

Virtual Private Cloud

User Guide

Issue 01
Date 2024-08-06



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

1 Permissions Management.....	1
1.1 Creating a User and Granting VPC Permissions.....	1
1.2 VPC Custom Policies.....	2
2 VPC and Subnet.....	5
2.1 VPC and Subnet Planning.....	5
2.2 VPC Connectivity Options.....	13
2.2.1 Overview.....	13
2.2.2 Connecting VPCs.....	16
2.2.3 Connecting VPCs to the Public Network.....	19
2.2.4 Connecting VPCs to On-Premises Data Centers.....	24
2.3 VPC.....	26
2.3.1 Creating a VPC and Subnet.....	26
2.3.2 Adding a Secondary IPv4 CIDR Block to a VPC.....	37
2.3.3 Obtaining a VPC ID.....	39
2.3.4 Modifying a VPC.....	40
2.3.5 Managing VPC Tags.....	41
2.3.6 Viewing a VPC Topology.....	43
2.3.7 Exporting VPC List.....	43
2.3.8 Deleting a Secondary IPv4 CIDR Block from a VPC.....	44
2.3.9 Deleting a VPC.....	44
2.4 Subnet.....	45
2.4.1 Creating a Subnet for the VPC.....	45
2.4.2 Modifying a Subnet.....	51
2.4.3 Managing Subnet Tags.....	53
2.4.4 Exporting Subnet List.....	55
2.4.5 Viewing and Deleting Resources in a Subnet.....	56
2.4.6 Viewing IP Addresses in a Subnet.....	58
2.4.7 Deleting a Subnet.....	59
3 Route Tables and Routes.....	61
3.1 Route Tables and Routes.....	61
3.2 Managing Route Tables.....	65
3.2.1 Creating a Custom Route Table.....	65

3.2.2 Associating a Route Table with a Subnet.....	66
3.2.3 Changing the Route Table Associated with a Subnet.....	67
3.2.4 Viewing the Route Table Associated with a Subnet.....	68
3.2.5 Viewing Route Table Information.....	68
3.2.6 Exporting Route Table Information.....	69
3.2.7 Deleting a Route Table.....	69
3.3 Managing Routes.....	70
3.3.1 Adding Routes to a Route Table.....	70
3.3.2 Modifying a Route.....	71
3.3.3 Replicating a Route.....	73
3.3.4 Deleting a Route.....	74
3.4 Route Configuration Examples.....	75
3.4.1 Configuring an SNAT Server to Enable ECSs to Share an EIP to Access the Internet.....	75
4 Virtual IP Address.....	79
4.1 Virtual IP Address Overview.....	79
4.2 Assigning a Virtual IP Address.....	83
4.3 Binding a Virtual IP Address to an Instance or EIP.....	83
4.4 Unbinding a Virtual IP Address from an Instance or EIP.....	90
4.5 Releasing a Virtual IP Address.....	91
4.6 Virtual IP Address Configuration Example.....	92
4.6.1 Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster.....	92
5 Elastic Network Interface and Supplementary Network Interface.....	109
5.1 Elastic Network Interface.....	109
5.1.1 Elastic Network Interface Overview.....	109
5.1.2 Creating a Network Interface.....	110
5.1.3 Viewing the Basic Information About a Network Interface.....	111
5.1.4 Attaching a Network Interface to a Cloud Server.....	112
5.1.5 Binding an EIP to a Network Interface.....	112
5.1.6 Binding a Network Interface to a Virtual IP Address.....	113
5.1.7 Detaching a Network Interface from an Instance or Unbinding an EIP from a Network Interface..	113
5.1.8 Changing Security Groups That Are Associated with a Network Interface.....	114
5.1.9 Deleting a Network Interface.....	115
5.2 Supplementary Network Interfaces.....	116
5.2.1 Supplementary Network Interface Overview.....	116
5.2.2 Creating a Supplementary Network Interface.....	117
5.2.3 Viewing the Basic Information About a Supplementary Network Interface.....	120
5.2.4 Binding or Unbinding an EIP to or from a Supplementary Network Interface.....	121
5.2.5 Changing Security Groups That Are Associated with a Supplementary Network Interface.....	122
5.2.6 Deleting a Supplementary Network Interface.....	123
5.3 Network Interface Configuration Examples.....	124
5.3.1 Binding an EIP to the Extension NIC of an ECS to Enable Internet Access.....	124
5.3.2 Configuring Policy-based Routes for an ECS with Multiple NICs.....	128

5.3.2.1 Overview.....	128
5.3.2.2 Collecting ECS Network Information.....	129
5.3.2.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs	131
5.3.2.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs.....	135
6 Access Control.....	138
6.1 Access Control Overview.....	138
6.2 Security Group.....	140
6.2.1 Security Groups and Security Group Rules.....	140
6.2.2 Default Security Groups.....	150
6.2.3 Security Group Examples.....	152
6.2.4 Common Ports Used by ECSs.....	157
6.2.5 Managing a Security Group.....	159
6.2.5.1 Creating a Security Group.....	160
6.2.5.2 Cloning a Security Group.....	164
6.2.5.3 Modifying a Security Group.....	165
6.2.5.4 Deleting a Security Group.....	166
6.2.6 Managing Security Group Rules.....	167
6.2.6.1 Adding a Security Group Rule.....	167
6.2.6.2 Fast-Adding Security Group Rules.....	173
6.2.6.3 Allowing Common Ports with a Few Clicks.....	177
6.2.6.4 Modifying a Security Group Rule.....	178
6.2.6.5 Replicating a Security Group Rule.....	179
6.2.6.6 Importing and Exporting Security Group Rules.....	180
6.2.6.7 Deleting a Security Group Rule.....	183
6.2.6.8 Querying Security Group Rule Changes.....	184
6.2.7 Managing Instances Added to a Security Group.....	187
6.2.7.1 Adding an Instance to or Removing an Instance from a Security Group.....	188
6.2.7.2 Changing the Security Group of an ECS.....	189
6.3 Network ACL.....	189
6.3.1 Network ACL Overview.....	190
6.3.2 Network ACL Configuration Examples.....	199
6.3.3 Managing Network ACLs.....	203
6.3.3.1 Creating a Network ACL.....	203
6.3.3.2 Modifying a Network ACL.....	204
6.3.3.3 Enabling or Disabling a Network ACL.....	204
6.3.3.4 Viewing a Network ACL.....	205
6.3.3.5 Deleting a Network ACL.....	206
6.3.4 Managing Network ACL Rules.....	206
6.3.4.1 Adding a Network ACL Rule (Default Priorities).....	206
6.3.4.2 Adding a Network ACL Rule (Custom Priorities).....	209
6.3.4.3 Modifying a Network ACL Rule.....	210
6.3.4.4 Enabling or Disabling a Network ACL Rule.....	212

6.3.4.5 Exporting and Importing Network ACL Rules.....	213
6.3.4.6 Deleting a Network ACL Rule.....	214
6.3.5 Managing Subnets Associated with a Network ACL.....	214
6.3.5.1 Associating Subnets with a Network ACL.....	214
6.3.5.2 Disassociating Subnets from a Network ACL.....	216
7 IP Address Group.....	218
7.1 IP Address Group.....	218
7.2 Managing an IP Address Group.....	219
7.2.1 Creating an IP Address Group.....	219
7.2.2 Associating an IP Address Group with Resources.....	221
7.2.3 Modifying an IP Address Group.....	222
7.2.4 Exporting IP Address Group Details.....	223
7.2.5 Viewing the Details of an IP Address Group.....	224
7.2.6 Deleting an IP Address Group.....	224
7.3 Managing IP Addresses in an IP Address Group.....	225
7.3.1 Adding IP Addresses to an IP Address Group.....	225
7.3.2 Modifying IP Addresses in an IP Address Group.....	226
7.3.3 Deleting IP Addresses from an IP Address Group.....	228
7.4 IP Address Group Configuration Examples.....	229
7.4.1 Using IP Address Groups to Reduce the Number of Security Group Rules.....	229
8 VPC Peering Connection.....	233
8.1 VPC Peering Connection.....	233
8.2 VPC Peering Connection Usage.....	235
8.2.1 VPC Peering Connection Usage Examples.....	235
8.2.2 Using a VPC Peering Connection to Connect Two VPCs.....	236
8.2.3 Using a VPC Peering Connection to Connect Subnets in Two VPCs.....	276
8.2.4 Using a VPC Peering Connection to Connect ECSs in Two VPCs.....	290
8.2.5 Unsupported VPC Peering Configurations.....	297
8.3 Creating a VPC Peering Connection to Connect Two VPCs in the Same Account.....	302
8.4 Creating a VPC Peering Connection Connect Two VPCs in Different Accounts.....	309
8.5 Obtaining the Peer Project ID of a VPC Peering Connection.....	317
8.6 Modifying a VPC Peering Connection.....	318
8.7 Viewing VPC Peering Connections.....	319
8.8 Deleting a VPC Peering Connection.....	319
8.9 Modifying Routes Configured for a VPC Peering Connection.....	320
8.10 Viewing Routes Configured for a VPC Peering Connection.....	321
8.11 Deleting Routes Configured for a VPC Peering Connection.....	323
9 Setting Up an IPv6 Network.....	325
10 VPC Flow Log.....	331
10.1 VPC Flow Log.....	331
10.2 Creating a VPC Flow Log.....	332

10.3 Viewing a VPC Flow Log.....	334
10.4 Enabling or Disabling VPC Flow Log.....	337
10.5 Deleting a VPC Flow Log.....	338
11 Elastic IP.....	339
11.1 EIP Overview.....	339
11.2 Assigning an EIP and Binding It to an ECS.....	340
11.3 Unbinding an EIP from an ECS and Releasing the EIP.....	344
11.4 Modifying an EIP Bandwidth.....	345
11.5 Exporting EIP Information.....	346
11.6 Managing EIP Tags.....	347
11.7 IPv6 EIP	348
12 Shared Bandwidth.....	354
12.1 Shared Bandwidth Overview.....	354
12.2 Assigning a Shared Bandwidth.....	355
12.3 Adding EIPs to a Shared Bandwidth.....	357
12.4 Removing EIPs from a Shared Bandwidth.....	357
12.5 Modifying a Shared Bandwidth.....	358
12.6 Deleting a Shared Bandwidth.....	359
13 Monitoring and Auditing.....	361
13.1 Cloud Eye Monitoring.....	361
13.1.1 Supported Metrics.....	361
13.1.2 Viewing Metrics.....	363
13.1.3 Creating an Alarm Rule.....	364
13.2 CTS Auditing.....	364
13.2.1 Key Operations Recorded by CTS.....	364
13.2.2 Viewing Traces.....	367
14 Managing Quotas.....	369

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 cloud account or cloud service to perform efficient O&M on your VPC resources.

If your cloud account 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

Before using VPCs and subnets to build cloud networks, determine how many VPCs and subnets do you need and plan the necessary CIDR blocks and connectivity options. If you need to connect different VPCs or connect a VPC to an on-premises data center, ensure that their CIDR blocks do not conflict. Properly plan your VPCs and subnets based on the guidelines provided here to avoid CIDR block conflicts, which will make future network expansion easier.

- [How Do I Determine How Many VPCs I Need?](#)
- [How Do I Determine How Many Subnets I Need?](#)
- [How Do I Plan CIDR Blocks for VPCs and Subnets?](#)
- [How Do I Plan How Many Route Tables I Need?](#)
- [How Do I Connect Two VPC or Connect a VPC to an On-premises Data Center?](#)

How Do I Determine How Many VPCs I Need?

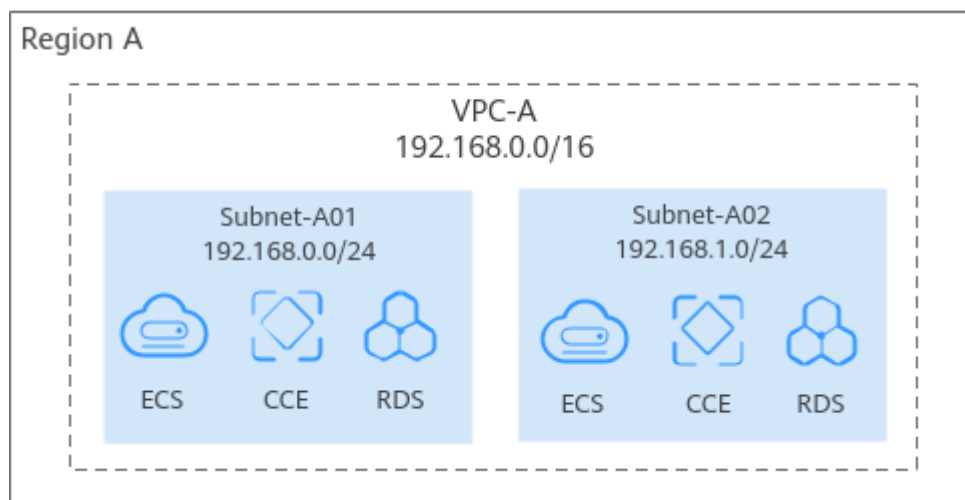
VPCs are region-specific. Cloud resources, such as ECSs, CCEs, and RDS instances, in a VPC must be in the same region as the VPC. By default, different VPCs are isolated from each other, but the subnets in a VPC can communicate with each other.

Planning a Single VPC

If your services are deployed in one region and do not have to handle a lot of traffic, you may not need network isolation. In this case, a single VPC should be enough.

You can create multiple subnets in a VPC for workloads with different requirements and associate route tables with these subnets to control traffic in and out of the subnets. In [Figure 2-1](#), services are deployed on different subnets in a VPC (VPC-A in this example).

Figure 2-1 Planning a single VPC



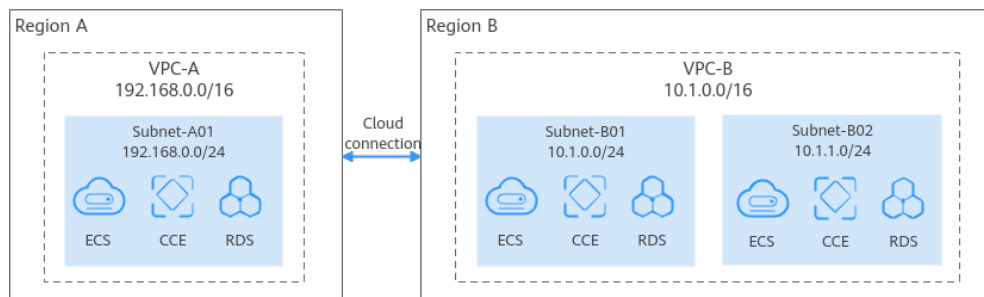
Planning Multiple VPCs

You need to plan multiple VPCs if you have:

- **Services that need to be deployed in different regions**


VPC is a region-specific service, so services cannot be deployed across regions in a VPC. If your services are deployed in multiple regions, plan at least one VPC in each region.

Figure 2-2 Planning multiple VPCs



- **Services that are deployed in the same region but need network isolation.**

If your services are deployed in the same region but need network isolation, you need to plan multiple VPCs in this region. Different VPCs are isolated from each other, so you can deploy different services in different VPCs, as shown in [Figure 2-3](#). In the figure, some services are deployed in VPC-A, and some are deployed in VPC-B. The two VPCs are isolated from each other.

Figure 2-3 Planning multiple VPCs**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](#).

How Do I Determine How Many Subnets I Need?

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

You can create different subnets for different services in a VPC. For example, you can create three subnets in a VPC, one subnet for web services, one for management services, and the third one for data services. Additionally, you can use network ACLs to control access to each subnet.

Note the following when selecting subnets and AZs for your resources:

- All instances in different subnets of the same VPC can communicate with each other by default, and the subnets can be located in different AZs. If you have a VPC with two subnets in it and they are located in different AZs, they can communicate with each other by default.
- 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.

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](#).

How Do I Plan CIDR Blocks for VPCs and Subnets?

After VPCs and subnets are created, their CIDR blocks cannot be changed. To ensure smooth service expansion and O&M, properly plan VPC and subnet CIDR blocks that best suit your service size and communication requirements.

 NOTE

Both IPv4 and IPv6 CIDR blocks can be assigned to a subnet. You can customize IPv4 CIDR blocks but not IPv6 CIDR blocks. The system assigns an IPv6 CIDR block with a 64-bit mask to each subnet, for example, 2407:c080:802:1b32::/64.

Planning VPC CIDR Blocks

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

- Reserve sufficient IP addresses for subsequent service expansion.
- Avoid CIDR block conflicts. To enable communications between VPCs or between a VPC and an on-premises data center, ensure their CIDR blocks do not overlap.

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.

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

Planning Subnet CIDR Blocks

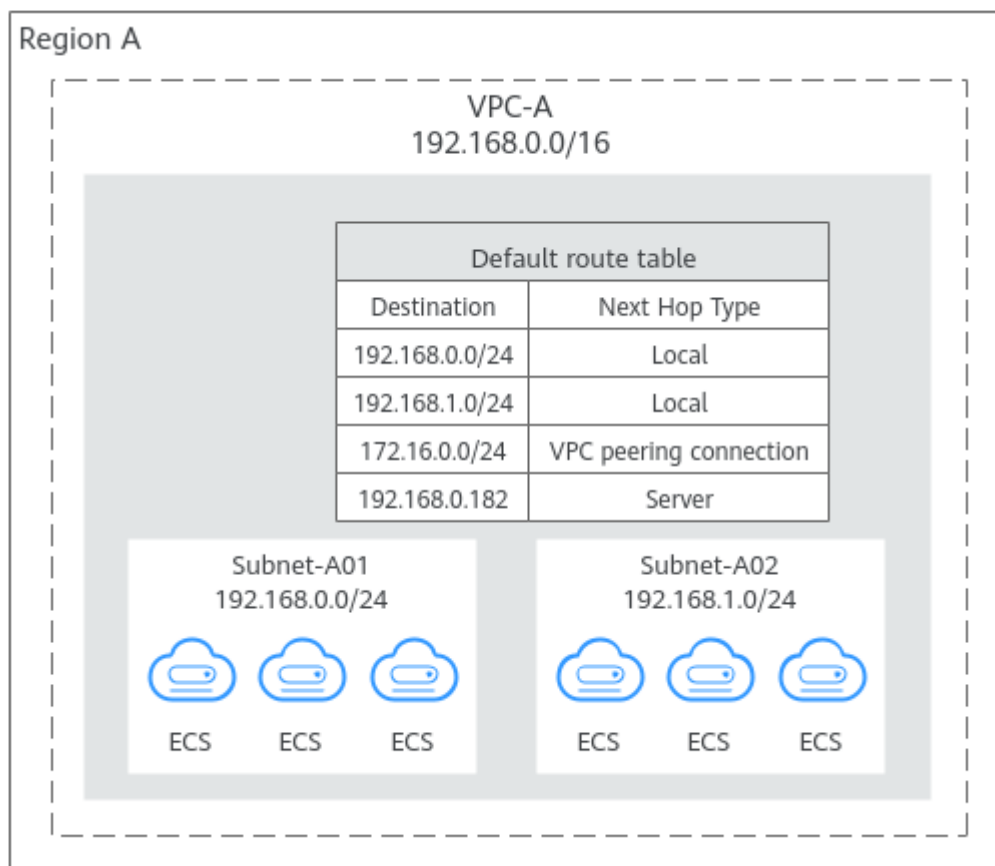
- Subnet mask planning: The subnet CIDR block must be within the VPC CIDR block. Subnet CIDR blocks in a VPC must be unique. A subnet mask can be between the netmask of its VPC CIDR block and a /29 netmask. If a VPC CIDR block is 10.0.0.0/16, its subnet mask can be anything from 16 to 29.
For example, if the CIDR block of a VPC is 10.0.0.0/16, you can specify 10.0.0.0/24 for a subnet in this VPC, 10.0.1.0/24 for the second subnet, and 10.0.2.0/24 for the third subnet.
- Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to properly plan the CIDR block in advance based on the number of IP addresses required by your service.
 - The subnet CIDR block size cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address.
 - The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks available later for new subnets, which can be a problem when you want to scale out services.
- Avoiding subnet CIDR block conflicts: Avoid CIDR block conflicts if you need to connect two VPCs or connect a VPC to an on-premises data center.
If the subnet CIDR blocks at both ends of the network conflict, [create a subnet](#).

How Do I Plan How Many Route Tables I Need?

A route table contains a set of routes that are used to control the traffic in and out of your subnets in a VPC. You can configure destination, next hop, and other information for each route. A VPC can have multiple route tables. Plan route tables based on the following sections.

Planning One Route Table

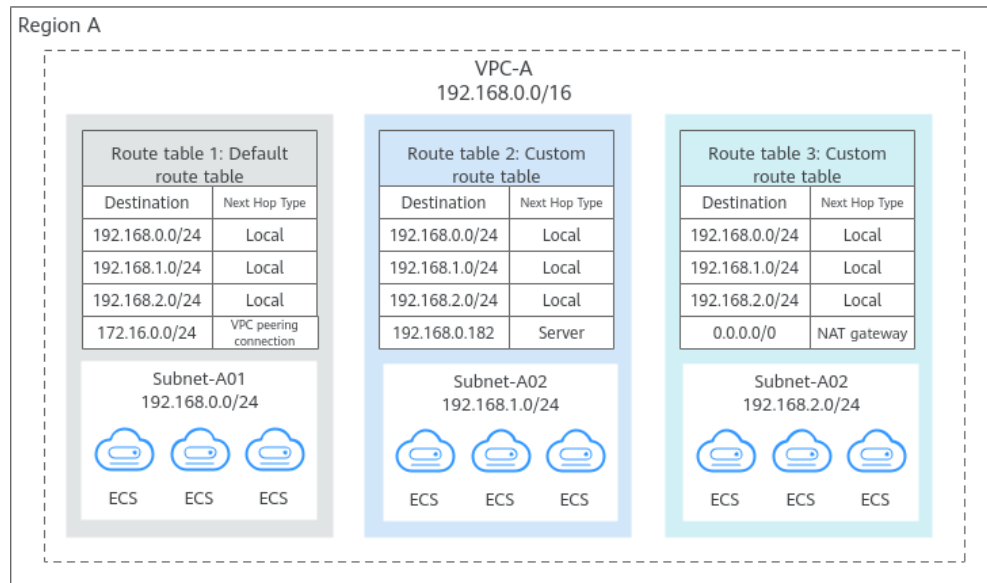
If you have the same or similar requirements for controlling the network traffic to and from subnets in a VPC, you can create one route table and associate it with these subnets in this VPC. Each VPC comes with a default route table. If you create a subnet in the VPC, the subnet is associated with the default route table. You can add routes to the default route table to control where the traffic is directed. In [Figure 2-4](#), VPC-A has only the default route table, and subnets Subnet-A01 and Subnet-A02 are associated with the default route table.

Figure 2-4 Planning one route table

Planning Multiple Route Tables

If you have different requirements for controlling the network traffic to and from subnets in a VPC, the default route table is not enough. You can create one or more custom route tables and associate them with these subnets in this VPC. In [Figure 2-5](#), VPC-A has three route tables. Subnet-A01 is associated with default route table 1, Subnet-A02 is associated with custom route table 2, and Subnet-A03 is associated with custom route table 3.

Figure 2-5 Planning multiple route tables




How Do I Connect Two VPC or Connect a VPC to an On-premises Data Center?

If you need to connect two VPCs or connect a VPC to an on-premises data center, ensure that their VPC CIDR blocks do not conflict.

Connecting Two VPCs

Connecting VPCs in the same region: In [Figure 2-6](#), there are three VPCs in region A: VPC-A, VPC-B, and VPC-X. If you want to connect VPC-A and VPC-B, but isolate VPC-C from other VPCs:


- Ensure that the CIDR blocks of VPC-A and VPC-B connected by a peering connection (Peering-AB in this example) must be unique.
- You do not need to worry about VPC CIDR block conflicts because VPC-X does not need to communicate with other VPCs. If VPC-X and VPC-B need to communicate with each other, you can specify different CIDR blocks for the subnets in the two VPCs and create a VPC peering connection to connect the subnets.

Figure 2-6 Connecting VPCs in the same region

Connecting a VPC to an On-premises Data Center

In [Figure 2-7](#), VPC-A and VPC-B in region A need to communicate with each other, and VPC-A needs to connect to on-premises data center IDC-A. In region C, VPC-C needs to connect to on-premises data center IDC-C.

- In region A, VPC-A and VPC-B have different CIDR blocks and can communicate with each other through a VPC peering connection. VPC-A and IDC-A have different CIDR blocks and are connected through a direct connection.
- In region C, VPC-C and IDC-C have different CIDR blocks and are connected through a VPC connection.

Figure 2-7 Connecting a VPC to an on-premises data center

Helpful Link

- You can create a VPC and an ECS to set up an IPv4 private network on the cloud and then bind an EIP to the ECS to allow the ECS to access the Internet. For details, see [Setting Up an IPv4 Network in a VPC](#).
- You can create a VPC with an IPv4 and IPv6 CIDR block and create an ECS with both IPv4 and IPv6 addresses in the VPC. You can bind an EIP and add the IPv6 address of the ECS to a shared bandwidth to enable the ECS to communicate with the Internet over both IPv4 and IPv6 networks. For details, see [Setting Up an IPv4/IPv6 Dual-Stack Network in a VPC](#).

2.2 VPC Connectivity Options

2.2.1 Overview

Huawei Cloud provides various network services for you to set up secure and scalable cloud networks. With these network services, you can connect VPCs in the same region or different regions, enable the instances (such as ECSs and RDS instances) in VPCs to access the public network, and enable on-premises data centers to access the VPCs. The following describes the function and highlights of each network service. You can flexibly configure VPC and other network services based on your network requirements:

- [Connecting VPCs](#)
- [Connecting VPCs to the Public Network](#)
- [Connecting VPCs to an On-Premises Data Center](#)

Connecting VPCs

With the network services described in [Table 2-3](#), you can flexibly connect VPCs in the same region, in different regions, or in different accounts.

Table 2-3 Network services that can connect VPCs

Network Service	Function	Highlights
VPC Peering	With VPC Peering, you can peer two VPCs in the same region. The VPCs can be in the same account or different accounts.	<ul style="list-style-type: none">• VPC Peering is free.• Routes can be configured on the console easily.
Enterprise Router	An enterprise router can connect multiple VPCs in the same account or different accounts to set up a hub-and-spoke network. Compared with VPC Peering, Enterprise Router is more suitable for complex networking where many VPCs need to be connected.	<ul style="list-style-type: none">• VPCs in the same region can be connected in minutes.• Routes can be automatically added.• Low latency and high speed• Simple network topology and high scalability
VPN	You can use VPN connect VPCs in different regions, so that they can communicate with each other over the Internet.	<ul style="list-style-type: none">• Low costs• Simple configuration• Immediate use• The network quality depends on the Internet.
Direct Connect	You can use Direct Connect to connect VPCs in different regions.	<ul style="list-style-type: none">• Dedicated connections with high security• Low latency and high speed

Connecting VPCs to the Public Network

With the network services described in [Table 2-4](#), you can connect VPCs to the public network so that instances in the VPCs can access the public network or provide services accessible on the public network.

Table 2-4 Network services that allow VPCs to communicate with the public network

Network Service	Function	Highlights
EIP	An EIP is an independent public IP address. You can bind it to an instance, such as an ECS, a NAT gateway, or a load balancer, so that the instance can access the public network or provide services accessible from the public network.	<ul style="list-style-type: none"> • EIPs can be bound to or unbound from instances if needed. • Shared bandwidths can be used to lower costs. • EIP bandwidth can be adjusted at any time.
NAT Gateway <ul style="list-style-type: none"> • SNAT • DNAT 	NAT Gateway supports both source NAT (SNAT) and destination NAT (DNAT). <ul style="list-style-type: none"> • SNAT enables multiple instances to share one or more EIPs to access the public network. <ul style="list-style-type: none"> - ECSs in the same VPC sharing an EIP - ECSs in different VPCs sharing an EIP • DNAT enables port 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. However, DNAT does not balance traffic. 	<ul style="list-style-type: none"> • Using shared EIPs to access the public network reduces the costs. • EIPs of ECSs are not exposed to the public network, which improves security. • Different specifications are available.
ELB	ELB evenly distributes incoming traffic to multiple backend servers. Together with EIPs, ELB allows a large number of users to access services deployed on cloud servers from the public network.	<ul style="list-style-type: none"> • ELB can process both Layer 4 and Layer 7 requests and supports advanced forwarding policies and multiple protocols. • ELB can eliminate single points of failure (SPOFs) for high availability.

Connecting VPCs to an On-Premises Data Center

If you have an on-premises data center and not all your workloads can be migrated to the cloud, you can use the network services described in [Table 2-5](#) to connect your on-premises data center to the VPCs.

Table 2-5 Network services that can connect VPCs to on-premises data centers

Network Service	Function	Highlights
VPN	VPN provides an encrypted, Internet-based channel that connects an on-premises data center and the cloud.	<ul style="list-style-type: none">• Low costs• Simple configuration• Immediate use• The network quality depends on the Internet.
Direct Connect	Direct Connect establishes a dedicated network connection between an on-premises data center and the cloud.	<ul style="list-style-type: none">• Dedicated connections with high security• Low latency and high speed
VPC Peering	With VPC Peering, you can peer two VPCs in the same region, no matter whether they are in the same account or different accounts. VPC Peering can work with Direct Connect or VPN to enable your on-premises data center to access multiple VPCs.	<ul style="list-style-type: none">• VPC Peering is free.• Routes can be configured on the console easily.

2.2.2 Connecting VPCs

Connecting VPCs in the Same Region

If the VPCs you want to connect are in the same region, you can use VPC Peering or Enterprise Router.

[Connecting VPCs](#) provides details about different network services.

NOTICE

Before connecting VPCs, you need to plan their CIDR blocks in advance. Overlapping CIDR blocks may cause communication failure.

VPC Peering

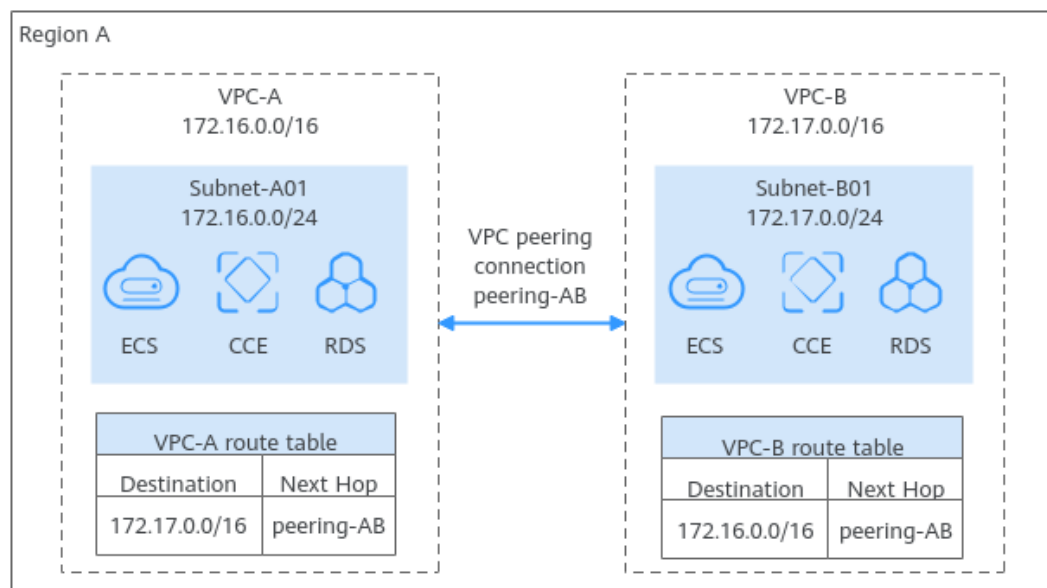
With VPC Peering, you can peer two VPCs in the same region. The VPCs can be in the same account or different accounts.

You can refer to the following topics:

- [Using a VPC Peering Connection to Connect Two VPCs in the Same Account](#)
- [Using a VPC Peering Connection to Connect Two VPCs in Different Accounts](#)

In [Figure 2-8](#), a VPC peering connection (Peering-AB) connects two VPCs (VPC-A and VPC-B) in a region.

Figure 2-8 Connecting VPCs in the same region over a VPC peering connection




Enterprise Router

An enterprise router can connect multiple VPCs in the same account or different accounts to set up a hub-and-spoke network. Compared with VPC Peering, Enterprise Router is more suitable for complex networking where many VPCs need to be connected.

For details, see [Using an Enterprise Router to Enable Communications Between VPCs in the Same Region](#).

In [Figure 2-9](#), an enterprise router connects multiple VPCs in the same region and forwards traffic among them. The routes are automatically configured for the VPCs and the enterprise router.

Figure 2-9 Connecting VPCs in the same region using an enterprise router



Connecting VPCs in Different Regions

If the VPCs to be connected are located in different regions, you can use Cloud Connect, Direct Connect, or VPN.

[Connecting VPCs](#) provides details about different network services.

NOTICE

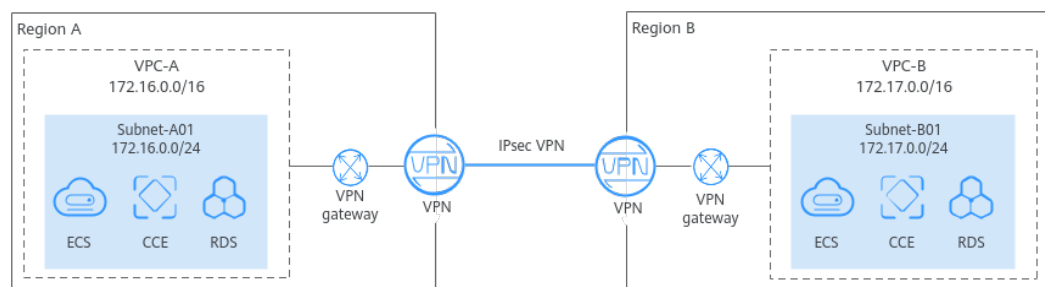
Before connecting VPCs, you need to plan their CIDR blocks in advance. Overlapping CIDR blocks may cause communication failure.

VPN

You can use **VPN** connect VPCs in different regions, so that they can communicate with each other over the Internet.

In **Figure 2-10**, there is a VPC in each region: VPC-A in region A and VPC-B in region B. Each VPC is connected to a VPN connection. The two VPCs can communicate with each other through an encrypted channel on the Internet. VPN can be enabled fast and is cost-effective.

Figure 2-10 Connecting VPCs in different regions using VPN

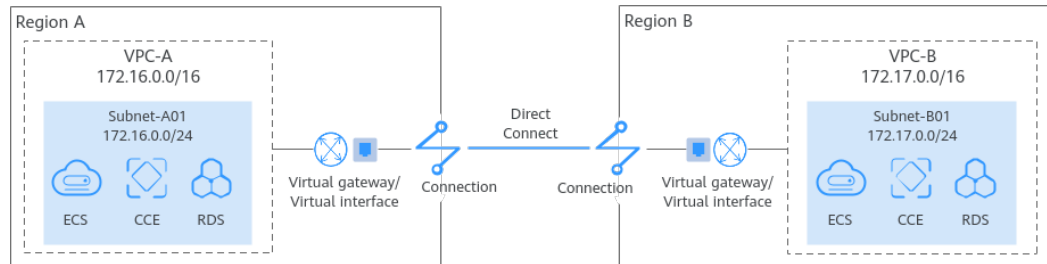


Direct Connect

You can use **Direct Connect** to connect VPCs in different regions.

In **Figure 2-11**, there is a VPC in each region: VPC-A in region A and VPC-B in region B. Each VPC is connected to a Direct Connect connection. The two VPCs can communicate with each other through a dedicated connection. Compared with VPN, Direct Connect enables faster, more stable data transmission.

Figure 2-11 Connecting VPCs in different regions using Direct Connect



2.2.3 Connecting VPCs to the Public Network

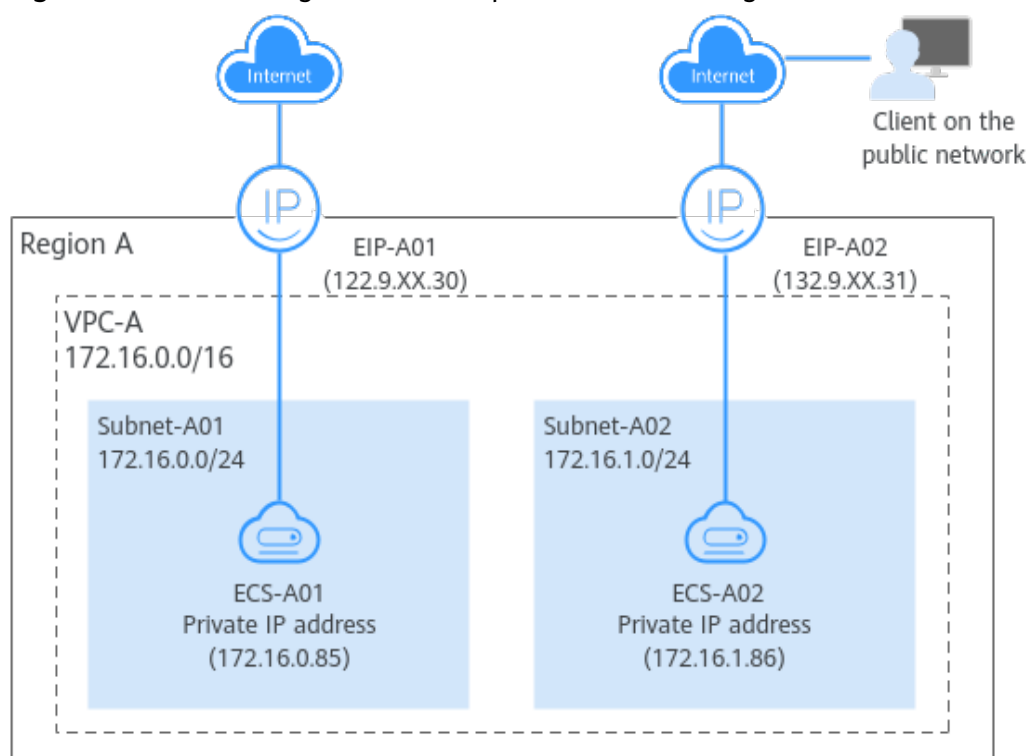
You can use EIP, NAT Gateway, or ELB to allow the resources in VPCs to access the public network.

EIP

An EIP is an independent public IP address. You can bind it to an instance, such as an ECS, a NAT gateway, or a load balancer, so that the instance can access the public network or provide services accessible from the public network.

- For details about EIPs in IPv4 networks, see [Setting Up an IPv4 Network in a VPC](#).
- For details about IPv4/IPv6 dual-stack networks, see [Quickly Setting Up an IPv4/IPv6 Dual-Stack Network in a VPC](#).

In **Figure 2-12**, there are two subnets (Subnet-A01 and Subnet-A02) in a region (region A), and there is an ECS on each subnet. The ECS (ECS-A01) on Subnet-A01 needs to access the public network, and the ECS (ECS-A02) on Subnet-A02 needs to provide web services for the public network. Two EIPs (EIP-A01 and EIP-A02) are required, with each bound to an ECS.

Figure 2-12 Connecting a VPC to the public network using EIP

NAT Gateway (SNAT)

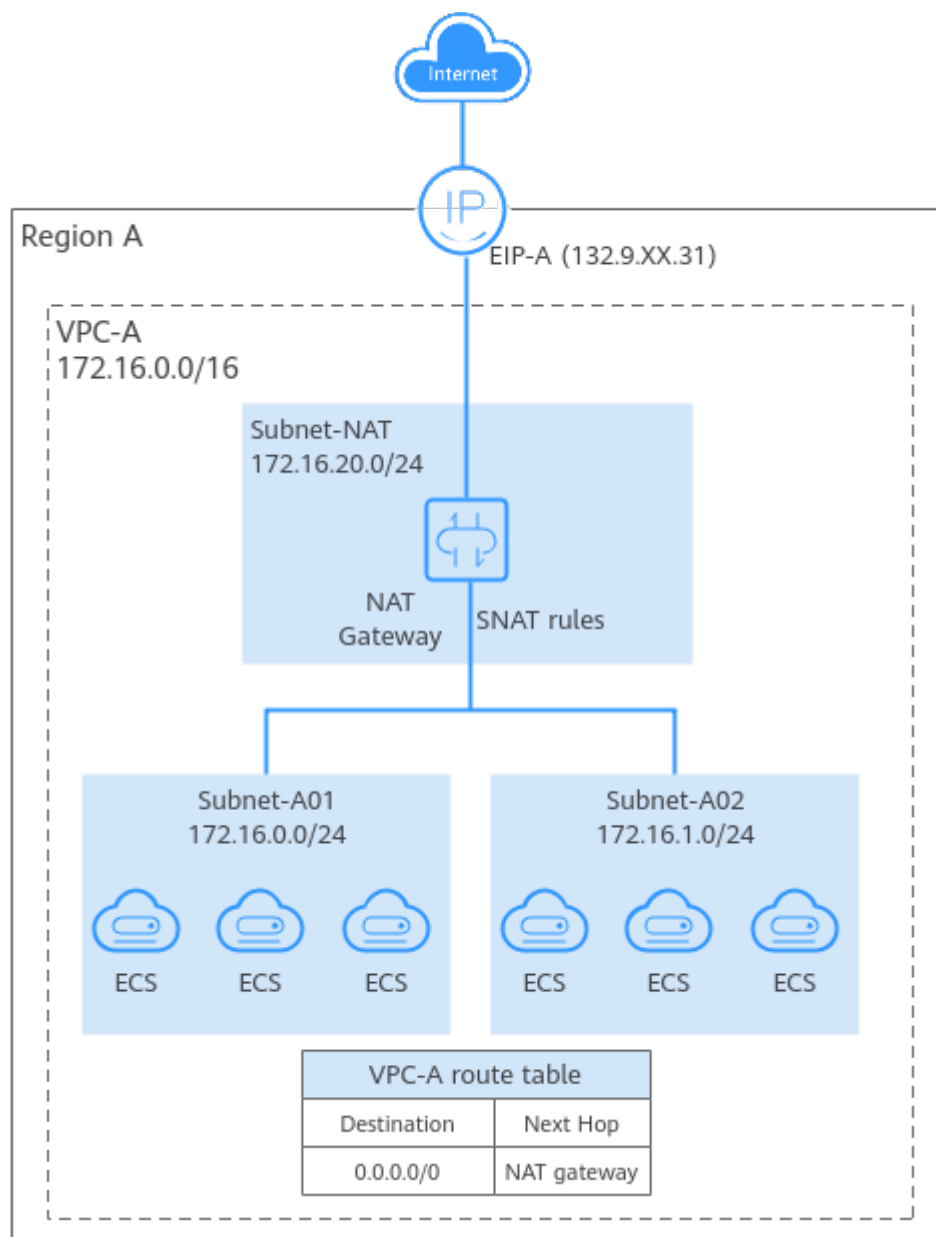
You can use a public network NAT gateway and configure SNAT rules to enable multiple ECSs in a VPC to share one or more EIPs to access the public network. If only SNAT rules are configured, the public network address of the NAT gateway cannot be directly accessed from the public network. This is more secure than using EIPs.

