 13-2** Assigning an EIP

The screenshot shows the 'Buy EIP' configuration interface. At the top right, there are links for 'Assured Purchase', 'Flexible Billing', and 'Documentation'. The main configuration area includes:

- Billing Mode:** Radio buttons for 'Yearly/Monthly' (selected) and 'Pay-per-use'.
- Region:** A dropdown menu with a search icon. Below it, a note states: 'Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.'
- EIP Type:** Radio buttons for 'Dynamic BGP' (selected) and 'Static BGP'. Below it, a note states: 'Greater than or equal to 99.95% service availability rate'.
- Bandwidth (Mbit/s):** A series of buttons for 1, 2, 5, 10, 100, 200, and 'Custom'. The '1' button is selected. A note next to the custom input says: 'The value ranges from 1 to 2,000 Mbit/s.'
- IPv6 EIP:** A checkbox for 'Enable IPv6 Internet access' which is unchecked. Below it, a note states: 'IPv6 EIP is free during the Open Beta Test. After IPv6 EIP is enabled, you need to configure security group rules to allow traffic to and from 198.19.0.0/16.'
- DDoS Protection:** A radio button for 'Cloud Native Anti-DDoS Basic' (selected). Below it, a note states: 'Provides up to 5 Gbit/s of DDoS mitigation capacity for free. If the attack to an EIP exceeds 5 Gbit/s, the EIP will be blocked.'
- EIP Name:** An empty text input field.
- Enterprise Project:** A dropdown menu with '-Select-' and a 'Create Enterprise Project' link.
- Advanced Settings:** A dropdown menu with 'Bandwidth Name' and 'Tag' options.

At the bottom, there is a price summary: 'EIP Reservation Price: Free + Bandwidth Price: [input field]' and a 'Next' button.

**Table 13-1** Parameter descriptions

Parameter	Description	Example Value
Billing Mode	The following billing modes are available: <ul style="list-style-type: none"> <li>• Yearly/Monthly</li> <li>• Pay-per-use</li> </ul>	Pay-per-use
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. <p><b>NOTE</b> The geographical location of an EIP purchased in CN North-Ulanqab1 is Beijing.</p>	CN-Hong Kong

Parameter	Description	Example Value
EIP Type	<ul style="list-style-type: none"><li>• <b>Dynamic BGP:</b> Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.</li><li>• <b>Static BGP:</b> Static BGP offers more routing control and protects against route flapping, but an optimal path cannot be selected in real time when a network connection fails.</li><li>• <b>Premium BGP:</b> Premium BGP chooses the optimal path and ensures low-latency and high-quality networks. BGP is used to interconnect with lines of multiple mainstream carriers. Public network connections that feature low latency and high quality are directly established between Chinese mainland and Hong Kong (China). (This parameter is available only in <b>CN-Hong Kong</b>.)</li></ul> <p>For details, see <a href="#">What Are the Differences Between Static BGP and Dynamic BGP?</a></p>	Dynamic BGP

Parameter	Description	Example Value
Billed By	<p>This parameter is available only when you set <b>Billing Mode</b> to <b>Pay-per-use</b>.</p> <ul style="list-style-type: none"><li>• <b>Bandwidth:</b> 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.</li><li>• <b>Traffic:</b> You specify a maximum bandwidth and pay for the total traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic.</li><li>• <b>Shared Bandwidth:</b> The bandwidth can be shared by multiple EIPs. This is suitable for scenarios with staggered traffic.</li></ul>	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100
DDoS Protection	Cloud Native Anti-DDoS Basic Cloud Native Anti-DDoS Basic provides up to 5 Gbit/s of DDoS mitigation capacity. If the attack to an EIP exceeds 5 Gbit/s, the EIP will be blocked.	-
EIP Name	The EIP name.	eip-test
Enterprise Project	<p>The enterprise project that the EIP belongs to.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is <b>default</b>.</p> <p>For details about creating and managing enterprise projects, see the <a href="#">Enterprise Management User Guide</a>.</p>	default
Advanced Settings	Click the drop-down arrow to configure parameters, including the bandwidth name and tag.	-
Bandwidth Name	The name of the bandwidth.	bandwidth

Parameter	Description	Example Value
Tag	<p>The EIP tags. Each tag contains a key and value pair.</p> <p>The tag key and value must meet the requirements listed in <a href="#">Table 13-2</a>.</p> <p><b>NOTE</b> If your organization has created a tag policy for EIP, you need to add tags for EIP based on the tag policy. If a tag does not comply with the tagging rules, the creation may fail. Contact the organization administrator to learn details about the tag policy.</p>	<ul style="list-style-type: none"> <li>• Key: lpv4_key1</li> <li>• Value: 3005eip</li> </ul>
Monitoring	<p>Used to monitor the EIP and enabled by default.</p> <p>You can use the management console or APIs provided by Cloud Eye to query the metrics and alarms generated for the EIP and bandwidth.</p>	-
Required Duration	<p>The duration for which the purchased EIP will use. The duration must be specified if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b>.</p>	1 month
Auto-renew	<p>Whether to select <b>Auto-renew</b>. You can select it if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b>. The auto-renewal period is determined by the required duration.</p> <ul style="list-style-type: none"> <li>• Monthly subscription: The subscription is renewed every month.</li> <li>• Yearly subscription: The subscription is renewed each year.</li> </ul>	-
Quantity	<p>The number of EIPs you want to purchase.</p> <p>The quantity must be specified if the <b>Billing Mode</b> is set to <b>Pay-per-use</b>.</p>	1

**Table 13-2** EIP tag requirements

Parameter	Requirement	Example Value
Key	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Must be unique for each EIP.</li><li>• Can contain a maximum of 36 characters.</li><li>• Can contain letters, digits, underscores (_), and hyphens (-).</li></ul>	Ipv4_key1
Value	<ul style="list-style-type: none"><li>• Can contain a maximum of 43 characters.</li><li>• Can contain letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	eip-01

 **NOTE**

- If you are buying an EIP billed on a pay-per-use basis and you want to use a shared bandwidth, you can only select an existing shared bandwidth from the **Bandwidth Name** drop-down list. If there are no shared bandwidths to select, purchase a shared bandwidth first.
  - A dedicated bandwidth cannot be changed to a shared bandwidth and vice versa. However, you can purchase a shared bandwidth for pay-per-use EIPs.
    - After an EIP is added to a shared bandwidth, the EIP will use the shared bandwidth.
    - After an EIP is removed from the shared bandwidth, the EIP will use the dedicated bandwidth.
3. Click **Next**.
  4. Click **Submit**.

## Binding an EIP

1. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
2. Select the instance that you want to bind the EIP to.  
Instances that can be bound with EIPs include ECSs, BMSs, virtual IP addresses, and supplementary network interfaces.
3. Click **OK**.