- If ECSs in the same VPC share an EIP, see [Using SNAT to Access the Internet](#).
- If ECSs in different VPCs share an EIP, see [Allowing VPCs to Share an EIP to Access the Internet Using Enterprise Router and NAT Gateway](#).

ECSs in a VPC Sharing an EIP

In [Figure 2-13](#), ECSs deployed on two subnets (Subnet-A01 and Subnet-A02) in a VPC (VPC-A) need to access the public network. For this to work, you first need to create a public NAT gateway in a third subnet (Subnet-NAT), and then configure SNAT rules on the public NAT gateway for Subnet-A01 and Subnet-A02. In this way, all ECSs in Subnet-A01 and Subnet-A02 can share an EIP to access the public network.

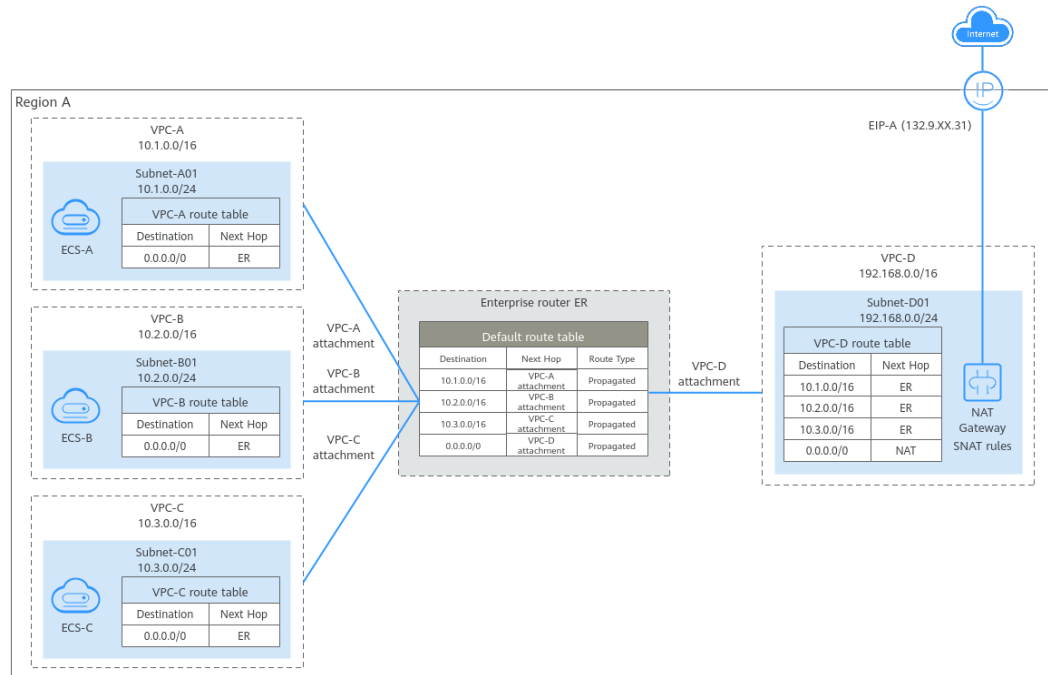
Figure 2-13 Enabling ECSs in a VPC to access the public network using a NAT gateway



ECSs in Different VPCs Sharing an EIP

In [Figure 2-14](#), three VPCs (VPC-A, VPC-B, and VPC-C) in a region need to communicate with each other and can use the NAT gateway deployed in another VPC (VPC-D) to access the public network. For this to work, you first need to attach the four VPCs to an enterprise router, then configure routes in the route tables of the VPCs and of the enterprise router, and configure SNAT rules on the public NAT gateway. In this way, the VPCs can communicate with each other and share an EIP to access the public network.

Figure 2-14 Enabling ECSs in different VPCs to access the public network using a NAT gateway



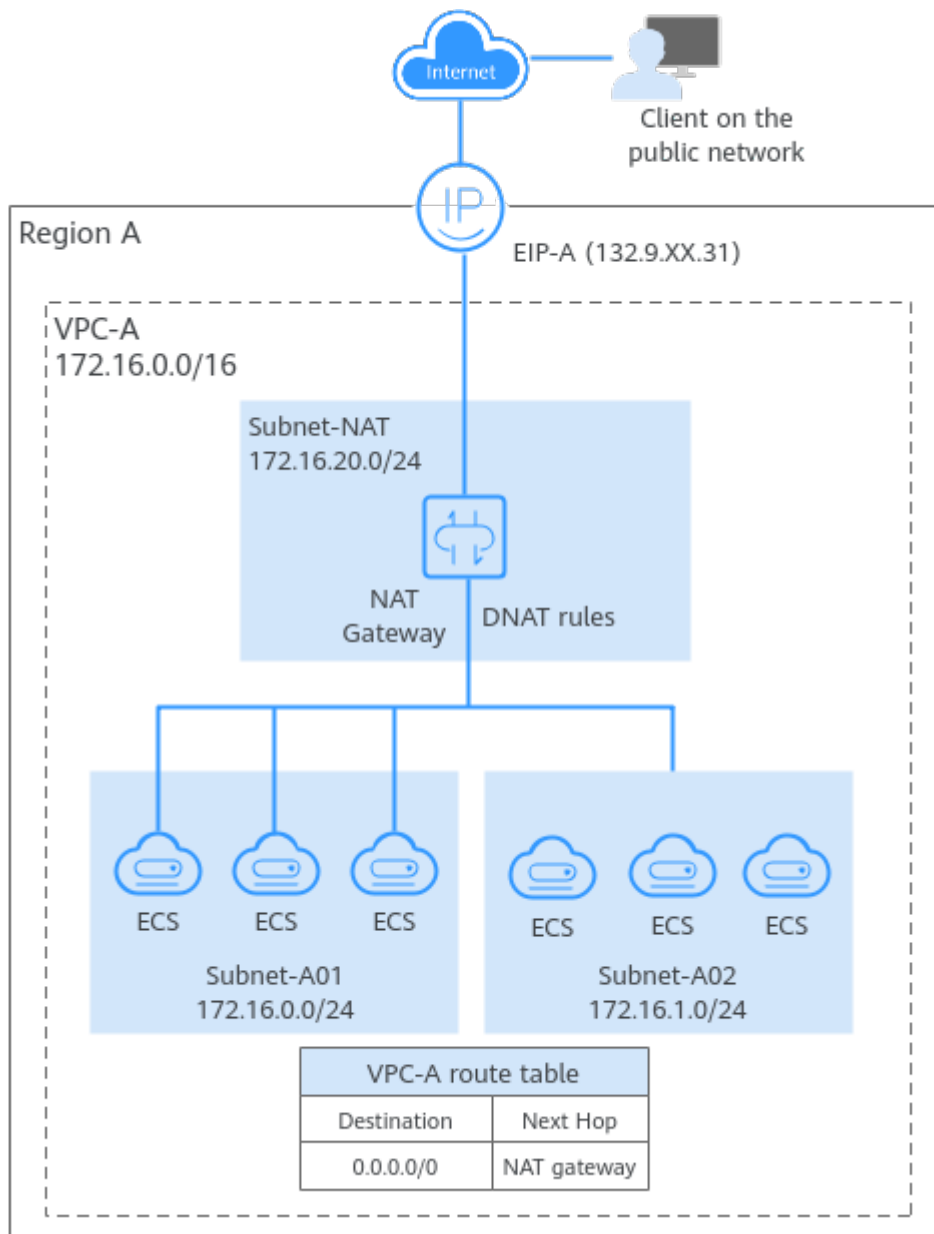
NAT Gateway (DNAT)

DNAT enables port forwarding. It maps EIP ports to ECS ports so that the ECSs in VPCs can share the same EIP and bandwidth to provide Internet-accessible services. However, DNAT does not balance traffic.

For details, see [Using DNAT to Provide Services for the Public Network](#).

In [Figure 2-15](#), ECSs deployed on two subnets (Subnet-A01 and Subnet-A02) in a VPC (VPC-A) need to provide web services for the public network. For this to work, you first need to create a public NAT gateway in a third subnet (Subnet-NAT in this example), and then configure DNAT rules on the public NAT gateway for Subnet-A01 and Subnet-A02. In this way, all ECSs in Subnet-A01 and Subnet-A02 can share an EIP to provide Internet-accessible services.

Figure 2-15 Enabling ECSs in a VPC to provide services for the public network using a NAT gateway



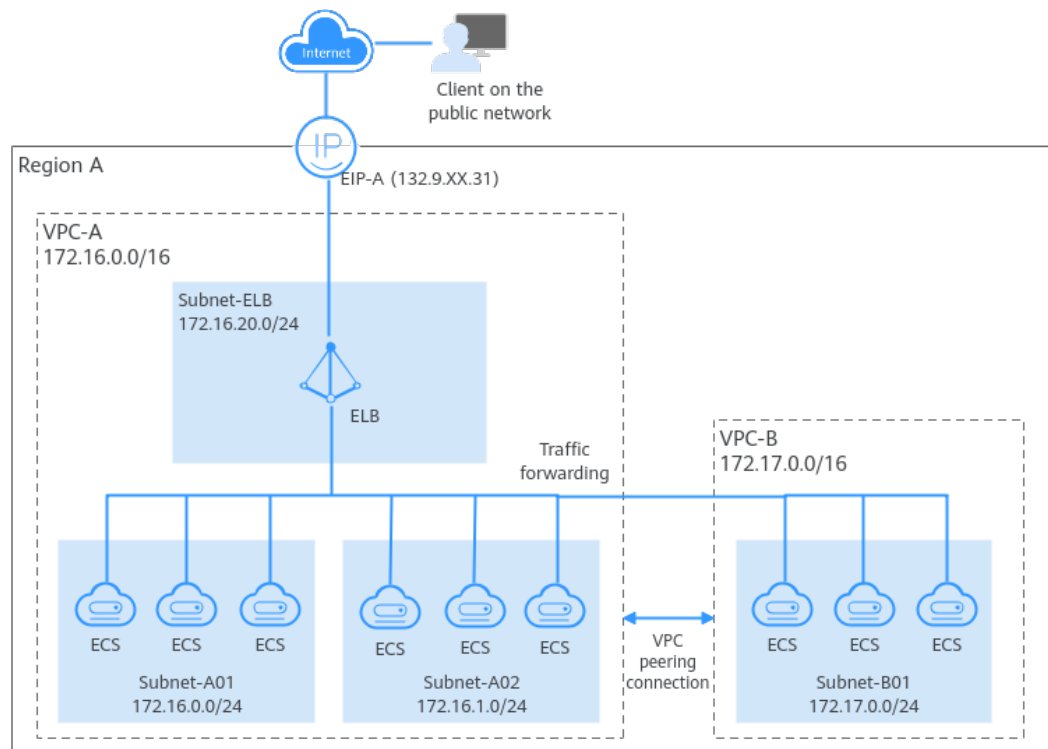
ELB

ELB evenly distributes incoming traffic to multiple backend servers. Together with EIPs, ELB allows a large number of users to access services deployed on cloud servers from the public network.

For details, see [Getting Started with ELB](#).

In [Figure 2-16](#), a web application is deployed on the ECSs in two VPCs (VPC-A and VPC-B) in a region. Because of the heavy incoming traffic, a load balancer is used to distribute the traffic across ECSs in different VPCs. For this to work, VPCs need to communicate with each other. In this example, a VPC peering connection is used to connect VPC-A and VPC-B.

Figure 2-16 ELB for evenly distributing incoming traffic from the public network



2.2.4 Connecting VPCs to On-Premises Data Centers

Connecting a Single VPC to an On-Premises Data Center

You can use Direct Connect or VPN to connect a VPC to an on-premises data center.

[Connecting VPCs to an On-Premises Data Center](#) provides details about different network services.

NOTICE

Before connecting a VPC to an on-premises data center, you need to plan their CIDR blocks in advance to ensure that the VPC CIDR block does not overlap with the on-premises CIDR block, or communications may fail.

VPN

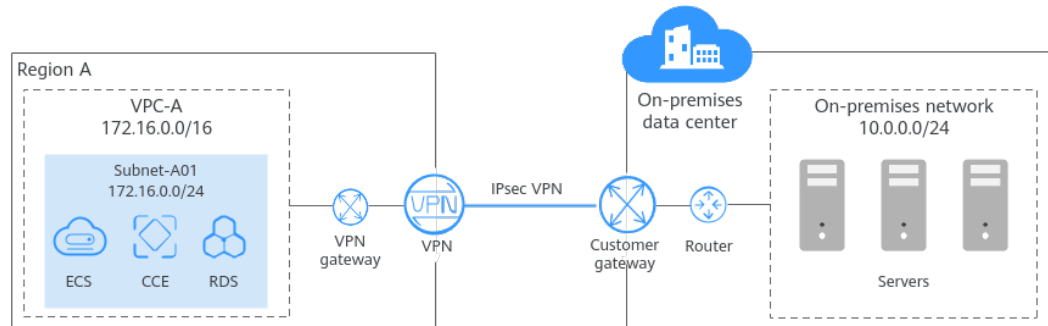
VPN provides an encrypted, Internet-based channel that connects an on-premises data center and the cloud.

For details, see [Configuring Enterprise Edition S2C VPN to Connect an On-premises Data Center to a VPC](#).

In [Figure 2-17](#), some workloads have been migrated to a VPC (VPC-A), and some workloads are still running on on-premises servers. With a VPN connection, on-

premises servers can quickly access the cloud resources in the VPC. Compared with Direct Connect, VPN is easier to configure and cost-effective.

Figure 2-17 Connecting a VPC to an on-premises data center using VPN



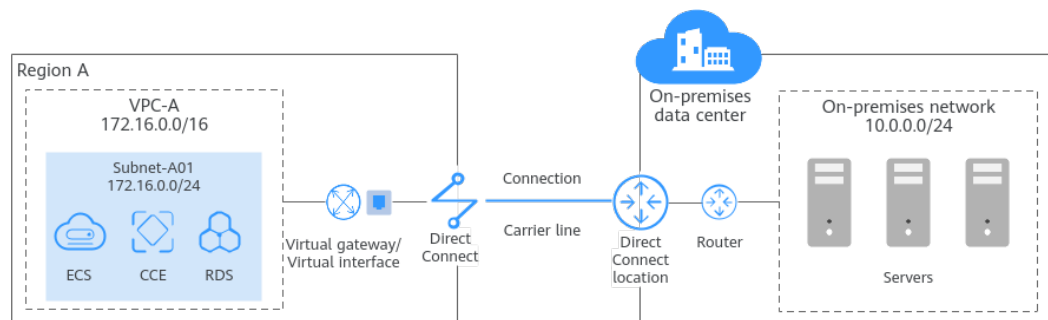
Direct Connect

Direct Connect establishes a dedicated network connection between an on-premises data center and the cloud.

For details, see [Accessing a VPC over a Direct Connect Connection and Using BGP to Route Traffic](#).

In [Figure 2-18](#), some workloads are running in a VPC (VPC-A) on the cloud, and some are running in the on-premises data center. A Direct Connect connection connects the on-premises data center to the cloud. Direct Connect connections are faster and more stable than VPN connections.

Figure 2-18 Connecting a VPC to an on-premises data center using Direct Connect



Connecting Multiple VPCs in the Same Region to an On-Premises Data Center

To connect multiple VPCs in a region to an on-premises data center, you can use Direct Connect or VPN to connect the data center to a VPC, and then use VPC Peering or Enterprise Router to connect all VPCs. In this way, the on-premises data center can access all the VPCs.

Compared with VPN, Direct Connect establishes a dedicated connection that enables faster, more secure data transmission. VPN is more cost-effective. To reduce network costs, you can use VPN instead of Direct Connect. [Connecting VPCs to an On-Premises Data Center](#) provides details about different network services.

NOTICE

To connect VPCs to an on-premises data center, you need to plan their CIDR blocks in advance. Note the following:

- Ensure that the VPC CIDR blocks do not overlap with the on-premises CIDR block, or communications may fail.
- Ensure that the VPC CIDR blocks do not overlap, or communications may fail.

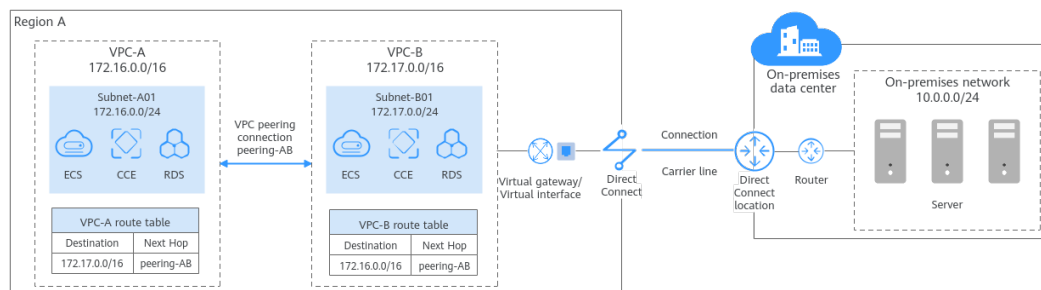
VPC Peering

With VPC Peering, you can peer two VPCs in the same region, no matter whether they are in the same account or different accounts. VPC Peering can work with Direct Connect or VPN to enable your on-premises data center to access multiple VPCs.

For details, see [Connecting an On-Premises Data Center to Multiple VPCs that Need to Communicate with Each Other](#).

In [Figure 2-19](#), some workloads are running in two VPCs (VPC-A and VPC-B) in a region, and some workloads are running in the on-premises data center. The on-premises data center connects to a VPC (VPC-B) over a Direct Connect connection, and VPC-A and VPC-B are connected over a VPC peering connection. In this way, the on-premises data center can access both VPC-A and VPC-B.

Figure 2-19 Connecting an on-premises data center to VPCs using Direct Connect and VPC Peering



2.3 VPC

2.3.1 Creating a VPC and Subnet

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 2-20 Creating a VPC and subnet

Basic Information

Region: EU-Dublin

Name: vpc-19e3

IPv4 CIDR Block: 192.168.0.0 / 16

Recommended: 10.0.0.0/8-24 (Select) 172.16.0.0/12-24 (Select) 192.168.0.0/16-24 (Select)

⚠ The CIDR block 192.168.0.0/16 overlaps with a CIDR block of another VPC in the current region. If you intend to enable communication between VPCs or between a VPC and an on-premises data center, change the CIDR block. View VPC CIDR blocks in current region

Enterprise Project: --Select--

Advanced Settings: Tag | Description

Default Subnet

AZ: AZ1

Name: subnet-1a18

IPv4 CIDR Block: 192.168.0.0 / 24

Available IP Addresses: 251
The CIDR block cannot be modified after the subnet has been created.


IPv6 CIDR Block: Enable

Associated Route Table: Default

Create Now

Table 2-6 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.	EU-Dublin
Name	The VPC name. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	vpc-test

Parameter	Description	Example Value
IPv4 CIDR Block	<p>The CIDR block of the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset of the CIDR block for the VPC (for multiple subnets in the VPC).</p> <p>The following CIDR blocks are supported:</p> <ul style="list-style-type: none"> • 10.0.0.0/8-24 • 172.16.0.0/12-24 • 192.168.0.0/16-24 <p>This parameter will be CIDR Block in regions where IPv4/IPv6 dual stack is not supported, and IPv4 CIDR Block if IPv4/IPv6 dual stack is supported.</p>	10.0.0.0/8
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 default.</p> <p>For details about creating and managing enterprise projects, see the Enterprise Management User Guide.</p>	default
Advanced Settings (Optional) > 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 Managing VPC Tags.</p>	<ul style="list-style-type: none"> • Key: vpc_key1 • Value: vpc-01


Parameter	Description	Example Value
Advanced Settings (Optional) > 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 (< or >).</p>	N/A


Table 2-7 Subnet parameter descriptions


Parameter	Description	Example Value
Name	<p>The subnet name. The name:</p> <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	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">• 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.• 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. <p>For details, see Region and AZ.</p>	AZ1



Parameter	Description	Example Value
IPv4 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>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"> • Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to properly plan the CIDR block in advance based on the number of IP addresses required by your service. <ul style="list-style-type: none"> - The subnet CIDR block size cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address. - The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks available later for new subnets, which can be a problem when you want to scale out services. • Avoiding subnet CIDR block conflicts: Avoid CIDR block conflicts if you need to 	10.0.0.0/24


Parameter	Description	Example Value
	<p>connect two VPCs or connect a VPC to an on-premises data center. If the subnet CIDR blocks at both ends of the network conflict, create a subnet.</p> <p>For details about subnet planning, see VPC and Subnet Planning.</p>	
IPv6 CIDR Block (Optional)	<p>The IPv6 CIDR block of the subnet. 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 IPv4 and IPv6 Dual-Stack Network.</p>	-

Parameter	Description	Example Value
Associated Route Table	<p>The default route table with which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. 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 Creating a Custom Route Table.</p>	-
Advanced Settings (Optional) > 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

Parameter	Description	Example Value
Advanced Settings (Optional) > DNS Server Address	<p>The DNS server addresses.</p> <p>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 Reset 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 (Optional) > 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>	test.com

Parameter	Description	Example Value
Advanced Settings (Optional) > IPv4 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">• Limited: Set the DHCP lease time. The unit can be day or hour.• Unlimited: The DHCP lease time does not expire. <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>	-
Advanced Settings (Optional) > 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 Managing Subnet Tags.</p>	<ul style="list-style-type: none">• Key: subnet_key1• Value: subnet-01

Parameter	Description	Example Value
Advanced Settings (Optional)> 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 (< or >).</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](#).

2.3.2 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 section "Updating VPC Information" in the *Virtual Private Cloud API Reference*.

Notes and Constraints

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

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

- The allowed secondary CIDR block size is between a /28 netmask and /3 netmask.
- **Table 2-8** lists the secondary CIDR blocks that are not supported.

Table 2-8 Restricted secondary CIDR blocks

Type	CIDR Block (Not Supported)
Primary CIDR blocks and existing CIDR blocks	<ul style="list-style-type: none">• 10.0.0.0/8• 172.16.0.0/12• 192.168.0.0/16
Reserved system CIDR blocks	<ul style="list-style-type: none">• 100.64.0.0/10• 214.0.0.0/7• 198.18.0.0/15• 169.254.0.0/16
Reserved public CIDR blocks	<ul style="list-style-type: none">• 0.0.0.0/8• 127.0.0.0/8• 240.0.0.0/4• 255.255.255.255/32

Procedure


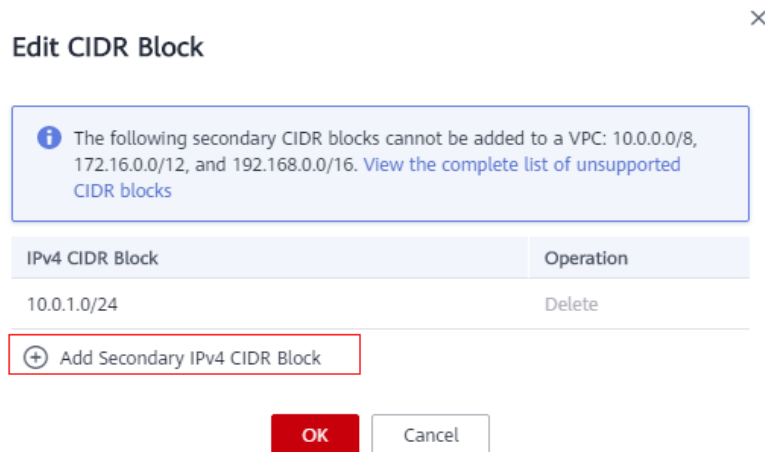
1. Log in to the management console.
2. Click  in the upper left corner and choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the VPC list, locate the row that contains the VPC and click **Edit CIDR Block** in the **Operation** column.
The **Edit CIDR Block** dialog box is displayed.
4. Click **Add Secondary IPv4 CIDR Block**.

Figure 2-21 Add Secondary IPv4 CIDR Block



5. Enter the secondary CIDR block and click **OK**.



2.3.3 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 2-22 VPC ID



2.3.4 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 section "Updating VPC Information" in the *Virtual Private Cloud API Reference*.



Modifying the Name and Description of a VPC

1. Log in to the management console.
2. Click  in the upper left corner and 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 a VPC name and click **OK**.
 - Method 2:
 - i. In the VPC list, locate the target VPC and click its name.
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**.

2.3.5 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 2-9](#) lists the tag key and value requirements.



Table 2-9 VPC tag key and value requirements

Parameter	Requirements	Example Value
Key	<ul style="list-style-type: none">• Cannot be left blank.• Must be unique for each VPC and can be the same for different VPCs.• Can contain a maximum of 36 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including hyphens (-) and underscores (_)	vpc_key1



Parameter	Requirements	Example Value
Value	<ul style="list-style-type: none">• Can contain a maximum of 43 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including periods (.), hyphens (-) and underscores (_)	vpc-01

Procedure

Search for VPCs by tag key or value on the VPC 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 upper right corner of the VPC list, click **Search by Tag**.
5. In the displayed area, enter the tag key and value of the VPC you are looking for.
Both the tag key and value must be specified. The system automatically displays the VPCs you are looking for if both the tag key and value are matched.
6. Click + to add more tag keys and values.
You can add multiple tag keys and values to refine your search results. If you add more than one tag to search for VPCs, the VPCs containing all specified tags will be displayed.
7. Click **Search**.
The system displays the VPCs you are looking for based on the entered tag keys and values.

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 **OK**.

2.3.6 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.



2.3.7 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.


2.3.8 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**.

2.3.9 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.

You can refer to [Why Can't I Delete My VPCs and Subnets?](#)

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and choose **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. Confirm the information and click **Yes**.

NOTICE

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

2.4 Subnet

2.4.1 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 2-10 Subnet parameter descriptions

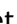

Parameter	Description	Example Value
Region	The region where VPC is located.	EU-Dublin
VPC	The VPC for which you want to create a subnet.	vpc-test
Subnet Name	The subnet name. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	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">• 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.• 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. <p>For details, see Region and AZ.</p>	AZ1

Parameter	Description	Example Value
IPv4 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>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">● Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to properly plan the CIDR block in advance based on the number of IP addresses required by your service.<ul style="list-style-type: none">– The subnet CIDR block size cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address.– The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks available later for new subnets, which can be a problem when you want to scale out services.● Avoid CIDR block conflicts if you need to connect two VPCs or connect a VPC to an on-premises data center. If the subnet CIDR blocks at both ends of the network conflict, create a subnet. <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

Parameter	Description	Example Value
IPv6 CIDR Block (Optional)	<p>The IPv6 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 IPv4 and IPv6 Dual-Stack Network.</p>	N/A
Associated Route Table	<p>The default route table with which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. 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 Creating a Custom Route Table.</p>	N/A
Advanced Settings (Optional) > 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

Parameter	Description	Example Value
Advanced Settings (Optional) > 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 Reset 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
Advanced Settings (Optional) > IPv4 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">● Limited: Set the DHCP lease time. The unit can be day or hour.● Unlimited: The DHCP lease time does not expire. <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>	N/A

Parameter	Description	Example Value
Advanced Settings (Optional) > Tag	The subnet tag. Click  to expand the configuration area and set this parameter. Add tags to help you quickly identify, classify, and search for your subnets. For details, see Managing Subnet Tags .	<ul style="list-style-type: none">• Key: subnet_key1• Value: subnet-01
Advanced Settings > Description	Supplementary information about the subnet. Click  to expand the configuration area and set this parameter. Enter the description about the subnet in the text box as required. The subnet description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

7. Click **Create Now**.
Return to the subnet list and view the new subnet.

2.4.2 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 2-11 Parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	Subnet
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 Reset 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
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">• Limited: Set the DHCP lease time. The unit can be day or hour.• Unlimited: The DHCP lease time does not expire. <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>	-
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 (< or >).</p>	-

2.4.3 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 2-12](#) lists the tag key and value requirements.

Table 2-12 Subnet tag key and value requirements

Parameter	Requirements	Example Value
Key	<ul style="list-style-type: none">• Cannot be left blank.• Must be unique for each subnet.• Can contain a maximum of 36 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including hyphens (-) and underscores (_)	subnet_key1
Value	<ul style="list-style-type: none">• Can contain a maximum of 43 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including hyphens (-) and underscores (_)	subnet-01

Procedure

Search for subnets by tag key or value on the subnet 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 > Subnets**.
The **Subnets** page is displayed.
5. In the upper right corner of the subnet list, click **Search by Tag**.
6. Enter the tag key of the subnet to be queried.
Both the tag key and value must be specified. The system automatically displays the subnets you are looking for if both the tag key and value are matched.
7. Click **+** to add another tag key and value.

You can add multiple tag keys and values to refine your search results. If you add more than one tag to search for subnets, the subnets containing all specified tags will be displayed.

8. Click **Search**.

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

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 **OK**.

2.4.4 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.

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

2.4.5 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

- 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.
- On the **Summary** page, view the resources in the subnet.
 - In the **Resources** area, view the quantities of resources, such as ECSs, BMSs, network interfaces, and load balancers, in the subnet. Click the number to the right of each resource to view the resources in the subnet.
 - In the **Networking Components** area on the right of the page, view the NAT gateways in the subnet.


Figure 2-23 Viewing resources in a subnet



7. Delete resources from the subnet.

Table 2-13 Viewing and deleting resources in a subnet

Resource	Reference
ECS	<p>You cannot jump to the target ECS from the current page. To delete an ECS from the subnet, you need to go to the ECS console, search for the target ECS in the ECS list, and delete it.</p> <ol style="list-style-type: none"> 1. In the ECS list, click the ECS name. The ECS details page is displayed. 2. In the NICs area on the Summary page, view the name of the subnet associated with the ECS. 3. Confirm the information and delete the ECS.
BMS	<p>You cannot jump to the target BMS from the current page. To delete an BMS from the subnet, you need to go to the BMS console, search for the target BMS in the BMS list, and delete it.</p> <ol style="list-style-type: none"> 1. In the BMS list, click the BMS name. The BMS details page is displayed. 2. In the NICs area on the Summary page, view the name of the subnet associated with the BMS. 3. Confirm the information and release the BMS.
Load balancer	<p>You can directly jump to the target load balancer page.</p> <ol style="list-style-type: none"> 1. Click the number to the right of Load Balancers. The load balancer list is displayed. 2. Confirm the load balancer that you want to delete and click Delete in the Operation column. For details, see Deleting a Load Balancer.

Resource	Reference
Network interface	<p>You can directly jump to the target network interface page.</p> <ol style="list-style-type: none">1. Click the number to the right of Network Interfaces. The Network Interfaces page is displayed.2. Confirm the network interface that you want to delete and choose More > Delete in the Operation column. For details, see Deleting a Network Interface.
NAT gateway	<p>You can directly jump to the target NAT gateway page.</p> <ol style="list-style-type: none">1. Click the NAT gateway name in the Networking Components area. The NAT gateway details page is displayed.2. Click  to return to the NAT gateway list.3. Locate the row that contains the NAT gateway and click Delete in the Operation column.<ul style="list-style-type: none">• Deleting a Public NAT Gateway

2.4.6 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.

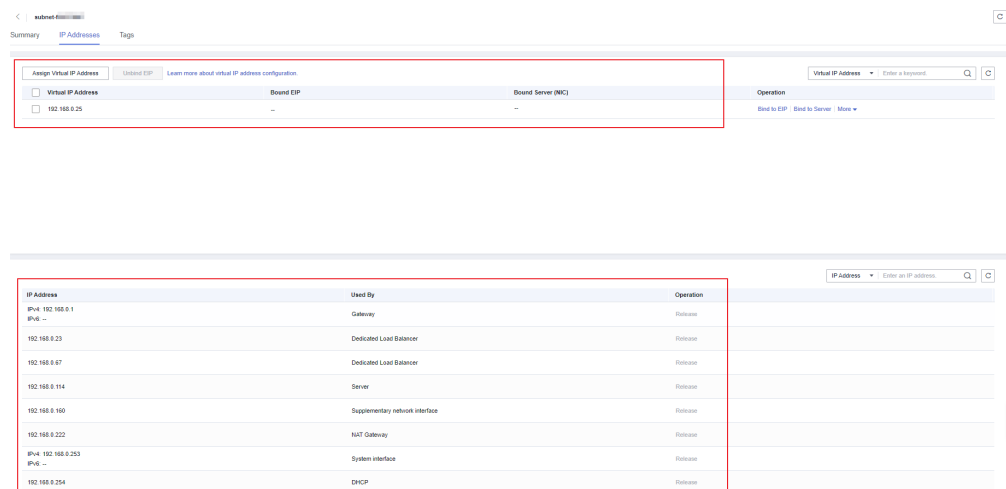
Notes and Constraints

- 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

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. Click the **IP Addresses** tab to view the IP addresses in the subnet.
 - a. In the virtual IP address list, you can view the virtual IP addresses assigned from the subnet.
 - b. In the private IP address list in the lower part of the page, you can view the private IP addresses and the resources that use the IP addresses of the subnet.

Figure 2-24 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?](#)

2.4.7 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.

You can refer to [Why Can't I Delete My VPCs and Subnets?](#)

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and 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 row that contains the subnet you want to delete and click **Delete** in the **Operation** column.
A confirmation dialog box is displayed.
6. Click **Yes**.

NOTICE

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

3 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 in and out of 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:** Each VPC comes with a default route table. If you create a subnet in the VPC, the subnet associates with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.

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

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

NOTE

By default, the quota for custom route tables is 0. To create custom route tables, [apply for a quota increase first](#).

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">• Default route table• Custom route table
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">• Default route table• Custom route table
BMS user-defined network	Traffic intended for the destination is forwarded to a BMS user-defined network.	<ul style="list-style-type: none">• Custom route table
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.	Custom route table
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	<ul style="list-style-type: none">• Default route table• Custom route table
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.	<ul style="list-style-type: none">• Default route table• Custom route table
Virtual IP address	Traffic intended for the destination is forwarded to a virtual IP address and then sent to active and standby ECSs to which the virtual IP address is bound.	<ul style="list-style-type: none">• Default route table• Custom route table
VPC endpoint	Traffic intended for the destination is forwarded to a VPC endpoint.	<ul style="list-style-type: none">• Default route table• Custom route table
Cloud container	Traffic intended for the destination is forwarded to a cloud container.	<ul style="list-style-type: none">• Default route table• Custom route table
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.	<ul style="list-style-type: none">• Default route table• Custom route table

Next Hop Type	Description	Supported Route Table
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.	<ul style="list-style-type: none"> • Default route table • Custom route table

 **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 added by the system 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.

Custom Route Table Configuration Process

Figure 3-2 Process for configuring a route table

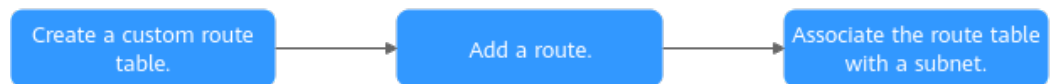


Table 3-2 Process for configuring a route table

No.	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.	Creating a Custom Route Table
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.	Adding Routes to a Route Table
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.	Associating a Route Table with a Subnet

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, the quota for custom route tables is 0. To create custom route tables, [apply for a quota increase first](#).

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">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	rtb-001
VPC	(Mandatory) The VPC that the route table is used to control traffic routing. The route table can be associated with the subnets in this VPC.	vpc-001
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 Adding Routes to a Route Table . 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. If you want to associate immediately, perform the following operations:

- a. Click **Associate Subnet**. The **Associated Subnets** 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 **Networking Components** area of the **Summary** 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 route 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 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 IP address . You can set an IP address or CIDR block.	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. NOTICE <ul style="list-style-type: none">The destination of each route in a route table must be unique.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.	IPv4: 192.168.0.0/16
Next Hop Type	Mandatory Set the type of the next hop. NOTE When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to VPN gateway , Direct Connect gateway , or Cloud connection .	VPC peering connection
Next Hop	Mandatory Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.	peer-AB
Description	Optional Enter the description of the route in the text box as required.	-

7. Click **OK**.

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 VPC endpoint, VPN or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.

Procedure



1. Log in to the management console.
2. Click  in the upper left corner and 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 IP address . You can set an IP address or CIDR block.	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. NOTICE <ul style="list-style-type: none">• The destination of each route in a route table must be unique.• 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.	IPv4: 192.168.0.0/16

Parameter	Description	Example Value
Next Hop Type	Mandatory Set the type of the next hop. NOTE When you add or modify a custom route in a default route table, the next hop type of the route cannot be set to VPN gateway , Direct Connect gateway , or Cloud connection .	VPC peering connection
Next Hop	Mandatory Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.	peer-AB
Description	Optional Enter the description of the route in the text box as required.	-

- 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 the Default Route Table	Can Be Replicated to a Custom Route Table
Local	No	No
Server	Yes	Yes
Extension NIC	Yes	Yes

Next Hop Type	Can Be Replicated to the Default Route Table	Can Be Replicated to a Custom Route Table
BMS user-defined network	No	Yes
VPN gateway	No	Yes
Direct Connect gateway	No	Yes
Cloud connection	No	Yes
Supplementary network interface	Yes	Yes
NAT gateway	Yes	Yes
VPC peering connection	Yes	Yes
Virtual IP address	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.
- The routes automatically delivered by VPN or Direct Connect to the default route table cannot be deleted. The next hop types of such routes are:
 - VPN gateway
 - Direct Connect gateway

To delete these routes, you need to delete the associated network instances first.

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. 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 go to the page showing ECS details.
5. On the displayed page, click the **Network Interfaces** tab.
6. Click the NIC IP address to view details and 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 [Binding an EIP to an Instance](#).
 - Bind an EIP to the virtual IP address of the ECS. For details, see [Binding a Virtual IP Address to an Instance or EIP](#).
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 allow all ECSs in the subnet (for example, 192.168.1.0/24) to access the Internet:

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

Figure 3-3 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 3-4](#) (for example, 192.168.1.0/24) is displayed, the configuration was successful.

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

Figure 3-4 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

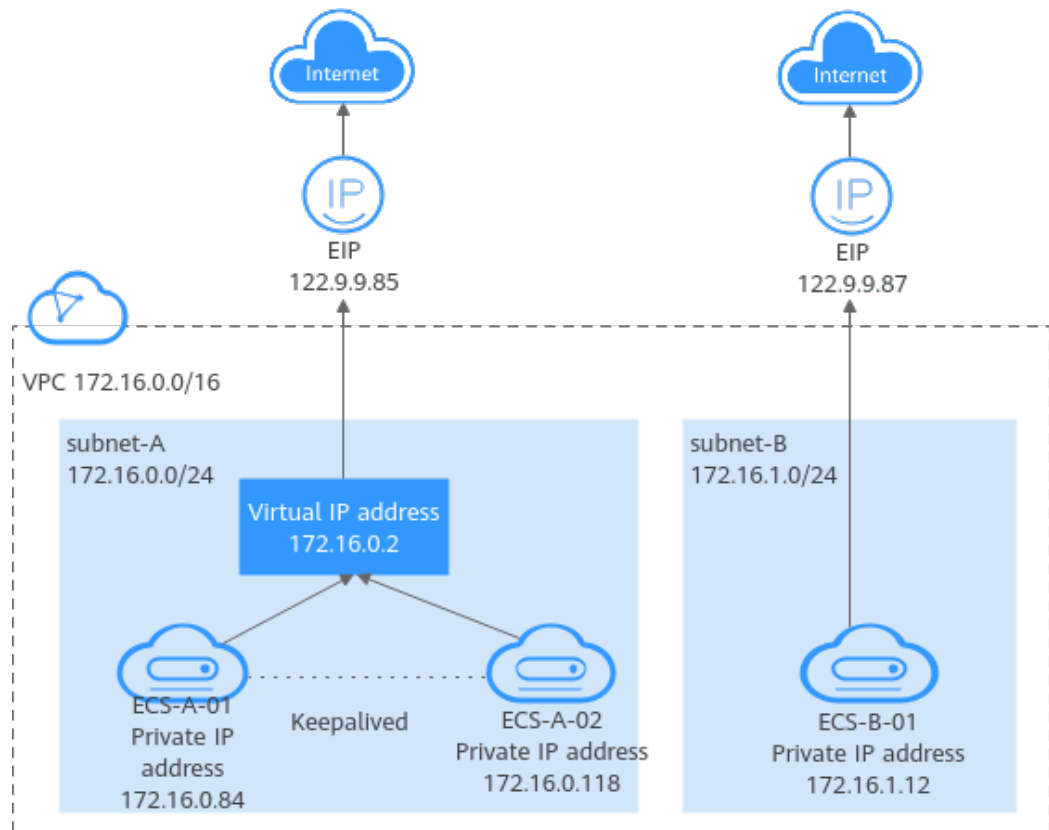
What Is a Virtual IP Address?

A virtual IP address is a private IP address that can be independently assigned from and released to a VPC subnet. You can:

- Bind one or more virtual IP addresses to an ECS so that you can use either the virtual IP address or private IP address to access the ECS. If you have multiple services running on an ECS, you can use different virtual IP addresses to access them.
- Bind a virtual IP address to multiple ECSs. You can use a virtual IP address and an HA software (such as Keepalived) to set up a high-availability active/standby cluster. If you want to improve service availability and eliminate single points of failure, you can deploy ECSs in the active/standby pair or deploy one active ECS and multiple standby ECSs. In this case, the ECSs can use the same virtual IP address. If the active ECS goes down, the standby ECS becomes the active ECS and continues to provide services.

Generally, ECSs use private IP addresses for internal network communication. A virtual IP address has the same network access capabilities as a private IP address. You can use either of them to enable layer 2 and layer 3 communications in a VPC, access a different VPC using a peering connection, enable Internet access through EIPs, and connect the cloud and the on-premises servers using VPN connections and Direct Connect connections. [Figure 4-1](#) describes how private IP addresses, the virtual IP address, and EIPs work together.

- Private IP addresses are used for internal network communication.
- The virtual IP address works with Keepalived to build an HA cluster. ECSs in this cluster can be accessed through one virtual IP address.
- EIPs are used for Internet communication.

Figure 4-1 Different types of IP addresses used by ECSs

Application Scenarios

You can use a virtual IP address and Keepalived to set up a high-availability active/standby cluster. If the active ECS goes down, the standby ECS becomes the active ECS and continues to provide services. The following describes the typical application scenarios of virtual IP addresses.

Using a Virtual IP Address and Keepalived to Set Up a High-Availability Cluster

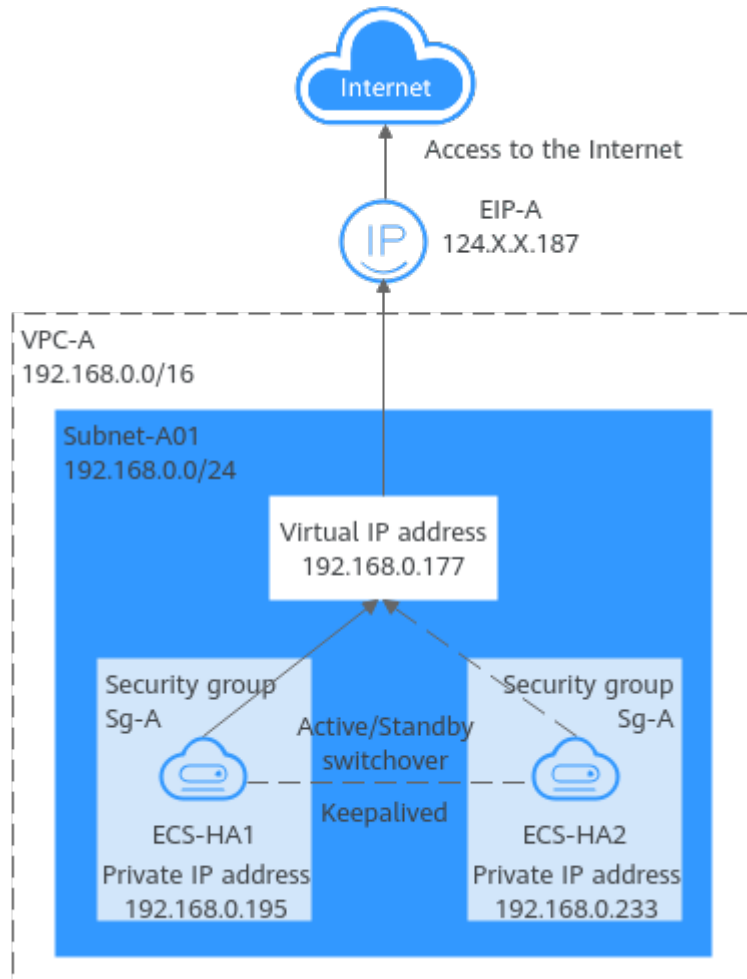
Figure 4-2 shows a high-availability cluster that is set up using a virtual IP address and Keepalived. They work as follows:

1. Virtual IP address **192.168.0.177** is bound to **ECS-HA1** and **ECS-HA2**. Keepalived is configured on the two ECSs.
2. EIP **EIP-A** is bound to the virtual IP address so that the ECSs can be accessed from the Internet.

In this cluster, **ECS-HA1** works as the active ECS and provides services accessible from the Internet using **EIP-A**. **ECS-HA2** works as the standby ECS, with no services deployed on it. If **ECS-HA1** goes down, **ECS-HA2** takes over services, ensuring service continuity.

For details about how to set up an HA cluster, see [Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster](#).

Figure 4-2 A high-availability cluster using a virtual IP address and Keepalived



Using a Virtual IP Address and Keepalived/LVS to Set Up a High-Availability Load Balancing Cluster

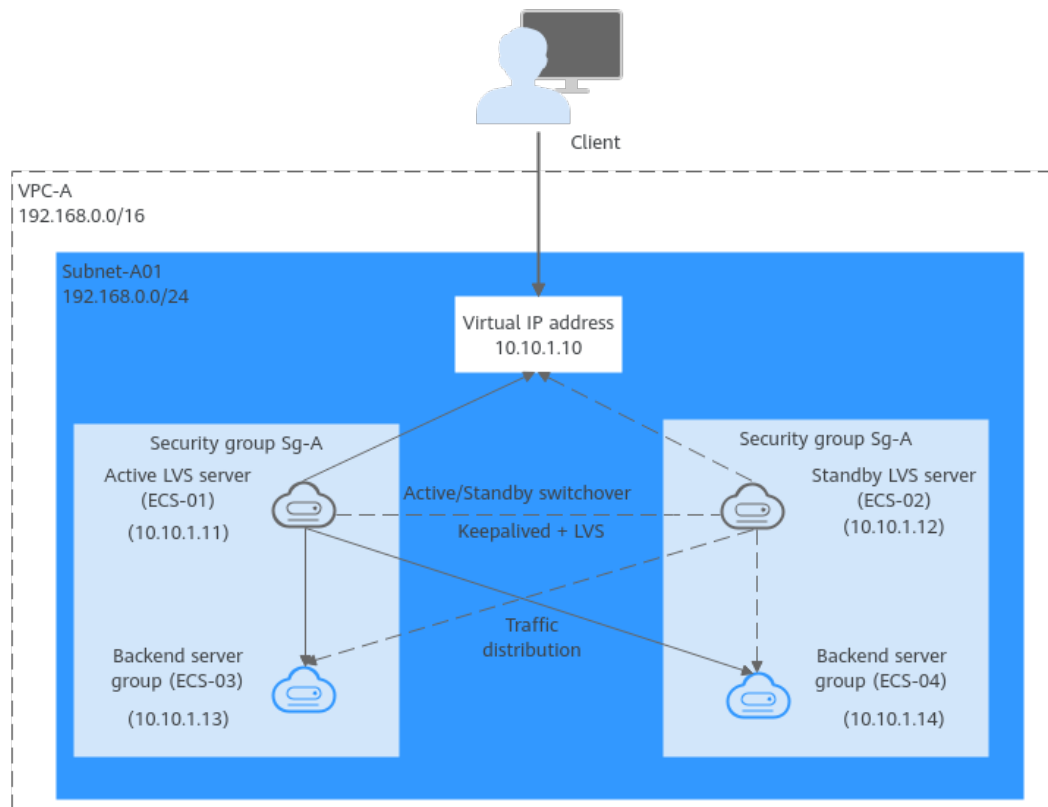
As shown in [Figure 4-3](#), a virtual IP address, Keepalived, and LVS are used to set up an HA load balancing cluster. LVS is used for load balancing, and Keepalived is used for high availability. They work as follows:

1. Virtual IP address **10.10.1.10** is bound to **ECS-01** and **ECS-02**. Keepalived and LVS (DR mode) are configured on **ECS-01** and **ECS-02** to set up the active/standby LVS servers. In this way, requests from clients can be evenly distributed to different backend servers.
2. **ECS-03** and **ECS-04** are configured as backend servers to handle service requests.
3. The source/destination check option needs to be disabled.

When you bind a virtual IP address to an ECS, the source/destination check option of the ECS NIC is automatically disabled. If the option is not disabled, disable it.

In this load balancing cluster, **ECS-01** works as the active LVS server to distribute requests from clients. If **ECS-01** is faulty, **ECS-02** takes over and distributes requests from clients, ensuring high availability of the LVS cluster.

Figure 4-3 A high-availability cluster using a virtual IP address and Keepalived/LVS



NOTE

For details about how to install and configure Keepalived and LVS services and how to configure backend servers, see the common practices in the industry.

Virtual IP Address Quotas

Table 4-1 lists the quotas about virtual IP addresses. Some default quotas can be increased.

Table 4-1 Virtual IP address quotas

Item	Default Quota	Adjustable
Maximum number of virtual IP addresses per region	2	Yes. For details, see Managing Quotas .
Maximum number of EIPs that a virtual IP address can be bound to.	1	No
Maximum number of instances (including ECSs and NICs) that a virtual IP address can be bound to.	10	No

Notes and Constraints



- If an ECS has multiple network interfaces that are in the same subnet, you are not advised to bind virtual IP addresses to the network interfaces. Using the virtual IP addresses may cause route conflicts on an ECS, which would lead to communication failures.
- A virtual IP address is assigned from a VPC subnet. They can only be bound to a cloud server in the same subnet as the virtual IP address.
- Virtual IP addresses and 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](#).

4.2 Assigning a Virtual IP Address

Scenarios

A virtual IP address is an IP address assigned from a VPC subnet. It can be assigned and released independently. You can follow the instructions in this section to assign a virtual IP address.

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.
The subnet details page is displayed.
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 Instance or EIP

Scenarios

You can bind a virtual IP address to an instance or EIP.



- Bind a virtual IP address to an instance. You can:

- Bind one or more virtual IP addresses to an instance.
- Bind a virtual IP address to multiple instances.
- Bind a virtual IP address to an EIP to enable public network communication.

Notes and Constraints

It is recommended that a maximum of eight virtual IP addresses be bound to an ECS. If an ECS has multiple virtual IP addresses, each virtual IP address is used by a specific service. If there are too many services, the ECS may become overloaded and compromise user experience.

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.
5. In the subnet list, locate the target subnet and click its name.
The subnet details page is displayed.
6. On the **IP Addresses** tab, bind an EIP to the virtual IP address:
 - a. Locate the row that contains the virtual IP address and click **Bind to EIP** in the **Operation** column.
The **Bind to EIP** dialog box is displayed.
 - b. Select an EIP and click **OK**.
In the virtual IP address list, you can view that the virtual IP address has an EIP bound.
7. On the **IP Addresses** tab, bind an instance to the virtual IP address:
 - a. Locate the row that contains the virtual IP address and click **Bind to Server** in the **Operation** column.
The **Bind to Server** dialog box is displayed.
 - b. Select an ECS and click **OK**.
In the virtual IP address list, you can view that the virtual IP address has an ECS bound.

NOTICE

- After you bind one or more virtual IP addresses to an ECS, you need to manually configure the virtual IP addresses on ECS. For details, see [Configuring a Virtual IP Address for an ECS](#).
 - If you want to bind a virtual IP address to multiple ECSs and use Keepalived to build an HA cluster, see [Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster](#).
-

Configuring a Virtual IP Address for an ECS

After you bind one or more virtual IP addresses to an ECS on the console, you must log in to the ECS to manually configure these virtual IP address.

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)

The following uses CentOS 7.2 64bit as an example.

1. Obtain the NIC that the virtual IP address is to be bound and the connection of the NIC:

nmcli connection

Information similar to the following is displayed:

```
[root@192.168.0.247 ~]# nmcli connection
NAME                UUID                                  TYPE      DEVICE
Wired connection 1  5e72ec5a-6165-3bd6-a34b-ce43981acb27  ethernet  eth0
docker0             cd351a91-c5eb-4b69-83eb-df892a2ccf6b  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@192.168.0.247 ~]# 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.

```
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:cd 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:a583:62c:7dd3: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)

The following uses Ubuntu 22.04 server 64bit as an example. If the 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
  ethernet:
    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**.

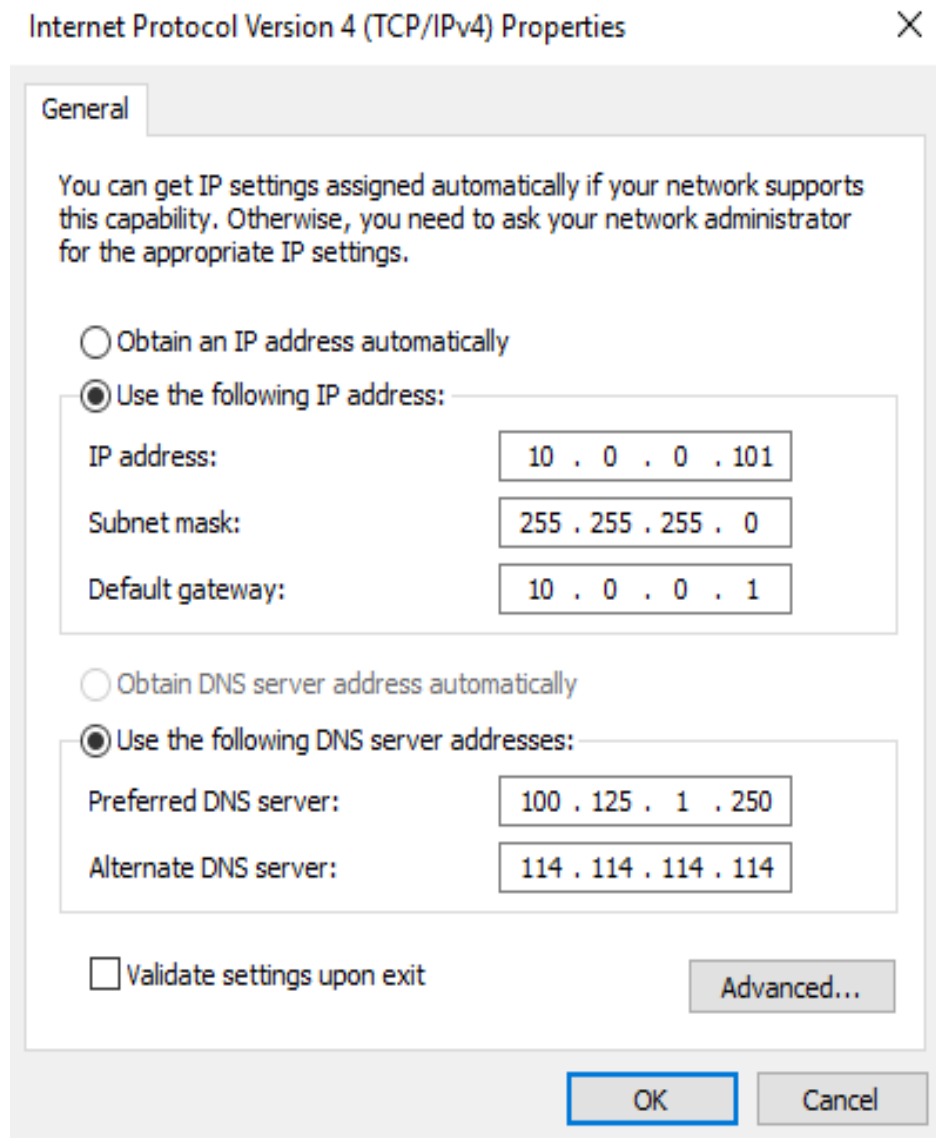
Windows OS

The following operations use Windows Server as an example.

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)**.

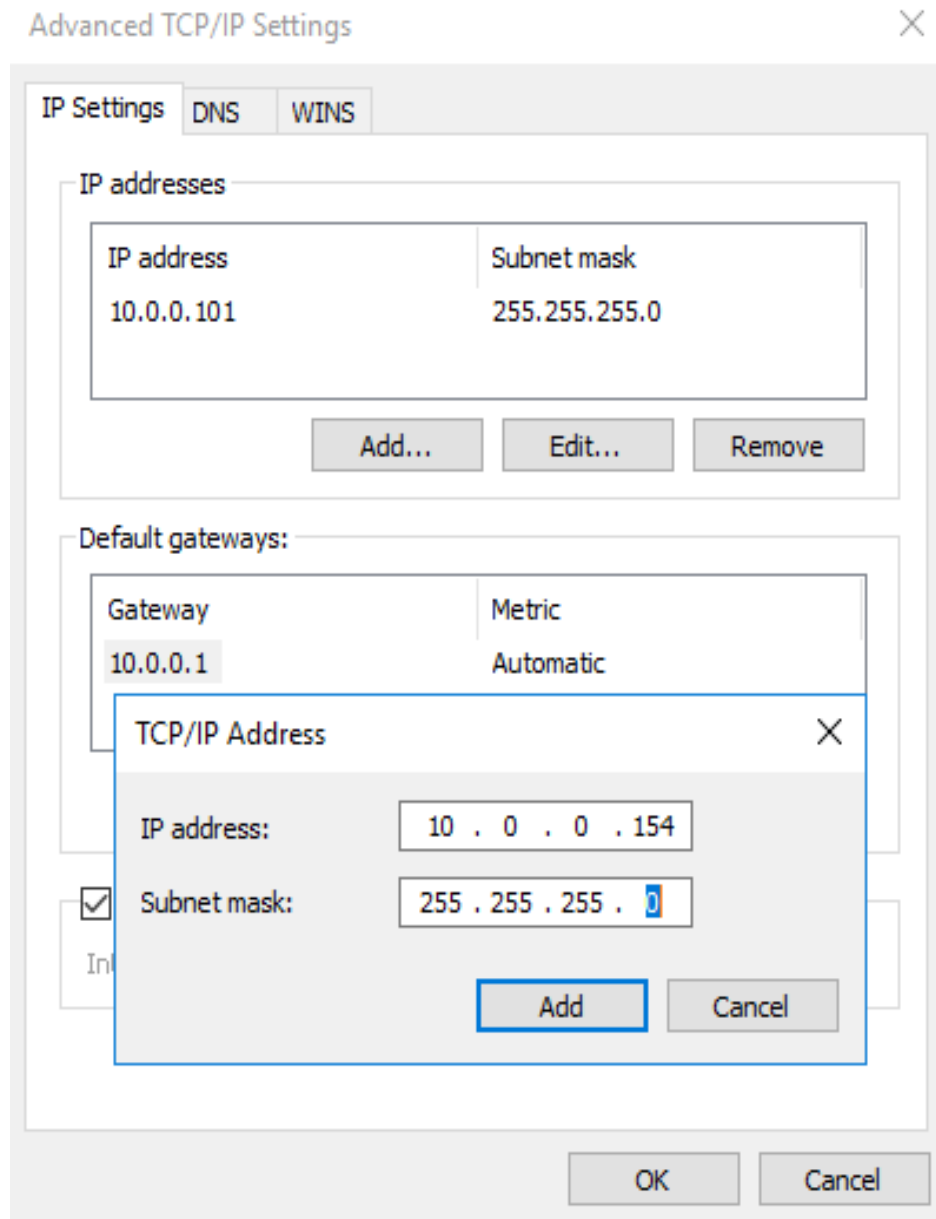
4. Click **Properties**.
5. Select **Use the following IP address** and set **IP address** to the private IP address of the ECS, for example, 10.0.0.101.

Figure 4-4 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 4-5 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?](#)



4.4 Unbinding a Virtual IP Address from an Instance or EIP

Scenarios



You can unbind a virtual IP address from an ECS or EIP:

- [Unbinding a Virtual IP Address from an Instance](#)
- [Unbinding a Virtual IP Address from an EIP](#)

Unbinding a Virtual IP Address from an Instance

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. 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**.
A confirmation dialog box is displayed.
8. In the displayed dialog box, perform the following operations to 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 **OK**.

Unbinding a Virtual IP Address 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 to display the service list 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. Click the **IP Addresses** tab.
The virtual IP address list is displayed.
7. Locate the target 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 **OK**.

4.5 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 4-2](#).

Table 4-2 Releasing a virtual IP address that is being used by a resource

Prompts	Cause Analysis and Solution
Scenario 1: This operation cannot be performed because the IP address is bound to an instance or an EIP. Unbind the IP address and try again.	This virtual IP address is being used by cloud resources such as an EIP or an ECS. For details, see Unbinding a Virtual IP Address from an Instance or EIP . Release the virtual IP address.
Scenario 2: This operation cannot be performed because the IP address is being used by a system component.	The virtual IP address is being used by an instance. Delete the instance, which will also release the virtual IP address. Search for the instance based on the instance information displayed on the virtual IP address console and delete the instance. <ul style="list-style-type: none">● RDS instance● CCE instance● API gateway

Figure 4-6 Scenario 1: Virtual IP address cannot be released**Figure 4-7** 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 **OK**.

4.6 Virtual IP Address Configuration Example

4.6.1 Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster

Scenarios

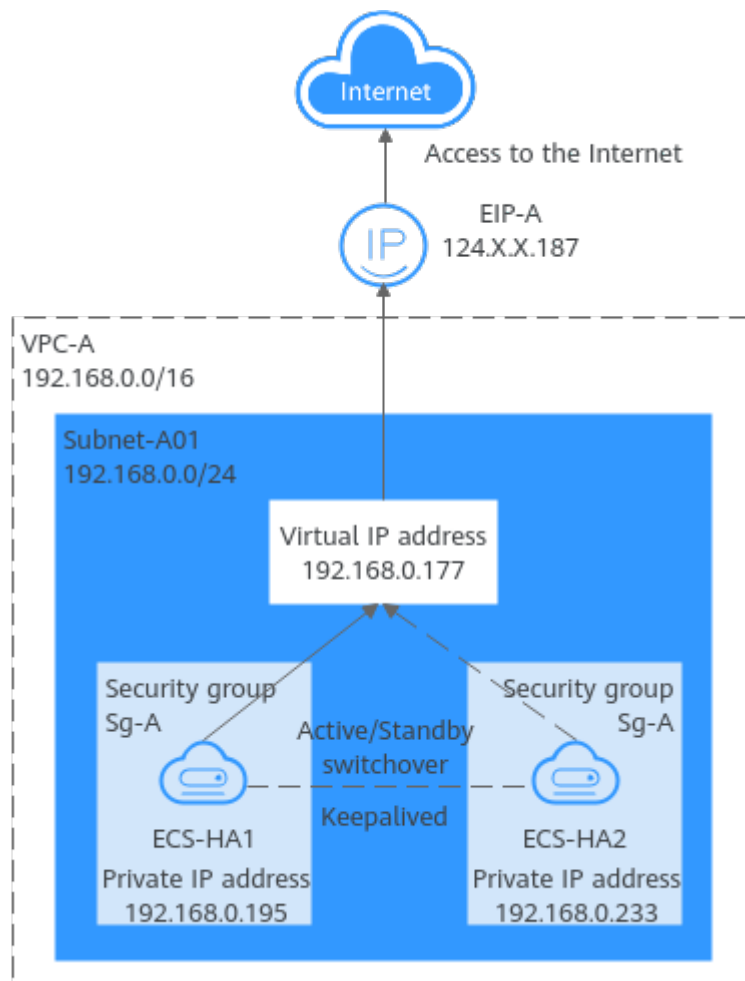
A virtual IP address is a private IP address assigned from a VPC subnet. You can use a virtual IP address and Keepalived to set up a high-availability active/standby web cluster. In such a cluster, if the active ECS goes down, the virtual IP address is bind to the standby ECS to provide services.

Architecture

Figure 4-8 shows a high-availability web cluster using Keepalived. In this architecture, virtual IP address **192.168.0.177** is bound to **ECS-HA1** and **ECS-HA2**. To allow **ECS-HA1** and **ECS-HA2** to access and be accessed from the Internet, an EIP (**EIP-A**) is bound to the virtual IP address. They work as follows:

1. **ECS-HA1** works as the active ECS and provides services accessible from the Internet using **EIP-A**. **ECS-HA2** works as the standby ECS, with no services deployed on it.
2. If **ECS-HA1** goes down, **ECS-HA2** takes over services, ensuring service continuity.

Figure 4-8 A high-availability web cluster using a virtual IP address and Keepalived



Advantages

A high-availability cluster can have one active ECS and one standby ECS or one active ECS and multiple standby ECSs. You can bind a virtual IP address to these ECSs. If the active ECS goes down, the standby ECS becomes the active ECS and continues to provide services.

Notes and Constraints

All servers of the HA cluster must be in the same subnet.

Resource Planning

In this example, the VPC, subnet, virtual IP address, EIP, and ECSs must be in the same region but can be in different AZs.

NOTE

The following resource details are only for your reference. You can modify them if needed.

Table 4-3 Resource planning

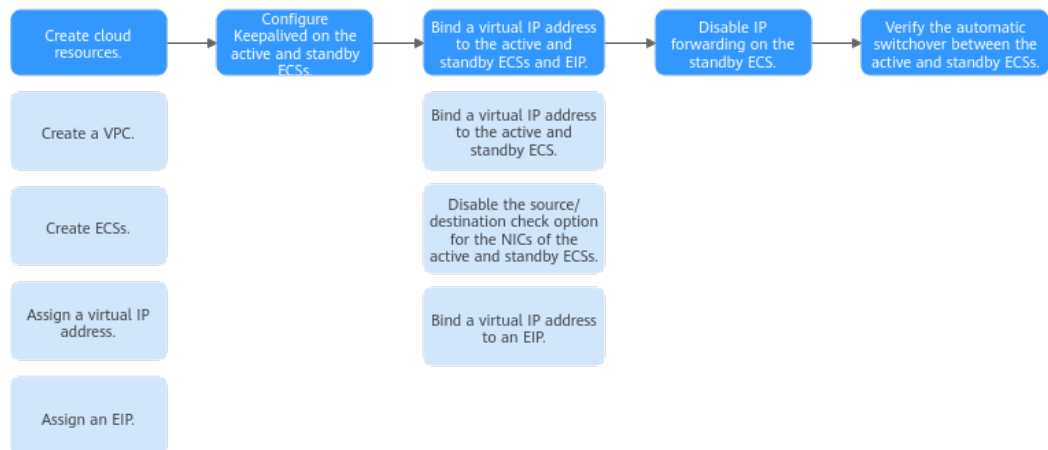
Resource Type	Quantity	Description
VPC and subnet	1	<ul style="list-style-type: none">• VPC name: Set it as needed. In this example, VPC-A is used.• VPC IPv4 CIDR block: Set it as needed. In this example, 192.168.0.0/16 is used.• Subnet name: Set it as needed. In this example, Subnet-A01 is used.• Subnet IPv4 CIDR block: Set it as needed. In this example, 192.168.0.0/24 is used.
ECS	2	<p>In this example, two ECSs are required for active/standby switchover. Configure the two ECSs as follows:</p> <ul style="list-style-type: none">• Name: Set this parameter as needed. In this example, the two ECSs are named ECS-HA1 and ECS-HA2.• Image: Select an image as needed. In this example, a public image (CentOS 7.8 64bit) is used.• System Disk: General Purpose SSD 40 GiB• Data Disk: In this example, no data disk is required. You can attach data disks based on service requirements and ensure data consistency between the two ECSs.• Network parameters<ul style="list-style-type: none">– VPC: Select a VPC. In this example, VPC-A is used.– Subnet: Select a subnet. In this example, Subnet-A01 is used.• Security Group: Select a security group as needed. In this example, ECS-HA1 and ECS-HA2 are associated with the same security group (Sg-A).• Private IP address: Specify 192.168.0.195 for ECS-HA1 and 192.168.0.233 for ECS-HA2.

Resource Type	Quantity	Description
Virtual IP address	1	Assign a virtual IP address from Subnet-A01 . <ul style="list-style-type: none"> • Assignment Mode: Set it as needed. In this example, Automatic is selected. • Virtual IP address: 192.168.0.177 is used in this example. • Instances: Bind 192.168.0.177 to ECS-HA1 and ECS-HA2. • EIP: Bind 192.168.0.177 to EIP-A.
EIP	1	<ul style="list-style-type: none"> • Billing Mode: Select a billing mode as needed. In this example, Pay-per-use is used. • EIP Name: Set it as needed. In this example, EIP-A is used. • EIP: The IP address is randomly assigned. In this example, 124.X.X.187 is used.

Procedure

You can follow the process in [Figure 4-9](#) to set up a high-availability web cluster using a virtual IP address and Keepalived

Figure 4-9 Process for setting up a high-availability web cluster



Step 1: Create Cloud Resources

1. Create a VPC and subnet.
For details, see [Creating a VPC and Subnet](#).
2. Create two ECSs, one as the active ECS and the other as the standby ECS.
For details, see [Purchasing an ECS](#).
Configure the ECSs as follows:

- **Network:** Select **VPC-A** and **Subnet-A01** you have created.
- **Security Group:** Create security group **Sg-A** and add inbound and outbound rules to it. Each security group comes with preset rules. You need to check and modify the rules as required.

Add rules in [Table 4-4](#) to **Sg-A** and associate **Sg-A** with **ECS-HA1** and **ECS-HA2**.

Table 4-4 Sg-A rules

Direction	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	Allow	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows remote logins to Linux ECSs over SSH port 22.
Inbound	Allow	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows remote logins to Windows ECSs over RDP port 3389.
Inbound	Allow	IPv4	TCP: 80	Source: 0.0.0.0/0	Allows external access to the website deployed on the ECSs over HTTP port 80.
Inbound	Allow	IPv4	All	Source: current security group (Sg-A)	Allows the ECSs in Sg-A to communicate with each other using IPv4 addresses.
Inbound	Allow	IPv6	All	Source: current security group (Sg-A)	Allows the ECSs in sg-A to communicate with each other using IPv6 addresses.
Outbound	Allow	IPv4	All	Destination: 0.0.0.0/0	Allows ECSs in Sg-A to access the Internet using IPv4 addresses.
Outbound	Allow	IPv6	All	Destination: ::/0	Allows ECSs in Sg-A to access the Internet using IPv6 addresses.

NOTICE

In this example, **Source** is set to **0.0.0.0/0**, which allows any external IP address to remotely log in to ECSs in **Sg-A**. To ensure security, you are advised to set **Source** to a specific IP address, for example, the IP address of your local PC.

If your ECSs are associated with different security groups, you need to add rules in [Table 4-5](#) to allow the ECSs in the two security groups to communicate with each other.

Table 4-5 Rules of security groups **Sg-A** and **Sg-B**

Security Group	Direction	Action	Type	Protocol & Port	Source/Destination	Description
Sg-A	Inbound	Allow	IPv4	All	Source: Sg-B	Allows ECSs in Sg-B to access those in Sg-A over any IPv4 protocol and port.
Sg-B	Inbound	Allow	IPv4	All	Source: Sg-A	Allows ECSs in Sg-A to access those in Sg-B over any IPv4 protocol and port.

- **EIP:** Select **Not required**.
- 3. Assign a virtual IP address from **Subnet-A01**.
For details, see [Assigning a Virtual IP Address](#).
- 4. Assign an EIP.
For details, see [Assigning an EIP](#).

Step 2: Configure Keepalived on ECS-HA1 and ECS-HA2.

1. Configure Keepalived on **ECS-HA1**.
 - a. Bind **EIP-A (124.X.X.187)** to **ECS-HA1**.
For details, see [Binding an EIP to an ECS](#).
 - b. Remotely log in to **ECS-HA1**.
For details, see [Logging In to an ECS](#).
 - c. Run the following command to install the Nginx and Keepalived packages and related dependency packages:

yum install nginx keepalived -y

If information similar to the following is displayed, the installation is complete:

```
[root@ecs-ha1 ~]# yum install nginx keepalived -y
Loaded plugins: fastestmirror
Determining fastest mirrors
base
| 3.6 kB 00:00:00
epel
| 4.3 kB 00:00:00
extras
| 2.9 kB 00:00:00
updates
| 2.9 kB 00:00:00
(1/7): epel/x86_64/
group
| 399 kB 00:00:00
```



```
(2/7): epel/x86_64/
updateinfo
| 1.0 MB 00:00:00
(3/7): base/7/x86_64/
primary_db
| 6.1 MB 00:00:00
(4/7): base/7/x86_64/
group_gz
| 153 kB 00:00:00
(5/7): epel/x86_64/
primary_db
| 8.7 MB 00:00:00
(6/7): extras/7/x86_64/
primary_db
| 253 kB 00:00:00
(7/7): updates/7/x86_64/primary_db

.....
Dependency Installed:
centos-indexhtml.noarch 0:7-9.el7.centos gperftools-libs.x86_64
0:2.6.1-1.el7 lm_sensors-libs.x86_64 0:3.4.0-8.20160601gitf9185e5.el7_9.1
net-snmp-agent-libs.x86_64 1:5.7.2-49.el7_9.4 net-snmp-libs.x86_64
1:5.7.2-49.el7_9.4 nginx-filessystem.noarch 1:1.20.1-10.el7
openssl11-libs.x86_64 1:1.1.1k-7.el7

Complete!
```

- d. Modify the Nginx configuration file.
 - i. Run the following command to open the `/etc/nginx/nginx.conf` file:
vim /etc/nginx/nginx.conf
 - ii. Press **i** to enter the editing mode.
 - iii. Replace the original content with the following:

```
user root;
worker_processes 1;
#error_log logs/error.log;
#error_log logs/error.log notice;
#error_log logs/error.log info;
#pid logs/nginx.pid;
events {
    worker_connections 1024;
}
http {
    include mime.types;
    default_type application/octet-stream;
    #log_format main '$remote_addr - $remote_user [$time_local] "$request" '
    # '$status $body_bytes_sent "$http_referer" '
    # '"$http_user_agent" "$http_x_forwarded_for"';
    #access_log logs/access.log main;
    sendfile on;
    #tcp_nopush on;
    #keepalive_timeout 0;
    keepalive_timeout 65;
    #gzip on;
    server {
        listen 80;
        server_name localhost;
        #charset koi8-r;
        #access_log logs/host.access.log main;
        location / {
            root html;
            index index.html index.htm;
        }
        #error_page 404 /404.html;
        # redirect server error pages to the static page /50x.html
        error_page 500 502 503 504 /50x.html;
        location = /50x.html {
            root html;
        }
    }
}
```

```
    }  
  }  
}
```

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- e. Modify the **index.html** file to verify whether the website is successfully accessed.
 - i. Run the following command to open the **/usr/share/nginx/html/index.html** file:
- ii. Press **i** to enter the editing mode.
- iii. Replace the original content with the following:
- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- f. Run the following commands to set the automatic startup of Nginx upon ECS startup:

```
systemctl enable nginx
```

```
systemctl start nginx.service
```


Information similar to the following is displayed:

```
[root@ecs-ha1 ~]# systemctl enable nginx  
Created symlink from /etc/systemd/system/multi-user.target.wants/nginx.service to /usr/lib/  
systemd/system/nginx.service.  
[root@ecs-ha1 ~]# systemctl start nginx.service
```

- g. Open a browser, enter the EIP address (**124.X.X.187**), and press **Enter** to verify the access to a single Nginx node.

If the web page shown in the following figure is displayed, Nginx is successfully configured for **ECS-HA1**.

Figure 4-10 ECS-HA1 accessed



- h. Modify the Keepalived configuration file.
 - i. Run the following command to open the **/etc/keepalived/keepalived.conf** file:
 - ii. Press **i** to enter the editing mode.
 - iii. Replace the IP parameters in the configuration file as follows:
 - **mcast_src_ip** and **unicast_src_ip**: Change their values to the private IP address of an ECS. In this example, private IP address **192.168.0.195** of **ECS-HA1** is used.
 - **virtual_ipaddress**: Change the value to a virtual IP address. In this example, **192.168.0.177** is used.

```
! Configuration File for keepalived  
global_defs {  
  router_id master-node  
}
```

```
vrrip_script chk_http_port {
    script "/etc/keepalived/chk_nginx.sh"
    interval 2
    weight -5
    fall 2
    rise 1
}
vrrip_instance VI_1 {
    state BACKUP
    interface eth0
    mcast_src_ip 192.168.0.195
    virtual_router_id 51
    priority 100
    advert_int 1
    authentication {
        auth_type PASS
        auth_pass 1111
    }
    unicast_src_ip 192.168.0.195
    virtual_ipaddress {
        192.168.0.177
    }
}
track_script {
    chk_http_port
}
```

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- i. Configure the Nginx monitoring script.

- i. Run the following command to open the **/etc/keepalived/chk_nginx.sh** file:

```
vim /etc/keepalived/chk_nginx.sh
```

- ii. Press **i** to enter the editing mode.
- iii. Replace the original content with the following:

```
#!/bin/bash
counter=$(ps -C nginx --no-heading|wc -l)
if [ "${counter}" = "0" ]; then
    systemctl start nginx.service
    sleep 2
    counter=$(ps -C nginx --no-heading|wc -l)
    if [ "${counter}" = "0" ]; then
        systemctl stop keepalived.service
    fi
fi
```

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
 - j. Run the following command to assign execute permissions to the **chk_nginx.sh** file:

```
chmod +x /etc/keepalived/chk_nginx.sh
```
 - k. Run the following commands to set the automatic startup of Keepalived upon ECS startup:

```
systemctl enable keepalived
systemctl start keepalived.service
```
 - l. Unbind **EIP-A** from **ECS-HA1**.
For details, see [Unbinding an EIP](#).
2. Configure Keepalived on **ECS-HA2**.
 - a. Bind **EIP-A (124.X.X.187)** to **ECS-HA2**.
For details, see [Binding an EIP to an ECS](#).

- b. Remotely log in to **ECS-HA2**.
For details, see [Logging In to an ECS](#).
- c. Run the following command to install the Nginx and Keepalived packages and related dependency packages:

yum install nginx keepalived -y

If information similar to the following is displayed, the installation is complete:

```
[root@ecs-ha2 ~]# yum install nginx keepalived -y
Loaded plugins: fastestmirror
Determining fastest mirrors
base
      | 3.6 kB  00:00:00
epel
      | 4.3 kB  00:00:00
extras
      | 2.9 kB  00:00:00
updates
      | 2.9 kB  00:00:00
(1/7): epel/x86_64/
group
      | 399 kB  00:00:00
(2/7): epel/x86_64/
updateinfo
      | 1.0 MB  00:00:00
(3/7): base/7/x86_64/
primary_db
      | 6.1 MB  00:00:00
(4/7): base/7/x86_64/
group_gz
      | 153 kB  00:00:00
(5/7): epel/x86_64/
primary_db
      | 8.7 MB  00:00:00
(6/7): extras/7/x86_64/
primary_db
      | 253 kB  00:00:00
(7/7): updates/7/x86_64/primary_db

.....
Dependency Installed:
  centos-indexhtml.noarch 0:7-9.el7.centos          gperftools-libs.x86_64
  0:2.6.1-1.el7            lm_sensors-libs.x86_64 0:3.4.0-8.20160601gitf9185e5.el7_9.1
  net-snmp-agent-libs.x86_64 1:5.7.2-49.el7_9.4      net-snmp-libs.x86_64
  1:5.7.2-49.el7_9.4      nginx-filessystem.noarch 1:1.20.1-10.el7
  openssl11-libs.x86_64 1:1.1.1k-7.el7

Complete!
```

- d. Modify the Nginx configuration file.
 - i. Run the following command to open the **/etc/nginx/nginx.conf** file:
vim /etc/nginx/nginx.conf
 - ii. Press **i** to enter the editing mode.
 - iii. Replace the original content with the following:

```
user root;
worker_processes 1;
#error_log logs/error.log;
#error_log logs/error.log notice;
#error_log logs/error.log info;
#pid logs/nginx.pid;
events {
    worker_connections 1024;
}
http {
```

```
include mime.types;
default_type application/octet-stream;
#log_format main '$remote_addr - $remote_user [$time_local] "$request" '
# '$status $body_bytes_sent "$http_referer" '
# '$http_user_agent' "$http_x_forwarded_for";
#access_log logs/access.log main;
sendfile on;
#tcp_nopush on;
#keepalive_timeout 0;
keepalive_timeout 65;
#gzip on;
server {
    listen 80;
    server_name localhost;
    #charset koi8-r;
    #access_log logs/host.access.log main;
    location / {
        root html;
        index index.html index.htm;
    }
    #error_page 404 /404.html;
    # redirect server error pages to the static page /50x.html
    error_page 500 502 503 504 /50x.html;
    location = /50x.html {
        root html;
    }
}
```

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- e. Modify the **index.html** file to verify whether the website is successfully accessed.
 - i. Run the following command to open the **/usr/share/nginx/html/index.html** file:
vim /usr/share/nginx/html/index.html
 - ii. Press **i** to enter the editing mode.
 - iii. Replace the original content with the following:
Welcome to ECS-HA2
 - iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- f. Run the following commands to set the automatic startup of Nginx upon ECS startup:
systemctl enable nginx
systemctl start nginx.service
Information similar to the following is displayed:
[root@ecs-ha2 ~]# systemctl enable nginx
Created symlink from /etc/systemd/system/multi-user.target.wants/nginx.service to /usr/lib/systemd/system/nginx.service.
[root@ecs-ha2 ~]# systemctl start nginx.service
- g. Open a browser, enter the EIP address (**124.X.X.187**), and press **Enter** to verify the access to a single Nginx node.
If the web page shown in the following figure is displayed, Nginx is successfully configured for **ECS-HA2**.