 **NOTE**

- An EIP and its bound cloud resource can use different billing modes. For example, a yearly/monthly EIP can be bound to a pay-per-use ECS.
- To bind an EIP to an ECS:
  - The ECS must be in the running or stopped status.
  - The ECS must be in the same region as that of the EIP.
  - The ECS has no EIP bound to it.

## Helpful Links

- [How Do I Assign or Retrieve a Specific EIP?](#)
- [How Do I Access an ECS with an EIP Bound from the Internet?](#)
- [Can I Bind an EIP of an ECS to Another ECS?](#)
- [How Do I Know If My EIP Bandwidth Limit Has Been Exceeded?](#)
- [Why Can't My ECS Access the Internet Even After an EIP Is Bound?](#)

## 13.3 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.

The price of a pay-per-use EIP includes **EIP reservation price** and **bandwidth price**.



- If a pay-per-use EIP is unbound from an instance and is not released, you need to pay for the **EIP reservation price** and **bandwidth price**.
- If a pay-per-use EIP is bound to an instance, you only need to pay for the **bandwidth price**.

### Notes and Constraints



- 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.
- You cannot buy an EIP that has been released if it is currently in use by another user.
- If an EIP is frozen due to account arrears or security reasons, it cannot be bound or unbound.

### Procedure

#### Unbinding a single EIP



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Elastic IP**.
4. On the displayed page, locate the row that contains the EIP, and click **Unbind**.
5. Click **Yes** in the displayed dialog box.

#### Releasing a single EIP



1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Elastic IP**.

4. On the displayed page, locate the row that contains the target EIP, click **More** and then **Release** in the **Operation** column.
5. Click **Yes** in the displayed dialog box.

#### Unbinding multiple EIPs at once

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. On the displayed page, select the EIPs to be unbound.
5. Click the **Unbind** button located above the EIP list.
6. 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. On the displayed page, select the EIPs to be released.
5. Click the **Release** button located above the EIP list.
6. Click **Yes** in the displayed dialog box.

## 13.4 Modifying an EIP Bandwidth

### Scenarios

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

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

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 13-3](#).

#### NOTE



Decreasing bandwidths may cause packet loss.

**Table 13-3** Impact on billing after bandwidth size change

Billing Mode	Billed By	Change	Impact
Yearly/ Monthly	Bandwidth	Increase bandwidth	The change will take effect immediately. The increased bandwidth will be billed accordingly.

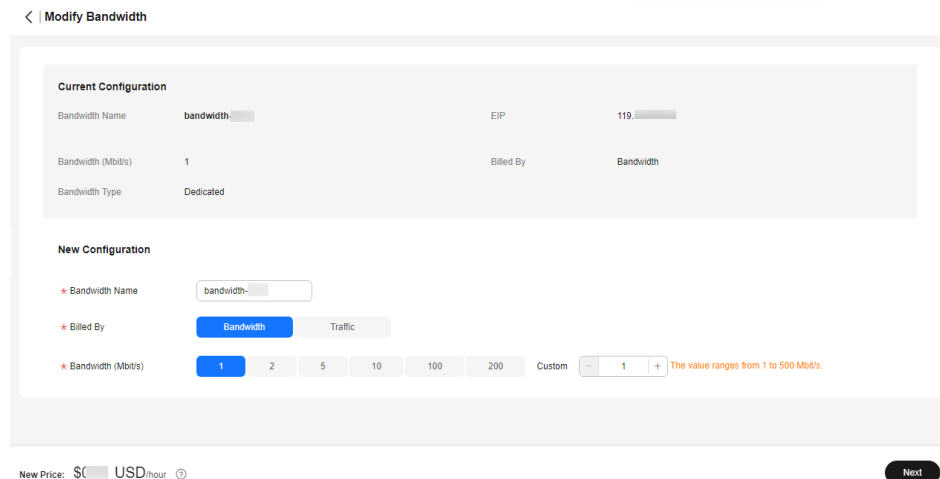
Billing Mode	Billed By	Change	Impact
	Bandwidth	Decrease bandwidth upon renewal	<p>The change will not take effect immediately.</p> <p>You need to select a new bandwidth size and a renewal duration. The change will take effect in the first billing cycle after a successful renewal.</p> <ul style="list-style-type: none"> <li>• The order can be unsubscribed before the bandwidth takes effect.</li> <li>• The bandwidth cannot be modified in the first billing cycle.</li> </ul>
Pay-per-use	Bandwidth	Increase or decrease the bandwidth	The change will take effect immediately.
	Traffic	Increase or decrease the bandwidth	<p>The change will take effect immediately.</p> <p>The bandwidth size you set is only used to limit the maximum data transfer rate.</p>

## Procedure

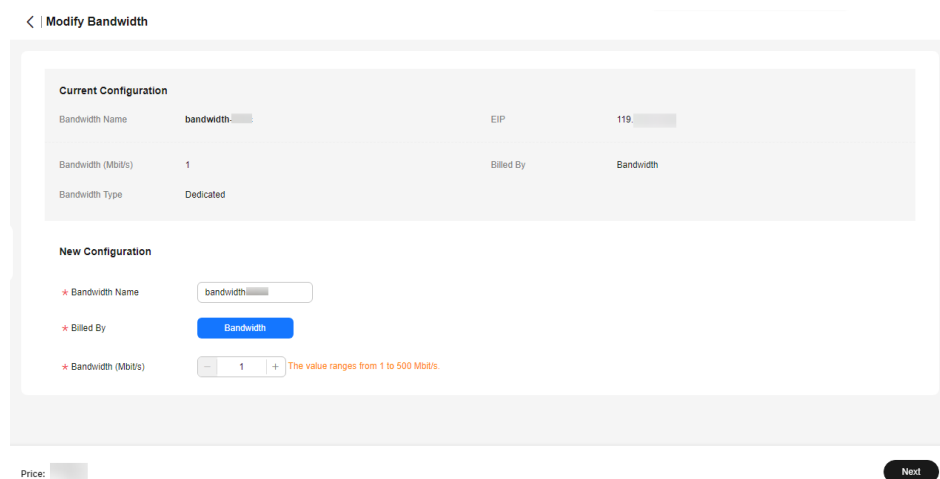
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. Locate the target EIP, click **More** in the **Operation** column, and select **Modify Bandwidth**.
  - If it is a pay-per-use EIP, the **Modify Bandwidth** page is displayed.
  - If it is a yearly/monthly EIP, select either of the following method to increase or decrease the bandwidth and click **Continue**.
    - Increase bandwidth
    - Decrease bandwidth
5. Modify the bandwidth parameters as prompted.



**Figure 13-3** Modifying the bandwidth of a pay-per-use EIP



**Figure 13-4** Modifying the bandwidth of a yearly/monthly EIP



6. Click **Next**.
7. 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 pay-per-use basis can be modified in batches.

## Helpful Links



- [How Do I Change the EIP Billing Option from Bandwidth to Traffic or from Traffic to Bandwidth?](#)
- [Can I Increase My Bandwidth Billed on Yearly/Monthly Basis and Then Decrease It?](#)

## 13.5 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. On the EIP list page, select one or more EIPs and click **Export** in the upper left corner.