Figure 4-11 ECS-HA2 accessed



Welcome to ECS-HA2

- h. Modify the Keepalived configuration file.
 - i. Run the following command to open the `/etc/keepalived/keepalived.conf` file:
vim /etc/keepalived/keepalived.conf
 - ii. Press **i** to enter the editing mode.
 - iii. Replace the IP parameters in the configuration file as follows:
 - **mcast_src_ip** and **unicast_src_ip**: Change their values to the private IP address of an ECS. In this example, private IP address of **ECS-HA2 (192.168.0.233)** is used.
 - **virtual_ipaddress**: Change the value to a virtual IP address. In this example, **192.168.0.177** is used.

```
! Configuration File for keepalived
global_defs {
    router_id master-node
}
vrrp_script chk_http_port {
    script "/etc/keepalived/chk_nginx.sh"
    interval 2
    weight -5
    fall 2
    rise 1
}
vrrp_instance VI_1 {
    state BACKUP
    interface eth0
    mcast_src_ip 192.168.0.233
    virtual_router_id 51
    priority 100
    advert_int 1
    authentication {
        auth_type PASS
        auth_pass 1111
    }
    unicast_src_ip 192.168.0.233
    virtual_ipaddress {
        192.168.0.177
    }
}
track_script {
    chk_http_port
}
```

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- i. Configure the Nginx monitoring script.
 - i. Run the following command to open the `/etc/keepalived/chk_nginx.sh` file:
vim /etc/keepalived/chk_nginx.sh

- ii. Press **i** to enter the editing mode.
- iii. Replace the original content with the following:

```
#!/bin/bash
counter=$(ps -C nginx --no-heading|wc -l)
if [ "${counter}" = "0" ]; then
    systemctl start nginx.service
    sleep 2
    counter=$(ps -C nginx --no-heading|wc -l)
    if [ "${counter}" = "0" ]; then
        systemctl stop keepalived.service
    fi
fi
```
- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- j. Run the following command to assign execute permissions to the **chk_nginx.sh** file:
chmod +x /etc/keepalived/chk_nginx.sh
- k. Run the following commands to set the automatic startup of Keepalived upon ECS startup:
systemctl enable keepalived
systemctl start keepalived.service
- l. Unbind **EIP-A** from **ECS-HA2**.
For details, see [Unbinding an EIP](#).

Step 3: Bind the Virtual IP Address to the Active and Standby ECSs and EIP


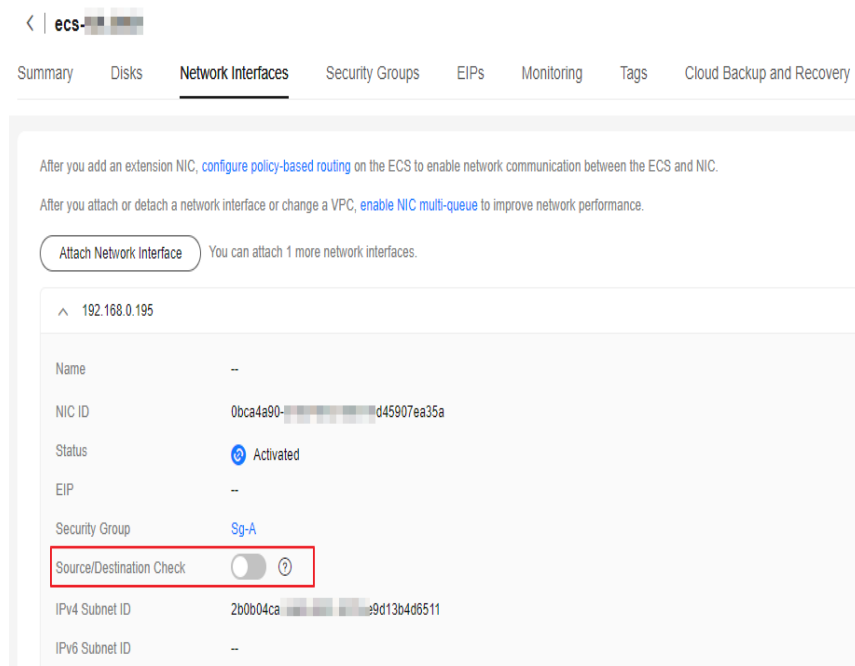
1. Bind virtual IP address **192.168.0.177** to **ECS-HA1** and **ECS-HA2**.
For details, see [Binding a Virtual IP Address to an Instance or EIP](#).
2. Disable **Source/Destination Check** for the network interfaces of the active and standby ECSs.
When you bind a virtual IP address to an ECS, **Source/Destination Check** is disabled by default. You can perform the following operations to check whether the function is disabled. If the function is not disabled, disable it.
 - a. In the ECS list, click the name of the target ECS.
The ECS details page is displayed.
 - b. On the **Network Interfaces** tab, click  to expand the details area and check whether **Source/Destination Check** is disabled.

Figure 4-12 Disabling Source/Destination Check

3. Bind virtual IP address **192.168.0.177** to **EIP-A**.
For details, see [Binding a Virtual IP Address to an Instance or EIP](#).


Step 4: Disable IP Forwarding on the Standby ECS

If a virtual IP address is bound to active/standby ECSs, you need to disable IP forwarding on the standby ECS. If an active/standby ECS switchover happens, ensure that IP forwarding of the new standby ECS is also disabled.

To make sure you do not miss any settings, it is better to disable IP forwarding on both of active and standby ECSs.

1. Open a browser, enter the EIP address (**124.X.X.187**), and press **Enter** to access the active ECS.

If the following page is displayed, the **ECS-HA1** is used as the active ECS.

Figure 4-13 The active ECS accessed

2. Remotely log in to the standby ECS (**ECS-HA2** in this example).
For details, see [Logging In to an ECS](#).
3. Disable IP forwarding by following the operations in [Table 4-6](#). In this example, the ECS runs the Linux OS.

Table 4-6 Disabling IP forwarding

OS	Operations
Linux	<ol style="list-style-type: none">1. Run the following command to switch to user root: su root2. Run the following command to check whether IP forwarding 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.<ul style="list-style-type: none">• If 0 is displayed, no further action is required.• If 1 is displayed, go to the next step.3. Use either of the following methods to modify the configuration file: Method 1<ol style="list-style-type: none">a. Run the following command to open the /etc/sysctl.conf file: vim /etc/sysctl.confb. Press i to enter the editing mode.c. Set net.ipv4.ip_forward to 0.d. Press ESC to exit and enter :wq! to save the configuration.Method 2 Run the sed command. An example command is as follows: sed -i '/net.ipv4.ip_forward/s/1/0/g' /etc/sysctl.conf4. Run the following command to apply the modification: sysctl -p /etc/sysctl.conf
Windows	<ol style="list-style-type: none">1. In the search box, enter cmd to open the command prompt window, and run the following command: ipconfig/all<ul style="list-style-type: none">• In the command output, if the value of IP Routing Enabled is No, IP forwarding is disabled.• If IP Routing Enabled is Yes, IP forwarding is not disabled. Go to the next step.2. Enter regedit in the search box to open the registry editor.3. Set the value of IPEnableRouter under HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to 0.<ul style="list-style-type: none">• If the value is set to 0, IP forwarding will be disabled.• If the value is set to 1, IP forwarding will be enabled.

Step 5: Verify the Automatic Switchover Between the Active and Standby ECSs

- Restart the active and standby ECSs.
 - Remotely log in to **ECS-HA1**.
For details, see [Logging In to an ECS](#).
 - Run the following command to restart **ECS-HA1**:
reboot
 - Repeat [1.a](#) to [1.b](#) to restart **ECS-HA2**.
- Check whether the website on the active ECS can be accessed.
 - Open a browser, enter the EIP address (**124.X.X.187**), and press **Enter**.
If the following page is displayed, **ECS-HA1** is used as the active ECS and the website can be accessed.

Figure 4-14 ECS-HA1 accessed



Welcome to ECS-HA1

- Remotely log in to **ECS-HA1** and run the following command to check whether the virtual IP address is bound to the eth0 NIC of **ECS-HA1**:

ip addr show

If information similar to the following is displayed, the virtual IP address (**192.168.0.177**) has been bound to the eth0 NIC of **ECS-HA1**, and this ECS is the active one.

```
[root@ecs-ha1 ~]# ip addr show
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 mq state UP group default qlen 1000
    link/ether fa:16:3e:fe:56:19 brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.195/24 brd 192.168.0.255 scope global noprefixroute dynamic eth0
        valid_lft 107898685sec preferred_lft 107898685sec
    inet 192.168.0.177/32 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe56:19/64 scope link
        valid_lft forever preferred_lft forever
```

- Run the following command to disable Keepalived on **ECS-HA1**:

systemctl stop keepalived.service

- Check whether **ECS-HA2** becomes the active ECS.
 - Remotely log in to **ECS-HA2** and run the following command to check whether the virtual IP address is bound to the eth0 NIC of **ECS-HA2**:
ip addr show

If information similar to the following is displayed, the virtual IP address (**192.168.0.177**) has been bound to the eth0 NIC of **ECS-HA2**, and this ECS becomes the active one.

```
[root@ecs-ha2 ~]# ip addr show
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 mq state UP group default
qlen 1000
    link/ether fa:16:3e:fe:56:3f brd ff:ff:ff:ff:ff:ff
    inet 192.168.0.233/24 brd 192.168.0.255 scope global noprefixroute dynamic eth0
        valid_lft 107898091sec preferred_lft 107898091sec
    inet 192.168.0.177/32 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fefe:563f/64 scope link
        valid_lft forever preferred_lft forever
```

- b. Open a browser, enter the EIP address (**124.X.X.187**), and press **Enter** to check whether the website on the active ECS (**ECS-HA2**) can be accessed. If the following page is displayed, **ECS-HA2** is used as the active ECS and the website can be accessed.

Figure 4-15 ECS-HA2 accessed



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.

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
Region	The region where the network interface is created. Select the region nearest to you to ensure the lowest latency possible.	EU-Dublin
Name	Enter the name of the network interface. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	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
Private IP Address	Select whether to automatically assign a private IP address.	-

Parameter	Parameter Description	Example Value
Security Group	Select the security group that will be associated with the network interface.	sg-001


4. Click **OK**.

5.1.3 Viewing the 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click the private IP address 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 an instance.
- Enable or disable **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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. 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.
5. Click **OK**.

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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. 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.

5. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. 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.
5. Locate the row that contains the target virtual IP address and click **Bind to Server** in the **Operation** column.
6. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. 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.
5. Click **OK**.
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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. In the network interface list, locate the row that contains the target network interface, and choose **More > Change Security Group** in the **Operation** column.
5. 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 choose **Networking > Virtual Private Cloud**.

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

3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. Click the private IP address of the target network interface.
5. Click the **Associated Security Groups** tab. Then, click **Change Security Group**.
6. On the **Change Security Group** page, select the security groups to be associated and click **OK**.

Other Operations

On the **Associated Security Groups** tab, locate the target security group and click **Manage Rule** to manage security group rules. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.

4. Locate the row that contains the network interface, click **More** in the **Operation** column, and click **Delete**.
A confirmation dialog box is displayed.
5. 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.
- 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. Go to the [supplementary network interface list page](#).
2. In the upper right corner of the page, click **Create Supplementary Network Interface**.
3. Configure the parameters based on [Table 5-2](#).

Table 5-2 Parameter descriptions

Parameter	Description	Example Value
Region	The region where the supplementary network interface is created. Select the region nearest to you to ensure the lowest latency possible.	EU-Dublin
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)

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

4. Click **Create Now**.

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 Viewing the Basic Information About a Supplementary Network Interface .
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.
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 the 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
5. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
5. Locate the row that contains the supplementary network interface, click **Bind EIP** in the **Operation** column, and select the EIP to be bound.
6. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
5. Locate the row that contains the supplementary network interface and click **Unbind EIP** in the **Operation** column.
6. Click **OK**.

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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click the **Supplementary Network Interfaces** tab.
5. Locate the row that contains the supplementary network interface and click **Change Security Group** in the **Operation** column.
6. On the **Change Security Group** page, select the security group to be associated.
7. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.

4. On the **Network Interfaces** page, click the **Supplementary Network Interfaces** tab.
5. Click the private IP address of the supplementary network interface whose security group is to be changed.
6. Click the **Associated Security Groups** tab. Then, click **Change Security Group**.
7. On the **Change Security Group** page, select the security group to be associated.
8. 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 choose **Networking > Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
3. In the navigation pane on the left, choose **Virtual Private Cloud > Network Interfaces**.
4. On the **Network Interfaces** page, click **Supplementary Network Interfaces** tab.
5. Locate the row that contains the supplementary network interface and click **Delete** in the **Operation** column.
A confirmation dialog box is displayed.
6. Confirm the information and click **OK**.
Deleting a supplementary network interface will also delete the VLAN sub-interfaces configured on the ECS.

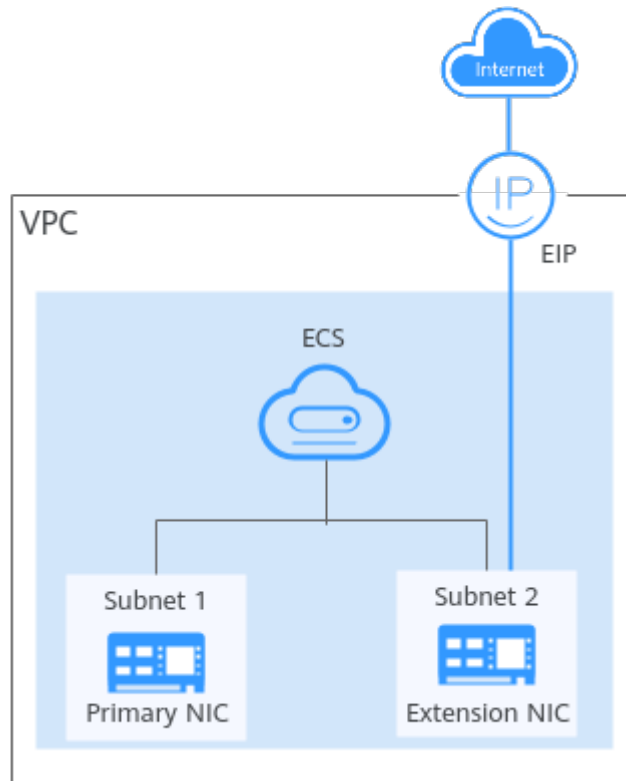
5.3 Network Interface Configuration Examples

5.3.1 Binding an EIP to the Extension NIC of an ECS to Enable Internet Access

Scenarios

As shown in [Figure 5-4](#), the ECS has two NICs, one primary NIC and one extension NIC. You can bind an EIP to the extension NIC of the ECS and configure policy-based routes to ensure that the ECS can access the Internet through the EIP.

Figure 5-4 Accessing the Internet through the EIP bound to the extension NIC



NOTE

This section uses a Linux ECS as an example.

Step 1: Create Cloud Resources and Attach an Extension NIC

1. Create a VPC and two subnets in the VPC.
In this example, the primary and extension NICs of the ECS are in different subnets.
For details, see [Creating a VPC and Subnet](#).
2. Create an ECS in the VPC subnet.



- For details about how to purchase an ECS, see [Purchasing an ECS](#).
3. Create a network interface and attach it to the ECS as an extension NIC.
When creating a network interface, select a different subnet from where the primary NIC is created. For details, see [Creating a Network Interface](#).
Attach the network interface to the ECS. For details, see section [Attaching a Network Interface to a Cloud Server](#).
 4. Assign an EIP and bind it to the extension NIC of the ECS.
For details, see [Assigning an EIP](#).
Bind the EIP to the extension NIC of the ECS. For details, see [Binding an EIP to a Network Interface](#).

Step 2: Obtain the ECS Network Information

Before configuring policy-based routes for the extension NIC, you need to obtain the network information in [Table 5-4](#).

Table 5-4 Required ECS network information

Item	Primary NIC	Extension NIC
Private IP address of the NIC	192.168.11.42	192.168.17.191
Subnet gateway address	192.168.11.1	192.168.17.1

1. Obtain the private IP addresses of the ECS NICs.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select the desired region and project.
 - c. Click **Service List** and choose **Compute > Elastic Cloud Server**.
 - d. In the ECS list, locate the target ECS and click its name.
The **Summary** tab page of the ECS is displayed.
 - e. Click the **Network Interfaces** tab and view the private IP addresses of the primary and extension NICs of the ECS.
2. Obtain the gateway address of the subnet.
 - a. Log in to the management console.
 - b. Click  in the upper left corner and select the desired region and project.
 - c. Click **Service List** and choose **Compute > Elastic Cloud Server**.
 - d. In the ECS list, locate the target ECS and click its name.
The **Summary** tab page of the ECS is displayed.
 - e. In the **ECS Information** area, click the VPC name.

- The **Virtual Private Cloud** page is displayed.
- f. In the VPC list and click the number in the **Subnets** column.
The **Subnets** page is displayed.
 - g. In the subnet list, click the subnet name.
The **Summary** page is displayed.
 - h. In the **Gateway and DNS Information** area, view the gateway address of the subnet.

Figure 5-5 Viewing the gateway address of the subnet

Gateway and DNS Information	
DHCP	Enabled
DNS Server Address	100.125.1.250, 100.125.129.250
IPv4 DHCP Lease Time	1250 days
Gateway	192.168.0.1
Domain Name	--
NTP Server Address	--

Step 3: Configure Policy-based Routes for the Extension NIC

1. ECS Remotely log in to the ECS.
For details, see [Logging In to an ECS](#).
2. Run the following command to query the route information of the NIC:

route -n

The following figure is displayed. In this figure:

- The destination of the route for the primary NIC is 192.168.11.0/24.
- The destination of the route for the extension NIC is 192.168.17.0/24.

```
[root@ecs-b926 ~]# route -n
Kernel IP routing table
Destination      Gateway         Genmask       Flags Metric Ref    Use Iface
0.0.0.0          192.168.11.1   0.0.0.0       UG    0      0      0 eth0
169.254.0.0     0.0.0.0        255.255.0.0   U      1002   0      0 eth0
169.254.0.0     0.0.0.0        255.255.0.0   U      1003   0      0 eth1
169.254.169.254 192.168.11.1   255.255.255.255 UGH    0      0      0 eth0
192.168.11.0    0.0.0.0        255.255.255.0 U      0      0      0 eth0
192.168.17.0    0.0.0.0        255.255.255.0 U      0      0      0 eth1
[root@ecs-b926 ~]#
```

3. Run the following command to query the NIC names of the ECS:

ifconfig

The following figure is displayed. Search for the NIC name based on the NIC address. In this figure:

- 192.168.11.42 is the IP address of the primary NIC, and the NIC name is eth0.
- 192.168.17.191 is the IP address of the extension NIC, and the NIC name is eth1.

```
[root@ecs-b926~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.11.42 netmask 255.255.255.0 broadcast 192.168.11.255
    inet6 fe80::f816:3eff:fef7:1c44 prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:f7:1c:44 txqueuelen 1000 (Ethernet)
    RX packets 127 bytes 21633 (21.1 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 258 bytes 22412 (21.8 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.17.191 netmask 255.255.255.0 broadcast 192.168.17.255
    inet6 fe80::f816:3eff:felc:b57f prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:1c:b5:7f txqueuelen 1000 (Ethernet)
    RX packets 11 bytes 1283 (1.2 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 12 bytes 1388 (1.3 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1 (Local Loopback)
    RX packets 51 bytes 12018 (11.7 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 51 bytes 12018 (11.7 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

4. Configure the default route for the ECS so that it can access the Internet through the extension NIC.

- a. Run the following command to delete the default route of the primary NIC:

```
route del -net 0.0.0.0 gw Subnet-gateway-IP-address dev NIC-name
```

The parameters are described as follows:

- 0.0.0.0: destination IP address, indicating that multiple IP addresses are matched. Do not change the value.
- Subnet gateway IP address: Enter the subnet gateway address of the primary NIC collected in section [Table 5-4](#).
- NIC name: Enter the name of the primary NIC obtained in [3](#).

Example command:

```
route del -net 0.0.0.0 gw 192.168.11.1 dev eth0
```

 **NOTE**

This operation will interrupt ECS traffic.

- b. Run the following command to configure the default route for the extension NIC:

```
route add default gw Subnet-gateway-IP-address
```

The parameters are described as follows:

Subnet gateway IP address: Enter the subnet gateway address of the extension NIC collected in section [Table 5-4](#).

Example command:

```
route add default gw 192.168.17.1
```

5. Verify network connectivity.

Run the following command to check whether the ECS can access the Internet:

ping *Public-IP-address-or-domain-name*

Example command:

ping support.huaweicloud.com

If information similar to the following is displayed, the ECS can communicate with the Internet.

```
[root@ecs-a01 ~]# ping support.huaweicloud.com
PING hcdnw.cbg-notzj.c.cdnhwc2.com (203.193.226.103) 56(84) bytes of data.
64 bytes from 203.193.226.103 (203.193.226.103): icmp_seq=1 ttl=51 time=2.17 ms
64 bytes from 203.193.226.103 (203.193.226.103): icmp_seq=2 ttl=51 time=2.13 ms
64 bytes from 203.193.226.103 (203.193.226.103): icmp_seq=3 ttl=51 time=2.10 ms
64 bytes from 203.193.226.103 (203.193.226.103): icmp_seq=4 ttl=51 time=2.09 ms
...
--- hcdnw.cbg-notzj.c.cdnhwc2.com ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3004ms
rtt min/avg/max/mdev = 2.092/2.119/2.165/0.063 ms
```

5.3.2 Configuring Policy-based Routes for an ECS with Multiple NICs

5.3.2.1 Overview

Background

If an ECS has multiple NICs, the primary NIC can communicate with external networks by default, but the extension NICs cannot. To enable extension NICs to communicate with external works either, you need to configure policy-based routes for these NICs.

Scenarios

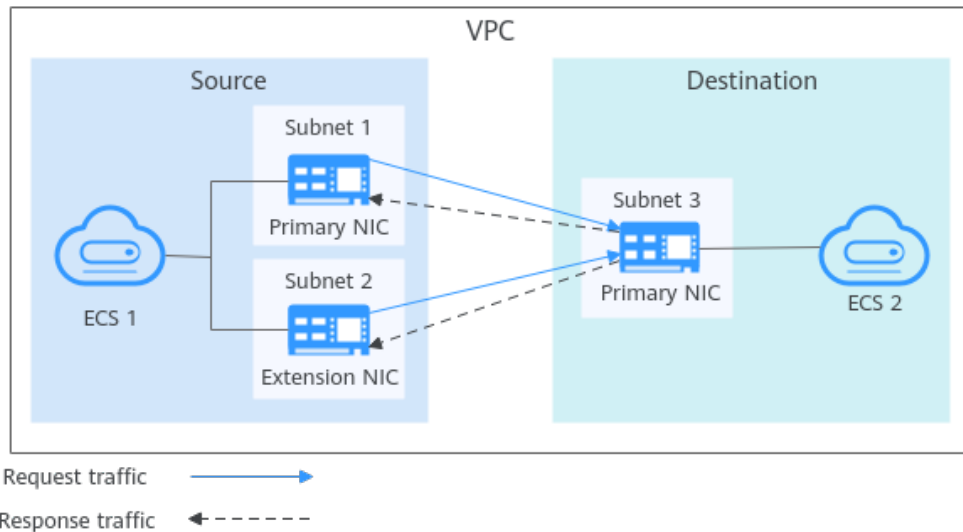
This example describes how to configure policy-based routes for an ECS with two NICs. [Figure 5-6](#) shows the networking. The details are as follows:

- The primary and extension NICs on the source ECS are in different subnets of the same VPC.
- The source and destination ECSs are in different subnets of the same VPC and the two ECSs can communicate with each other through primary NICs without configuring policy-based routes.
- After policy-based routes are configured for the two NICs of the source ECS, both the primary and extension NICs can communicate with the destination ECS.

NOTICE

You can select a destination IP address based on service requirements. Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

Figure 5-6 Dual-NIC ECS networking



5.3.2.2 Collecting ECS Network Information

Scenarios

Before configuring policy-based routes for a multi-NIC ECS, you need to collect network information about the ECS.

- [Table 5-5](#) lists the information to be collected for a Linux ECS using IPv4.

Table 5-5 Linux ECS using IPv4


EC S	Primary NIC	Extension NIC	How to Obtain
So ur ce	<ul style="list-style-type: none"> • NIC address: 10.0.0.115 • Subnet: 10.0.0.0/24 • Subnet gateway: 10.0.0.1 	<ul style="list-style-type: none"> • NIC address: 10.0.1.183 • Subnet: 10.0.1.0/24 • Subnet gateway: 10.0.1.1 	<ul style="list-style-type: none"> • Obtaining ECS NIC Addresses • Obtaining Subnet CIDR Blocks and Gateway Addresses
De sti na tio n	NIC address: 10.0.2.12	N/A	

- [Table 5-6](#) lists the information to be collected for a Windows ECS using IPv4.


Table 5-6 Windows ECS using IPv4

EC S	Primary NIC	Extension NIC	How to Obtain
Source	<ul style="list-style-type: none">NIC address: 10.0.0.59Subnet gateway: 10.0.0.1	<ul style="list-style-type: none">NIC address: 10.0.1.104Subnet gateway: 10.0.1.1	<ul style="list-style-type: none">Obtaining ECS NIC AddressesObtaining Subnet CIDR Blocks and Gateway Addresses
Destination	NIC address: 10.0.2.12	N/A	

Obtaining ECS NIC Addresses

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click **Service List** and choose **Compute > Elastic Cloud Server**.
4. In the ECS list, click the target ECS name.
The **Summary** tab page of the ECS is displayed.
5. In the **NICs** area, view the IP addresses of the primary and extension NICs.

Obtaining Subnet CIDR Blocks and Gateway Addresses

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click **Service List** and choose **Compute > Elastic Cloud Server**.
4. In the ECS list, click the target ECS name.
The **Summary** tab page of the ECS is displayed.
5. In the **ECS Information** area, click the VPC hyperlink.
The **Virtual Private Cloud** page is displayed.
6. Locate the target VPC and click the number in the **Subnets** column.
The **Subnets** page is displayed.
7. In the subnet list, view the CIDR blocks of the subnets.
8. In the subnet list, click the subnet name.
The **Summary** page is displayed.
9. Click the **IP Addresses** tab and view the gateway addresses of the subnet.

5.3.2.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs

Scenarios

This section describes how to configure policy-based routes for a dual-NIC ECS running CentOS 8.0 (64-bit).

For details about the background knowledge and networking of dual-NIC ECSs, see [Overview](#).

Procedure (Linux ECS Using IPv4)

1. Collect the ECS network information required for configuring policy-based routes.
For details, see [Collecting ECS Network Information](#).
2. Log in to an ECS.
3. Check whether the source ECS can use its primary NIC to communicate with the destination ECS:

ping -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

In this example, run the following command:

ping -I 10.0.0.115 10.0.2.12

If information similar to the following is displayed, the source ECS can use its primary NIC to communicate with the destination ECS.

```
[root@ecs-resource ~]# ping -I 10.0.0.115 10.0.2.12
PING 10.0.2.12 (10.0.2.12) from 10.0.0.115 : 56(84) bytes of data.
64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=0.775 ms
64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.268 ms
64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.220 ms
64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.167 ms
^C
--- 10.0.2.12 ping statistics ---
```

NOTE

Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

4. Query the NIC names of the ECS:

ifconfig

Search for the NIC name based on the NIC address.

- 10.0.0.115 is the IP address of the primary NIC, and the NIC name is eth0.
- 10.0.1.183 is the IP address of the extension NIC, and the NIC name is eth1.

```
[root@ecs-resource ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.0.0.115 netmask 255.255.255.0 broadcast 10.0.0.255
    inet6 fe80::f816:3eff:fe92:6e0e prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:92:6e:0e txqueuelen 1000 (Ethernet)
    RX packets 432288 bytes 135762012 (129.4 MiB)
    RX errors 0 dropped 0 overruns 0 frame 1655
    TX packets 423744 bytes 106716932 (101.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 10.0.1.183 netmask 255.255.255.0 broadcast 10.0.1.255
```

```
inet6 fe80::f816:3eff:febf:5818 prefixlen 64 scopeid 0x20<link>  
ether fa:16:3e:bf:58:18 txqueuelen 1000 (Ethernet)  
RX packets 9028 bytes 536972 (524.3 KiB)  
RX errors 0 dropped 0 overruns 0 frame 1915  
TX packets 6290 bytes 272473 (266.0 KiB)  
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

5. Configure temporary routes for the ECS.

NOTICE

Temporary routes take effect immediately after being configured and will be lost after the ECS is restarted. To prevent network interruptions after the ECS is restarted, perform [6](#) after this step to configure persistent routes.

- a. Configure policy-based routes for both the primary and extension NICs:
 - Primary NIC
ip route add default via *Subnet gateway* dev *NIC name* table *Route table name*
ip route add *Subnet CIDR block* dev *NIC name* table *Route table name*
ip rule add from *NIC address* table *Route table name*
 - Extension NIC
ip route add default via *Subnet gateway* dev *NIC name* table *Route table name*
ip route add *Subnet CIDR block* dev *NIC name* table *Route table name*
ip rule add from *NIC address* table *Route table name*

Configure the parameters as follows:

- NIC name: Enter the name obtained in [4](#).
- Route table name: Customize a route table name using a number.
- Other network information: Enter the IP addresses collected in [1](#).

In this example, run the following commands:

- Primary NIC
ip route add default via 10.0.0.1 dev eth0 table 10
ip route add 10.0.0.0/24 dev eth0 table 10
ip rule add from 10.0.0.115 table 10
- Extension NIC
ip route add default via 10.0.1.1 dev eth1 table 20
ip route add 10.0.1.0/24 dev eth1 table 20
ip rule add from 10.0.1.183 table 20

 NOTE

If the ECS has multiple NICs, configure policy-based routes for all NICs one by one.

- b. Check whether the policy-based routes are successfully added.

ip rule

ip route show table *Route table name of the primary NIC*

ip route show table *Route table name of the extension NIC*

The route table name is customized in [5.a](#).

In this example, run the following commands:

ip rule

ip route show table 10

ip route show table 20

If information similar to the following is displayed, the policy-based routes have been added.

```
[root@ecs-resource ~]# ip rule
0:    from all lookup local
32764: from 10.0.1.183 lookup 20
32765: from 10.0.0.115 lookup 10
32766: from all lookup main
32767: from all lookup default
[root@ecs-resource ~]# ip route show table 10
default via 10.0.0.1 dev eth0
10.0.0.0/24 dev eth0 scope link
[root@ecs-resource ~]# ip route show table 20
default via 10.0.1.1 dev eth1
10.0.1.0/24 dev eth1 scope link
```

- c. Check whether the source ECS and the destination ECS can communicate with each other.

ping -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

ping -I *IP address of the extension NIC on the source ECS IP address of the destination ECS*

In this example, run the following commands:

ping -I 10.0.0.115 10.0.2.12

ping -I 10.0.1.183 10.0.2.12

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS.

```
[root@ecs-resource ~]# ping -I 10.0.0.115 10.0.2.12
PING 10.0.2.12 (10.0.2.12) from 10.0.0.115 : 56(84) bytes of data.
64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=0.775 ms
64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.268 ms
64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.220 ms
64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.167 ms
^C
--- 10.0.2.12 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 102ms
rtt min/avg/max/mdev = 0.167/0.357/0.775/0.244 ms
[root@ecs-resource ~]# ping -I 10.0.1.183 10.0.2.12
PING 10.0.2.12 (10.0.2.12) from 10.0.1.183 : 56(84) bytes of data.
64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=2.84 ms
64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.258 ms
64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.234 ms
64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.153 ms
^C
```

```
--- 10.0.2.12 ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 92ms  
rtt min/avg/max/mdev = 0.153/0.871/2.840/1.137 ms
```

6. Configure persistent routes for the ECS.

- a. Run the following command to open the **/etc/rc.local** file:

```
vi /etc/rc.local
```

- b. Press **i** to enter the editing mode.

- c. Add the following content to the end of the file:

```
# wait for nics up  
sleep 5  
# Add v4 routes for eth0  
ip route flush table 10  
ip route add default via 10.0.0.1 dev eth0 table 10  
ip route add 10.0.0.0/24 dev eth0 table 10  
ip rule add from 10.0.0.115 table 10  
# Add v4 routes for eth1  
ip route flush table 20  
ip route add default via 10.0.1.1 dev eth1 table 20  
ip route add 10.0.1.0/24 dev eth1 table 20  
ip rule add from 10.0.1.183 table 20  
# Add v4 routes for cloud-init  
ip rule add to 169.254.169.254 table main
```

Parameters are described as follows:

- wait for nics up: file startup time. Set the value to be the same as that in the preceding configurations.
 - Add v4 routes for eth0: policy-based routes of the primary NIC. Set the value to be the same as that configured in [5.a](#).
 - Add v4 routes for eth1: policy-based routes of the extension NIC. Set the value to be the same as that configured in [5.a](#).
 - Add v4 routes for cloud-init: Configure the Cloud-Init address. Set the value to be the same as that in the preceding configurations.
- d. Press **ESC** to exit and enter **:wq!** to save the configuration.
- e. Run the following command to assign execute permissions to the **/etc/rc.local** file:

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

 **NOTE**

If your operating system is Red Hat or EulerOS, run the following command after you perform [6.e](#):

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

- f. Run the following command to restart the ECS:

```
reboot
```

NOTICE

Policy-based routes added to the **/etc/rc.local** file take effect only after the ECS is restarted. Ensure that workloads on the ECS will not be affected before restarting the ECS.

- g. Repeat [5.b](#) to [5.c](#) to check whether the policy-based routes are added and whether the source ECS and the destination ECS can communicate with each other.

5.3.2.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs

Scenarios

This section describes how to configure policy-based routes for a dual-NIC ECS running Windows Server 2012 (64-bit).

For details about the background knowledge and networking of dual-NIC ECSs, see [Overview](#).

Procedure (Windows ECS Using IPv4)

1. Collect the ECS network information required for configuring policy-based routes.
For details, see [Collecting ECS Network Information](#).
2. Log in to an ECS.
3. Check whether the source ECS can use its primary NIC to communicate with the destination ECS:

ping -S *IP address of the primary NIC on the source ECS IP address of the destination ECS*

In this example, run the following command:

ping -S 10.0.0.59 10.0.2.12

If information similar to the following is displayed, the source ECS can use its primary NIC to communicate with the destination ECS.

```
C:\Users\Administrator>ping -S 10.0.0.59 10.0.2.12
Pinging 10.0.2.12 from 10.0.0.59 with 32 bytes of data:
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
```

NOTE

Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

4. Configure a policy-based route for the extension NIC.

route add -p 0.0.0.0 mask 0.0.0.0 *Subnet gateway of the extension NIC*
metric *Route priority*

Configure the parameters as follows:

- **0.0.0.0/0**: Default route. Do not change it.
- Subnet gateway of the extension NIC: Enter the IP address collected in [1](#).
- Route priority: Set its value to 261. The priority of the extension NIC must be lower than that of the primary NIC. A larger value indicates a lower priority.

In this example, run the following command:

```
route add -p 0.0.0.0 mask 0.0.0.0 10.0.1.1 metric 261
```

NOTE

- The primary NIC already has policy-based routes and you do not need to configure again.
- If the ECS has multiple extension NICs, configure policy-based routes for all extension NICs one by one.

5. Check whether the policy-based route is successfully added.

route print

If information similar to the following is displayed, the policy-based route has been added. The route is persistent and will not be lost after the ECS is restarted.

```
C:\Users\Administrator>route print
=====
Interface List
19...fa 16 3e fc 7b 76 .....Red Hat VirtIO Ethernet Adapter #3
14...fa 16 3e 5d 3e b6 .....Red Hat VirtIO Ethernet Adapter
1.....Software Loopback Interface 1
16...00 00 00 00 00 00 e0 Microsoft ISA/ATP Adapter #2
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0          10.0.1.1         10.0.1.104       266
0.0.0.0                    0.0.0.0          10.0.0.1         10.0.0.59        5
10.0.0.0                   255.255.255.0   On-link         10.0.0.59        261
10.0.0.59                 255.255.255.255 On-link         10.0.0.59        261
10.0.0.255                255.255.255.255 On-link         10.0.0.59        261
10.0.1.0                  255.255.255.0   On-link         10.0.1.104       261
10.0.1.104               255.255.255.255 On-link         10.0.1.104       261
10.0.1.255                255.255.255.255 On-link         10.0.1.104       261
127.0.0.0                 255.0.0.0       On-link         127.0.0.1        306
127.0.0.1                 255.255.255.255 On-link         127.0.0.1        306
127.255.255.255          255.255.255.255 On-link         127.0.0.1        306
169.254.169.254          255.255.255.255 10.0.0.254      10.0.0.59        6
224.0.0.0                 240.0.0.0       On-link         127.0.0.1        306
224.0.0.0                 240.0.0.0       On-link         10.0.0.59        261
224.0.0.0                 240.0.0.0       On-link         10.0.1.104       261
255.255.255.255          255.255.255.255 On-link         127.0.0.1        306
255.255.255.255          255.255.255.255 On-link         10.0.0.59        261
255.255.255.255          255.255.255.255 On-link         10.0.1.104       261
=====
Persistent Routes:
Network Address            Netmask          Gateway Address   Metric
0.0.0.0                    0.0.0.0          10.0.1.1         261
=====

IPv6 Route Table
=====
Active Routes:
If Metric Network Destination      Gateway
1       306   ::1/128                    On-link
14      261   fe80::/64                  On-link
19      261   fe80::/64                  On-link
19      261   fe80::197b:3504:e05:5a4d/128 On-link
14      261   fe80::e115:8e6a:5dcc:6715/128 On-link
1       306   ff00::/8                   On-link
14      261   ff00::/8                   On-link
19      261   ff00::/8                   On-link
=====
Persistent Routes:
None
```

6. Check whether the source ECS and the destination ECS can communicate with each other.

```
ping -S IP address of the primary NIC on the source ECS IP address of the destination ECS
```

```
ping -S IP address of the extension NIC on the source ECS IP address of the destination ECS
```

In this example, run the following commands:

```
ping -S 10.0.0.59 10.0.2.12
```

```
ping -S 10.0.1.104 10.0.2.12
```

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS.

```
C:\Users\Administrator>ping -S 10.0.0.59 10.0.2.12

Pinging 10.0.2.12 from 10.0.0.59 with 32 bytes of data:
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64

Ping statistics for 10.0.2.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Users\Administrator>ping -S 10.0.1.104 10.0.2.12

Pinging 10.0.2.12 from 10.0.1.104 with 32 bytes of data:
Reply from 10.0.2.12: bytes=32 time=4ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64
Reply from 10.0.2.12: bytes=32 time<1ms TTL=64

Ping statistics for 10.0.2.12:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 4ms, Average = 1ms
```


6 Access Control


6.1 Access Control Overview

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	Does not support Allow or Deny rules.	Supports both Allow and Deny rules.
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	<p>If an instance is associated with multiple security groups that have multiple rules:</p> <ol style="list-style-type: none"> 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. Rules are then matched by priority in that security group. Rules with lower values have higher priorities than those with higher values. Deny rules take precedence over allow rules if the rules have the same priority. 	<p>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.</p>

Item	Security Group	Network ACL
Usage	<ul style="list-style-type: none">• 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.• After creating an instance, you can:<ul style="list-style-type: none">– Add or remove the instance to or from the security group on the security group console.– Associate or disassociate a security group with or from the instance on the instance console.	Selecting a network ACL is not allowed when you create a subnet. You must create a network ACL, add inbound and outbound rules, associate subnets with it, and enable network ACL. The network ACL then protects the associated subnets and instances in the subnets.

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.