The system will automatically export all EIPs to an Excel file and download the file to a local directory.

## 13.6 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 20 tags can be added to each EIP.

If your organization has created a tag policy for EIP, you need to add tags for EIP based on the tag policy. If a tag does not comply with the tagging rules, the EIP may fail to be created or the tag may fail to be added. Contact the organization administrator to learn more about the tag policy.

#### NOTE

The Organizations service is in open beta test (OBT). To use organization rules, apply for OBT.



A tag consists of a key and value pair. [Table 13-4](#) lists the tag key and value requirements.

**Table 13-4** EIP tag requirements



Parameter	Requirement	Example Value
Key	<ul style="list-style-type: none"><li>• Cannot be left blank.</li><li>• Must be unique for each EIP.</li><li>• Can contain a maximum of 36 characters.</li><li>• Can contain letters, digits, underscores (_), and hyphens (-).</li></ul>	Ipv4_key1
Value	<ul style="list-style-type: none"><li>• Can contain a maximum of 43 characters.</li><li>• Can contain letters, digits, underscores (_), periods (.), and hyphens (-).</li></ul>	eip-01

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

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

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Elastic IP**.
4. On the displayed page, locate the EIP whose tags you want to manage and click the EIP name.
5. On the page showing EIP details, click the **Tags** tab and perform desired operations on tags.
  - View tags.

On the **Tags** tab, you can view details about tags added to the current EIP, including the number of tags and the key and value of each tag.
  - Add a tag.

Click **Add Tag** in the upper left corner. In the displayed **Add Tag** dialog box, enter the tag key and value, and click **OK**.
  - Edit a tag.

Locate the row that contains the tag you want to edit, and click **Edit** in the **Operation** column. Enter the new tag value, and click **OK**.  
The tag key cannot be modified.

- Delete a tag.  
Locate the row that contains the tag you want to delete, and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

## 13.7 IPv6 EIP

### Overview

Both IPv4 and IPv6 EIPs are available. You can assign an IPv6 EIP or map an existing IPv4 EIP to an IPv6 EIP.

After the IPv6 EIP function is enabled, you will obtain both an IPv4 EIP and its corresponding IPv6 EIP. External IPv6 addresses can access cloud resources through this IPv6 EIP.

IPv4 EIPs are billed. IPv6 EIPs are currently free, but will be billed at a later date (price yet to be determined).

### Application Scenarios of IPv4/IPv6 Dual Stack

If your ECS supports IPv6, you can use the IPv4/IPv6 dual stack. [Table 13-5](#) shows the example application scenarios.

**Table 13-5** Application scenarios of IPv4/IPv6 dual stack

Application Scenario	Description	Requirement	IPv4 or IPv6 Subnet	ECS
Private IPv4 communication	Your applications on ECSs need to communicate with other systems (such as databases) through private networks using IPv4 addresses.	<ul style="list-style-type: none"><li>No EIPs have been bound to the ECSs.</li></ul>	IPv4 CIDR Block	<b>Private IPv4 address:</b> used for private IPv4 communication.

Application Scenario	Description	Requirement	IPv4 or IPv6 Subnet	ECS
Public IPv4 communication	Your applications on ECSs need to communicate with other systems (such as databases) through public IPv4 addresses.	<ul style="list-style-type: none"> <li>EIPs have been bound to the ECSs.</li> </ul>	IPv4 CIDR Block	<ul style="list-style-type: none"> <li><b>Private IPv4 address:</b> used for private IPv4 communication.</li> <li><b>Public IPv4 address:</b> used for public IPv4 communication.</li> </ul>

Application Scenario	Description	Requirement	IPv4 or IPv6 Subnet	ECS
Private IPv6 communication	Your applications on ECSs need to communicate with other systems (such as databases) through private IPv6 addresses.	<ul style="list-style-type: none"> <li>• IPv6 has been enabled for the VPC subnet.</li> <li>• The network has been configured for the ECSs as follows:               <ul style="list-style-type: none"> <li>– <b>Flavor:</b> Any ECS flavor that supports the IPv6 network. For details, see section "x86 ECS Specifications and Types" in the <a href="#">Elastic Cloud Server User Guide</a>.</li> <li>– <b>VPC and Subnet:</b> IPv6-enabled subnet and VPC.</li> <li>– <b>Self-assigned IPv6 address:</b> Selected.</li> <li>– <b>Shared Bandwidth:</b> Selected <b>Do not configure</b>.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>• IPv4 CIDR Block</li> <li>• IPv6 CIDR block</li> </ul>	<ul style="list-style-type: none"> <li>• <b>Private IPv4 address + IPv4 EIP:</b> Bind an IPv4 EIP to the instance to allow public IPv4 communication.</li> <li>• <b>Private IPv4 address:</b> Do not bind any IPv4 EIP to the instance and use only the private IPv4 address to allow private IPv4 communication.</li> <li>• <b>IPv6 address:</b> Do not configure shared bandwidth for the IPv6 address to allow private IPv6 communication.</li> </ul>

Application Scenario	Description	Requirement	IPv4 or IPv6 Subnet	ECS
Public IPv6 communication	An IPv6 network is required for the ECS to access the IPv6 service on the Internet.	<ul style="list-style-type: none"> <li>IPv6 has been enabled for the VPC subnet.</li> <li>The network has been configured for the ECSs as follows:                             <ul style="list-style-type: none"> <li><b>Flavor:</b> Any ECS flavor that supports the IPv6 network. For details about the ECS flavor that support the IPv6 network, see section "x86 ECS Specifications and Types" in the <a href="#">Elastic Cloud Server User Guide</a>.</li> <li><b>VPC and Subnet:</b> IPv6-enabled subnet and VPC.</li> <li><b>Self-assigned IPv6 address:</b> Selected.</li> <li><b>Shared Bandwidth:</b> Selected a shared bandwidth.</li> </ul> </li> </ul> <p><b>NOTE</b> For details, see <a href="#">Setting Up an IPv6 Network</a>.</p>	<ul style="list-style-type: none"> <li>IPv4 CIDR Block</li> <li>IPv6 CIDR block</li> </ul>	<ul style="list-style-type: none"> <li><b>Private IPv4 address + IPv4 EIP:</b> Bind an IPv4 EIP to the instance to allow public IPv4 communication.</li> <li><b>Private IPv4 address:</b> Do not bind any IPv4 EIP to the instance and use only the private IPv4 address to allow private IPv4 communication.</li> <li><b>IPv6 address + shared bandwidth:</b> Allow both private IPv6 communication and public IPv6 communication.</li> </ul>

For details, see [IPv4 and IPv6 Dual-Stack Network](#).

## Application Scenarios of IPv6 EIP

If you want an ECS to provide IPv6 services but the ECS does not support IPv6 networks or you do not want to build an IPv6 network, you can use IPv6 EIP to

quickly address your requirements. For details about application scenarios and resource planning, see [Table 13-6](#).