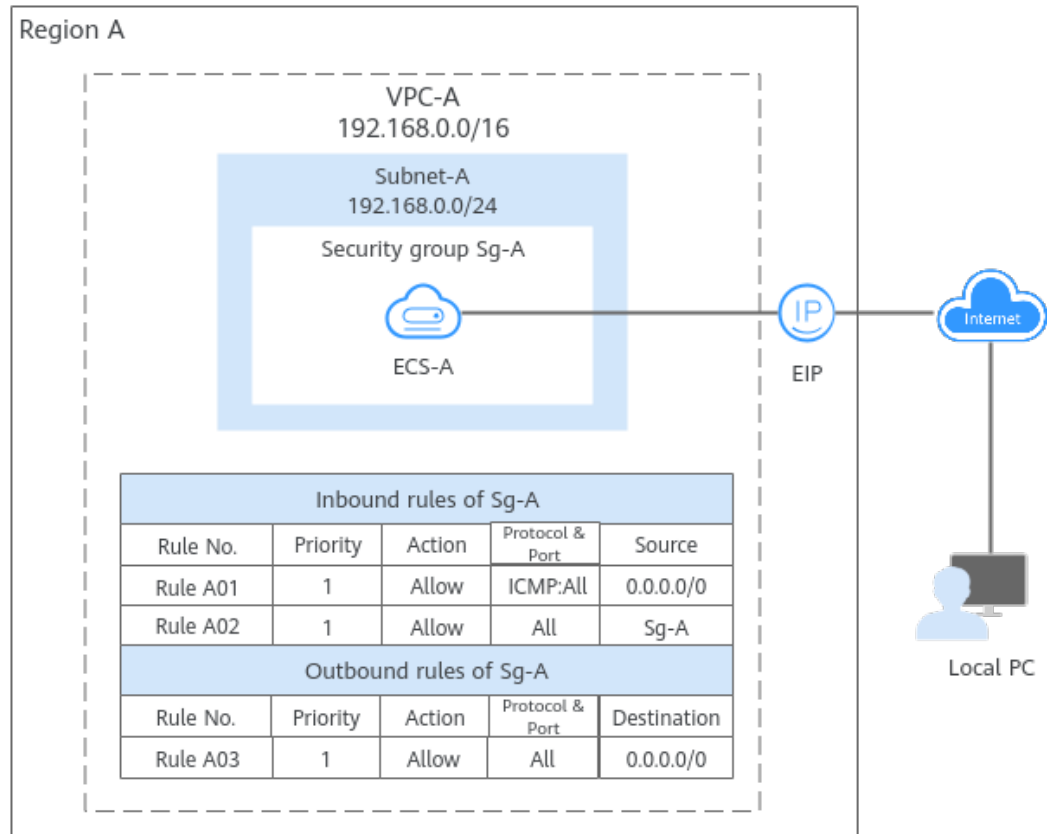
When creating an instance (for example, an ECS), you must associate it with a security group. If there are no security groups yet, a **default security group** will be automatically created and associated with the instance. You can also create a security group based on service requirements and associate 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 6-2**, 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 have 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 6-2 A 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 is 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 to access 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 an easier 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 work 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 6-2](#) 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 6-2](#) 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.

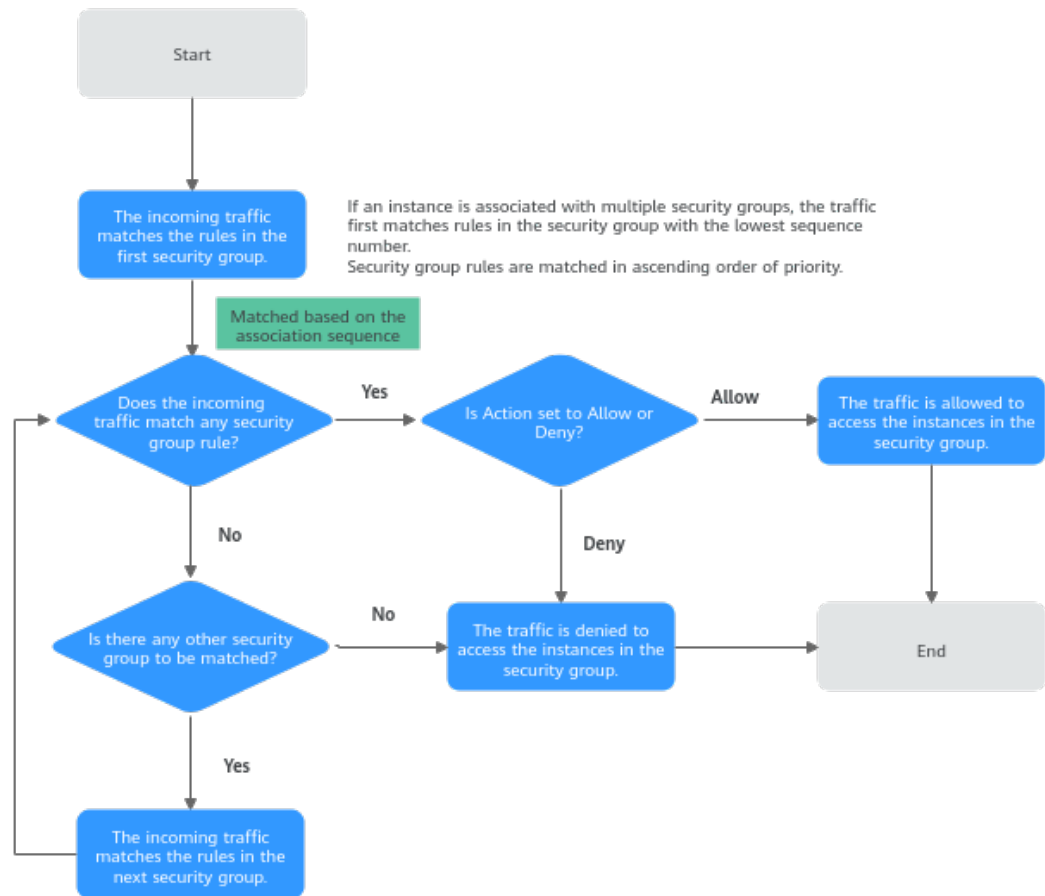
Table 6-2 Security group rules

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 6-3 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 Security Groups

In [Figure 6-4](#), 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**.

- 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.
- Inbound rule A02 of **Sg-A** allows the ECSs in this security group to communicate with each other using any protocol and port.
- 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**.
- 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 6-4 Allowing traffic from given IP addresses and security groups



NOTE

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

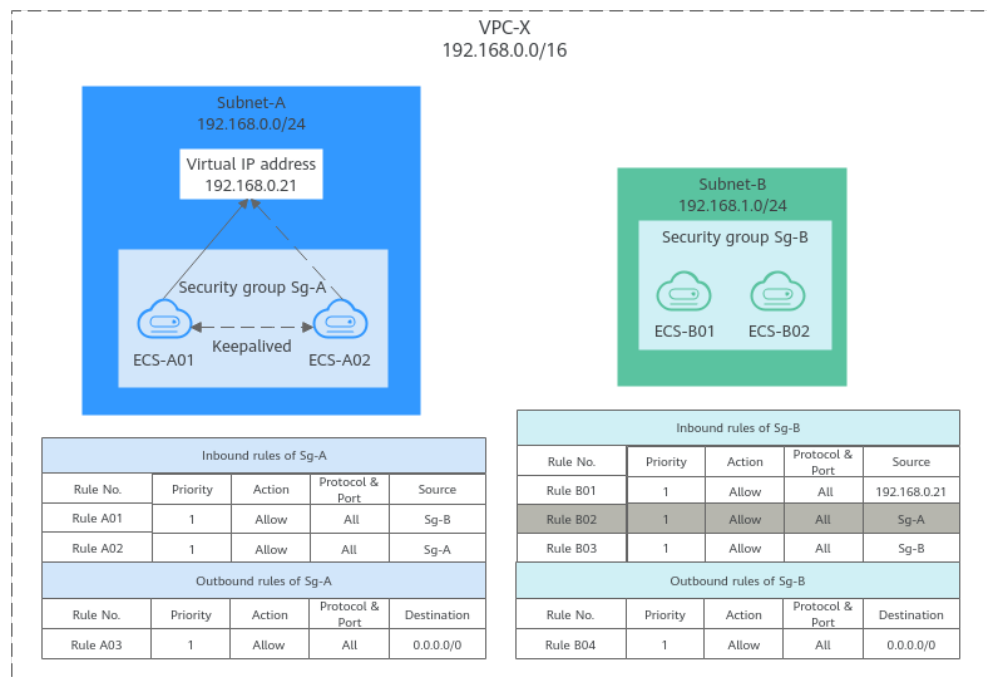
Allowing Traffic from a Virtual IP Address

In [Figure 6-5](#), 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 6-5](#), **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 6-5 Allowing traffic from a virtual IP address



NOTE

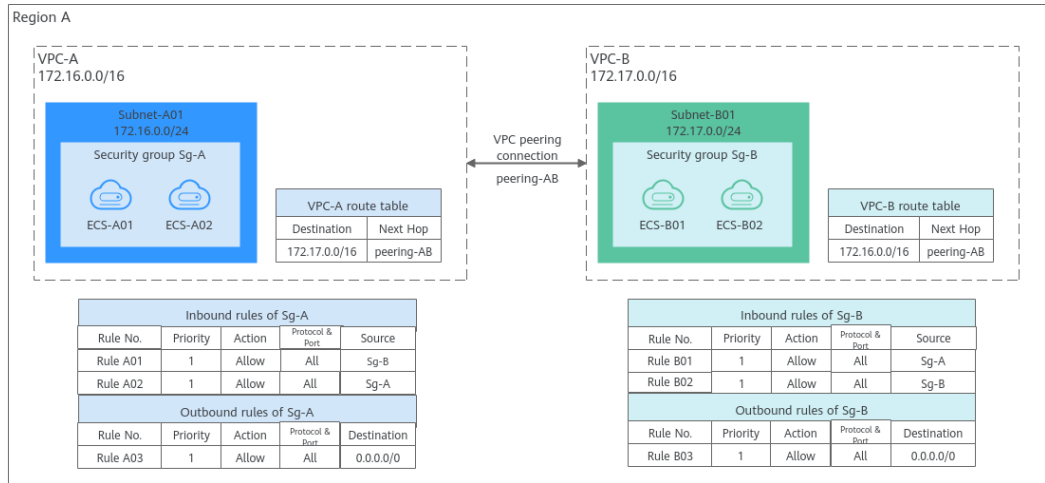
[Security Group Examples](#) lists more security group rule configuration examples.

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

In [Figure 6-6](#), 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 6-6 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 6-7 Process of using a security group

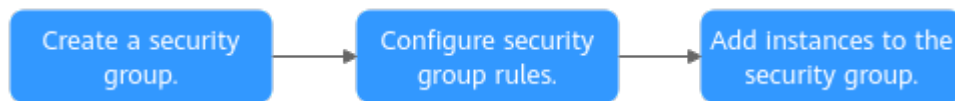


Table 6-3 Security group configuration process description

No.	Step	Description	Reference
1	Create a security group.	When creating a security group, you can use the preset rules.	Creating a Security Group
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.	Adding a Security Group Rule Fast-Adding Security Group Rules
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.	Adding an Instance to or Removing an Instance from a Security Group

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.
- 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 an easier 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 6-8 Default security group

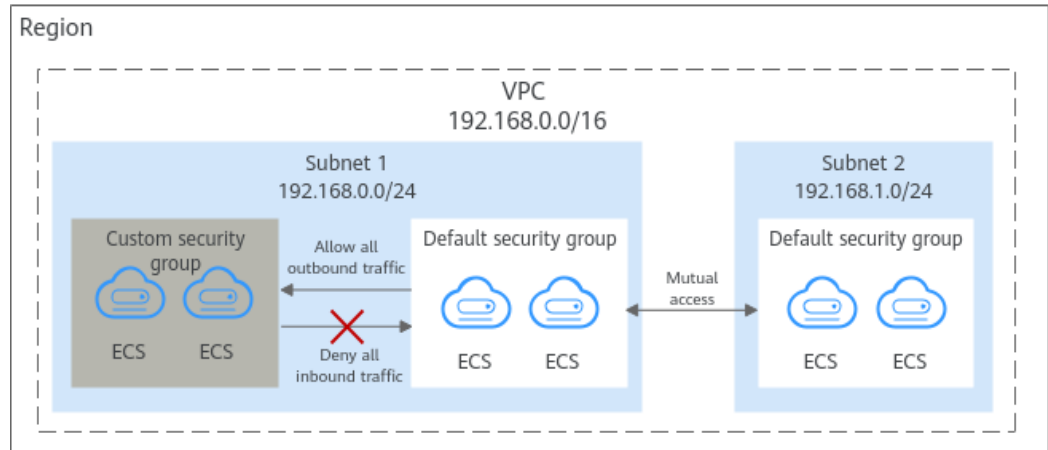


Table 6-4 describes the default rules for the default security group.

Table 6-4 Default security group rules

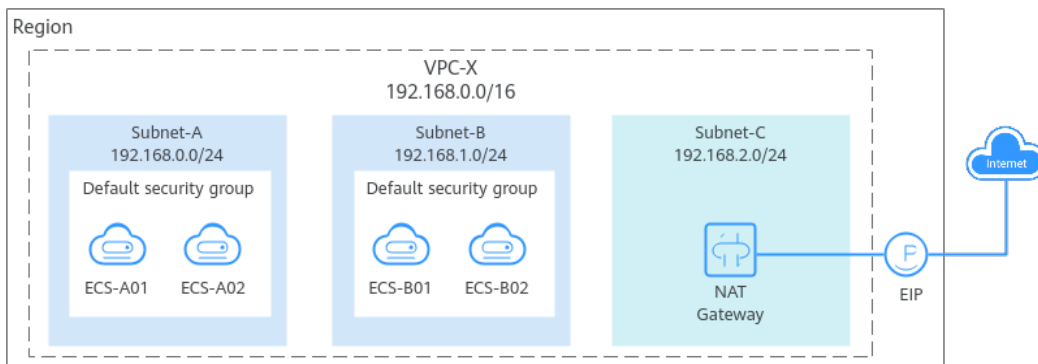
Direction	Action	Type	Protocol & Port	Source/Destination	Description
Inbound	Allow	IPv4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inbound	Allow	IPv6	All	Source: default security group (default)	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 6-9**, 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 6-9 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 6-5](#).

Table 6-5 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 6-6](#).
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see [Table 6-7](#).

Table 6-6 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 6-7 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 6-8](#).

Table 6-8 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 6-9 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 6-10 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 6-11 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 6-12 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 6-13 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 sg-A to access the MySQL database service.
Inbound	1	Allow	IPv4	TCP: 1521	Security group: sg-B	Allows the ECSs in security group sg-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 ipGroup-A 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 6-15](#) 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 6-14 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 6-15 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 6-16](#) 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 6-16 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 6-17 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 6-17 Common ports used by ECSs

Port	Protocol	Description
21	FTP	Used by FTP services for uploading and downloading files. For configuration examples, see Remotely Connecting to an ECS from a Local Server to Upload or Download FTP Files .
22	SSH	Used to remotely connect to Linux ECSs. For configuration examples, see Remotely Logging In to an ECS from a Local Server .
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 Why Is Outbound Access Through TCP Port 25 Restricted?
80	HTTP	Used to access websites over HTTP. For configuration examples, see Setting Up a Website on an ECS to Provide Services Externally .
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 Setting Up a Website on an ECS to Provide Services Externally .
1433	SQL Server	A TCP port of the SQL Server for providing services. For configuration examples, see Allowing External Instances to Access the Database Deployed on an ECS .
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 Allowing External Instances to Access the Database Deployed on an ECS .
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 Allowing External Instances to Access the Database Deployed on an ECS .
3306	MySQL	Used by MySQL databases to provide services. For configuration examples, see Allowing External Instances to Access the Database Deployed on an ECS .
3389	Windows Server Remote Desktop Services	Used to connect to Windows ECSs. For configuration examples, see Remotely Logging In to an ECS from a Local Server . For details about how to log in to a Windows ECS, see .
8080	Proxy	Used by the WWW proxy service for web browsing, like port 80. If you use port 8080, you need to add :8080 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">• Ports 137 and 138: UDP ports that are used when files are transferred using Network Neighborhood (My Network Places).• Port 139: Connections from this port try to access the NetBIOS/SMB service.

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 6-18](#) describes the security group templates.

Table 6-18 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"> Remotely log in to an instance (such as an ECS) in a security group from an external network. Enable external servers to ping the instances in a security group to verify network connectivity. Use instances in a security group as web servers to provide website services accessible from the Internet.
		TCP: 3389 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 3389 (RDP) for remotely logging in to Windows instances.	
		TCP: 80 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 80 (HTTP) for visiting websites.	
		TCP: 443 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over port 443 (HTTPS) for visiting websites.	
		ICMP: All (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in a security group over any port for using the ping command to test connectivity.	
		All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network over any protocol and port.	

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

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

Procedure

1. 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 6-19 Parameter description

Parameter	Description	Example Value
Name	Mandatory The name of the security group. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.). NOTE You can change the security group name after a security group is created. It is recommended that you give each security group a different name.	sg-AB
Enterprise Project	Mandatory When creating a security group, you can add the security group to an enabled enterprise project. An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is default . For details about creating and managing enterprise projects, see the <i>Enterprise Management User Guide</i> .	default

Parameter	Description	Example Value
Template	Mandatory 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. Table 6-18 describes the security group templates.	General-purpose web server
Description (Optional)	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.

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 Deleting a Security Group

Scenarios

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

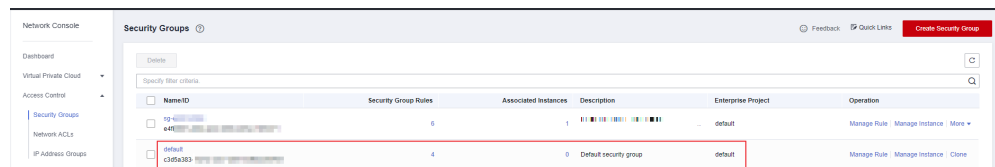
NOTE

Both default and custom security groups are free.

Notes and Constraints

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

Figure 6-10 Default security group





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

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

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

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and 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 **OK**.

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 6-20](#) shows the inbound rules.


Table 6-20 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 target security group and click **Manage Rules** 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 6-11 Add Inbound Rule

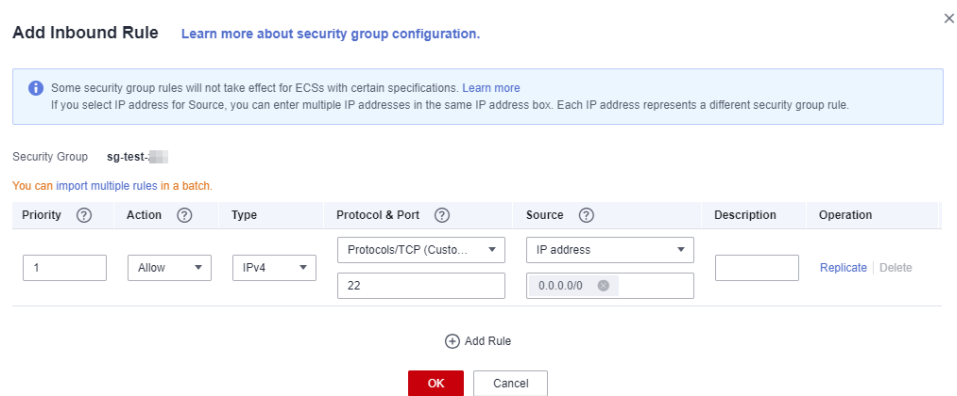


Table 6-21 Inbound rule parameter description

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

Parameter	Description	Example Value
Action	<p>The value can be Allow or Deny.</p> <ul style="list-style-type: none">If the Action is set to Allow, access from the source is allowed to ECSs in the security group over specified ports.If the Action is set to Deny, access from the source is denied to ECSs in the security group over specified ports. <p>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.</p>	Allow
Type	<p>Source IP address version. You can select:</p> <ul style="list-style-type: none">IPv4IPv6	IPv4
Protocol & Port	<p>The network protocol used to match traffic in a security group rule. The protocol can be All, TCP, UDP, GRE, or ICMP.</p>	TCP
	<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">Individual port: Enter a port, such as 22.Consecutive ports: Enter a port range, such as 22-30.All ports: Leave it empty or enter 1-65535.	22, or 22-30
Source	<p>Source of the security group rule. The value can be an IP address or a security group, to allow access from IP addresses or instances in the security group.</p> <ul style="list-style-type: none">IP address:<ul style="list-style-type: none">Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) <p>If the source is a security group, this rule will apply to all instances associated with the selected security group.</p>	0.0.0.0/0

Parameter	Description	Example Value
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	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 6-12 Add Outbound Rule

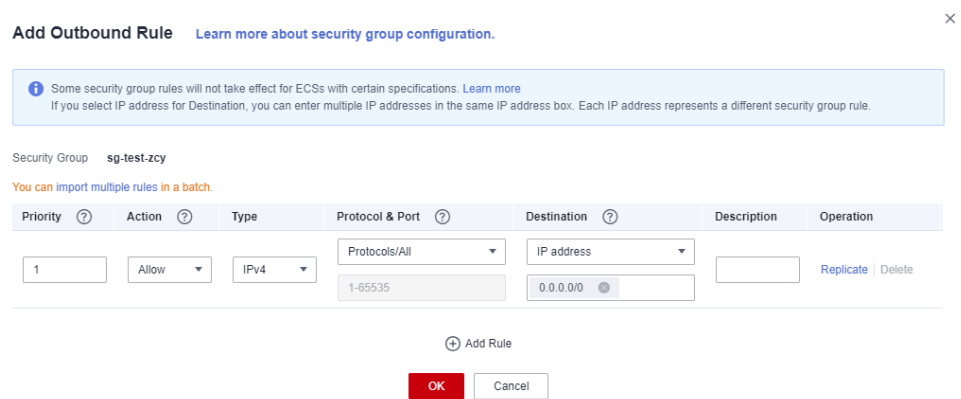


Table 6-22 Outbound rule parameter description

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

Parameter	Description	Example Value
Action	<p>The value can be Allow or Deny.</p> <ul style="list-style-type: none"> If the Action is set to Allow, access from ECSs in the security group is allowed to the destination over specified ports. If the Action is set to Deny, access from ECSs in the security group is denied to the destination over specified ports. <p>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.</p>	Allow
Type	<p>Destination IP address version. You can select:</p> <ul style="list-style-type: none"> IPv4 IPv6 	IPv4
Protocol & Port	<p>The network protocol used to match traffic in a security group rule. The protocol can be All, TCP, UDP, GRE, or ICMP.</p>	TCP
	<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. Specify one of the following:</p> <ul style="list-style-type: none"> Individual port: Enter a port, such as 22. Consecutive ports: Enter a port range, such as 22-30. All ports: Leave it empty or enter 1-65535. 	22, or 22-30
Destination	<p>Destination of the security group rule. The value can be an IP address or a security group, to allow access to IP addresses or instances in the security group.</p> <ul style="list-style-type: none"> IP address: <ul style="list-style-type: none"> Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	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 (< or >).</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 6-23](#) shows the rule.

Table 6-23 Security group rule

Direction	Protocol & Port	Source
Inbound	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 6-13 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 6-14 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



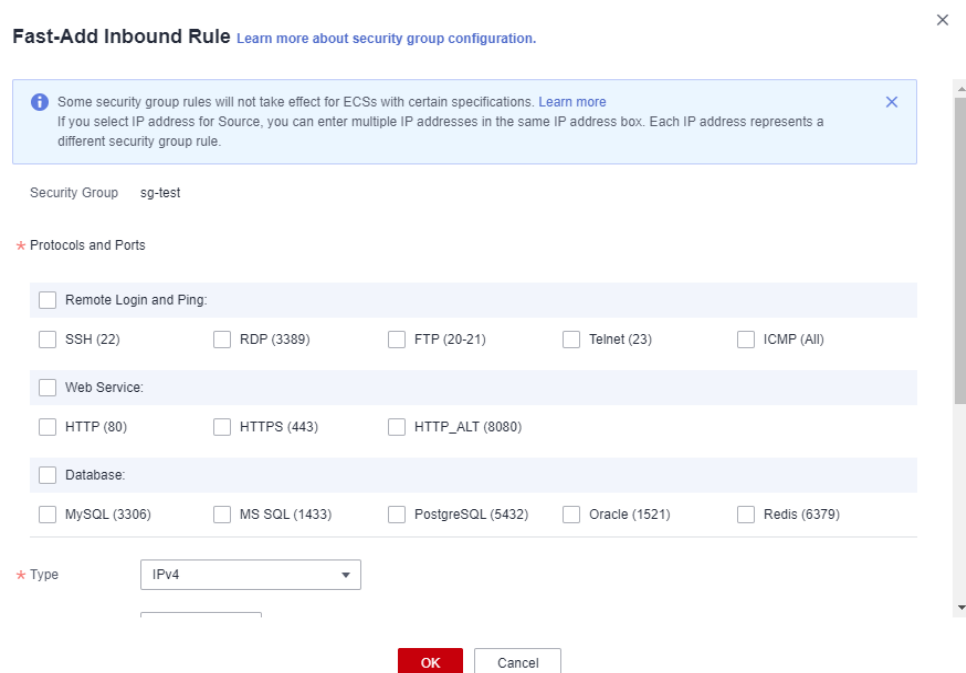
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Click  in the upper left corner and choose **Networking** > **Virtual Private Cloud**.
The **Virtual Private Cloud** page is displayed.
4. In the navigation pane on the left, choose **Access Control** > **Security Groups**.
The security group list is displayed.
5. Locate the row that contains the target security group and click **Manage Rules** 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 6-15 Fast-Add Inbound Rule



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

Security Group sg-test

*** 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

OK **Cancel**

Table 6-24 Inbound rule parameter description

Parameter	Description	Example Value
Protocols and Ports	Common protocols and ports are provided for: <ul style="list-style-type: none">• Remote login and ping• Web services• Databases	SSH (22)
Type	Source IP address version. You can select: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Source	Source of the security group rule. The value can be an IP address or a security group to allow access from IP addresses or instances in the security group. You can specify: <ul style="list-style-type: none">• Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6)• IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)• All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)• Security group: sg-abc If the source is a security group, this rule will apply to all instances associated with the selected security group.	0.0.0.0/0
Action	The value can be Allow or Deny . <ul style="list-style-type: none">• If the Action is set to Allow, access from the source is allowed to ECSs in the security group over specified ports.• If the Action is set to Deny, access from the source is denied to ECSs in the security group over specified ports. 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 .	Allow
Priority	Security group rule priority. 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.	1

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

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 6-16 Fast-Add Outbound Rule

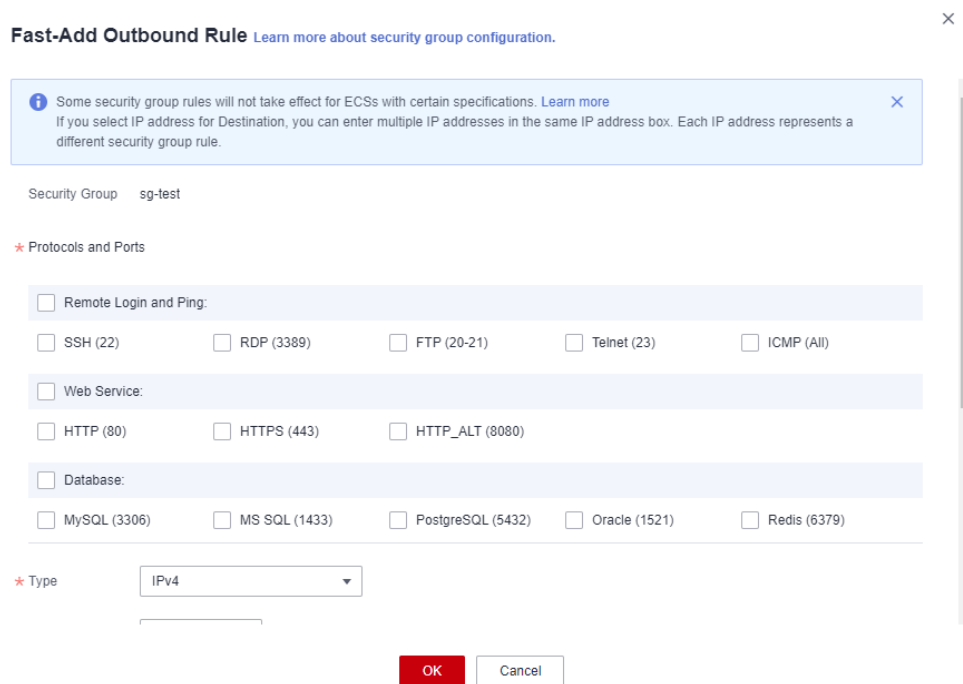


Table 6-25 Outbound rule parameter description

Parameter	Description	Example Value
Protocols and Ports	Common protocols and ports are provided for: <ul style="list-style-type: none"> • Remote login and ping • Web services • Databases 	SSH (22)

Parameter	Description	Example Value
Type	Source IP address version. You can select: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Destination	Destination of the security group rule. The value can be an IP address or a security group to allow access to IP addresses or instances in the security group. You can specify: <ul style="list-style-type: none">• xxx.xxx.xxx.xxx/32 (IPv4 address)• xxx.xxx.xxx.0/24 (IPv4 address range)• 0.0.0.0/0 (all IPv4 addresses)• sg-abc (security group)• Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6)• IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)• All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)• Security group: sg-abc	0.0.0.0/0
Priority	Security group rule priority. 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.	1
Action	The value can be Allow or Deny . <ul style="list-style-type: none">• If the Action is set to Allow, access from ECSs in the security group is allowed to the destination over specified ports.• If the Action is set to Deny, access from ECSs in the security group is denied to the destination over specified ports. 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 .	Allow
Description	(Optional) Supplementary information about the security group rule. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

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 6-26](#) describes the common ports that can be opened with a few clicks.

Table 6-26 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.
	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 6-27](#) ensure that instances in the security group can communicate with each other. Do not modify these rules.
- The outbound rules in [Table 6-27](#) 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 6-27 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 target 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 target rule and click **Replicate** in the **Operation** column.
The **Replicate Inbound Rule** or **Replicate Outbound 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 a 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 a security group rule to be imported is the same as an existing one, the security group rule cannot be imported. You can delete the rule and try again.
- Do not import two security group rules with the same **Direction, Type, Protocol & Port**, and **Source/Destination**, but different **Action** configurations. [Table 6-28](#) 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 6-28 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 6-29 describes the parameters in the template for importing rules.

Table 6-29 Template parameters

Parameter	Description	Example Value
Direction	The direction in which the security group rule takes effect. <ul style="list-style-type: none"> • Inbound: Inbound rules control incoming traffic to instances in the security group. • Outbound: Outbound rules control outgoing traffic from instances in the security group. 	Inbound

Parameter	Description	Example Value
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	The value can be Allow or Deny . <ul style="list-style-type: none">If the Action is set to Allow, access from the source is allowed to ECSs in the security group over specified ports.If the Action is set to Deny, access from the source is denied to ECSs in the security group over specified ports. 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 .	Allow
Protocol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	TCP
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Inbound rules control incoming traffic over specific ports to instances in the security group. Outbound rules control outgoing traffic over specific ports from instances in the security group. Specify one of the following: <ul style="list-style-type: none">Individual port: Enter a port, such as 22.Consecutive ports: Enter a port range, such as 22-30.All ports: Leave it empty or enter 1-65535.	22, or 22-30
Type	Source IP address version. You can select: <ul style="list-style-type: none">IPv4IPv6	IPv4

Parameter	Description	Example Value
Source	Source of the security group rule. The value can be an IP address or a security group, to allow access from IP addresses or instances in the security group. <ul style="list-style-type: none">IP address:<ul style="list-style-type: none">Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	sg-test[96a8a93f-XXX-d7872990c314]
Destination	Destination of the security group rule. The value can be an IP address or a security group, to allow access to IP addresses or instances in the security group.	sg-test[96a8a93f-XXX-d7872990c314]
Description	Supplementary information about the security group rule. This parameter is optional. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Last 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 6-30](#) ensure that instances in the security group can communicate with each other. Do not delete these rules.
- The outbound rules in [Table 6-30](#) 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 6-30 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**.

6.2.6.8 Querying Security Group Rule Changes

Scenarios

CTS records the changes made to security group rules. You can query the change details of:

- New security group rules
- Modified security group rules
- Deleted security group rules

Precautions

- To use CTS to record security group rule changes, you need to [enable CTS](#) first.
- CTS records operations performed on each cloud service. You can query specific operations by trace name, resource type, or operation time. [Table 6-31](#) lists the operations on security group rules supported by CTS.

Table 6-31 Operations on security group rules supported by CTS

Operation	Trace Name	Resource Type
Adding a security group rule	createSecurity-group-rule	security-group-rules
Modifying a security group rule	updateSecurity-group-rule	security-group-rules
Deleting a security group rule	deleteSecurity-group-rule	security-group-rules

Procedure

The following describes how to view the rule described in [Table 6-32](#) that is added to security group **Sg-A**.


Table 6-32 The new security group rule

Direction	Action	Type	Protocol & Port	Source	Last Modified
Inbound	Allow	IPv4	TCP: 23	10.0.0.0/16	June 19, 2024 10:46:07 GMT+08:00

- Log in to the CTS console, search for the operations by trace name (**createSecurity-group-rule** in this example) and locate the specific operation by operation time.

For details, see [Querying Real-Time Traces](#).

Figure 6-17 The trace list for new security group rules



- In the trace list, locate the target trace and click its name.

On the **Trace Overview** page, you can view the details about the operation. **Table 6-33** provides the detailed information about the operation, including operator ID and details about the security group rules.

 **NOTE**

The trace details in **Table 6-33** are only for your reference. The actual information may vary.

Table 6-33 The trace details for the new security group rule

Example Command Output	Description
<pre>"source_ip": "124.71.XX.146",</pre>	<p>IP address of the client that performs the operation. If this parameter is left blank, the operation is performed by the system. In this example, the IP address is 124.71.XX.146.</p>
<pre>"user": { "access_key_id": "HSTA205XXXXXC4MHAE", "account_id": "3c24f6f885294XXXXX93ce075fbd", "user_name": "cts-test-01", "domain": { "name": "cts-test", "id": "3c24f6f885294XXXXX93ce075fbd" }, "name": "cts-test-01", "principal_is_root_user": "false", "id": "a26ee7e7224XXXXXe4a28a9ce503",</pre>	<p>Account of the operator who performs the operation. Key parameters are described as follows:</p> <ul style="list-style-type: none"> name under domain: indicates the account name. In this example, the account name is cts-test. id under domain: indicates the account ID. In this example, the ID is 3c24f6f885294XXXXX93ce075fbd. name: IAM username. In this example, the username is cts-test-01, which is an IAM user under account cts-test. id: IAM user ID. In this example, the ID is a26ee7e7224XXXXXe4a28a9ce503. <p>For details about more parameters of CTS traces, see the response parameter description in Trace Structure.</p>

Example Command Output	Description
<pre> "response": "{\request_id \": \"8d2d1111cafaXXX9b49d53e2da38f \", \"security_group_rules\": [{\id\": \"b6acda6e-0976- XXX-82bc-a8093cbd591d\", \"project_id \": \"15289aca74eXXXa37dea0315d99\", \"security_group _id\": \"3730d371-3111-4ace-XXX- b00b7259e178\", \"remote_group_id\": null, \"direction \": \"ingress\", \"protocol\": \"tcp\", \"description \": \"\", \"created_at\": \"2024-06-19T02:46:07Z \", \"updated_at\": \"2024-06-19T02:46:07Z \", \"ethertype\": \"IPv4\", \"remote_ip_prefix \": \"10.0.0/16\", \"multiport \": \"23\", \"remote_address_group_id\": null, \"action \": \"allow\", \"priority\": 1}]}", </pre>	<p>Details about the security group rule in response. Key parameters are described as follows:</p> <ul style="list-style-type: none"> • direction: indicates the direction of the security group rule. ingress indicates the inbound direction, and egress indicates the outbound direction. In this example, ingress is returned, indicating an inbound rule is added. • protocol: indicates the protocol of the security group rule. In this example, the protocol is TCP. • ethertype: indicates the source IP address version. In this example, the version is IPv4. • remote_ip_prefix: indicates the source or destination of the security group rule. In this example, an inbound rule is added, so this parameter indicates IP address range 10.0.0/16. • multiport: indicates the port used to filter traffic. In this example, the port is 23. • action: indicates whether to allow or deny traffic. allow indicates traffic is allowed, while deny indicates traffic is denied. In this example, the action is allow. • priority: indicates the priority of the security group rule. In this example, the priority is 1.

6.2.7 Managing Instances Added to 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.
7. Click the **Servers** tab and click **Add**.
The **Add Server** dialog box is displayed.
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.
6. Click the required instance type tab.
The following operations use **Servers** as an example.
7. 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.
8. Confirm the information and click **OK**.


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

1. Log in to the management console.
2. Click . Under **Compute**, click **Elastic Cloud Server**.
3. 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.
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

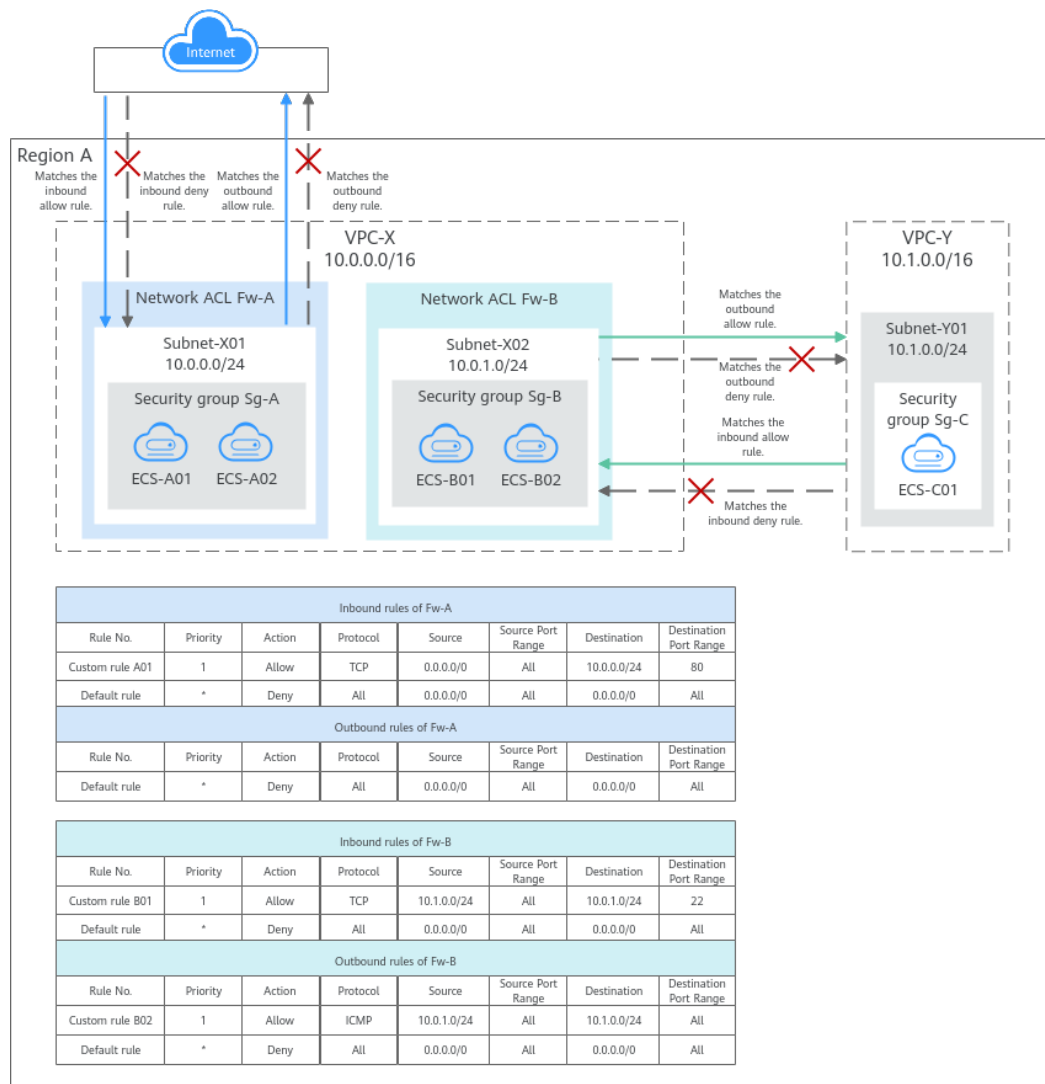
A network ACL is an optional layer of security for your subnets. After you add inbound and outbound rules to a network ACL and associate subnets with it, you can control traffic in and out of the subnets.

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 the entire subnet. 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 6-18](#). **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**:
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**:
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.
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 6-18 Network ACL rules



NOTE

Figure 6-18 shows how network ACLs control traffic in and out of subnets. In actual services, the security groups control traffic from and to the instances associated with it. For details about network ACLs and security groups, 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.
- **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 6-34](#). If no custom rules are configured, 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 6-34 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 6-35](#).