**Table 13-6** Application scenarios and resource planning of an IPv6 EIP network (with IPv6 EIP enabled)

Application Scenario	Description	Requirement	IPv4 or IPv6 Subnet	ECS
Public IPv6 communication	You want to allow an ECS to provide IPv6 services for clients on the Internet without setting up an IPv6 network.	<ul style="list-style-type: none"><li>An EIP has been bound to the ECS.</li><li>IPv6 EIP has been enabled.</li></ul>	IPv4 CIDR Block	<ul style="list-style-type: none"><li><b>Private IPv4 address:</b> used for private IPv4 communication.</li><li><b>IPv4 EIP (with IPv6 EIP enabled):</b> used for public network communication through IPv4 and IPv6 addresses.</li></ul>



## Application Scenarios and Resource Planning of IPv6 Networks

Figure 13-5 Application scenarios and resource planning of IPv6 networks



### Enabling IPv6 (Assigning IPv6 EIPs)

- Method 1:

Select the **IPv6 EIP** option when you assign an EIP by referring to [Assigning an EIP and Binding It to an ECS](#) so that you can obtain both an IPv4 and an IPv6 EIP.

External IPv6 addresses can access cloud resources through this IPv6 EIP.

- Method 2:

If you want an IPv6 EIP in addition to an existing IPv4 EIP, locate the row that contains the target IPv4 EIP, click **More** in the **Operation** column, and select **Enable IPv6 EIP**. Then, a corresponding IPv6 EIP will be assigned.

After the IPv6 EIP is enabled, you will obtain both an IPv4 EIP and an IPv6 EIP. External IPv6 addresses can access cloud resources through this IPv6 EIP.

#### NOTE

There is no adverse impact on the cloud resources bound with existing IPv4 EIPs.

## Configuring Security Groups

After IPv6 EIP is enabled, add inbound and outbound security group rules to allow packets to and from the IP address range **198.19.0.0/16**. **Table 13-7** shows the security group rules. IPv6 EIP uses NAT64 to convert the source IP address in the inbound direction to an IPv4 address in the IP address range 198.19.0.0/16. The source port can be a random one, the destination IP address is the private IPv4 address of your local server, and the destination port remains unchanged.

For details, see [Virtual Private Cloud User Guide](#).

**Table 13-7** Security group rules

Direction	Protocol	Source or Destination
Inbound	All	Source: 198.19.0.0/16
Outbound	All	Destination: 198.19.0.0/16

## Disabling IPv6 EIP

If you do not need the IPv6 EIP, locate the row that contains its corresponding IPv4 EIP, click **More** in the **Operation** column, and select **Disable IPv6 EIP**. Then, the IPv6 EIP will be released. You will only have the IPv4 EIP.

# 14 Shared Bandwidth

---

## 14.1 Shared Bandwidth Overview

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

### 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, which significantly increases bandwidth costs. If all EIPs share the same bandwidth, you can lower bandwidth costs and easily perform system O&M.

- Lowered Bandwidth Costs  
Region-level bandwidth sharing and multiplexing reduce bandwidth usage and O&M costs.
- Flexible Operations  
You can add pay-per-use EIPs (except for **5\_gray** EIPs of dedicated load balancers) to or remove them from a shared bandwidth regardless of the type of instances that they are bound to.
- Flexible Billing Modes  
The yearly/monthly and pay-per-use billing modes are provided.

You can use a shared bandwidth in either of the following ways:

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

- [Assigning an EIP and Binding It to an ECS](#)

## Shared Bandwidth Quotas

- Each account can have a maximum of 5 shared bandwidths. If you need more shared bandwidths, submit a service ticket to request a quota increase.
- If you want to increase a pay-per-use shared bandwidth that is greater than 1 Gbit/s, the minimum increase is 500 Mbit/s.

## Notes and Constraints

- The minimum size of a shared bandwidth that can be purchased is 5 Mbit/s. You can only add pay-per-use EIPs to a shared bandwidth.
- If a yearly/monthly shared bandwidth is deleted upon expiration, EIPs sharing the bandwidth will be removed from the bandwidth and be billed based on the mode before they are added to the shared bandwidth.
- A shared bandwidth can only be used by resources from its same account.

### NOTE

- A dedicated bandwidth cannot be changed to a shared bandwidth and vice versa. However, you can purchase a shared bandwidth for pay-per-use EIPs.
  - Add an EIP to a shared bandwidth and then the EIP will use the shared bandwidth.
  - Remove the EIP from the shared bandwidth and then the EIP will use the dedicated bandwidth.
- If you want to submit a service ticket, refer to [Submitting a Service Ticket](#).

# 14.2 Assigning a Shared Bandwidth

## Scenarios

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

Assign a shared bandwidth for use with EIPs.

## Procedure

1. Go to the [Buy Shared Bandwidth](#) page.
2. Set the parameters as prompted.

**Table 14-1** Parameter descriptions

Parameter	Description	Example Value
Billing Mode	<p>A shared bandwidth can be billed on a yearly/monthly or pay-per-use basis.</p> <ul style="list-style-type: none"><li>● <b>Yearly/Monthly:</b> You pay for the bandwidth by year or month before using it. No other charges apply during the validity period of the bandwidth.</li><li>● <b>Pay-per-use:</b> You pay for the bandwidth based on the amount of time you use the bandwidth.</li></ul>	Yearly/Monthly
Region	<p>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.</p>	CN-Hong Kong
Bandwidth Type	<p>Select a type of the shared bandwidth based on your EIP type.</p> <ul style="list-style-type: none"><li>● <b>Standard:</b> Dynamic BGP and premium BGP EIPs can be added to a shared bandwidth of this type.</li><li>● <b>Premium BGP:</b> Premium BGP EIPs can be added to a shared bandwidth of this type.</li></ul> <p><b>NOTE</b> In the CN-Hong Kong region, only dynamic BGP EIPs can be added to standard shared bandwidths.</p>	Standard
Billed By	<p>The billing method for the shared bandwidth.</p> <p>You can specify a shared bandwidth to be billed by bandwidth.</p>	Bandwidth
Bandwidth	<p>The bandwidth size in Mbit/s. The minimum value is 5 Mbit/s.</p>	10
Enterprise Project	<p>The enterprise project that the EIP belongs to.</p> <p>An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is <b>default</b>.</p>	default
Name	<p>The name of the shared bandwidth.</p>	Bandwidth-001

Parameter	Description	Example Value
Required Duration	The duration for which the purchased EIP will use. The duration must be specified if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> .	2 months
Auto-renew	Whether to select <b>Auto-renew</b> . You can select it if the <b>Billing Mode</b> is set to <b>Yearly/Monthly</b> . The auto-renewal period is determined by the required duration. <ul style="list-style-type: none"><li>• Monthly subscription: The subscription is renewed every month.</li><li>• Yearly subscription: The subscription is renewed each year.</li></ul>	-

3. Click **Next**.
4. Confirm the configurations.
  - If you set **Billing Mode** to **Pay-per-Use**, click **Submit**.
  - If you set **Billing Mode** to **Yearly/Monthly**, click **Pay Now**.On the payment page, confirm the order information and click **Confirm**.