Table 6-35 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

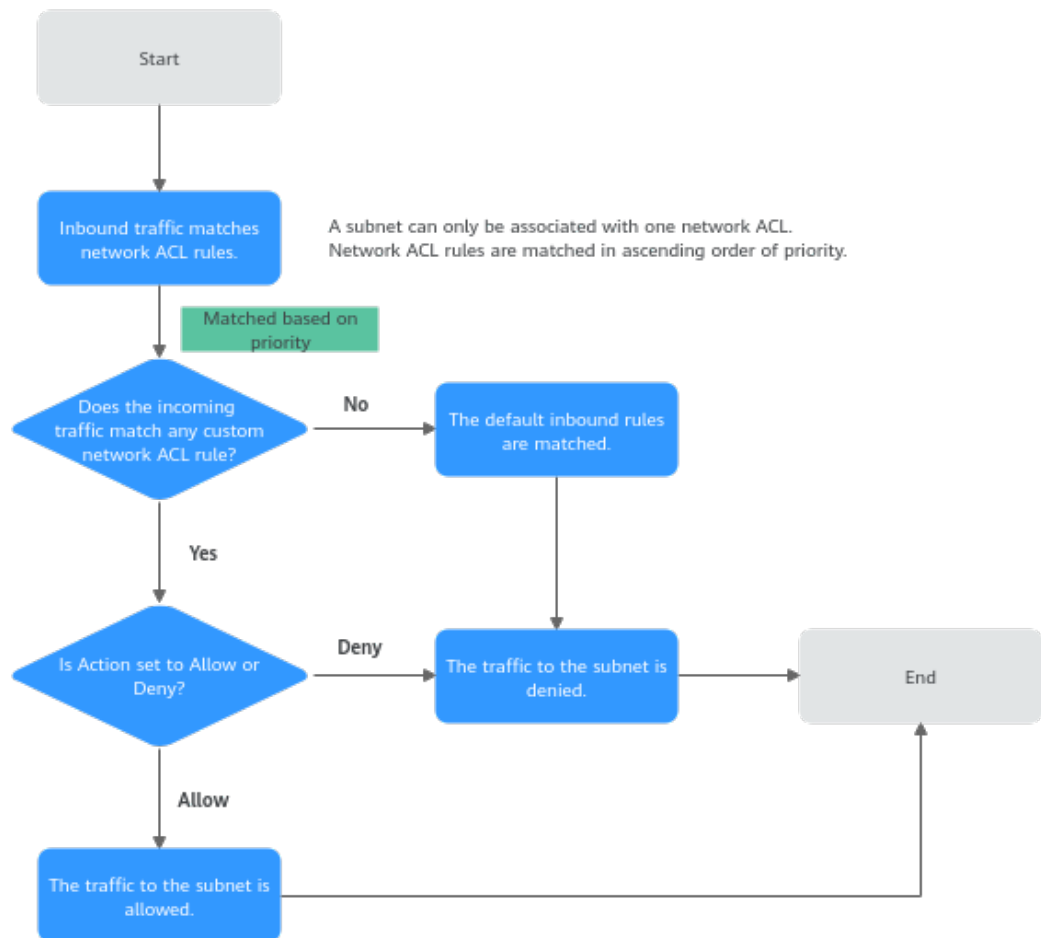
Direction	Description
	Broadcast traffic to 255.255.255.255/32
	Multicast traffic to 224.0.0.0/24
	TCP metadata traffic to 169.254.169.254/32 over port 80
	Traffic to 100.125.0.0/16 that is reserved for public services on the cloud, such as the DNS server address and NTP server address.

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 6-19 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.

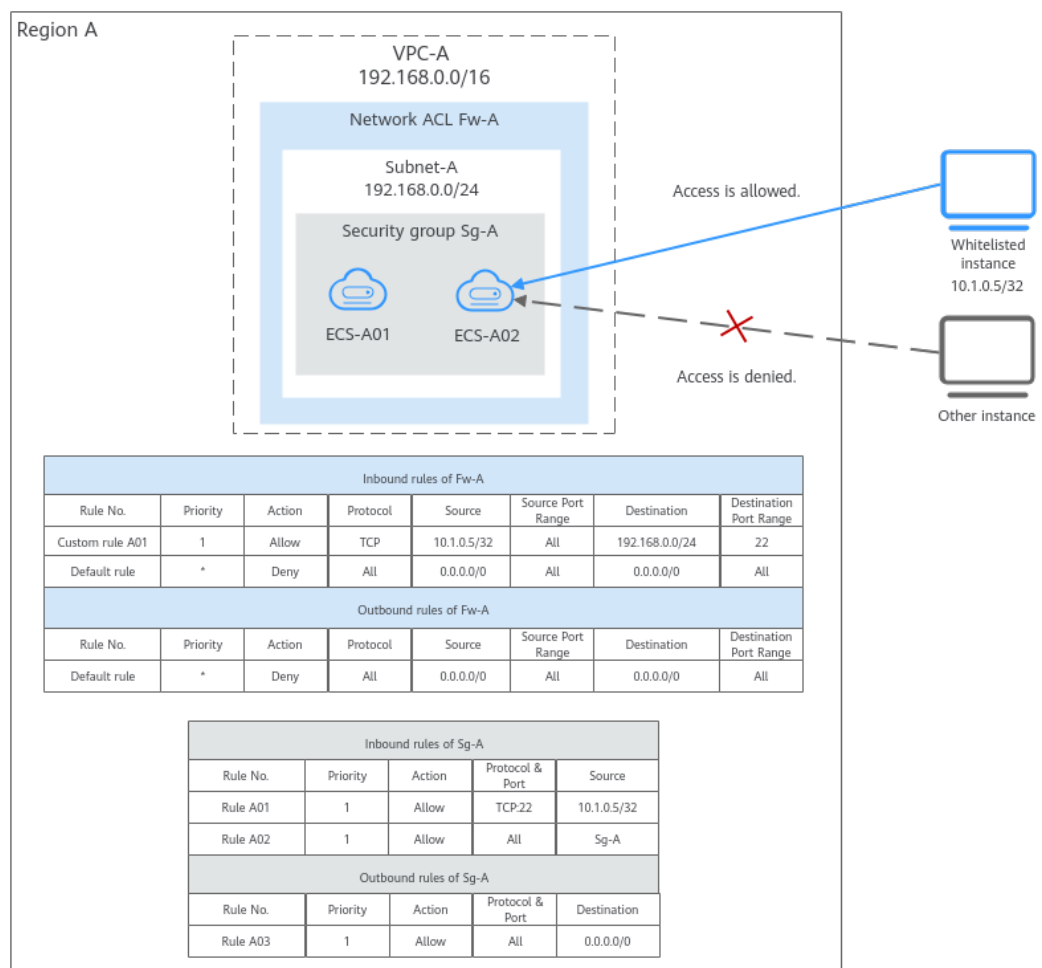
Controlling External Access to Instances in a Subnet

As shown in [Figure 6-20](#), 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 6-20 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 6-36](#), 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 6-36 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

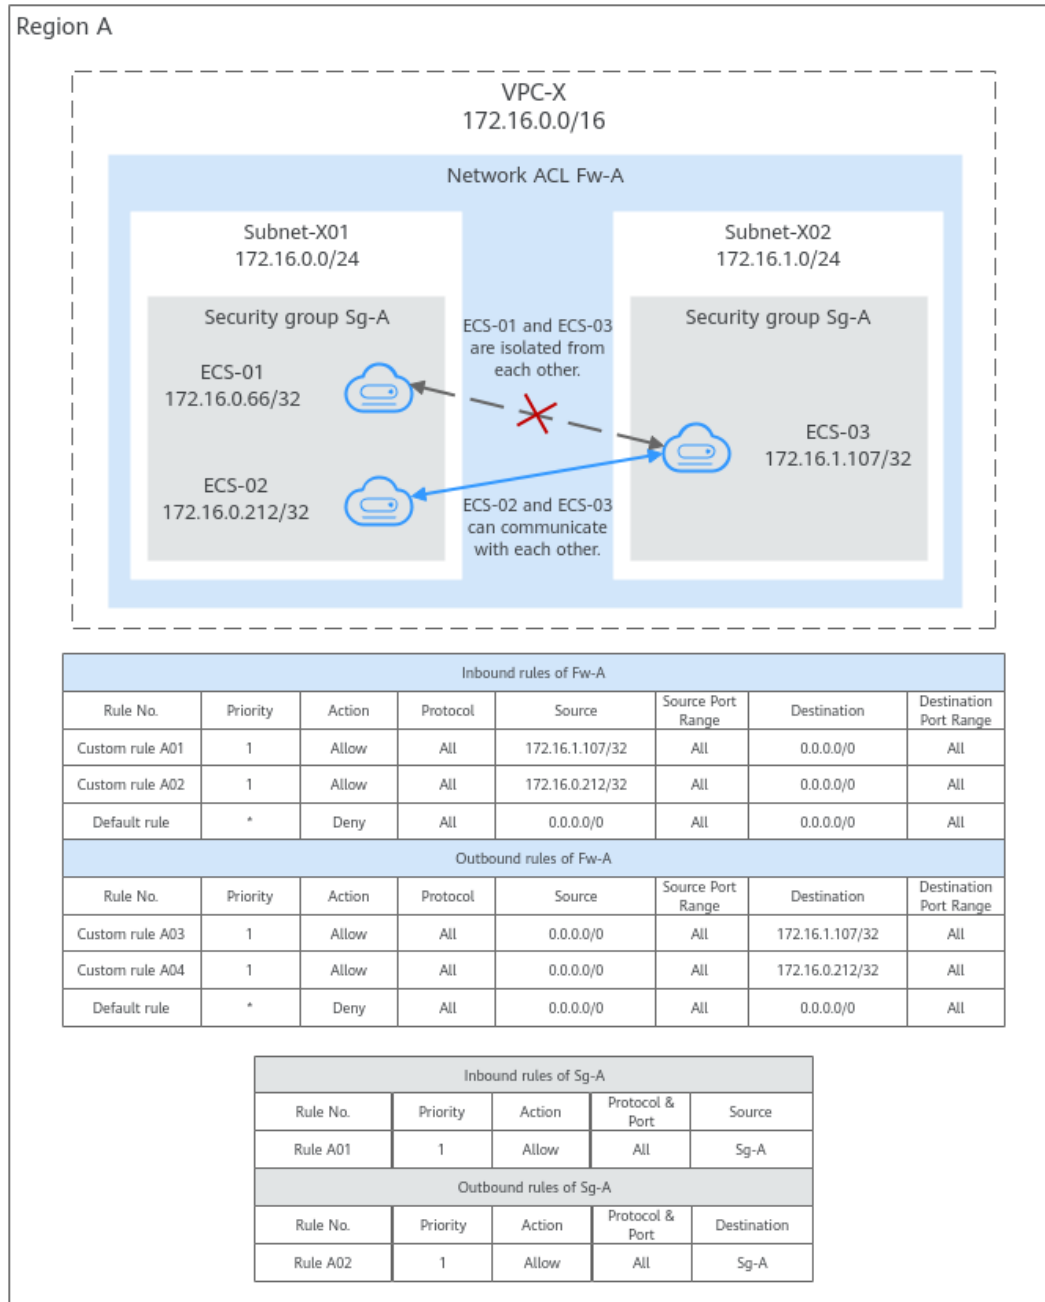
As shown in [Figure 6-21](#), VPC-X has two subnets: **Subnet-X01** and **Subnet-X02**. **ECS-01** and **ECS-02** work in **Subnet-X01**, and **ECS-03** works in **Subnet-X02**. Suppose you want to:

- Connect **ECS-02** to **ECS-03**.
- Isolate **ECS-01** from **ECS-03**.

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 6-21 Controlling communications between instances in different subnets



NOTE

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

Network ACL Configuration Procedure

Figure 6-22 Procedure for configuring a network ACL




Table 6-37 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.	Creating a Network ACL
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.	Adding a Network ACL Rule (Default Priorities) Adding a Network ACL Rule (Custom Priorities)
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.	Associating Subnets with a Network ACL

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 6-38](#). If a network ACL has no custom rules, the default rule is applied, denying all traffic in and out of a subnet.

Table 6-38 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 6-39](#) 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 6-39 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 6-40](#) 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 6-40 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 6-41](#).
 - 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 6-41 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 6-42](#).
 - 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 6-42 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 6-43 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
Description (Optional)	Supplementary information about the network ACL. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	N/A

4. 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 **OK**.
 - 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 **OK**.



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**.

6.3.4 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 6-44](#), 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 6-44 Default priorities

Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
Custom rule A	1	Custom rule A	1
--	--	Custom rule B	2
Custom rule B	2	Custom rule C	3
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 6-45 Parameter descriptions

Parameter	Description	Example Value
Type	Network ACL type. There are two options: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Action	The action in the network ACL. There are two options: <ul style="list-style-type: none">• Allow: allows matched traffic in and out of a subnet.• Deny: denies matched traffic in and out of a subnet.	Allow
Protocol	The protocol supported by the network ACL to match traffic. The value can be TCP, UDP, or ICMP .	TCP
Source	The source from which the traffic is allowed. The source can be an IP address or IP address range. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)- All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	192.168.0.0/24
Source Port Range	The source port or port range used to match traffic. The value ranges from 1 to 65535.	22-30
Destination	The destination to which the traffic is allowed. The destination can be an IP address or IP address range. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)- All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	0.0.0.0/0
Destination Port Range	The destination port or port range used to match traffic. The value ranges from 1 to 65535.	22-30

Parameter	Description	Example Value
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 (< or >).</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 6-46](#), 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 6-46 Custom priorities

Priorities (Rules A and B)		Priorities (Rules A, B, and C)	
Custom rule A	1	Custom rule A	1
--	--	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.
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, locate the target rule, click **Modify** in the **Operation** column, and modify parameters based on [Table 6-47](#).

Table 6-47 Parameter descriptions

Parameter	Description	Example Value
Type	Network ACL type. There are two options: <ul style="list-style-type: none">• IPv4• IPv6	IPv4
Action	The action in the network ACL. There are two options: <ul style="list-style-type: none">• Allow: allows matched traffic in and out of a subnet.• Deny: denies matched traffic in and out of a subnet.	Allow
Protocol	The protocol supported by the network ACL to match traffic. The value can be TCP, UDP, or ICMP .	TCP
Source	The source from which the traffic is allowed. The source can be an IP address or IP address range. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)- All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	192.168.0.0/24
Source Port Range	The source port or port range used to match traffic. The value ranges from 1 to 65535.	22-30
Destination	The destination to which the traffic is allowed. The destination can be an IP address or IP address range. <ul style="list-style-type: none">• IP address:<ul style="list-style-type: none">- Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6)- All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6)- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	0.0.0.0/0
Destination Port Range	The destination port or port range used to match traffic. The value ranges from 1 to 65535.	22-30

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

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 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 target subnet and click its name.
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 target subnet and click the name of the network ACL 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** in the **Operation** column.
A confirmation dialog box is displayed.
 - iv. Click **OK** 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 target 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 **OK** 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

An IP address group is a collection of IP addresses that can use the same security group rules or network ACL rules. You can use an IP address group to manage IP addresses that have the same security requirements or whose security requirements change frequently.

You can create an IP address group and add IP addresses that need to be managed in a unified manner to the group. Then, you can select this IP address group when configuring a security group rule. The rule will take effect for all IP addresses in the 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 an easier 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 7-1](#).

Table 7-1 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">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.). You can customize the name of an IP address group that is uniquely identified by its ID.	ipGroup-A

Parameter	Description	Example Value
Max. IP Addresses	<p>Mandatory</p> <p>Set the number of IP addresses that can be added to an IP address group. By default, the system displays the maximum number of IP addresses that can be added to an IP address group. You can change number as required.</p> <p>If you want to increase the maximum number of IP addresses in a group, submit a service ticket.</p>	20
IP Address Version	<p>Mandatory</p> <p>Select the type of IP addresses that can be added to an IP address group.</p> <ul style="list-style-type: none"> • IPv4 • IPv6 	IPv4
IP Addresses (Optional)	<p>Optional</p> <p>Enter an IP address or IP address range on each line, and press Enter. The format is "IP address Description". Description is optional, can contain 0 to 255 characters, and cannot contain angle brackets (< or >). You can enter:</p> <ul style="list-style-type: none"> • An IPv4 address range, for example, 192.168.0.0/16 or 192.168.0.0/16 ECS01 • A single IPv4 address, for example, 192.168.10.10 or 192.168.10.10 ECS01 • An IPv6 address range, for example, 2001:db8:a583:6e::/64 or 2001:db8:a583:6e::/64 ECS01 • A single IPv6 address, for example, 2001:db8:a583:6e::5c or 2001:db8:a583:6e::5c ECS01 	<ul style="list-style-type: none"> • Without description: 192.168.0.0/16 • With description: 192.168.0.0/16 ECS01
Description (Optional)	<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. 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 7-2](#).

Table 7-2 Associating an IP address group with resources

Resource	Description	Reference
Security group	The Source or Destination of a security group rule can be set to IP address group .	Adding a Security Group Rule <ul style="list-style-type: none">• Inbound rule: Set Source to an IP address group.• Outbound rule: Set Destination to an IP address group.

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

7.2.3 Modifying an IP Address Group

Scenarios

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

- Name
- Max. IP Addresses
- Description

Procedure


1. Go to the [IP address group list page](#).
2. In the IP address group list, click the name of the target IP address group.
The basic information page of the IP address group is displayed.
3. 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 7-3](#).

Table 7-3 IP address group parameters

Parameter	Description	Example Value
Name	<p>Mandatory</p> <p>Enter the name of the IP address group. The name:</p> <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.). <p>You can customize the name of an IP address group that is uniquely identified by its ID.</p>	ipGroup-A
Max. IP Addresses	<p>Mandatory</p> <p>Set the number of IP addresses that can be added to an IP address group. By default, the system displays the maximum number of IP addresses that can be added to an IP address group. You can change number as required.</p> <p>If you want to increase the maximum number of IP addresses in a group, submit a service ticket.</p>	20
Description	<p>Optional</p> <p>Enter the description of the IP address group in the text box as required.</p>	-

4. Click .

7.2.4 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. Go to the [IP address group list page](#).
2. 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.

7.2.5 Viewing the 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. Go to the [IP address group list page](#).
2. 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.
3. 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.

7.2.6 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. Go to the [IP address group list page](#).
2. 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**.

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. Go to the [IP address group list page](#).
2. In the IP address group list, click the name of the target IP address group.
The basic information page of the IP address group is displayed.
3. In the left corner above the IP address list, click **Add**.
The **Add IP Address** dialog box is displayed.
4. Add IP addresses to the IP address group as prompted.

Table 7-4 IP address group parameters

Parameter	Description	Example Value
Name	The name of the IP address group.	ipGroup-A
Max. IP Addresses	<p>Mandatory</p> <p>Set the number of IP addresses that can be added to an IP address group. By default, the system displays the maximum number of IP addresses that can be added to an IP address group. You can change number as required.</p> <p>If you want to increase the maximum number of IP addresses in a group, submit a service ticket.</p>	20

Parameter	Description	Example Value
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">• IPv4• IPv6	IPv4
IP Addresses	Mandatory Enter an IP address or IP address range on each line, and press Enter . You can enter: <ul style="list-style-type: none">• An IPv4 address range, for example, 192.168.0.0/16• A single IPv4 address, for example, 192.168.10.10/32• An IPv6 address range, for example, 2001:db8:a583:6e::/64• A single IPv6 address, for example, 2001:db8:a583:6e::5c/128	192.168.0.0/16 192.168.10.10/32

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. Go to the [IP address group list page](#).
2. In the IP address group list, click the name of the target IP address group. The basic information page of the IP address group is displayed.
3. In the left corner above the IP address list, click **Modify**. The **Modify IP Address** dialog box is displayed.

4. Modify the information as prompted.
For details, see [Table 7-5](#).

Table 7-5 Parameters for modifying IP addresses

Parameter	Description	Example Value
Name	The name of the IP address group.	ipGroup-A
Max. IP Addresses	<p>Mandatory</p> <p>Set the number of IP addresses that can be added to an IP address group. By default, the system displays the maximum number of IP addresses that can be added to an IP address group. You can change number as required.</p> <p>If you want to increase the maximum number of IP addresses in a group, submit a service ticket.</p>	20
IP Address Version	<p>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:</p> <ul style="list-style-type: none">• IPv4• IPv6	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 (< or >). You can enter:</p> <ul style="list-style-type: none">• An IPv4 address range, for example, 192.168.0.0/16 or 192.168.0.0/16 ECS01• A single IPv4 address, for example, 192.168.10.10 or 192.168.10.10 ECS01• An IPv6 address range, for example, 2001:db8:a583:6e::/64 or 2001:db8:a583:6e::/64 ECS01• A single IPv6 address, for example, 2001:db8:a583:6e::5c or 2001:db8:a583:6e::5c ECS01	<ul style="list-style-type: none">• Without description: 192.168.0.0/16• With description: 192.168.0.0/16 ECS01

5. Click **OK**.

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

7.3.3 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. Go to the [IP address group list page](#).
2. In the IP address group list, click the name of the target IP address group. The basic information page of the IP address group is displayed.
3. Delete IP addresses:
 - Delete a single IP address.

- i. In the IP address list, locate the target 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 above the IP address list.
A confirmation dialog box is displayed.
 - iii. Confirm the information and click **OK**.

7.4 IP Address Group Configuration Examples

7.4.1 Using IP Address Groups to Reduce the Number of Security Group Rules

Scenarios

An IP address group is a collection of one or more IP addresses. You can use IP address groups when configuring security group rules. If you change the IP addresses in an IP address group, the security group rules are changed accordingly. You do not need to modify the security group rules one by one.

Finance and securities enterprises have high security requirements when planning cloud networks. Access to instances is often controlled based on IP addresses. To simplify security group rule configuration and control access based on IP addresses, you can use IP address groups to manage IP address ranges and IP addresses with the same security requirements. For more information about IP address groups, see [IP Address Group Overview](#).

Suppose your enterprise has an online office system deployed on the cloud. To provide services for different departments, you associate office servers with different security groups based on security levels. These servers are accessed from a large number of IP addresses that may change from time to time.

- If IP address groups are not used, you need to configure multiple rules to control access from different sources. Once the IP addresses change, you need to adjust the rules in each security group one by one. The management workload increases with the number of security groups and rules.
- If IP address groups are used, you can add the IP addresses with the same security requirements to an IP address group and add rules with source set to this IP address group. When an IP address changes, you only need to change it in the IP address group. Then, the security group rules using the IP address group change accordingly. You do not need to modify the security group rules one by one. This simplifies security group management and improves efficiency.

Solution Architecture

In this practice, the instances are associated with three security groups based on different security requirements. In addition, these instances need to be accessed by

specific IP addresses over SSH port 22. To simplify management, you can use IP address groups.

1. Create an IP address group and add IP addresses that need to access the instances.
2. Add inbound rules to allow traffic from the IP address group to the instances in the three security groups.

Table 7-6 Inbound rules

Direction	Action	Type	Protocol & Port	Source
Inbound	Allow	IPv4	TCP:22	IP address group

3. Change the IP addresses in the IP address group if any IP addresses change. Then, the rules using the IP address group change accordingly.

Constraints

Security group rules using IP address groups do not take effect for the following instances:

- 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)

Resource Planning

In this practice, the IP address group and security groups must be in the same region. For details, see [Table 7-7](#). The following resource details are only examples. You can modify them as required.

Table 7-7 Resource planning

Resource	Quantity	Description
IP address group	1	Create an IP address group and add IP addresses that need to access the instances. <ul style="list-style-type: none">• Name: ipGroup-A• Max. IP Addresses: Set it as required. In this practice, 20 is used.• IP Address Version: Set it as required. In this practice, IPv4 is used.• IP Addresses:<ul style="list-style-type: none">- 11.xx.xx.64/32- 116.xx.xx.252/30- 113.xx.xx.0/25- 183.xx.xx.208/28
Security group	3	Add inbound rules to allow traffic from ipGroup-A to the instances in the three security groups, as shown in Table 7-8 .

Table 7-8 Inbound rules

Direction	Action	Type	Protocol & Port	Source
Inbound	Allow	IPv4	TCP:22	ipGroup-A

Procedure

Step 1 Create IP address group **ipGroup-A** and add IP addresses that need to access the instances.

For details, see [Creating an IP Address Group](#).

Step 2 Add inbound rules to allow traffic from **ipGroup-A** to the instances in the three security groups.

For details, see [Adding a Security Group Rule](#).

After the rules are added, traffic from 11.xx.xx.64/32, 116.xx.xx.252/30, 113.xx.xx.0/25, and 183.xx.xx.208/28 are allowed to the Linux ECSs over SSH port 22.

Step 3 Change IP addresses in the IP address group.

After security group rules are added, you can add IP addresses to **ipGroup-A**. For example, you can add 117.xx.xx.0/25 to **ipGroup-A**, and the security groups rule is applied automatically, allowing traffic from 117.xx.xx.0/25 over SSH port 22.

For details, see [Managing IP Addresses in an IP Address Group](#).

----End

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.

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 Two VPCs connected by a VPC peering connection



NOTICE

Currently, VPC peering connections are free.

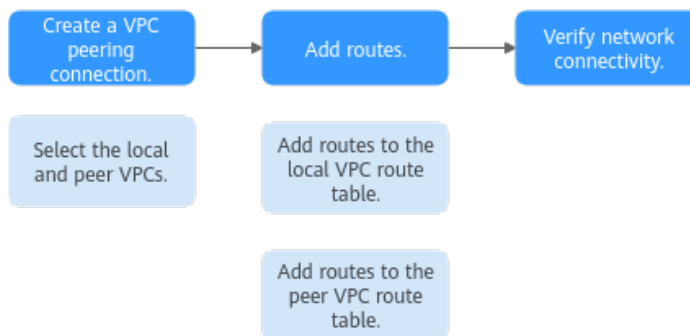
VPC Peering Connection Creation Process

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 to Connect Two VPCs in the Same Account](#).

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

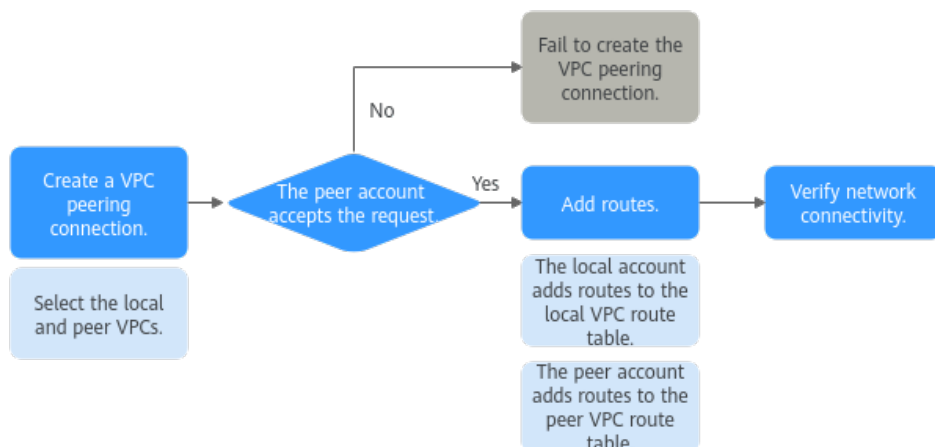


- If two VPCs are in different accounts, the process of creating a VPC peering connection is shown in [Figure 8-3](#).

For details about how to create a VPC peering connection, see [Creating a VPC Peering Connection Connect Two 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.

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



Notes and Constraints

- A VPC peering connection can only connect VPCs in the same region.
 - If 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 be usable.
- 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

8.2.1 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	Example
VPCs in the same region	<ul style="list-style-type: none">• VPC CIDR blocks do not overlap.• Subnet CIDR blocks of VPCs do not overlap.	You can create VPC peering connections to connect entire CIDR blocks of VPCs. Then, all resources in the VPCs can communicate with each other.	Using a VPC Peering Connection to Connect Two VPCs
VPCs in the same region	<ul style="list-style-type: none">• VPC CIDR blocks overlap.• Some subnet CIDR blocks overlap.	You can create VPC peering connections to connect specific subnets or ECSs from different VPCs. <ul style="list-style-type: none">• To connect specific subnets from two VPCs, the subnet CIDR blocks cannot overlap.• To connect specific ECSs from two VPCs, each ECS must have a unique private IP address.	Using a VPC Peering Connection to Connect Subnets in Two VPCs
			Using a VPC Peering Connection to Connect ECSs in Two VPCs

Location	CIDR Block	Description	Example
VPCs in the same region	<ul style="list-style-type: none">VPC CIDR blocks overlap.All subnet CIDR blocks overlap.	VPC peering connections are not usable.	Unsupported VPC Peering Configurations

NOTICE

Alternatively, you can use enterprise routers to connect VPCs in the same region. **Enterprise Router** is more suitable for complex networking that needs to connect multiple VPCs. With enterprise routers, you do not have to create a large number of VPC peering connections or add too many routes. This makes your network topology simpler and more scalable.

All route tables in a VPC can have a maximum of 1,000 routes. If you want to create VPC peering connections to connect multiple VPCs, consider this restriction when planning the networking.

8.2.2 Using a VPC Peering Connection to Connect Two VPCs

You can configure a VPC peering connection and set the destination of the routes added to VPC route tables to the peer VPC CIDR block. In this way, all resources in the two VPCs are connected. [Table 8-2](#) shows example scenarios.

Table 8-2 Scenario description

Scenario	Scenario Description	IP Address Version	Example
Two VPCs peered together	You have two VPCs that require full access to each other's resources. For example, your company has VPC-A for the human resource department, and VPC-B for the finance department. The two departments require full access to each other's resources.	IPv4	Two VPCs Peered Together (IPv4)
		IPv6	Two VPCs Peered Together (IPv6)

Scenario	Scenario Description	IP Address Version	Example
Multiple VPCs peered together	You have multiple VPCs that require access to each other's resources.	IPv4	Multiple VPCs Peered Together (IPv4)
	For example, your company has VPC-A for the human resource department, VPC-B for the finance department, and VPC-C for the marketing department. These departments require full access to each other's resources.	IPv4	Multiple VPCs Peered Together Through Transitive Peering Connections (IPv4)
		IPv6	Multiple VPCs Peered Together (IPv6)
One central VPC peered with two VPCs	You have a central VPC that requires access to two peer VPCs, and similarly, the peer VPCs require access to the central VPC. However, the two peer VPCs need to be isolated from each other.	IPv4	One Central VPC Peered with Two VPCs (IPv4)
	For example, public services (such as databases) are deployed on VPC-A. Both VPC-B and VPC-C need to access the databases, but they do not need to access each other.	IPv6	One Central VPC Peered with Two VPCs (IPv6)
One central VPC with primary and secondary CIDR blocks peered with two VPCs	You have a central VPC that has both primary and secondary CIDR blocks. The central VPC needs to communicate with two peer VPCs, but the peer VPCs need to be isolated from each other.	IPv4	One Central VPC with Primary and Secondary CIDR Blocks Peered with Two VPCs (IPv4)

Scenario	Scenario Description	IP Address Version	Example
One central VPC peered with multiple VPCs	You have a central VPC that requires access to the multiple peer VPCs, and similarly, the peer VPCs require access to the central VPC. However, the peer VPCs need to be isolated from each other. For example, public services (such as databases) are deployed on your central VPC-A. VPC-B, VPC-C, VPC-D, VPC-E, VPC-F, and VPC-G need to access the databases, but these VPCs do not need to access each other.	IPv4	One Central VPC Peered with Multiple VPCs (IPv4)
		IPv6	One Central VPC Peered with Multiple VPCs (IPv6)

Notes and Constraints

If you create a VPC peering connection that connects entire CIDR blocks of two VPCs, the VPC CIDR blocks cannot overlap. Otherwise, the VPC peering connection does not take effect. For details, see [Invalid VPC Peering for Overlapping VPC CIDR Blocks](#).

Even if you intend to use the VPC peering connection for IPv6 communication only, you cannot create a VPC peering connection if the VPCs have matching or overlapping IPv4 CIDR blocks. In all examples in this section, the IPv4 CIDR blocks of any VPCs connected by a VPC peering connection do not overlap.

Two VPCs Peered Together (IPv4)

Create Peering-AB between VPC-A and VPC-B. The CIDR blocks of VPC-A and VPC-B do not overlap.

- For details about resource planning, see [Table 8-3](#).
- For details about VPC peering relationships, see [Table 8-4](#).

Figure 8-4 Networking diagram (IPv4)

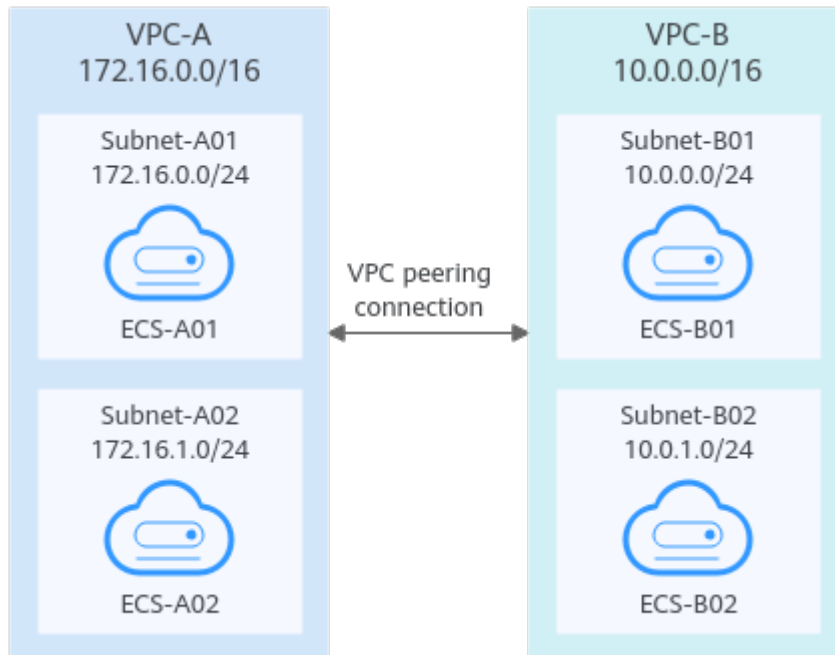


Table 8-3 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167

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

After the VPC peering connection is created, add the following routes to the route tables of the local and peer VPCs:

Table 8-5 VPC route tables (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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	
	172.16.0.0/16 (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.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

Two VPCs Peered Together (IPv6)

Create Peering-AB between VPC-A and VPC-B. The subnets of VPC-A and VPC-B have both IPv4 and IPv6 CIDR blocks and their IPv4 CIDR blocks do not overlap.

- For details about resource planning, see [Table 8-6](#).
- For details about VPC peering relationships, see [Table 8-7](#).

Figure 8-5 Networking diagram (IPv6)




Table 8-6 Resource planning details (IPv6)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	<ul style="list-style-type: none"> IPv4: 172.16.0.0/24 IPv6: 2407:c080:802:c34::/64 	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	<ul style="list-style-type: none"> IPv4: 172.16.0.111 IPv6: 2407:c080:802:c34:a925:f12e:cfa0:8edb
		Subnet-A02	<ul style="list-style-type: none"> IPv4: 172.16.1.0/24 IPv6: 2407:c080:802:c37::/64 	rtb-VPC-A	ECS-A02		<ul style="list-style-type: none"> IPv4: 172.16.1.91 IPv6: 2407:c080:802:c37:594b:4c0f:2fcd:8b72

VP C Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-B	10.0.0.0/16	Subnet-B01	<ul style="list-style-type: none"> IPv4: 10.0.0.0/24 IPv6: 2407:c080:802:c35::/64 	rtb-VPC-B	ECS-B01		<ul style="list-style-type: none"> IPv4: 10.0.0.139 IPv6: 2407:c080:802:c35:493:33f4:4531:5162
		Subnet-B02	<ul style="list-style-type: none"> IPv4: 10.0.1.0/24 IPv6: 2407:c080:802:c38::/64 	rtb-VPC-B	ECS-B02		<ul style="list-style-type: none"> IPv4: 10.0.1.167 IPv6: 2407:c080:802:c38:b9a9:aa03:2700:c1cf

Table 8-7 Peering relationships (IPv6)

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

After the VPC peering connection is created, add the following routes to the route tables of the local and peer VPCs:

Table 8-8 VPC route tables (IPv6)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	Local	System	
	2407:c080:802:c37::/64	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 for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c38::/64 (Subnet-B02)	Peering-AB	Custom	
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c35::/64	Local	System	
	10.0.1.0/24	Local	System	
	2407:c080:802:c38::/64	Local	System	
	172.16.0.0/16 (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 for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AB	Custom	

NOTE

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

Multiple VPCs Peered Together (IPv4)

If multiple VPCs need to communicate with each other, their CIDR blocks cannot overlap and you need to create a VPC peering connection between every two VPCs.

- For details about resource planning, see [Table 8-9](#).
- For details about VPC peering relationships, see [Table 8-10](#).

Figure 8-6 Networking diagram (IPv4)




Table 8-9 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167
VPC-C	192.168.0/16	Subnet-C01	192.168.0.0/24	rtb-VPC-C	ECS-C01		192.168.0.194
		Subnet-C02	192.168.1.0/24	rtb-VPC-C	ECS-C02		192.168.1.200

Table 8-10 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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-11 VPC route tables (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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.

Route Table	Destination	Next Hop	Route Type	Description
	192.168.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	172.16.0.0/16 (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.
	192.168.0.0/16 (VPC-C)	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	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	192.168.1.0/24	Local	System	
	172.16.0.0/16 (VPC-A)	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 (VPC-B)	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 the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

Multiple VPCs Peered Together Through Transitive Peering Connections (IPv4)

VPC peering connections are transitive. As shown in [Figure 8-7](#), there is a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. To enable communication between VPC-B and VPC-C, you can use either of the following methods:

- Create a VPC peering connection between VPC-B and VPC-C. For details, see [Multiple VPCs Peered Together \(IPv4\)](#).

- Add routes to direct traffic between VPC-B and VPC-C based on VPC-A. For details, see [Table 8-14](#).

Figure 8-7 Transitive VPC peering connections




Table 8-12 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167
VPC-C	192.168.0.0/16	Subnet-C01	192.168.0.0/24	rtb-VPC-C	ECS-C01		192.168.0.194
		Subnet-C02	192.168.1.0/24	rtb-VPC-C	ECS-C02		192.168.1.200

Table 8-13 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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-14 VPC route tables (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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.
	192.168.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	172.16.0.0/16 (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.
	192.168.0.0/16 (VPC-C)	Peering-AB	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-AB as the next hop.
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	192.168.1.0/24	Local	System	
	172.16.0.0/16 (VPC-A)	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 (VPC-B)	Peering-AC	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-AC as the next hop.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

Multiple VPCs Peered Together (IPv6)

If multiple VPCs need to communicate with each other, you need to create a VPC peering connection between every two VPCs. In this example, subnets in VPC-A, VPC-B, and VPC-C have IPv6 CIDR blocks and the IPv4 CIDR blocks of VPC-A, VPC-B, and VPC-C cannot overlap.

- For details about resource planning, see [Table 8-15](#).
- For details about VPC peering relationships, see [Table 8-16](#).

Figure 8-8 Networking diagram (IPv6)




Table 8-15 Resource planning details (IPv6)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC -A	172.16.0.0/16	Subnet-A01	<ul style="list-style-type: none"> IPv4: 172.16.0.0/24 IPv6: 2407:c080:802:c34::/64 	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	<ul style="list-style-type: none"> IPv4: 172.16.0.111 IPv6: 2407:c080:802:c34:a925:f12e:cfa0:8edb
		Subnet-A02	<ul style="list-style-type: none"> IPv4: 172.16.1.0/24 IPv6: 2407:c080:802:c37::/64 	rtb-VPC-A	ECS-A02		<ul style="list-style-type: none"> IPv4: 172.16.1.91 IPv6: 2407:c080:802:c37:594b:4c0f:2fcd:8b72
VPC -B	10.0.0.0/16	Subnet-B01	<ul style="list-style-type: none"> IPv4: 10.0.0.0/24 IPv6: 2407:c080:802:c35::/64 	rtb-VPC-B	ECS-B01		<ul style="list-style-type: none"> IPv4: 10.0.0.139 IPv6: 2407:c080:802:c35:493:33f4:4531:5162
		Subnet-B02	<ul style="list-style-type: none"> IPv4: 10.0.1.0/24 IPv6: 2407:c080:802:c38::/64 	rtb-VPC-B	ECS-B02		<ul style="list-style-type: none"> IPv4: 10.0.1.167 IPv6: 2407:c080:802:c38:b9a9:aa03:2700:c1cf

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-C	192.168.0/16	Subnet-C01	<ul style="list-style-type: none"> IPv4: 192.168.0.0/24 IPv6: 2407:c080:802:c3c::/64 	rtb-VPC-C	ECS-C01		<ul style="list-style-type: none"> IPv4: 192.168.0.194 IPv6: 2407:c080:802:c3c:d2f3:d891:24f5:f4af
		Subnet-C02	<ul style="list-style-type: none"> IPv4: 192.168.1.0/24 IPv6: 2407:c080:802:c3d::/64 	rtb-VPC-C	ECS-C02		<ul style="list-style-type: none"> IPv4: 192.168.1.200 IPv6: 2407:c080:802:c3d:e9ca:169a:390c:74d1

Table 8-16 Peering relationships (IPv6)

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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-17 VPC route tables (IPv6)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	Local	System	
	2407:c080:802:c37::/64	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 for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c38::/64 (Subnet-B02)	Peering-AB	Custom	
	192.168.0.0/16 (VPC-C)	Peering-AC	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c3c::/64 (Subnet-C01)	Peering-AC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the next hop for IPv6 communication.
	2407:c080:802:c3d::/64 (Subnet-C02)	Peering-AC	Custom	
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c35::/64	Local	System	
	10.0.1.0/24	Local	System	
	2407:c080:802:c38::/64	Local	System	

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/16 (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 for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AB	Custom	
	192.168.0.0/16 (VPC-C)	Peering-BC	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-BC as the next hop for IPv4 communication.
	2407:c080:802:c3c::/64 (Subnet-C01)	Peering-BC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-BC as the next hop for IPv6 communication.
	2407:c080:802:c3d::/64 (Subnet-C02)	Peering-BC	Custom	
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c3c::/64	Local	System	
	192.168.1.0/24	Local	System	
	2407:c080:802:c3d::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AC	Custom	

Route Table	Destination	Next Hop	Route Type	Description
	10.0.0.0/16 (VPC-B)	Peering-BC	Custom	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-BC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-BC as the next hop for IPv6 communication.
	2407:c080:802:c38::/64 (Subnet-B02)	Peering-BC	Custom	

 **NOTE**

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

One Central VPC Peered with Two VPCs (IPv4)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. The three VPCs do not have overlapping CIDR blocks.