## 14.3 Adding EIPs to a Shared Bandwidth



### Scenarios

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

### Notes and Constraints

- To add a yearly/monthly EIP to a shared bandwidth, you need to first change its billing mode to pay-per-use.
- 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 select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.
5. In the shared bandwidth list, locate the target shared bandwidth that you want to add EIPs to. In the **Operation** column, choose **Add Public IP Address**, and select the EIPs to be added.

 **NOTE**

- After an EIP is added to a shared bandwidth, the dedicated bandwidth used by the EIP will become invalid and the EIP will start to use the shared bandwidth. The EIP's dedicated bandwidth will be deleted and will no longer be billed.
  - An EIP cannot be configured for two shared bandwidths at the same time, so if you attempt to add an EIP to a second shared bandwidth, it will be automatically removed from the original shared bandwidth.
6. Click **OK**.

## Helpful Links

[What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth? Can a Dedicated Bandwidth Be Changed to a Shared Bandwidth or the Other Way Around?](#)

# 14.4 Removing EIPs from a Shared Bandwidth



## Scenarios

Remove EIPs that are no longer required from a shared bandwidth if needed.

## Notes and Constraints

A yearly/monthly EIP cannot be removed from a shared bandwidth purchased during OBT.

## Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the target shared bandwidth from which EIPs are to be removed, choose **More > Remove Public IP Address** in the **Operation** column, and select the EIPs to be removed in the displayed dialog box.
6. Set the EIP bandwidth after the EIP is removed. You can configure the EIP billing mode and bandwidth size.
7. Click **OK**.

# 14.5 Modifying a Shared Bandwidth



## Scenarios

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



- If a shared bandwidth is billed on a pay-per-use basis, the modification will take effect immediately. For details, see [Modifying a Shared Bandwidth \(Pay-per-Use\)](#).
- If a shared bandwidth is billed on a yearly/monthly basis:
  - **You can increase the bandwidth.** The increased bandwidth size will take effect immediately and the price difference will be billed accordingly.
  - **You can decrease the bandwidth.** The decreased bandwidth size will take effect in the first billing cycle after a successful renewal.

If you want to change the billing mode of a shared bandwidth, see [How Do I Change My EIP Billing Mode from Pay-per-Use to Yearly/Monthly?](#)

## Modifying a Shared Bandwidth (Pay-per-Use)



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

## Increasing a Shared Bandwidth (Yearly/Monthly)

1. Log in to the management console.
  2. Click  in the upper left corner and select the desired region and project.
  3. Click  in the upper left corner and choose **Networking > Elastic IP**.
  4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
  5. In the shared bandwidth list, locate the row that contains the target shared bandwidth, and click **Modify Bandwidth** in the **Operation** column.
  6. Select **Increase bandwidth** and click **Continue**.
  7. In the **New Configuration** area on the **Modify Bandwidth** page, change the bandwidth name and size.
  8. Click **Next**.
  9. Confirm the information and click **Pay Now**.
- After you complete the payment, the increased bandwidth will take effect immediately.



## Decreasing a Shared Bandwidth (Yearly/Monthly)

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the row that contains the target shared bandwidth, and click **Modify Bandwidth** in the **Operation** column.
6. Select **Decrease bandwidth** and click **Continue**.
7. In the **New Configuration** area on the **Modify Bandwidth** page, change the bandwidth name and size.
8. Click **Next**.
9. Confirm the information and click **Pay Now**.

After you complete the payment, the decreased bandwidth will take effect in the first billing cycle after the current subscription ends.

## 14.6 Deleting a Shared Bandwidth

### Scenarios

Delete a shared bandwidth billed on a pay-per-use basis if it is no longer required.



### Notes and Constraints

- A yearly/monthly shared bandwidth cannot be directly deleted. It can only be unsubscribed.
- If you want to delete a shared bandwidth with EIPs added, you have to remove the EIPs from the shared bandwidth first.

### 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  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking > Elastic IP**.
4. In the navigation pane on the left, choose **Elastic IP and Bandwidth > Shared Bandwidths**.
5. In the shared bandwidth list, locate the row that contains the pay-per-use shared bandwidth you want to delete, click **More** in the **Operation** column, and then click **Delete**.

6. In the displayed dialog box, click **OK**.

# 15 Shared Data Package

---

## 15.1 Shared Data Package Overview

Shared data package provides a quota for data usage. Such packages are cost-effective and easy to use. Shared data packages take effect immediately after your purchase. If you have subscribed to pay-per-use EIPs billed by traffic in a region and buy a shared data package in the same region, the EIPs will use the shared data package. After the package quota is used up or the package expires, the EIPs will continue to be billed on a pay-per-use basis. For billing details, see [Product Pricing Details](#).

- Two types of packages are available: dynamic BGP and static BGP. Dynamic BGP data packages will be used by pay-per-use EIPs (billed by traffic) of the dynamic BGP type, and static BGP data packages will be used by pay-per-use EIPs (billed by traffic) of the static BGP type.
- Shared data packages can be purchased yearly or monthly. Packages purchased for a year are more cost effective. If you have multiple shared data packages, the data package with the shortest validity period will be used first.
- If your usage exceeds your shared data package quota within its validity, you will be billed on a pay-per-use basis for the additional traffic usage.
- If a shared data package expires, make sure your account balance is sufficient and your EIP will be billed on a pay-per-use basis.

### Notes and Constraints

- Shared data packages require a one-off payment and take effect immediately after purchase. You cannot specify the effective date.
- Shared data packages cannot be unsubscribed from nor be modified once purchased and cannot be renewed upon expiration.
- Shared data packages are billed by month or year. Once expired, remaining package quota cannot be used any more.
- Shared data packages can only be used by pay-per-use dedicated bandwidth billed by traffic. Two types of shared data packages are available: static BGP (for static BGP bandwidth) and dynamic BGP (for dynamic BGP bandwidth).
- A shared data package cannot be used for bandwidth of a specific EIP.

- A shared data package cannot be used for a shared bandwidth.
- A shared data package cannot be used by EIPs of the premium BGP type.

## 15.2 Buying a Shared Data Package

### Scenarios

This section describes how to buy a shared data package. Shared data packages take effect immediately after your purchase. If you have subscribed to pay-per-use EIPs billed by traffic in a region and buy a shared data package in the same region, the EIPs will use the shared data package. After the package quota is used up or the package expires, the EIPs will continue to be billed on a pay-per-use basis.

### Notes and Constraints

- Shared data packages require a one-off payment and take effect immediately after purchase. You cannot specify the effective date.
- Shared data packages cannot be unsubscribed from nor be modified once purchased and cannot be renewed upon expiration.
- Shared data packages are billed by month or year. Once expired, remaining package quota cannot be used any more.
- Shared data packages can only be used by pay-per-use dedicated bandwidth billed by traffic. Two types of shared data packages are available: static BGP (for static BGP bandwidth) and dynamic BGP (for dynamic BGP bandwidth).
- A shared data package cannot be used for bandwidth of a specific EIP.
- A shared data package cannot be used for a shared bandwidth.
- A shared data package cannot be used by EIPs of the premium BGP type.
- If you have an order that has not been paid within the payment period, you need to cancel or pay for the order first. Then, you can purchase a shared data package.

### Procedure

1. Go to the [Buy Shared Data Package](#) page.
2. Set the parameters as prompted.

**Table 15-1** Parameter descriptions

Parameter	Description	Example Value
Region	A shared data package can only be used by resources in its same region. Select the region based on your requirements.	CN-Hong Kong

Parameter	Description	Example Value
Type	The shared data package type. Set this parameter based on the bandwidth type of the EIP. The following two types of packages are available: <ul style="list-style-type: none"><li>• Dynamic BGP: A dynamic BGP data package can only be used by dynamic BGP EIPs billed by traffic on a pay-per-use basis.</li><li>• Static BGP: A static BGP data package can only be used by static BGP EIPs billed by traffic on a pay-per-use basis.</li></ul>	Static BGP
Package Validity	The validity period of the shared data package. Select a validity period based on service requirements. A shared data package cannot be unsubscribed and takes effect immediately after you purchase it. Expired shared data packages will longer be available for use.	1 month
Specification	The size of the shared data package in GB.	10 GB
Usage Duration	The validity period of the shared data package.	Default

3. Click **Next**.
4. Confirm the configurations and click **Submit**.
5. On the payment page, confirm the order information and click **Confirm**.

# 16 Monitoring and Auditing

## 16.1 Monitoring

### 16.1.1 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 16-1 EIP and bandwidth metrics

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
upstream_bandwidth	Outbound Bandwidth	Network rate of outbound traffic (Previously called "Upstream Bandwidth") Unit: bit/s	$\geq 0$ bit/s	Bandwidth or EIP	1 minute

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
downstream_bandwidth	Inbound Bandwidth	Network rate of inbound traffic (Previously called "Downstream Bandwidth") Unit: bit/s	$\geq 0$ bit/s	Bandwidth or EIP	1 minute
upstream_bandwidth_usage	Outbound Bandwidth 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
upstream	Outbound Traffic	Network traffic going out of the cloud platform in a minute (Previously called "Upstream Traffic") Unit: byte	$\geq 0$ bytes	Bandwidth or EIP	1 minute
downstream	Inbound Traffic	Network traffic going into the cloud platform in a minute (Previously called "Downstream Traffic") Unit: byte	$\geq 0$ bytes	Bandwidth or EIP	1 minute

## Dimensions

Key	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=publicip\_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"  
  }  
],



## 16.1.2 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. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click  to open the service list and choose **Management & Deployment > Cloud Eye**.
4. Click **Cloud Service Monitoring** on the left of the page, and choose **Elastic IP and Bandwidth**.
5. Locate the target bandwidth or EIP and click **View Metric** in the **Operation** column to check the bandwidth or EIP monitoring information.


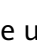


## 16.1.3 Creating an Alarm Rule

### 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. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click  to open the service list and choose **Management & Deployment > Cloud Eye**.
4. In the left navigation pane on the left, choose **Alarm Management > Alarm Rules**.
5. On the **Alarm Rules** page, click **Create Alarm Rule** and set required parameters, or modify an existing alarm rule.
6. 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 [Cloud Eye User Guide](#).

## 16.2 Interconnecting with CTS

### 16.2.1 Key Operations Recorded by CTS

With CTS, you can record operations performed on the VPC service for further query, audit, and backtracking purposes.

**Table 16-2** lists the VPC operations that can be recorded by CTS.

**Table 16-2** 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

Operation	Resource Type	Trace
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
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

Operation	Resource Type	Trace
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
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

Operation	Resource Type	Trace
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
Updating a flow log	flowlogs	updateFlowLog
Deleting a flow log	flowlogs	deleteFlowLog

## 16.2.2 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.



---

#### NOTICE

CTS only retains traces for seven days. To store traces for a longer time, configure your tracker to transfer traces to OBS buckets. For details, see [Configuring a Tracker](#).

---

### Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click  to go to the service list. Under **Management & Governance**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Trace List**.
5. Specify filters as needed. The following filters are available:
  - **Trace Type**: Set it to **Management** or **Data**.
  - **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.
6. Click arrow on the left of the required trace to expand its details.
  7. Locate the required trace and click **View Trace** in the **Operation** column.  
A dialog box is displayed, showing the trace content.

# A Change History

Release Date	What's New
2024-06-05	<p>This issue is the fifty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Modified the suggestions for selecting VPC CIDR blocks in <a href="#">VPC and Subnet Planning Suggestions</a>.</li><li>• Added cloud resources that can be created in a shared VPC subnet in <a href="#">VPC Sharing</a>.</li></ul>
2024-05-28	<p>This issue is the fiftieth official release, which incorporates the following changes:</p> <p>Added network ACL architecture, working principles, and examples, in <a href="#">Network ACL Overview</a>, <a href="#">Network ACL Configuration Examples</a>, and <a href="#">Adding a Network ACL Rule (Custom Priorities)</a>.</p>
2024-04-26	<p>This issue is the forty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Changed the screenshots based on the console style changes.</li><li>• Added descriptions indicating that associated resources of the VPC and subnet to be deleted are displayed on the console in <a href="#">Deleting a VPC</a> and <a href="#">Deleting a Subnet</a>.</li><li>• Added description about security group tags in <a href="#">Creating a Security Group</a> and <a href="#">Managing Security Group Tags</a>.</li><li>• Added the description about network ACL tags in <a href="#">Creating a Network ACL</a> and <a href="#">Managing Network ACL Tags</a>.</li></ul>

Release Date	What's New
2024-03-15	<p>This issue is the forty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added the descriptions about the VPC with subnets shared with other accounts from section <a href="#">VPC Sharing</a> to section <a href="#">Stopping Sharing a Subnet</a>.</li><li>• Added the descriptions about tag policies in <a href="#">Creating a VPC</a> and <a href="#">Creating a Subnet for the VPC</a>.</li><li>• Added constraints on importing security group rules in <a href="#">Importing and Exporting Security Group Rules</a>.</li></ul>
2023-12-05	<p>This issue is the forty-seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added description of associating or disassociating subnets on the subnet list page in <a href="#">Associating Subnets with a Network ACL</a> and <a href="#">Disassociating Subnets from a Network ACL</a>.</li><li>• Added description about IP addresses in <a href="#">Creating an IP Address Group</a> and <a href="#">Adding IP Addresses to an IP Address Group</a>.</li><li>• Added description about IP address group operations in <a href="#">Exporting IP Address Group Details</a>, <a href="#">Modifying IP Addresses in an IP Address Group</a>, and <a href="#">Importing IP Addresses to an IP Address Group in Batches</a>.</li><li>• Added supported secondary CIDR blocks in <a href="#">Adding a Secondary IPv4 CIDR Block to a VPC</a>.</li></ul>
2023-11-02	<p>This issue is the forty-sixth official release, which incorporates the following changes:</p> <p>Added the traffic mirroring function in <a href="#">Traffic Mirroring</a> to <a href="#">Deleting a Mirror Session</a>.</p>