- For details about resource planning, see [Table 8-18](#).
- For details about VPC peering relationships, see [Table 8-19](#).

Figure 8-9 Networking diagram (IPv4)




Table 8-18 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167
VPC-C	192.168.0.0/16	Subnet-C01	192.168.0.0/24	rtb-VPC-C	ECS-C01		192.168.0.194
		Subnet-C02	192.168.1.0/24	rtb-VPC-C	ECS-C02		192.168.1.200

Table 8-19 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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-20 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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.
	192.168.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	172.16.0.0/16 (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.
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	192.168.1.0/24	Local	System	

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

One Central VPC Peered with Two VPCs (IPv6)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. Each VPC has IPv6 subnets. The IPv4 CIDR blocks of the three VPCs do not overlap with each other.

- For details about resource planning, see [Table 8-21](#).
- For details about VPC peering relationships, see [Table 8-22](#).

Figure 8-10 Networking diagram (IPv6)




Table 8-21 Resource planning details (IPv6)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC -A	172.16.0.0/16	Subnet-A01	<ul style="list-style-type: none"> IPv4: 172.16.0.0/24 IPv6: 2407:c08:0:802:c34::/64 	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	<ul style="list-style-type: none"> IPv4: 172.16.0.111 IPv6: 2407:c08:0:802:c34:a925:f12e:cfa0:8edb
		Subnet-A02	<ul style="list-style-type: none"> IPv4: 172.16.1.0/24 IPv6: 2407:c08:0:802:c37::/64 	rtb-VPC-A	ECS-A02		<ul style="list-style-type: none"> IPv4: 172.16.1.91 IPv6: 2407:c08:0:802:c37:594b:4c0f:2fcd:8b72
VPC -B	10.0.0.0/16	Subnet-B01	<ul style="list-style-type: none"> IPv4: 10.0.0.0/24 IPv6: 2407:c08:0:802:c35::/64 	rtb-VPC-B	ECS-B01		<ul style="list-style-type: none"> IPv4: 10.0.0.139 IPv6: 2407:c08:0:802:c35:493:33f4:4531:5162
		Subnet-B02	<ul style="list-style-type: none"> IPv4: 10.0.1.0/24 IPv6: 2407:c08:0:802:c38::/64 	rtb-VPC-B	ECS-B02		<ul style="list-style-type: none"> IPv4: 10.0.1.167 IPv6: 2407:c08:0:802:c38:b9a9:aa03:2700:c1cf

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-C	192.168.0/16	Subnet-C01	<ul style="list-style-type: none"> IPv4: 192.168.0.0/24 IPv6: 2407:c08:0:802:c3c::/64 	rtb-VPC-C	ECS-C01		<ul style="list-style-type: none"> IPv4: 192.168.0.194 IPv6: 2407:c08:0:802:c3c:d2f3:d891:24f5:f4af
		Subnet-C02	<ul style="list-style-type: none"> IPv4: 192.168.1.0/24 IPv6: 2407:c08:0:802:c3d::/64 	rtb-VPC-C	ECS-C02		<ul style="list-style-type: none"> IPv4: 192.168.1.200 IPv6: 2407:c08:0:802:c3d:e9ca:169a:390c:74d1

Table 8-22 Peering relationships (IPv6)

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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-23 VPC route table details (IPv6)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	Local	System	
	2407:c080:802:c37::/64	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 for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c38::/64 (Subnet-B02)	Peering-AB	Custom	
	192.168.0.0/16 (VPC-C)	Peering-AC	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c3c::/64 (Subnet-C01)	Peering-AC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the next hop for IPv6 communication.
	2407:c080:802:c3d::/64 (Subnet-C02)	Peering-AC	Custom	
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c35::/64	Local	System	
	10.0.1.0/24	Local	System	
	2407:c080:802:c38::/64	Local	System	

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/16 (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 for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AB	Custom	
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c3c::/64	Local	System	
	192.168.1.0/24	Local	System	
	2407:c080:802:c3d::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AC	Custom	

 **NOTE**

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

One Central VPC with Primary and Secondary CIDR Blocks Peered with Two VPCs (IPv4)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. VPC-A has both primary and secondary CIDR blocks. The three VPCs do not have overlapping CIDR blocks.

- For details about resource planning, see [Table 8-24](#).
- For details about VPC peering relationships, see [Table 8-25](#).

Figure 8-11 Networking diagram (IPv4)




Table 8-24 Resource planning details

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	Primary CIDR block: 172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
	Secondary CIDR block: 192.167.0.0/16	Subnet-A-Extended01	192.167.0.0/24	rtb-VPC-A	ECS-A-Extended01		192.167.0.100
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
VPC-C	192.168.0.0/16	Subnet-C01	192.168.0.0/24	rtb-VPC-C	ECS-C01		192.168.0.194

Table 8-25 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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-26 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	192.167.0.0/24	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.
	192.168.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peering-AB	Custom	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AB as the next hop.
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peering-AB	Custom	
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peering-AC	Custom	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AC as the next hop.
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peering-AC	Custom	

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

One Central VPC Peered with Multiple VPCs (IPv4)

Create a VPC peering connection between VPC-A and VPC-B, between VPC-A and VPC-C, between VPC-A and VPC-D, between VPC-A and VPC-E, between VPC-A

and VPC-F, and between VPC-A and VPC-G. The CIDR blocks of these VPCs do not overlap.

- For details about resource planning, see [Table 8-27](#).
- For details about VPC peering relationships, see [Table 8-28](#).

Figure 8-12 Networking diagram (IPv4)

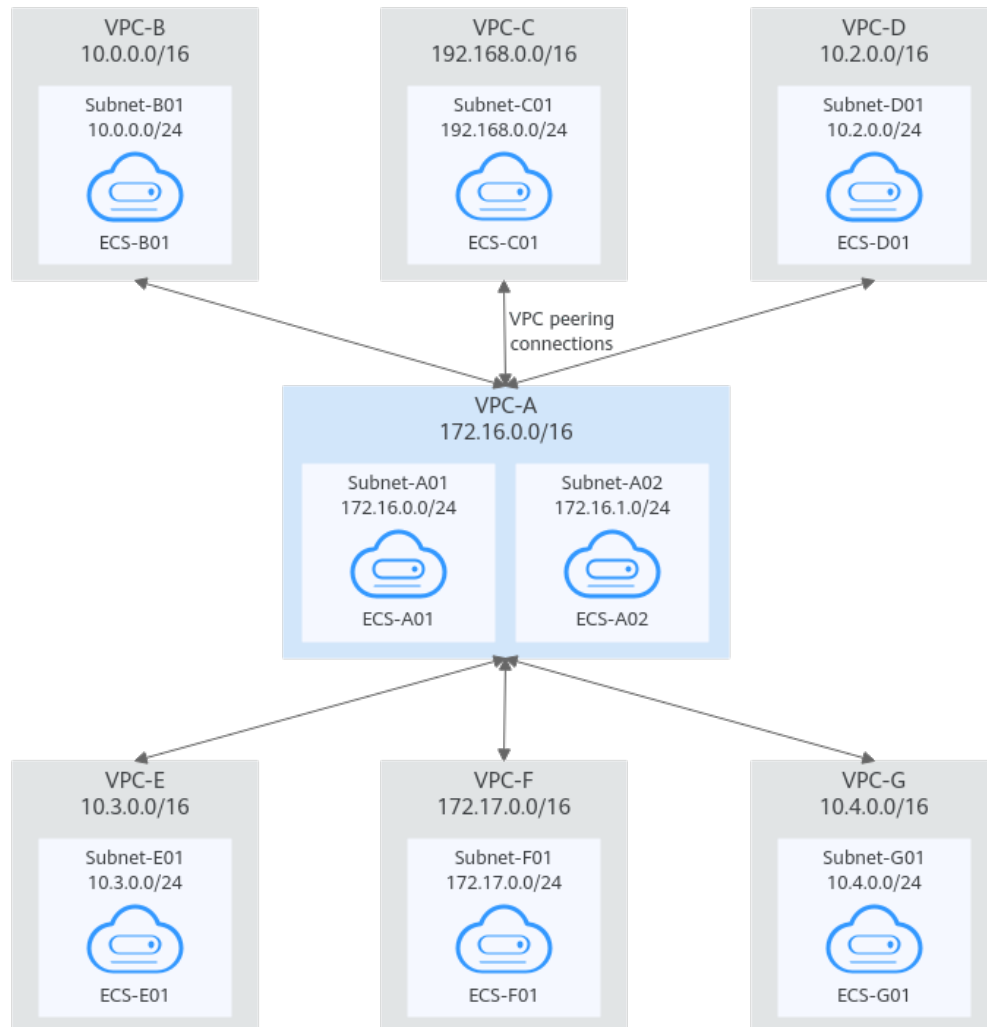


Table 8-27 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
VPC-C	192.168.0.0/16	Subnet-C01	192.168.0.0/24	rtb-VPC-C	ECS-C01		192.168.0.194
VPC-D	10.2.0.0/16	Subnet-D01	10.2.0.0/24	rtb-VPC-D	ECS-D01		10.2.0.237
VPC-E	10.3.0.0/16	Subnet-E01	10.3.0.0/24	rtb-VPC-E	ECS-E01		10.3.0.87
VPC-F	172.17.0.0/16	Subnet-F01	172.17.0.0/24	rtb-VPC-F	ECS-F01		172.17.0.103
VPC-G	10.4.0.0/16	Subnet-G01	10.4.0.0/24	rtb-VPC-G	ECS-G01		10.4.0.10

Table 8-28 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
VPC-A is peered with VPC-G.	Peering-AG	VPC-A	VPC-G

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-29 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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.
	192.168.0.0/16 (VPC-C)	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 (VPC-D)	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 (VPC-E)	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 (VPC-F)	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 (VPC-G)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (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.
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	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	10.2.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/16 (VPC-A)	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	10.3.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	Peering-AE	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.
rtb-VPC-F	172.17.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	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	10.4.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	Peering-AG	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each other's resources.

One Central VPC Peered with Multiple VPCs (IPv6)

Create a VPC peering connection between VPC-A and VPC-B, between VPC-A and VPC-C, between VPC-A and VPC-D, between VPC-A and VPC-E, between VPC-A and VPC-F, and between VPC-A and VPC-G. Each VPC has IPv6 subnets. The IPv4 CIDR blocks of these VPCs do not overlap.

- For details about resource planning, see [Table 8-30](#).
- For details about VPC peering relationships, see [Table 8-31](#).

Figure 8-13 Networking diagram (IPv6)

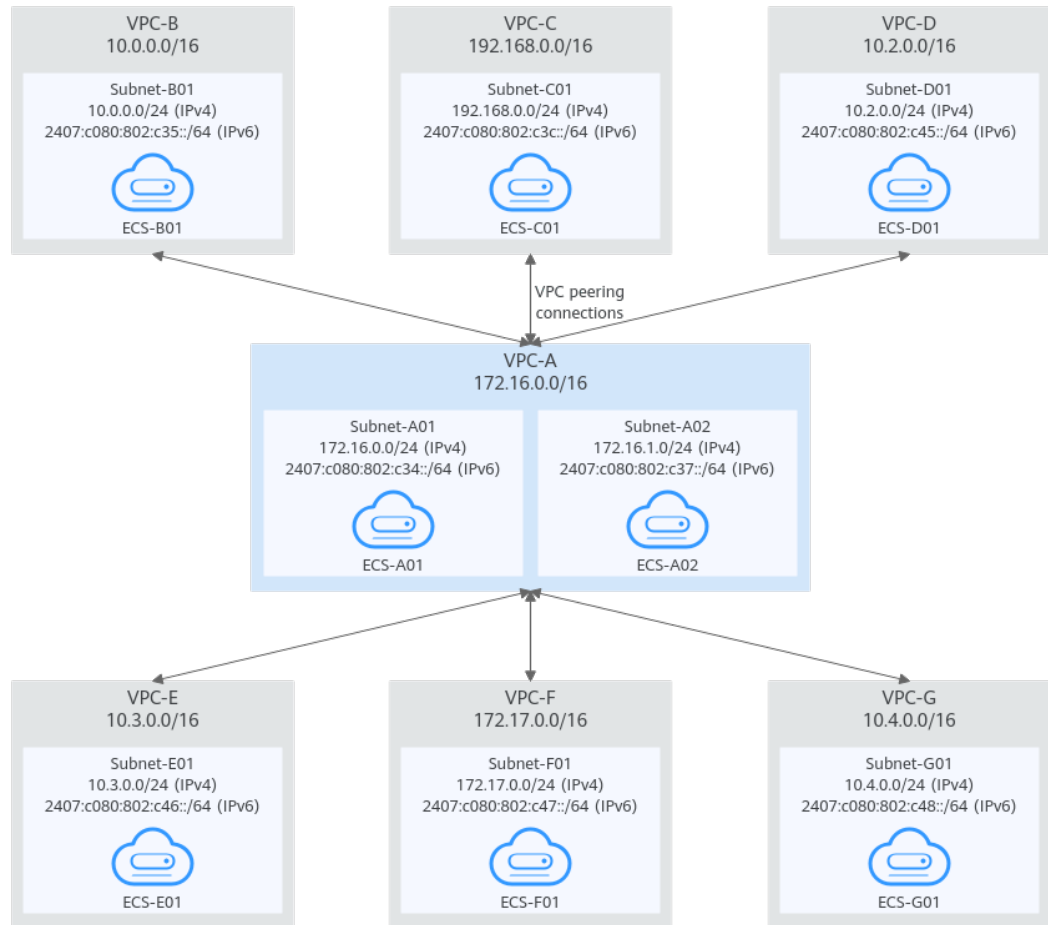


Table 8-30 Resource planning details (IPv6)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	<ul style="list-style-type: none"> IPv4: 172.16.0.0/24 IPv6: 2407:c080:802:c34::/64 	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	<ul style="list-style-type: none"> IPv4: 172.16.0.111 IPv6: 2407:c080:802:c34:a925:f12e:cfa0:8edb

VP C Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
		Subnet-A02	<ul style="list-style-type: none"> IPv4: 172.16.1.0/24 IPv6: 2407:c080:802:c37::/64 	rtb-VPC-A	ECS-A02		<ul style="list-style-type: none"> IPv4: 172.16.1.91 IPv6: 2407:c080:802:c37:594b:4c0f:2fcd:8b72
VPC-B	10.0.0/16	Subnet-B01	<ul style="list-style-type: none"> IPv4: 10.0.0/24 IPv6: 2407:c080:802:c35::/64 	rtb-VPC-B	ECS-B01		<ul style="list-style-type: none"> IPv4: 10.0.0.139 IPv6: 2407:c080:802:c35:493:33f4:4531:5162
VPC-C	192.168.0/16	Subnet-C01	<ul style="list-style-type: none"> IPv4: 192.168.0/24 IPv6: 2407:c080:802:c3c::/64 	rtb-VPC-C	ECS-C01		<ul style="list-style-type: none"> IPv4: 192.168.0.194 IPv6: 2407:c080:802:c3c:d2f3:d891:24f5:f4af
VPC-D	10.0.0/16	Subnet-D01	<ul style="list-style-type: none"> IPv4: 10.2.0/24 IPv6: 2407:c080:802:c45::/64 	rtb-VPC-D	ECS-D01		<ul style="list-style-type: none"> IPv4: 10.2.0.237 IPv6: 2407:c080:802:c45:6bb7:f161:3596:6e4c

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Subnet Route Table	ECS Name	Security Group	Private IP Address
VPC-E	10.3.0.0/16	Subnet-E01	<ul style="list-style-type: none"> IPv4: 10.3.0.0/24 IPv6: 2407:c080:802:c46::/64 	rtb-VPC-E	ECS-E01		<ul style="list-style-type: none"> IPv4: 10.3.0.87 IPv6: 2407:c080:802:c46:2a2f:558a:85da:4c70
VPC-F	172.17.0.0/16	Subnet-F01	<ul style="list-style-type: none"> IPv4: 172.17.0.0/24 IPv6: 2407:c080:802:c47::/64 	rtb-VPC-F	ECS-F01		<ul style="list-style-type: none"> IPv4: 172.17.0.103 IPv6: 2407:c080:802:c47:b5e2:e6f0:c42b:44fd
VPC-G	10.4.0.0/16	Subnet-G01	<ul style="list-style-type: none"> IPv4: 10.4.0.0/24 IPv6: 2407:c080:802:c48::/64 	rtb-VPC-G	ECS-G01		<ul style="list-style-type: none"> IPv4: 10.4.0.10 IPv6: 2407:c080:802:c48:3020:f48c:4e54:aa17

Table 8-31 Peering relationships (IPv6)

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

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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-32 VPC route table details (IPv6)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	Local	System	
	2407:c080:802:c37::/64	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 for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-AB	Custom	Add a route with the IPv6 CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop for IPv6 communication.

Route Table	Destination	Next Hop	Route Type	Description
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c3c::/64 (Subnet-C01)	Peerin g-AC	Cust om	Add a route with the IPv6 CIDR block of Subnet-C01 as the destination and Peering-AC as the next hop for IPv6 communication.
	10.2.0.0/16 (VPC-D)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-D as the destination and Peering-AD as the next hop for IPv4 communication.
	2407:c080:802:c45::/64 (Subnet-D01)	Peerin g-AD	Cust om	Add a route with the IPv6 CIDR block of Subnet-D01 as the destination and Peering-AD as the next hop for IPv6 communication.
	10.3.0.0/16 (VPC-E)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-E as the destination and Peering-AE as the next hop for IPv4 communication.
	2407:c080:802:c46::/64 (Subnet-E01)	Peerin g-AE	Cust om	Add a route with the IPv6 CIDR block of Subnet-E01 as the destination and Peering-AE as the next hop for IPv6 communication.
	172.17.0.0/16 (VPC-F)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-F as the destination and Peering-AF as the next hop for IPv4 communication.
	2407:c080:802:c47::/64 (Subnet-F01)	Peerin g-AF	Cust om	Add a route with the IPv6 CIDR block of Subnet-F01 as the destination and Peering-AF as the next hop for IPv6 communication.
	10.4.0.0/16 (VPC-G)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-G as the destination and Peering-AG as the next hop for IPv4 communication.
	2407:c080:802:c48::/64 (Subnet-G01)	Peerin g-AG	Cust om	Add a route with the IPv6 CIDR block of Subnet-G01 as the destination and Peering-AG as the next hop for IPv6 communication.

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c35::/64	Local	System	
	172.16.0.0/16 (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 for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AB	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AB	Custom	
rtb-VPC-C	192.168.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c3c::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AC	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AC	Custom	
rtb-VPC-D	10.2.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c45::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AD	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop for IPv4 communication.

Route Table	Destination	Next Hop	Route Type	Description
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AD	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AD as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AD	Custom	
rtb-VPC-E	10.3.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c46::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AE	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AE	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AE as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AE	Custom	
rtb-VPC-F	172.17.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c47::/64	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AF	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AF	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AF as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AF	Custom	
rtb-VPC-G	10.4.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c48::/64	Local	System	

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/16 (VPC-A)	Peering-AG	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AG	Custom	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AG as the next hop for IPv6 communication.
	2407:c080:802:c37::/64 (Subnet-A02)	Peering-AG	Custom	

 NOTE

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

8.2.3 Using a VPC Peering Connection to Connect Subnets in Two VPCs

You can configure a VPC peering connection and set the destination of the routes added to VPC route tables to the subnet CIDR block of the peer VPC. In this way, all resources in the VPC subnets are connected. [Table 8-33](#) shows example scenarios.

Table 8-33 Scenario description

Scenario	Scenario Description	IP Address Version	Example
Two VPCs peered to two subnets in a central VPC	<p>You have a central VPC that requires access to the multiple other VPCs. The other VPCs need to be isolated from each other.</p> <ul style="list-style-type: none"> The central VPC has separate sets of resources in different subnets. The other VPCs require access to some of the resources, but not all of them. 	IPv4	Two VPCs Peered to Two Subnets in a Central VPC (IPv4)
		IPv6/IPv4	Two VPCs Peered to Two Subnets in a Central VPC (IPv6/IPv4)

Scenario	Scenario Description	IP Address Version	Example
One central VPC peered to specific subnets in two VPCs	<p>You have a central VPC that requires access to two other VPCs. The other VPCs need to be isolated from each other.</p> <ul style="list-style-type: none">• The central VPC has public resources deployed and the other VPCs require access to all resources in the central VPC.• Other VPCs have multiple subnets and only one in each VPC is used for accessing resources in the central VPC.	IPv4	One Central VPC Peered to Specific Subnets in Two VPCs (IPv4)
One central VPC peered to overlapping subnets from two VPCs	<p>This scenario is similar to the preceding one. If two VPCs with overlapping subnets need to peer with the central VPC, traffic may fail to be forwarded to the required destination. To prevent this, plan the network according to this example.</p>	IPv4	One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)

Two VPCs Peered to Two Subnets in a Central VPC (IPv4)

The central VPC-A has two subnets, Subnet-A01 and Subnet-A02. The subnets are associated with different route tables. You need to create Peering-AB between Subnet-A01 and VPC-B, and Peering-AC between Subnet-A02 and VPC-C. VPC-B and VPC-C have the same CIDR block. However, there will be no route conflicts because the two subnets in VPC-A are associated with different route tables.

- For details about resource planning, see [Table 8-34](#).
- For details about VPC peering relationships, see [Table 8-35](#).

Figure 8-14 Networking diagram (IPv4)




Table 8-34 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A01	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A02	ECS-A02		172.16.1.91
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
VPC-C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71

 NOTE

VPC-A has two route tables. Route table rtb-VPC-A01 is associated with Subnet-A01, and route table rtb-VPC-A02 is associated with Subnet-A02. The two subnets can communicate with each other.

Table 8-35 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
Subnet-A01 of VPC-A is peered to VPC-B.	Peering-AB	VPC-A	VPC-B
Subnet-A02 of VPC-A is peered to VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-36 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A01	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	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-A02	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	Local	System	
	10.0.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.0/24 (Subnet-A01)	Peering-AB	Custom	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop.
rtb-VPC-C	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24 (Subnet-A02)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AC as the next hop.

Two VPCs Peered to Two Subnets in a Central VPC (IPv6/IPv4)

The central VPC-A has two subnets, Subnet-A01 and Subnet-A02. The subnets are associated with different route tables. You need to create Peering-AB between Subnet-A01 and VPC-B for IPv6 communication, and Peering-AC between Subnet-A02 and VPC-C for IPv4 communication. VPC-B and VPC-C have the same CIDR block. However, there will be no route conflicts because the two subnets in VPC-A are associated with different route tables.

- For details about resource planning, see [Table 8-37](#).
- For details about VPC peering relationships, see [Table 8-38](#).

Figure 8-15 Networking diagram (IPv6/IPv4)

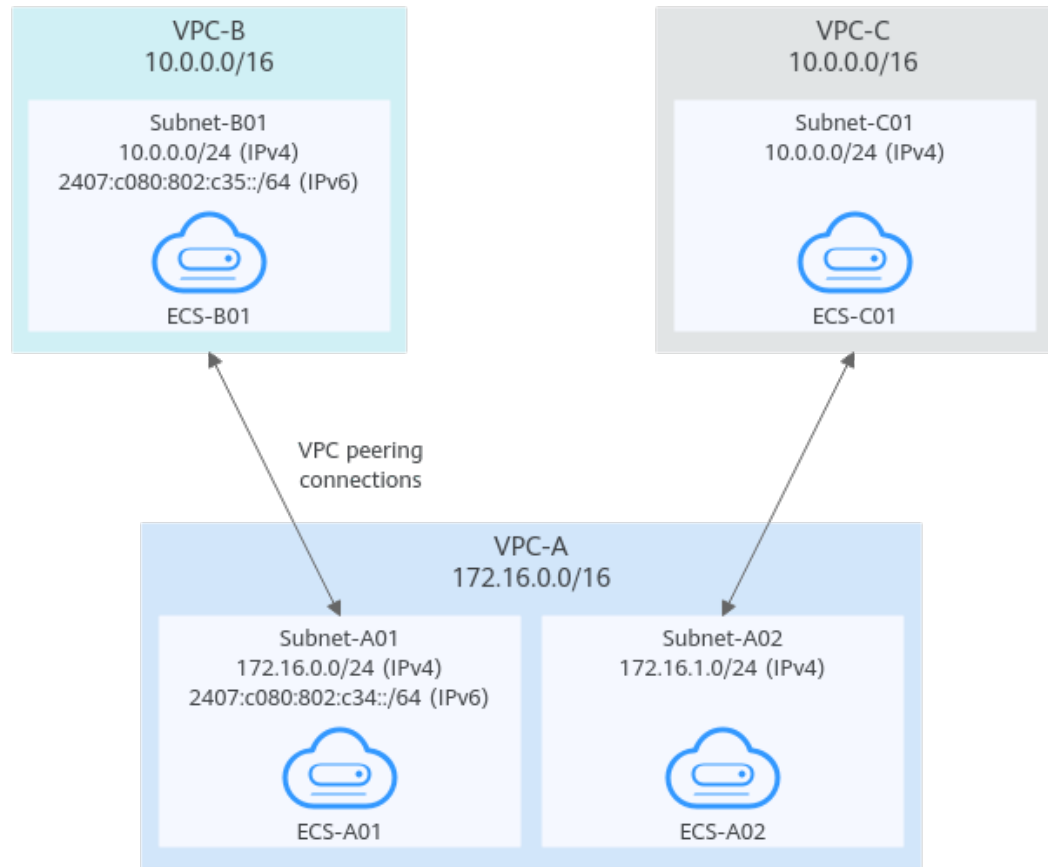


Table 8-37 Resource planning details (IPv6/IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	<ul style="list-style-type: none"> IPv4: 172.16.0.0/24 IPv6: 2407:c080:802:c34::/64 	rtb-VPC-A01	ECS-A01	sg-web: general-purpose web server	<ul style="list-style-type: none"> IPv4: 172.16.0.111 IPv6: 2407:c080:802:c34:a925:f12e:cfa0:8edb
		Subnet-A02	172.16.1.0/24	rtb-VPC-A02	ECS-A02		172.16.1.91

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-B	10.0.0.0/16	Subnet-B01	<ul style="list-style-type: none"> IPv4: 10.0.0.0/24 IPv6: 2407:c080:802:c35::/64 	rtb-VPC-B	ECS-B01		<ul style="list-style-type: none"> IPv4: 10.0.0.139 IPv6: 2407:c080:802:c35:493:33f4:4531:5162
VPC-C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71

 **NOTE**

VPC-A has two route tables. Route table rtb-VPC-A01 is associated with Subnet-A01, and route table rtb-VPC-A02 is associated with Subnet-A02. The two subnets can communicate with each other.

Table 8-38 Peering relationships (IPv6/IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
Subnet-A01 of VPC-A is peered to VPC-B. (IPv6)	Peering-AB	VPC-A	VPC-B
Subnet-A02 of VPC-A is peered to VPC-C. (IPv4)	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-39 VPC route table details (IPv6/IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A01	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	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 for IPv4 communication.
	2407:c080:802:c35::/64 (Subnet-B01)	Peering-AB	Custom	Add a route with the IPv6 CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop for IPv6 communication.
rtb-VPC-A02	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c34::/64	Local	System	
	172.16.1.0/24	Local	System	
	10.0.0.0/16 (VPC-C)	Peering-AC	Custom	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.
rtb-VPC-B	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	2407:c080:802:c35::/64	Local	System	
	172.16.0.0/24 (Subnet-A01)	Peering-AB	Custom	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop for IPv4 communication.
	2407:c080:802:c34::/64 (Subnet-A01)	Peering-AB	Custom	Add a route with the IPv6 CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop for IPv6 communication.

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-C	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24 (Subnet-A02)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AC as the next hop for IPv4 communication.

One Central VPC Peered to Specific Subnets in Two VPCs (IPv4)

You need to create Peering-AB between central VPC-A and Subnet-B01 in VPC-B, and Peering-AC between central VPC-A and Subnet-C02 in VPC-C. VPC-B and VPC-C have the same CIDR block, but the CIDR blocks of Subnet-B01 and Subnet-C02 do not overlap. Therefore, there will be no route conflicts.

- For details about resource planning, see [Table 8-40](#).
- For details about VPC peering relationships, see [Table 8-41](#).

Figure 8-16 Networking diagram (IPv4)

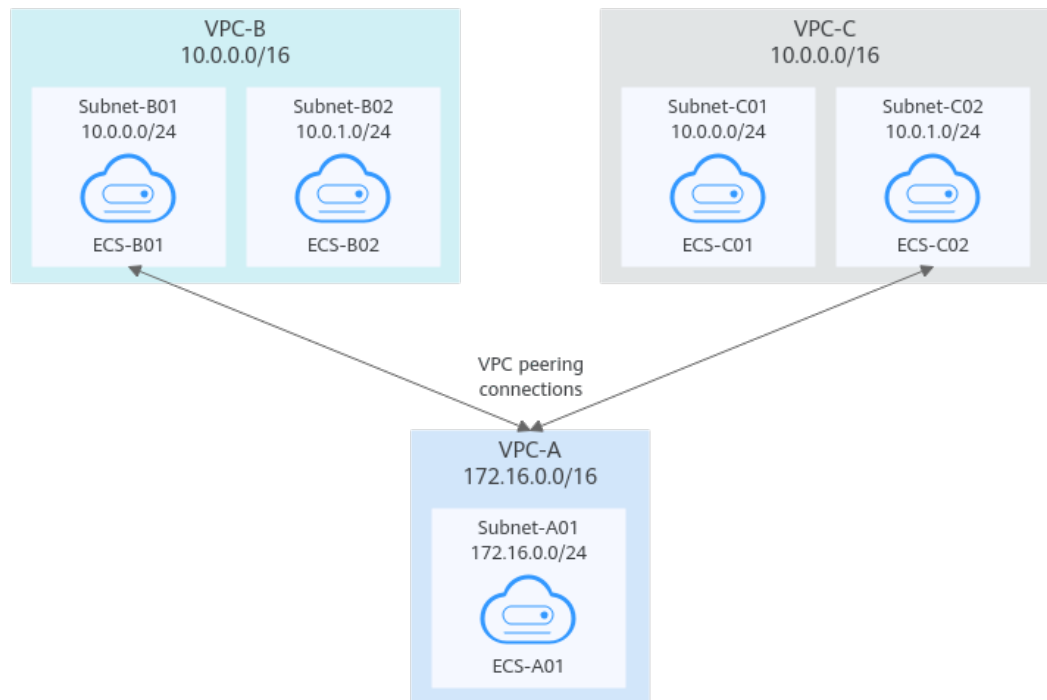


Table 8-40 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167
VPC-C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71
		Subnet-C02	10.0.1.0/24	rtb-VPC-C	ECS-C02		10.0.1.116

Table 8-41 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered to Subnet-B01 of VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered to Subnet-C02 of VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-42 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	10.0.0.0/24 (Subnet-B01)	Peering-AB	Custom	Add a route with the CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop.
	10.0.1.0/24 (Subnet-C02)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC 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	
	172.16.0.0/16 (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.
rtb-VPC-C	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)

If you want to create VPC peering connections between a VPC and multiple overlapping subnets from different VPCs, ensure that the destinations of the routes added for the peering connections do not conflict and traffic can be correctly forwarded.

In this example, you need to create Peering-AB between central VPC-A and Subnet-B02 in VPC-B, and Peering-AC between central VPC-A and Subnet-C02 in VPC-C. Subnet-B02 and Subnet-C02 have the same CIDR block, and ECS-B02 and ECS-C02 have the same private IP address (10.0.1.167/32).

- For details about resource planning, see [Table 8-43](#).
- For details about VPC peering relationships, see [Table 8-44](#).

Figure 8-17 Networking diagram (IPv4)




Table 8-43 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
		Subnet-B02	10.0.1.0/24	rtb-VPC-B	ECS-B02		10.0.1.167
VPC-C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71
		Subnet-C02	10.0.1.0/24	rtb-VPC-C	ECS-C02		10.0.1.167

Table 8-44 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered to Subnet-B02 of VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered to Subnet-C02 of VPC-C.	Peering-AC	VPC-A	VPC-C

If you add routes to the route tables of the local and peer VPCs according to [Table 8-45](#), the response traffic cannot be correctly forwarded. The details are as follows:

1. ECS-B02 in Subnet-B02 of VPC-B sends request traffic to VPC-A through the route with Peering-AB as the next hop in the rtb-VPC-B route table.
2. VPC-A receives the request traffic from ECS-B02 and expects to send the response traffic to ECS-B02. The rtb-VPC-A route table has the route with 10.0.1.167/32 as the destination, but its next hop is Peering-AC. The response traffic is incorrectly sent to VPC-C.
3. ECS-C02 in Subnet-C02 of VPC-C has the same private IP address (10.0.1.167/32) as ECS-B02. The response traffic is incorrectly sent to ECS-C02, and ECS-B02 cannot receive the response traffic.

Table 8-45 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24 (Subnet-C02)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC 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	
	172.16.0.0/16 (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.

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-C	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24	Local	System	
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

If there are overlapping subnets, configure routes as follows to prevent traffic from being incorrectly forwarded:

- Suggestion 1: In the rtb-VPC-A route table, add a route with Peering-AB as the next hop and the private IP address of ECS-B02 (10.0.1.167/32) as the destination. The route with 10.0.1.167/32 as the destination is preferentially matched based on the longest prefix match rule to ensure that VPC-A sends the response traffic to ECS-B02. For more configurations, see [Using a VPC Peering Connection to Connect ECSs in Two VPCs](#).

Table 8-46 VPC route table details

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.167/32 (ECS-B02)	Peering-AB	Custom	Add a route with the private IP address of ECS-B02 as the destination and Peering-AB as the next hop.
	10.0.1.0/24 (Subnet-C02)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.

- Suggestion 2: In the rtb-VPC-A route table, change the destination of the route with Peering-AC as the next hop from Subnet-C02 to Subnet-C01. Add a route with Peering-AB as the next hop and Subnet-B02 as the destination to ensure that VPC-A can send the response traffic to Subnet-B02 in VPC-B.

Table 8-47 VPC route table details

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24 (Subnet-B02)	Peering-AB	Custom	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
	10.0.0.0/24 (Subnet-C01)	Peering-AC	Custom	Add a route with the CIDR block of Subnet-C01 as the destination and Peering-AC as the next hop.

8.2.4 Using a VPC Peering Connection to Connect ECSs in Two VPCs

You can configure a VPC peering connection and set the destination of the routes added to VPC route tables to the private IP address of ECS in the peer VPC. In this way, the two ECS are connected.

To enable traffic forwarding among these ECSs, you need to add routes with private IP addresses of these ECSs as the destinations and a VPC peering connection as the next hop to VPC route tables. [Table 8-48](#) shows example scenarios.

Table 8-48 Scenario description

Scenario	Scenario Description	IP Address Version	Example
ECS in a central VPC peered to ECSs in two other VPCs	You want a central VPC to communicate with the other two VPCs. However, you do not want the other two VPCs to communicate with each other. The other two VPCs have the same CIDR block and also include subnets that overlap. To prevent route conflicts in the central VPC, you can configure VPC peering connections to connect to specific ECSs in the other two VPCs.	IPv4	ECS in a Central VPC Peered to ECSs in Two Other VPCs (IPv4)

Scenario	Scenario Description	IP Address Version	Example
A central VPC peered with two other VPCs using longest prefix match	<p>This scenario is similar to the preceding one. In addition to peering specific ECSs, you can create the following VPC peering connections based on the longest prefix match rule:</p> <ul style="list-style-type: none">• Create a VPC peering connection between the central VPC and an ECS in VPC-B• Create a VPC peering connection between the central VPC and a subnet in VPC-C <p>This configuration expands the communication scope.</p>	IPv4	A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)

ECS in a Central VPC Peered to ECSs in Two Other VPCs (IPv4)

You want to create a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. VPC-B and VPC-C have matching CIDR blocks. You can set the destinations of routes to private IP addresses of specific ECSs to limit traffic to these ECSs. If the destination of a route is not properly planned, traffic cannot be correctly forwarded. For details, see [One Central VPC Peered to Overlapping Subnets from Two VPCs \(IPv4\)](#).

In this example, you need to create Peering-AB between ECS-A01-1 in VPC-A and ECS-B01 in VPC-B, and Peering-AC between ECS-A01-2 in VPC-A and ECS-C01 in VPC-C. Subnet-B01 and Subnet-C01 have matching CIDR blocks. The private IP addresses of ECS-B01 and ECS-C01 must be different. Otherwise, there will be route conflicts because the route table of VPC-A will have routes with the same destination.

- For details about resource planning, see [Table 8-49](#).
- For details about VPC peering relationships, see [Table 8-50](#).

Figure 8-18 Networking diagram (IPv4)

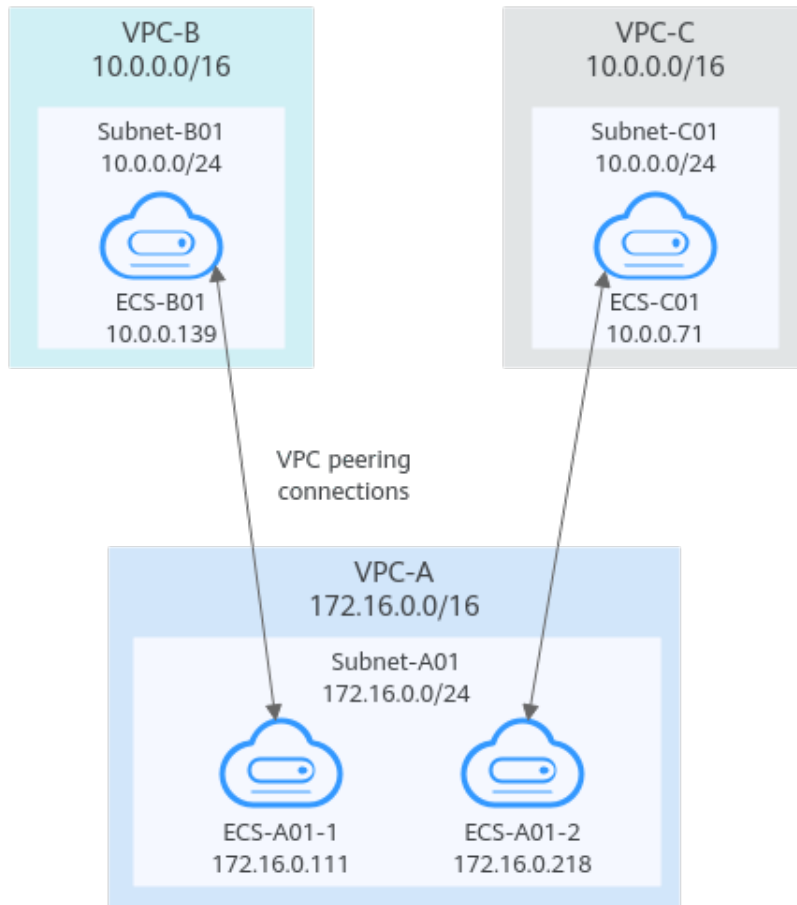


Table 8-49 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC-A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01-1	sg-web: general-purpose web server	172.16.0.111
					ECS-A01-2		172.16.0.218
VPC-B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
VPC-C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71

Table 8-50 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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-51 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.0.139/32 (ECS-B01)	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 (ECS-C01)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.111/32 (ECS-A01-1)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.

Route Table	Destination	Next Hop	Route Type	Description
	172.16.0.218/32 (ECS-A01-2)	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.

A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)

You want to create a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. VPC-B and VPC-C have matching CIDR blocks. You can set the destinations of routes to private IP addresses of specific ECSs to limit traffic to these ECSs. If the destination of a route is not properly planned, traffic cannot be correctly forwarded. For details, see [One Central VPC Peered to Overlapping Subnets from Two VPCs \(IPv4\)](#).

In this example, you need to create Peering-AB between central VPC-A and ECS-B01 in VPC-B, and Peering-AC between central VPC-A and VPC-C. Subnet-B01 and Subnet-C01 have matching CIDR blocks. You can use the longest prefix match rule to control traffic forwarding.

- For details about resource planning, see [Table 8-52](#).
- For details about VPC peering relationships, see [Table 8-53](#).

Figure 8-19 Networking diagram (IPv4)




Table 8-52 Resource planning details (IPv4)

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	VPC Route Table	ECS Name	Security Group	Private IP Address
VPC -A	172.16.0.0/16	Subnet-A01	172.16.0.0/24	rtb-VPC-A	ECS-A01	sg-web: general-purpose web server	172.16.0.111
		Subnet-A02	172.16.1.0/24	rtb-VPC-A	ECS-A02		172.16.1.91
VPC -B	10.0.0.0/16	Subnet-B01	10.0.0.0/24	rtb-VPC-B	ECS-B01		10.0.0.139
VPC -C	10.0.0.0/16	Subnet-C01	10.0.0.0/24	rtb-VPC-C	ECS-C01		10.0.0.71

Table 8-53 Peering relationships (IPv4)

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

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Table 8-54 VPC route table details (IPv4)

Route Table	Destination	Next Hop	Route Type	Description
rtb-VPC-A	172.16.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	Local	System	
	10.0.0.139/32 (ECS-B01)	Peering-AB	Custom	Add a route with the private IP address of ECS-B01 as the destination and Peering-AB as the next hop.
	10.0.0.0/16 (VPC-C)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (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.
rtb-VPC-C	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	Peering-AC	Custom	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

8.2.5 Unsupported VPC Peering Configurations

Scenarios

The VPC peering connection configurations are not supported in [Table 8-55](#).