Release Date	What's New
2023-10-09	<p>This issue is the forty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added description about associated resource deletion in <a href="#">Deleting an IP Address Group</a>.</li><li>• Added description about deleting network interfaces in <a href="#">Deleting a Network Interface</a> and <a href="#">Deleting a Supplementary Network Interface</a>.</li><li>• Added description about <b>Allow Common Ports</b> in <a href="#">Allowing Common Ports with A Few Clicks</a>.</li><li>• Modified description about adding routes for a VPC peering connection in <a href="#">Creating a VPC Peering Connection with Another VPC in Your Account</a>, <a href="#">Creating a VPC Peering Connection with a VPC in Another Account</a>, <a href="#">Modifying Routes Configured for a VPC Peering Connection</a>, <a href="#">Viewing Routes Configured for a VPC Peering Connection</a>, and <a href="#">Deleting Routes Configured for a VPC Peering Connection</a>.</li><li>• Added description about security group rules and rule configuration examples in <a href="#">Security Groups and Security Group Rules</a> and <a href="#">Security Group Examples</a>.</li></ul>
2023-08-07	<p>This issue is the forty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added content about elastic network interfaces in <a href="#">Elastic Network Interface</a>.</li><li>• Added content about supplementary network interfaces in <a href="#">Supplementary Network Interfaces</a>.</li></ul>
2023-07-07	<p>This issue is the forty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added description that multiple IP addresses can be configured for <b>Source</b> and <b>Destination</b> in <a href="#">Adding a Security Group Rule</a> and <a href="#">Fast-Adding Security Group Rules</a>.</li><li>• Added description about security group templates for creating security groups quickly in <a href="#">Creating a Security Group</a>.</li></ul>



Release Date	What's New
2023-06-08	<p>This issue is the forty-second official release, which incorporates the following change:</p> <ul style="list-style-type: none"><li>• Added description about security group rule examples in <a href="#">Security Groups and Security Group Rules</a>.</li><li>• Added description about security group rule templates in <a href="#">Creating a Security Group</a>.</li><li>• Modified description about source and destination in <a href="#">Adding a Security Group Rule</a>, <a href="#">Fast-Adding Security Group Rules</a>, and <a href="#">Importing and Exporting Security Group Rules</a>.</li><li>• Added support for enterprise projects in <a href="#">Creating a Network ACL</a>.</li><li>• Updated screenshots and parameters in <a href="#">Creating a VPC Peering Connection with Another VPC in Your Account</a> and <a href="#">Creating a VPC Peering Connection with a VPC in Another Account</a>.</li><li>• Added IP address group functions to <a href="#">IP Address Group</a> and <a href="#">Deleting an IP Address Group</a>.</li><li>• Added parameter <b>Domain Name</b> in <a href="#">Creating a VPC</a>, <a href="#">Creating a Subnet for the VPC</a>, and <a href="#">Modifying a Subnet</a>.</li></ul>
2023-02-25	<p>This issue is the forty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">VPC Peering Connection Usage Examples</a>.</li><li>• Added <a href="#">Obtaining the Peer Project ID of a VPC Peering Connection</a>.</li><li>• Added <a href="#">Modifying Routes Configured for a VPC Peering Connection</a>.</li></ul>
2022-12-23	<p>This issue is the fortieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Viewing and Deleting Resources in a Subnet</a>.</li><li>• Added <a href="#">Viewing IP Addresses in a Subnet</a>.</li><li>• Modified <a href="#">Deleting a VPC</a> and <a href="#">Deleting a Subnet</a>.</li><li>• Updated the document based on the navigation pane changes of subnets, route tables, VPC peering connections, and elastic network interfaces.</li></ul>

Release Date	What's New
2022-11-15	<p>This issue is the thirty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added the link to the price calculator in <a href="#">Shared Data Package Overview</a>.</li><li>• Added <a href="#">Unbinding a Virtual IP Address from an Instance</a> and <a href="#">Unbinding a Virtual IP Address from an EIP</a>.</li><li>• Added notes and constraints in <a href="#">Releasing a Virtual IP Address</a>.</li><li>• Added description that security groups are free of charge in <a href="#">Deleting a Security Group</a>.</li></ul>
2022-11-01	<p>This issue is the thirty-eighth official release, which incorporates the following change:</p> <ul style="list-style-type: none"><li>• Changed the number of instances that can be associated with a security group in <a href="#">Security Groups and Security Group Rules</a>.</li></ul>
2022-07-26	<p>This issue is the thirty-seventh official release, which incorporates the following changes:</p> <p>Added descriptions that EIPs cannot be used across regions in the following section:</p> <p><a href="#">EIP Overview</a></p>
2022-07-14	<p>This issue is the thirty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Optimized description about premium BGP in <a href="#">Assigning an EIP and Binding It to an ECS</a>.</li><li>• Put VPC flow logs into commercial use.</li></ul>
2022-06-15	<p>This issue is the thirty-fifth official release, which incorporates the following changes:</p> <p>Deleted the content about L2CGs. For the latest document about L2CGs, see <a href="#">Enterprise Switch User Guide</a>.</p>
2022-05-15	<p>This issue is the thirty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added description about different types of routes that can be added to default and custom route tables in <a href="#">Route Tables and Routes</a>.</li><li>• Added notes and constraints about whether different types of routes can be replicated in <a href="#">Replicating a Route</a>.</li></ul>
2021-12-30	<p>This issue is the thirty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added application scenarios of IPv6 networks in <a href="#">Setting Up an IPv6 Network</a>.</li><li>• Added <a href="#">Viewing a VPC Topology</a>.</li></ul>

Release Date	What's New
2021-10-30	This issue is the thirty-second official release, which incorporates the following change: Added section "Permissions Management".
2021-05-20	This issue is the thirty-first official release, which incorporates the following changes: Added FAQ "Why Am I Stilled Be Billed After All VPCs Are Deleted?"
2021-03-24	This issue is the thirtieth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Added "Network Service Overview" in section "Service Overview".</li><li>• Added FAQ "Why Cannot I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?"</li></ul>
2021-03-01	This issue is the twenty-ninth official release, which incorporates the following changes: Added sections "Adding a Secondary CIDR Block to a VPC" and "Removing a Secondary CIDR Block from a VPC".
2020-12-17	This issue is the twenty-eighth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Added restrictions in section "Notes and Constraints".</li><li>• Added a figure to illustrate the VPC peering connection.</li></ul>
2020-11-03	This issue is the twenty-seventh official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Adjusted sections under <b>VPC and Subnet</b>.</li><li>• Added "Denying Access from a Specific IP Address" to section "Network ACL".</li><li>• Deleted FAQ "Will the EIP Bound to an ECS Be Changed After the ECS Is Stopped and Then Started?"</li></ul>
2020-10-23	This issue is the twenty-sixth official release, which incorporates the following change: Optimized sections under "Security Group".
2020-09-18	This issue is the twenty-fifth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Modified the steps in section "Changing the Security Group of an ECS".</li><li>• Added the L2CG function.</li></ul>

Release Date	What's New
2020-09-07	<p>This issue is the twenty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "What Should I Do If My Security Group Rules Do Not Take Effect?"</li><li>• Deleted FAQ "How Do I Handle the VPC Peering Connection Failure?"</li><li>• Modified FAQ "Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?"</li><li>• Modified FAQ "Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?"</li></ul>
2019-07-23	<p>This issue is the twenty-third official release, which incorporates the following change:</p> <p>Added parameter <b>IP address group</b> to security group sections and added section "IP Address Group".</p>
2020-06-09	<p>This issue is the twenty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "Why Can't I Ping an ECS with Two NICs Configured?"</li><li>• Added IPv4/IPv6 dual stack function.</li></ul>
2020-05-20	<p>This issue is the twenty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "Why Does My Server Can Be Accessed from the Internet But Cannot Access the Internet?"</li><li>• Added section "Cloning a Security Group".</li><li>• Modified FAQ "How Can I Delete a Subnet That Is Being Used by Other Resources?"</li></ul>
2020-04-15	<p>This issue is the twentieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "Can an EIP Be Bound to a Cloud Resource in Another Region?"</li><li>• Added FAQ "How Do I Query the Region of My EIPs?"</li><li>• Added FAQ "Can I Transfer an EIP to Another Account?"</li><li>• Added FAQ "Can I Assign a Specific EIP?"</li><li>• Added FAQ "Will an EIP Be Changed After I Assign It?"</li><li>• Added FAQ "How Do I Change an EIP for an Instance?"</li><li>• Added FAQ "How Do I Switch to a Private DNS Server?"</li></ul>

Release Date	What's New
2020-03-30	<p>This issue is the nineteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added basic information to sections "Security Group Overview" and "Network ACL Overview".</li><li>• Added FAQ "Does a Security Group Rule or a Network ACL Rule Immediately Take Effect for Its Original Traffic After It Is Modified?"</li><li>• Added section "Billing".</li><li>• Added category "Billing and Payments" to FAQs.</li></ul>
2020-03-20	<p>This issue is the eighteenth official release, which incorporates the following change:</p> <p>Deleted FAQ "Which Security Group Rule Has Priority When Multiple Security Group Rules Conflict?"</p>
2020-02-18	<p>This issue is the seventeenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "How Is an EIP Charged?"</li><li>• Optimized FAQ "Can I Bind One EIP to Multiple ECSs?"</li></ul>
2019-12-23	<p>This issue is the sixteenth official release, which incorporates the following changes:</p> <p>Updated the document based on the navigation path and function changes of <b>Subnets</b> and <b>Route Tables</b>.</p>
2019-12-03	<p>This issue is the fifteenth official release, which incorporates the following change:</p> <p>Optimized description and figures in section "Service Overview".</p>
2019-11-20	<p>This issue is the fourteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "What Is the EIP Assignment Policy?"</li><li>• Added FAQ "What Are the Differences Between Static BGP and Dynamic BGP?"</li></ul>
2019-10-15	<p>This issue is the thirteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added section "VPC Flow Log (OBT)".</li><li>• Updated the screenshots of adding security group rules.</li><li>• Added FAQ "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?"</li></ul>

Release Date	What's New
2019-10-09	<p>This issue is the twelfth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added FAQ "Can I Bind an EIP to an ECS, to Another ECS?"</li><li>• Added FAQ "Will the EIP Bound to an ECS Be Changed After the ECS Is Stopped and Then Started?"</li></ul>
2019-09-26	<p>This issue is the eleventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Optimized the sections in "VPC Peering Connection".</li><li>• Added section "Common Ports Used by ECSs".</li><li>• Added FAQ "Why Are Some Ports Inaccessible?"</li></ul>
2019-09-12	<p>This issue is the tenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Deleted section "Deleting a VPN".</li><li>• Added FAQ "What Is the Relationship Between Bandwidth and Upload or Download Rate?"</li><li>• Added content to FAQ "What Are the Restrictions on Deleting a Security Group?"</li></ul>
2019-08-15	<p>This issue is the ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added the example of allowing external access to a specified port in the section "Security Group Configuration Examples".</li><li>• Added billing information in section "Modifying EIP Bandwidth".</li><li>• Added FAQ "How Do I Change the Billing Mode?"</li><li>• Added FAQ "How Do I Change the Bandwidth Billing Option from <b>Bandwidth</b> to <b>Traffic</b> or from <b>Traffic</b> to <b>Bandwidth</b>?"</li><li>• Added FAQ "Can I Increase My Yearly/Monthly Bandwidth Size and Then Decrease It?"</li></ul>
2019-02-28	<p>This issue is the eighth official release, which incorporates the following change:</p> <ul style="list-style-type: none"><li>• Added section <b>Shared Data Package</b>.</li></ul>
2018-12-30	<p>This issue is the seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Modified the description about how to switch to the security group, network ACL, EIP, and shared bandwidth pages based on the changes made on the management console.</li><li>• Added section "Network ACL Overview".</li><li>• Added section "Network ACL Configuration Examples".</li></ul>

Release Date	What's New
2018-11-30	<p>This issue is the sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Updated the document based on changes made to the network ACL GUI.<ul style="list-style-type: none"><li>– Added description about how to delete multiple network ACL rules at a time and how to disassociate multiple subnets from a network ACL at a time.</li><li>– Changed parameter <b>Any</b> to <b>All</b>.</li></ul></li><li>• Modified FAQ "Why Does My ECS Fail to Obtain an IP Address?"</li></ul>
2018-09-30	<p>This issue is the fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added description about how to create multiple subnets at a time to section "Creating a VPC".</li><li>• Added description about how to add multiple network ACL rules at a time and parameter <b>Description</b> to section "Adding a Network ACL Rule".</li></ul>
2018-08-15	<p>This issue is the fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added sections under "Shared Bandwidth".</li><li>• Optimized sections under "Service Overview."</li><li>• Optimized sections under "Security Group".</li></ul>
2018-05-30	<p>This issue is the third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added section "Monitoring".</li><li>• Optimized the words on button labels to ensure consistency.</li></ul>
2018-04-28	<p>This issue is the second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added section <b>Exporting VPC Information</b>.</li><li>• Modified EIP description.</li></ul>
2017-12-31	<p>This issue is the first official release.</p>