Table 8-55 Scenarios that VPC peering connections are invalid

Scenario	Example
<ul style="list-style-type: none">• If VPCs with the same CIDR block also include subnets that overlap, VPC peering connections are not usable.• If two VPCs have overlapping CIDR blocks but some of their subnets do not overlap, you cannot create a VPC peering connection to connect specific subnets that do not overlap.	Invalid VPC Peering for Overlapping VPC CIDR Blocks <ul style="list-style-type: none">• VPCs with the same CIDR block also include subnets that overlap.• Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.
VPC peering connections cannot enable ECSs in their VPCs to share an EIP to access the Internet. If VPC-A and VPC-B are peered and ECS-A01 in VPC-A has an EIP, ECS-B01 in VPC-B cannot access the Internet using the EIP bound to ECS-A01.	Invalid VPC Peering for Sharing an EIP

Invalid VPC Peering for Overlapping VPC CIDR Blocks

If two VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect due to route conflicts. The following describes the reasons and configuration suggestions.

- VPCs with the same CIDR block also include subnets that overlap.
VPC peering connections are not usable. As shown in [Table 8-56](#), VPC-A and VPC-B, and their subnets have the same CIDR block. If you create a VPC peering connection between VPC-A and VPC-B, their route tables are shown in [Table 8-56](#).
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 8-20 Networking diagram (IPv4)





Table 8-56 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.

If two VPCs want to use their IPv6 CIDR blocks for communication by a VPC peering connection but the IPv4 CIDR blocks of the VPCs or subnets overlap, the connection is not usable.

Figure 8-21 Networking diagram (IPv6)



- Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.

VPC peering connections will not take effect in the following scenarios:

- Connecting overlapping CIDR blocks of VPCs

As shown in [Figure 8-22](#), if you create a VPC peering connection between VPC-A and VPC-B, the VPC peering connection will not take effect because the two VPCs have the same CIDR block.

- Connecting overlapping subnets from different VPCs

If you create a VPC peering connection between Subnet-A01 and Subnet-B02, the route tables are shown in [Table 8-57](#). In the `rtb-VPC-B` route table, the custom route for routing traffic from Subnet-B02 to Subnet-A01 and the local route have overlapping destinations. The local route has a higher priority and traffic will be forwarded within Subnet-B02 and cannot reach Subnet-A01.

Figure 8-22 Networking diagram (IPv4)

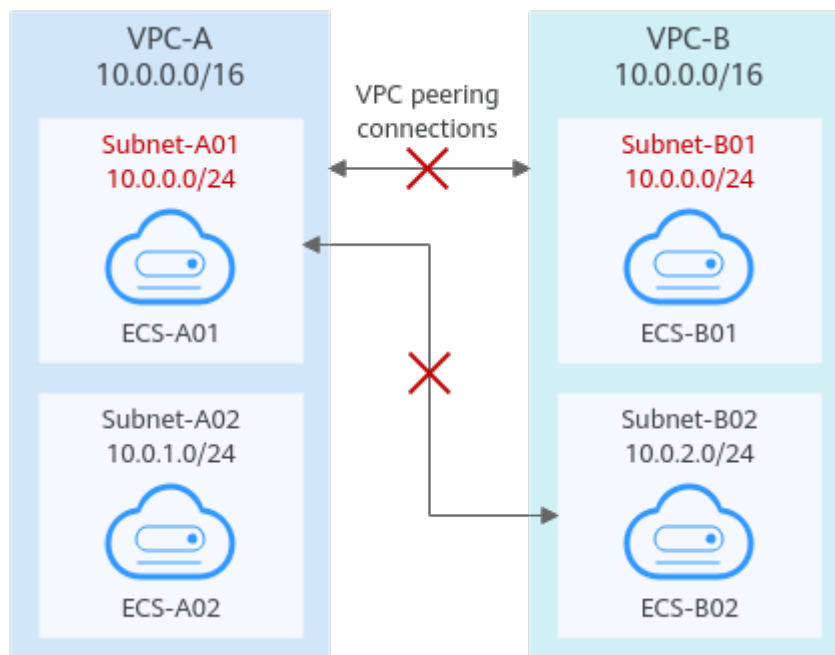



Table 8-57 VPC route table details

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	
	10.0.2.0/24 (Subnet-B02)	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	10.0.0.0/24	Local	System	Local routes are automatically added for communications within a VPC.
	10.0.2.0/24	Local	System	
	10.0.0.0/24 (Subnet-A01)	Peering-AB	Custom	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop.


If the subnets connected by a VPC peering connection do not overlap, the connection will take effect. As shown in [Figure 8-23](#), you can create a VPC peering connection between Subnet-A02 and Subnet-B02. In this case, the routes do not conflict and the VPC peering connection takes effect.

Figure 8-23 Networking diagram (IPv4)



If two VPCs want to use their IPv6 CIDR blocks for communication by a VPC peering connection but the IPv4 CIDR blocks of the VPCs or subnets overlap, the connection is not usable.


Figure 8-24 Networking diagram (IPv6)



Invalid VPC Peering for Sharing an EIP

As shown in [Figure 8-25](#), although VPC-A and VPC-B are peered and ECS-A01 in VPC-A has an EIP, ECS-B01 in VPC-B cannot access the Internet using the EIP bound to ECS-A01.

Figure 8-25 Networking diagram



8.3 Creating a VPC Peering Connection to Connect Two VPCs in the Same Account

Scenarios

Two VPCs from the same region cannot communicate with each other by default, but you can use a VPC peering connection to connect them.

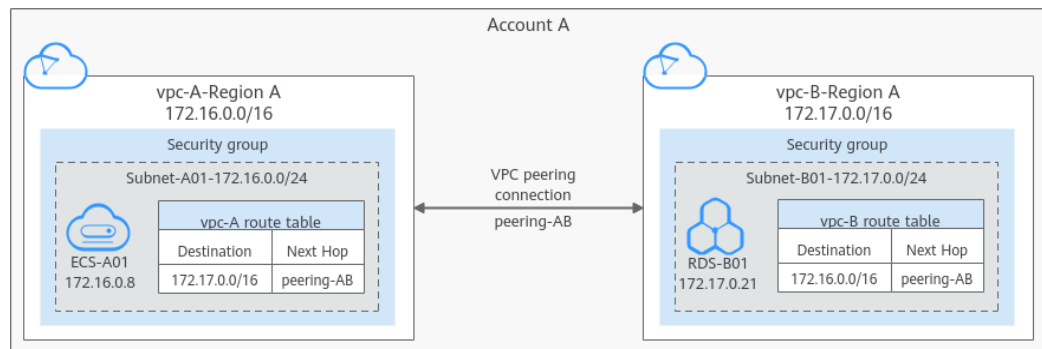
The following describes how to create a VPC peering connection to connect two VPCs (**vpc-A** and **vpc-B** in this example) in the same account. In this way, instances (**ECS-A01** and **RDS-B01** in this example) in the two VPCs can communicate with each other.

The procedure is as follows:

Step 1: Create a VPC Peering Connection

Step 2: Add Routes for the VPC Peering Connection

Step 3: Verify Network Connectivity

Figure 8-26 Connecting two VPCs in an account using a VPC peering connection**NOTICE**

Currently, VPC peering connections are free.

Notes and Constraints

- Only one VPC peering connection can be created between two VPCs at the same time.
- A VPC peering connection can only connect VPCs in the same region.
 - If 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 be usable.

Prerequisites

You have two VPCs from the same account in the same region. If you want to create one, see [Creating a VPC and Subnet](#).

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 8-58](#).

Figure 8-27 Creating a VPC peering connection

✕

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](#)

* VPC Peering Connection Name

Local VPC Settings

* Local VPC C

Local VPC CIDR Block

Peer VPC Settings

* Account My account Another account ?

* Peer Project

If you select **My account**, the project is filled in by default.

* Peer VPC

If the local and peer VPCs overlap, your VPC peering connection may not be usable. [Learn more](#)

Table 8-58 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
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

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">Options: My account and Another accountSelect My account.	My account
Peer Project	The project is selected in by default if Account is set to My account . In this example, vpc-A and vpc-B are created in region A, and the corresponding project of the account in region A is selected by default.	ab-cdef-1
Peer VPC	This parameter is mandatory if Account is set to My account . 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 be usable. For details, see VPC Peering Connection Usage Examples .	172.17.0.0/16

Parameter	Description	Example Value
Description	Optional Enter a description of the VPC peering connection in the text box as required.	peering-AB connects vpc-A and vpc-B .

4. Click **OK**.
A dialog box for adding routes is displayed.
5. 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

1. In the lower part of the VPC peering connection details page, click **Add Route**.
The **Add Route** dialog box is displayed.
2. Add routes to the route tables as prompted.
[Table 8-59](#) describes the parameters.

Table 8-59 Parameter description

Parameter	Description	Example Value
VPC	Select a VPC that is connected by the VPC peering connection.	vpc-A
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">• If there is only the default route table in the drop-down list, select the default route table.• 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.	rtb-vpc-A (Default)

Parameter	Description	Example Value
Destination	An IP address or address range in the VPC being 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 VPC Peering Connection Usage Examples .	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
Add a route for the other VPC	If you select this option, you can also add a route for the other VPC connected by the VPC peering connection. 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 VPC Peering Connection Usage Examples .	Selected
VPC	By default, the system selects the VPC connected by the VPC peering connection. You do not need to specify this parameter.	vpc-B

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">• If there is only the default route table in the drop-down list, select the default route table.• 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.	rtb-vpc-B (Default)
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 VPC Peering Connection Usage Examples.</p>	vpc-A CIDR block: 172.16.0.0/16
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 (< or >).</p>	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 communications between the local and peer VPCs.

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

ping *RDS-B01-IP-address*

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-A01 ~]# 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 Connect Two VPCs in Different Accounts

Scenarios

Two VPCs from the same region cannot communicate with each other by default, but you can use a VPC peering connection to connect them.

The following describes how to create a VPC peering connection to connect two VPCs, **vpc-A** in one account and **vpc-B** in another account. In this way, instances (**ECS-A01** and **RDS-B01** in this example) in the two VPCs can communicate with each other.

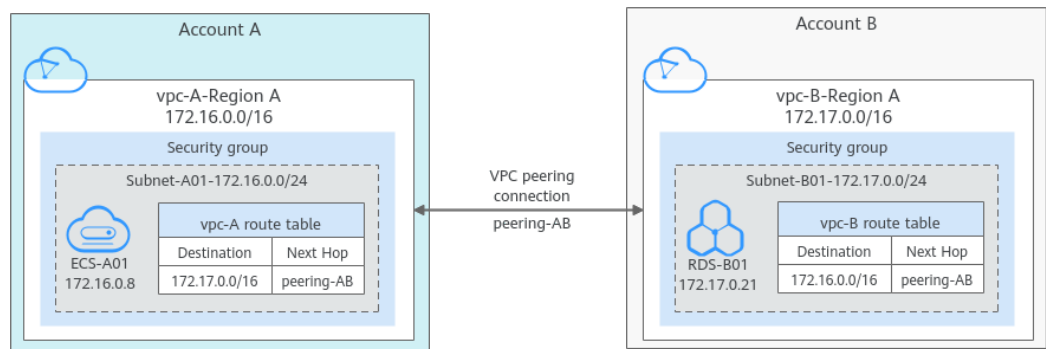
The procedure is as follows:

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 8-28 Connecting two VPCs in different accounts using a VPC peering connection**NOTICE**

Currently, VPC peering connections are free.

Notes and Constraints

- Only one VPC peering connection can be created between two VPCs at the same time.
- A VPC peering connection can only connect VPCs in the same region.
 - If 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 be usable.
- 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 and Subnet](#).

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 8-60](#).

Figure 8-29 Creating a VPC peering connection

✕

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

* VPC Peering Connection Name

Local VPC Settings

* Local VPC C

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

OK
Cancel

Table 8-60 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
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

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">Options: My account and Another accountSelect Another account.	Another account
Peer Project ID	This parameter is mandatory if Account is set to Another account . The project ID of the region that the peer VPC resides. For details about how to obtain the project ID, see Obtaining the Peer Project ID of a VPC Peering Connection .	Project ID of vpc-B in region A: 067cf8aecf3XXX08322f13b
Peer VPC ID	This parameter is mandatory if Account is set to Another account . ID of the VPC at the other end of the VPC peering connection. For details about how to obtain the ID, see Obtaining a VPC ID .	vpc-B ID: 17cd7278-XXX-530c952dcf35
Description (Optional)	Optional Enter a 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 8-30 Awaiting acceptance

NameID	Status	Local VPC	Local VPC CIDR Block	Peer Project ID	Peer VPC	Peer VPC CIDR Block	Descr...	Operation
peering-aB 04f2930-caef-44ef-8992- 67a5a953a9e8	Awaiting acceptance	vpc-A	172.16.0.0/16	accB 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.
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 both 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.

- c. Add routes to the route tables as prompted.

[Table 8-61](#) describes the parameters.

Table 8-61 Parameter description

Parameter	Description	Example Value
VPC	By default, the VPC in the current account is selected. You do not need to select a VPC.	vpc-A
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">• If there is only the default route table in the drop-down list, select the default route table.• 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.	rtb-vpc-A (Default route table)
Destination	An IP address or address range in the VPC being 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 VPC Peering Connection Usage Examples .	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

Parameter	Description	Example Value
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.
 - c. Add routes to the route table as prompted.
Table 8-62 describes the parameters.

Table 8-62 Parameter description

Parameter	Description	Example Value
VPC	By default, the VPC in the current account is selected. 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">• If there is only the default route table in the drop-down list, select the default route table.• 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.	rtb-vpc-B (Default route table)
Destination	An IP address or address range in the VPC being 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 VPC Peering Connection Usage Examples .	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 routes in the route list.

Step 4: Verify Network Connectivity

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

1. Log in to ECS-A01 in the local VPC.

2. Check whether ECS-A01 can communicate with RDS-B01.

ping *RDS-B01-IP-address*

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-A01 ~]# 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

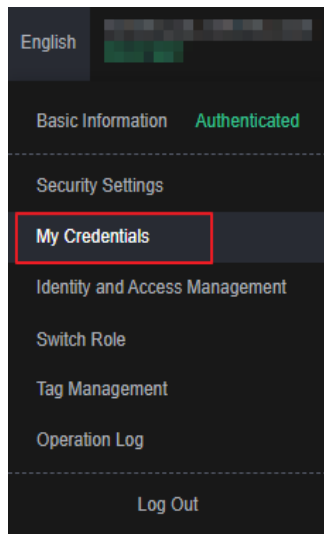
Scenarios

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 8-31 My Credentials



- 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 8-32 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

- 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 > 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. Confirm the information and enter **DELETE**.
7. Click **OK**.

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 name of the target route table in the **Route Table** column.
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 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 9-1 shows how an IPv4 and IPv6 dual-stack network works.

Figure 9-1 An IPv4 and IPv6 dual-stack network

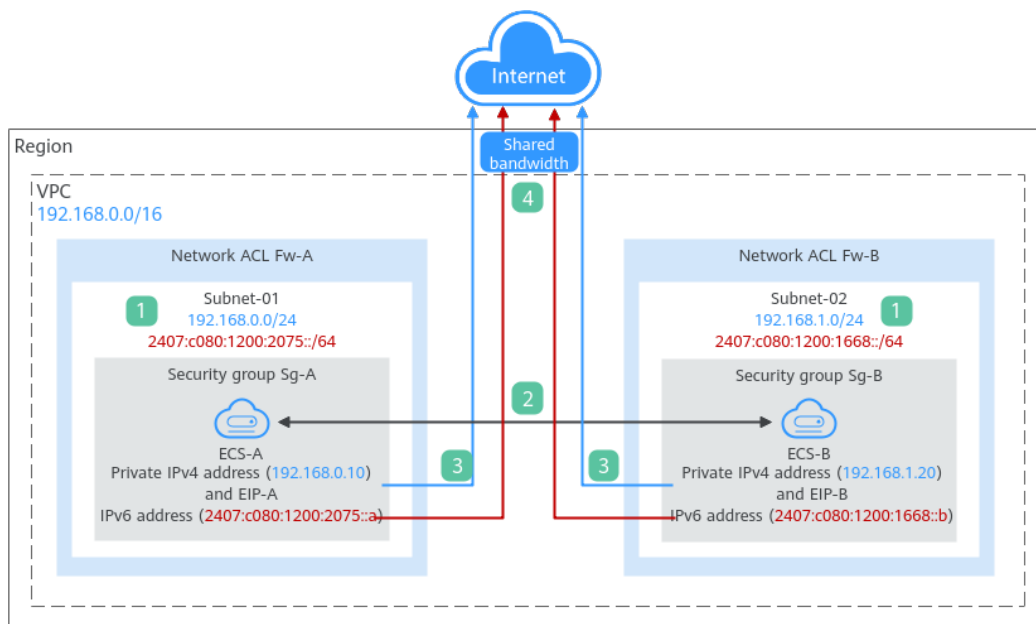


Table 9-1 Steps for deploying a dual-stack network

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">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.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. <p>As shown in Figure 9-1, if allow rules are configured for network ACLs Fw-A and Fw-B and security groups Sg-A and Sg-B, ECS-A and ECS-B can communicate with each other:</p> <ul style="list-style-type: none">Using private IPv4 addresses (192.168.0.10 and 192.168.1.20).Using IPv6 addresses (2407:c080:1200:2075::a and 2407:c080:1200:1668::b).
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 Figure 9-1, you can bind EIP-A to ECS-A and EIP-B to ECS-B so that ECS-A and ECS-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 Figure 9-1, you can add the IPv6 addresses of ECS-A and ECS-B to a shared bandwidth so that ECS-A and ECS-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).
- The IPv6 function is now available for open beta test in [certain regions](#). You can use the IPv6 function only after obtaining the OBT permission.
- Only certain ECS specifications support IPv6 networks. You need to select such ECSs in supported regions.

On the ECS console, click **Buy ECS**. On the displayed page, check the ECS specifications. If **Yes** is shown in the **IPv6** column, the ECS with this specification supports IPv6.

IPv4 and IPv6 Dual-Stack Application Scenarios

If your ECS supports IPv6, you can build an IPv4 and IPv6 dual-stack network. [Table 9-2](#) shows where IPv4 and IPv6 dual-stack networks can be used.

Table 9-2 Application scenarios of IPv4 and IPv6 dual-stack networks

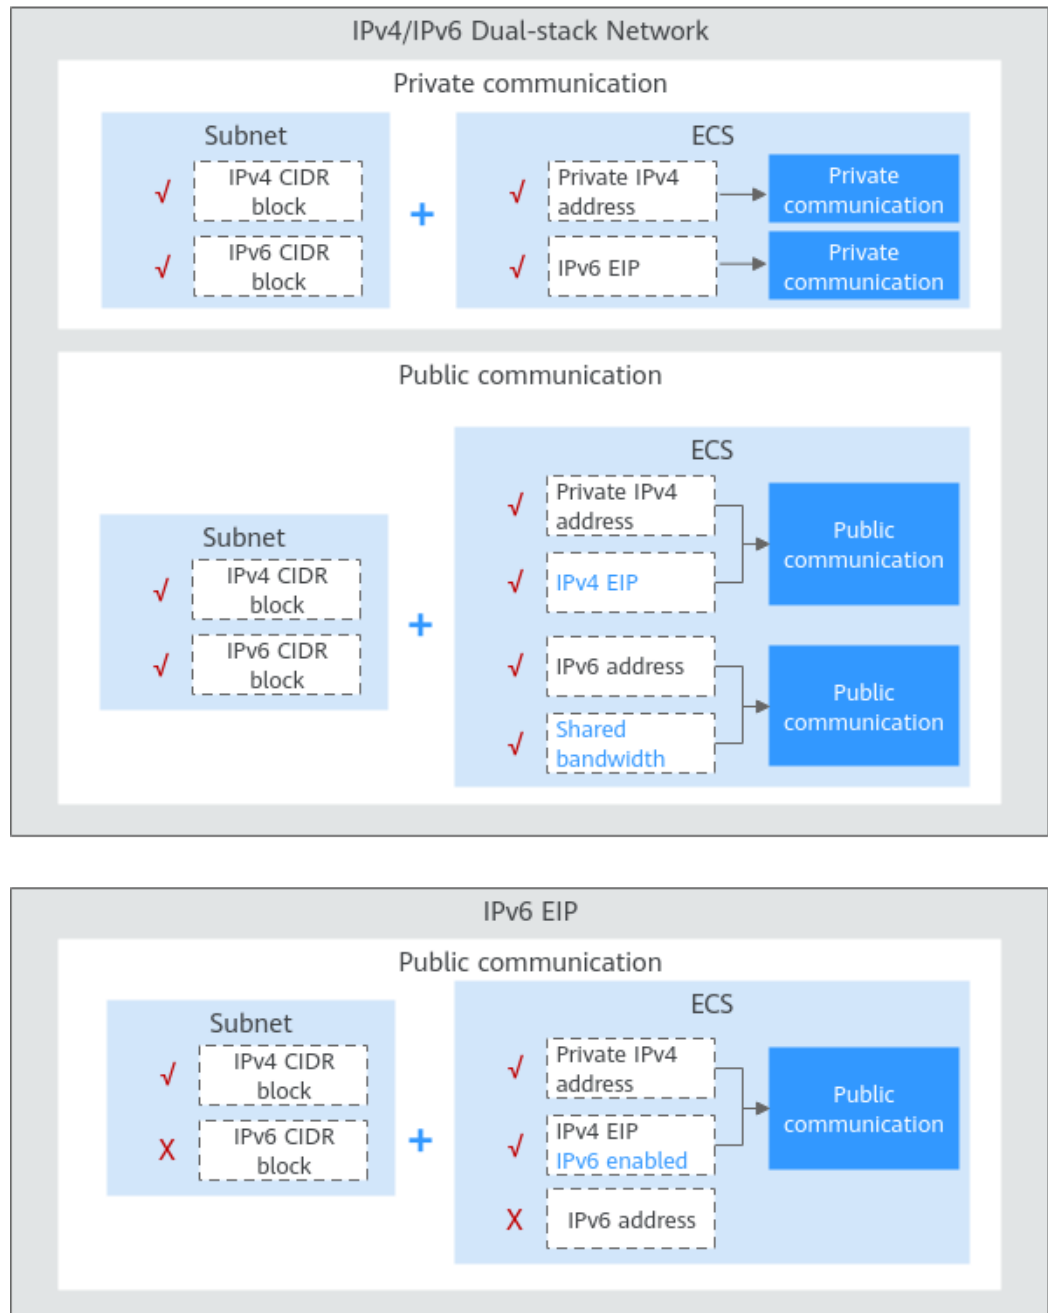
Application Scenario	Scenario	Subnet	ECS
Private communication using IPv6 addresses	Your applications deployed on ECSs need to communicate with other systems (such as databases) through private networks using IPv6 addresses.	<ul style="list-style-type: none"> IPv4 CIDR block IPv6 CIDR block 	<ul style="list-style-type: none"> Private IPv4 address: used for private communication IPv6 address: used for private communication.
Public communication using IPv6 addresses	Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.	<ul style="list-style-type: none"> IPv4 CIDR block IPv6 CIDR block 	<ul style="list-style-type: none"> Private IPv4 address + IPv4 EIP: used for public network communication IPv6 address + shared bandwidth: used for public network communication
	Your applications deployed on ECSs need to both provide services accessible from the Internet and analyze the access request data using IPv6 addresses.		

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 9-3](#).

Table 9-3 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"> Private IPv4 address IPv4 EIP (with IPv6 function enabled): used for public communication using IPv4 and IPv6 EIPs

Figure 9-2 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 9-4](#) describes how you can build and use an IPv6 network.

Table 9-4 Operation guide on IPv6 networks

Scenario	Description	Reference
Creating an IPv6 subnet	Select Enable for IPv6 CIDR Block when creating a subnet. An IPv6 CIDR block will be automatically assigned to the subnet. <ul style="list-style-type: none"> You cannot customize an IPv6 CIDR block. IPv6 cannot be disabled after the subnet is created. You can enable IPv6 for existing subnets. 	Creating a Subnet for the VPC
Viewing in-use IPv6 addresses	In the subnet list, click the subnet name. On the displayed page, view in-use IPv4 and IPv6 addresses on the IP Addresses tab.	Viewing IP Addresses in a Subnet
Adding a security group rule (IPv6)	Add a security group rule with Type set to IPv6 and Source or Destination set to an IPv6 address or IPv6 CIDR block.	Adding a Security Group Rule
Adding a network ACL rule (IPv6)	Add a network ACL rule with Type set to IPv6 and Source or Destination set to an IPv6 address or IPv6 CIDR block.	Adding a Network ACL Rule (Default Priorities)
Purchasing an EIP (IPv6)	When purchasing an EIP, select Enable IPv6 Internet access , or choose More > Enable IPv6 EIP in the Operation column of an existing IPv4 EIP. After IPv6 EIP is enabled, both IPv4 and IPv6 EIPs are assigned.	IPv6 EIP
Adding an IPv6 EIP or IPv6 address to a shared bandwidth	After purchasing a shared bandwidth, you can add IPv6 EIPs or IPv6 addresses to it.	Adding EIPs to a Shared Bandwidth
Adding an IPv6 route to the VPC route table	Add a route with Destination and Next Hop set to an IPv4 or IPv6 CIDR block. <ul style="list-style-type: none"> 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. 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. 	Adding Routes to a Route Table

Scenario	Description	Reference
Assigning a virtual IPv6 address	If IPv6 is enabled for a VPC subnet, you can set IP Address Type to IPv6 when assigning for a virtual IP address.	Assigning a Virtual IP Address

10 VPC Flow Log

10.1 VPC Flow Log

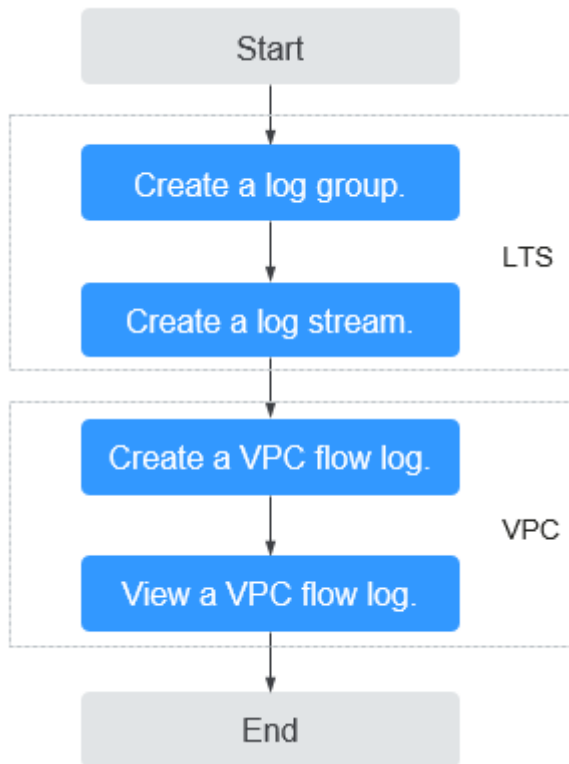
What Is a VPC Flow Log?

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.

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 10-1 shows the process for configuring VPC flow logs.

Figure 10-1 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.

10.2 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 10-1 Parameter descriptions

Parameter	Description	Example Value
Name	The VPC flow log name. The name: <ul style="list-style-type: none">• Can contain 1 to 64 characters.• Can contain letters, digits, underscores (_), hyphens (-), and periods (.).	flowlog-495d
Resource Type	The type of resources whose traffic is to be logged. You can select NIC , Subnet , or VPC .	NIC
Resource	The specific NIC whose traffic is to be logged. NOTE We recommend that you select an ECS that is in the running state. If an ECS in the stopped state is selected, restart the ECS after creating the VPC flow log for accurately recording the information about the traffic going to and from the ECS NIC.	N/A
Filter	<ul style="list-style-type: none">• All traffic: specifies that both accepted and rejected traffic of the specified resource will be logged.• Accepted traffic: specifies that only accepted traffic of the specified resource will be logged. Accepted traffic refers to the traffic permitted by the security group or network ACL.• Rejected traffic: specifies that only rejected traffic of the specified resource will be logged. Rejected traffic refers to the traffic denied by the network ACL.	All
Log Group	The log group created in LTS.	lts-group-abc
Log Stream	The log stream created in LTS.	lts-topic-abc
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.

10.3 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 an ECS is in the stopped state, its flow log records will not be displayed.

Procedure

1. Go to the [VPC flow log list page](#).
2. 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>
```

Table 10-2 describes the fields of a flow log record.

Table 10-2 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

Field	Description	Example Value
protocol	The Internet Assigned Numbers Authority (IANA) protocol number of the traffic. For details, see Assigned Internet Protocol Numbers .	17
packets	The number of packets transferred during the capture window.	1
bytes	The number of bytes transferred during the capture window.	96
start	The time, in Unix seconds, of the start of the capture window.	1548752136
end	The time, in Unix seconds, of the end of the capture window.	1548752736
action	The action associated with the traffic: <ul style="list-style-type: none">• ACCEPT: The recorded traffic was allowed by the security groups or network ACLs.• REJECT: The recorded traffic was denied by the security groups or network ACLs.	ACCEPT

Field	Description	Example Value
log-status	<p>The logging status of the VPC flow log:</p> <ul style="list-style-type: none"> • OK: Data is logging normally to the chosen destinations. • NODATA: There was no traffic of the Filter setting to or from the NIC during the capture window. • SKIPDATA: Some flow log records were skipped during the capture window. This may be caused by an internal capacity constraint or an internal error. <p>Example:</p> <p>When Filter is set to Accepted traffic, if there is accepted traffic, the value of log-status is OK. If there is no accepted traffic, the value of log-status is NODATA regardless of whether there is rejected traffic. If some accepted traffic is abnormally skipped, the value of log-status is SKIPDATA.</p>	OK

Table 10-3 provides you with flow log examples.

Table 10-3 Flow log examples

Scenario	Example Value
A flow log record in which data was recorded during the capture window	<p>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.</p> <pre>1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983- bd6512aed4bd 192.168.0.154 192.168.3.25 38929 53 17 1 96 1548752136 1548752736 ACCEPT OK</pre>

Scenario	Example Value
A flow log record in which no data was recorded during the capture window	1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983-bd6512aed4bd - - - - - 1431280876 1431280934 - NODATA
A flow log record in which data was skipped during the capture window	1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983-bd6512aed4bd - - - - - 1431280876 1431280934 - SKIPDATA

You can enter a keyword on the log stream details page on the LTS console to search for flow log records.

10.4 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. Go to the [VPC flow log list page](#).
2. Locate the target flow log and click **Enable** or **Disable** in the **Operation** column.
A confirmation dialog box is displayed.
3. Confirm the information and click **OK**.

10.5 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. Go to the [VPC flow log list page](#).
2. Locate the target flow log and click **Delete** in the **Operation** column.
A confirmation dialog box is displayed.
3. Confirm the information and click **OK**.

11 Elastic IP


11.1 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 11-1 Accessing the Internet using an EIP



EIP Quotas

You can log in to the console to query your EIP quotas by referring to [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.
- 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

- 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.

11.2 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.

Table 11-1 Parameter descriptions

Parameter	Description	Example Value
Billing Mode	The following billing modes are available: <ul style="list-style-type: none">• Yearly/Monthly• Pay-per-use	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.	N/A
EIP Type	<ul style="list-style-type: none">• Dynamic BGP: Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.	Dynamic BGP
Billed By	This parameter is available only when you set Billing Mode to Pay-per-use . <ul style="list-style-type: none">• Bandwidth: You specify a maximum bandwidth and pay for the amount of time you use the bandwidth. This is suitable for scenarios with heavy or stable traffic.• Traffic: You specify a maximum bandwidth and pay for the total traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic.• Shared Bandwidth: The bandwidth can be shared by multiple EIPs. This is suitable for scenarios with staggered traffic.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	100

Parameter	Description	Example Value
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.	N/A
EIP Name	The EIP name.	eip-test
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is default . For details about creating and managing enterprise projects, see the Enterprise Management User Guide .	default
Advanced Settings	Click the drop-down arrow to configure parameters, including the bandwidth name and tag.	N/A
Bandwidth Name	The name of the bandwidth.	bandwidth
Tag	The EIP tags. Each tag contains a key and value pair. The tag key and value must meet the requirements listed in Table 11-2 .	<ul style="list-style-type: none">Key: lpv4_key1Value: 3005eip
Monitoring	Used to monitor the EIP and enabled by default. You can use the management console or APIs provided by Cloud Eye to query the metrics and alarms generated for the EIP and bandwidth.	N/A
Required Duration	The duration for which the purchased EIP will use. The duration must be specified if the Billing Mode is set to Yearly/Monthly .	1 month

Parameter	Description	Example Value
Auto-renew	<p>Whether to select Auto-renew. You can select it if the Billing Mode is set to Yearly/Monthly. The auto-renewal period is determined by the required duration.</p> <ul style="list-style-type: none"> • Monthly subscription: The subscription is renewed every month. • Yearly subscription: The subscription is renewed each year. 	N/A
Quantity	<p>The number of EIPs you want to purchase.</p> <p>The quantity must be specified if the Billing Mode is set to Pay-per-use.</p>	1

Table 11-2 EIP tag requirements

Parameter	Requirement	Example Value
Key	<ul style="list-style-type: none"> • Cannot be left blank. • Must be unique for each EIP. • Can contain a maximum of 36 characters. • Can contain only the following character types: <ul style="list-style-type: none"> - Uppercase letters - Lowercase letters - Digits - Special characters, including hyphens (-) and underscores (_) 	Ipv4_key1
Value	<ul style="list-style-type: none"> • Can contain a maximum of 43 characters. • Can contain only the following character types: <ul style="list-style-type: none"> - Uppercase letters - Lowercase letters - Digits - Special characters, including hyphens (-) and underscores (_) 	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.
3. Click **OK**.

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?](#)

11.3 Unbinding an EIP from an ECS and Releasing the EIP

Scenarios



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

Notes and Constraints



- 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.



11.4 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.

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.
6. Click **Next**.
7. Click **Submit**.

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?](#)

11.5 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 displayed page, click  in the upper right corner of the EIP list.

The system will automatically export all EIPs in the current region of your account to an Excel file and download the file to a local directory.

11.6 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.



A tag consists of a key and value pair. [Table 11-3](#) lists the tag key and value requirements.

Table 11-3 EIP tag requirements

Parameter	Requirement	Example Value
Key	<ul style="list-style-type: none">• Cannot be left blank.• Must be unique for each EIP.• Can contain a maximum of 36 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including hyphens (-) and underscores (_)	Ipv4_key1
Value	<ul style="list-style-type: none">• Can contain a maximum of 43 characters.• Can contain only the following character types:<ul style="list-style-type: none">- Uppercase letters- Lowercase letters- Digits- Special characters, including hyphens (-) and underscores (_)	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. Click the search box above the EIP list.
5. Select the tag key and value of the EIP.
You can add multiple tag keys and values to refine your search results. If you add more than one tag to search for EIPs, the system will display only the EIPs that contain all of the tags you specified.
6. Click **OK**.
The system displays the EIPs you are looking for based on the entered tag keys and values.

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

1. Log in to the management console.
2. Click  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**.

11.7 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.

Application Scenarios of IPv4/IPv6 Dual Stack

If your ECS supports IPv6, you can use the IPv4/IPv6 dual stack. [Table 11-4](#) shows the example application scenarios.

Table 11-4 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"> No EIPs have been bound to the ECSs. 	IPv4 CIDR Block	Private IPv4 address: used for private IPv4 communication.
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"> EIPs have been bound to the ECSs. 	IPv4 CIDR Block	<ul style="list-style-type: none"> Private IPv4 address: used for private IPv4 communication. Public IPv4 address: used for public IPv4 communication.

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"> • IPv6 has been enabled for the VPC subnet. • The network has been configured for the ECSs as follows: <ul style="list-style-type: none"> – Flavor: Any ECS flavor that supports the IPv6 network. For details, see section "x86 ECS Specifications and Types" in the Elastic Cloud Server User Guide. – VPC and Subnet: IPv6-enabled subnet and VPC. – Self-assigned IPv6 address: Selected. – Shared Bandwidth: Selected Do not configure. 	<ul style="list-style-type: none"> • IPv4 CIDR Block • IPv6 CIDR block 	<ul style="list-style-type: none"> • Private IPv4 address + IPv4 EIP: Bind an IPv4 EIP to the instance to allow public IPv4 communication. • Private IPv4 address: Do not bind any IPv4 EIP to the instance and use only the private IPv4 address to allow private IPv4 communication. • IPv6 address: Do not configure shared bandwidth for the IPv6 address to allow private IPv6 communication.

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"> IPv6 has been enabled for the VPC subnet. The network has been configured for the ECSs as follows: <ul style="list-style-type: none"> VPC and Subnet: IPv6-enabled subnet and VPC. Self-assigned IPv6 address: Selected. Shared Bandwidth: Selected a shared bandwidth. <p>NOTE For details, see Setting Up an IPv6 Network.</p>	<ul style="list-style-type: none"> IPv4 CIDR Block IPv6 CIDR block 	<ul style="list-style-type: none"> Private IPv4 address + IPv4 EIP: Bind an IPv4 EIP to the instance to allow public IPv4 communication. Private IPv4 address: Do not bind any IPv4 EIP to the instance and use only the private IPv4 address to allow private IPv4 communication. IPv6 address + shared bandwidth: Allow both private IPv6 communication and public IPv6 communication.

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 11-5](#).

Table 11-5 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"> An EIP has been bound to the ECS. IPv6 EIP has been enabled. 	IPv4 CIDR Block	<ul style="list-style-type: none"> Private IPv4 address: used for private IPv4 communication. IPv4 EIP (with IPv6 EIP enabled): used for public network communication through IPv4 and IPv6 addresses.

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 11-6](#) 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 11-6 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.

12 Shared Bandwidth

12.1 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](#)

Notes and Constraints

- 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](#).

12.2 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 12-1 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">• Yearly/Monthly: You pay for the bandwidth by year or month before using it. No other charges apply during the validity period of the bandwidth.• Pay-per-use: You pay for the bandwidth based on the amount of time you use the bandwidth.	Yearly/Monthly

Parameter	Description	Example Value
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across different regions. For lower network latency and faster access to your resources, select the region nearest you.	N/A
Bandwidth Type	Select a type of the shared bandwidth based on your EIP type. <ul style="list-style-type: none">● Standard: Dynamic BGP and premium BGP EIPs can be added to a shared bandwidth of this type. NOTE In the CN-Hong Kong region, only dynamic BGP EIPs can be added to standard shared bandwidths.	Standard
Billed By	The billing method for the shared bandwidth. You can specify a shared bandwidth to be billed by bandwidth.	Bandwidth
Bandwidth	The bandwidth size in Mbit/s.	10
Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The name of the default project is default .	default
Name	The name of the shared bandwidth.	Bandwidth-001
Required Duration	The duration for which the purchased EIP will use. The duration must be specified if the Billing Mode is set to Yearly/Monthly .	2 months
Auto-renew	Whether to select Auto-renew . You can select it if the Billing Mode is set to Yearly/Monthly . The auto-renewal period is determined by the required duration. <ul style="list-style-type: none">● Monthly subscription: The subscription is renewed every month.● Yearly subscription: The subscription is renewed each year.	N/A

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**.

12.3 Adding EIPs to a Shared Bandwidth



Scenarios

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

Notes and Constraints

- Currently, yearly/monthly EIPs cannot be added to a shared bandwidth.
- 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.
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?](#)

12.4 Removing EIPs from a Shared Bandwidth

Scenarios

Remove EIPs that are no longer required from a shared bandwidth if needed.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and 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.
7. Click **OK**.

12.5 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 its bandwidth.** The change will be applied immediately and the price difference will be billed accordingly.
 - **You can decrease its bandwidth.** The change will be applied 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.

12.6 Deleting a Shared Bandwidth

Scenarios

Delete a shared bandwidth when it is no longer required.

Notes and Constraints



- A yearly/monthly shared bandwidth cannot be directly deleted. It can only be unsubscribed from.

- 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. Confirm the information and enter **DELETE**.
7. In the displayed dialog box, click **OK**.

13 Monitoring and Auditing

13.1 Cloud Eye Monitoring

13.1.1 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 13-1 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	≥ 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	≥ 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	≥ 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	≥ 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"
 }
],



13.1.2 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)


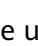
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click  to open the service list and choose **Management & Deployment > Cloud Eye**.
4. Click **Cloud Service Monitoring** on the left of the page, 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.

13.1.3 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](#).

13.2 CTS Auditing

13.2.1 Key Operations Recorded by CTS

With CTS, you can record operations performed on the VPC service for further query, audit, and backtracking purposes.

Table 13-2 lists the VPC operations that can be recorded by CTS.

Table 13-2 VPC operations that can be recorded by CTS

Operation	Resource Type	Trace Name
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 Name
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 Name
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 Name
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

13.2.2 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 & Deployment**, 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.


14 Managing Quotas

What Is a Quota?

A quota limits the quantity of a resource available to users, thereby preventing spikes in the usage of the resource. For example, a VPC quota limits the number of VPCs that can be created.

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

How Do I View My Quotas?

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

How Do I Apply for a Higher Quota?

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