Virtual Private Cloud

User Guide

 Issue
 01

 Date
 2025-07-22





HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

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

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

Trademarks and Permissions

NUAWEI 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 an IAM 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	17
2.2.3 Connecting VPCs to the Public Network	20
2.2.4 Connecting VPCs to On-Premises Data Centers	25
2.3 VPC	29
2.3.1 Creating a VPC with a Subnet	29
2.3.2 Extending a VPC CIDR Block Using a Secondary IPv4 CIDR Block	40
2.3.3 Obtaining a VPC ID	43
2.3.4 Modifying a VPC	43
2.3.5 Viewing a VPC Topology	45
2.3.6 Exporting VPCs	46
2.3.7 Managing VPC Tags	46
2.3.8 Deleting a VPC	48
2.4 Subnet	49
2.4.1 Creating a Subnet for an Existing VPC	49
2.4.2 Modifying a Subnet	55
2.4.3 Exporting Subnets	57
2.4.4 Viewing and Deleting Resources in a Subnet	58
2.4.5 Viewing IP Addresses in a Subnet	60
2.4.6 Managing Subnet Tags	61
2.4.7 Deleting a Subnet	63
3 Route Table and Route	65
3.1 Route Table and Route Overview	65
3.2 Managing Route Tables	73
3.2.1 Creating a Custom Route Table	74
3.2.2 Associating a Route Table with a Subnet	75

3.2.3 Changing the Poute Table Associated with a Subnet	76
3.2.4 Viewing the Poute Table Associated with a Subnet	70
3.2.5 Viewing Route Table Information	70
3.2.6 Deleting a Route Table	77
3.3 Managing Routes	
3.3.1 Adding Routes to a Route Table	
3.3.2 Modifying a Route	80
3.3.3 Replicating a Route	82
3.3.4 Deleting a Route	83
3.4 Route Configuration Examples	84
3.4.1 Configuring an SNAT Server to Enable ECSs to Share an EIP to Access the Internet	84
4 Virtual IP Address	88
4.1 Virtual IP Address Overview	88
4.2 Assigning a Virtual IP Address	92
4.3 Binding a Virtual IP Address to an Instance or EIP	93
4.4 Unbinding a Virtual IP Address from an Instance or EIP	100
4.5 Releasing a Virtual IP Address	101
4.6 Virtual IP Address Configuration Example	102
4.6.1 Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster	102
5 Elastic Network Interface and Supplementary Network Interface	. 119
5.1 Elastic Network Interface	119
5.1.1 Elastic Network Interface Overview	119
5.1.2 Creating a Network Interface	120
5.1.3 Managing Network Interfaces	121
5.1.4 Deleting a Network Interface	125
5.2 Supplementary Network Interfaces	126
5.2.1 Supplementary Network Interface Overview	126
5.2.2 Creating a Supplementary Network Interface	127
5.2.3 Viewing the Basic Information About a Supplementary Network Interface	160
5.2.4 Binding or Unbinding an EIP to or from a Supplementary Network Interface	161
5.2.5 Changing Security Groups That Are Associated with a Supplementary Network Interface	162
5.2.6 Deleting a Supplementary Network Interface	163
5.3 Network Interface Configuration Examples	164
5.3.1 Binding an EIP to the Extended Network Interface of an ECS to Enable Internet Access	164
5.3.2 Configuring Policy-based Routes for an ECS with Multiple Network Interfaces	169
5.3.2.1 Overview	169
5.3.2.2 Collecting ECS Network Information	170
5.3.2.3 Automatically Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Network Interfaces (Huawei Cloud EulerOS 2.0/CentOS 8.0 or Later)	171
5.3.2.4 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Netwo Interfaces (CentOS)	ork 191

5.3.2.5 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Netw Interfaces (Ubuntu)	/ork 203
5.3.2.6 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Windows ECS with Multiple Network Interfaces	215
6 Access Control	221
6.1 Access Control Overview	221
6.2 Security Group	223
6.2.1 Security Group and Security Group Rule Overview	223
6.2.2 Default Security Groups	234
6.2.3 Security Group Examples	236
6.2.4 Common ECS Ports	242
6.2.5 Managing a Security Group	244
6.2.5.1 Creating a Security Group	244
6.2.5.2 Cloning a Security Group	248
6.2.5.3 Modifying a Security Group	249
6.2.5.4 Deleting a Security Group	250
6.2.6 Managing Security Group Rules	251
6.2.6.1 Adding a Security Group Rule	251
6.2.6.2 Fast-Adding Security Group Rules	263
6.2.6.3 Modifying a Security Group Rule	268
6.2.6.4 Replicating a Security Group Rule	269
6.2.6.5 Importing and Exporting Security Group Rules	269
6.2.6.6 Deleting One or More Security Group Rules	272
6.2.6.7 Querying Security Group Rule Changes	274
6.2.7 Managing Instances Added to a Security Group	277
6.2.7.1 Adding an Instance to or Removing an Instance from a Security Group	278
6.2.7.2 Changing the Security Group of an ECS	279
6.3 Network ACL	280
6.3.1 Network ACL Overview	280
6.3.2 Network ACL Configuration Examples	290
6.3.3 Managing Network ACLs	294
6.3.3.1 Creating a Network ACL	294
6.3.3.2 Modifying a Network ACL	295
6.3.3.3 Enabling or Disabling a Network ACL	295
6.3.3.4 Viewing a Network ACL	296
6.3.3.5 Deleting a Network ACL	296
6.3.4 Managing Network ACL Rules	297
6.3.4.1 Adding a Network ACL Rule	297
6.3.4.2 Modifying a Network ACL Rule	300
6.3.4.3 Enabling or Disabling One or More Network ACL Rules	302
6.3.4.4 Exporting and Importing Network ACL Rules	303
6.3.4.5 Deleting One or More Network ACL Rules	304
6.3.5 Managing Subnets Associated with a Network ACL	305

6.3.5.1 Associating Subnets with a Network ACL	305
6.3.5.2 Disassociating Subnets from a Network ACL	307
7 IP Address Group	309
7.1 IP Address Group Overview	309
7.2 Managing an IP Address Group	311
7.2.1 Creating an IP Address Group	311
7.2.2 Associating an IP Address Group with Resources	313
7.2.3 Modifying an IP Address Group	314
7.2.4 Exporting IP Address Group Details	315
7.2.5 Viewing the Details of an IP Address Group	315
7.2.6 Deleting an IP Address Group	316
7.3 Managing IP Address Entries in an IP Address Group	316
7.3.1 Adding IP Address Entries to an IP Address Group	316
7.3.2 Modifying IP Address Entries in an IP Address Group	318
7.3.3 Deleting IP Address Entries from an IP Address Group	320
7.4 IP Address Group Configuration Examples	320
7.4.1 Using IP Address Groups to Reduce the Number of Security Group Rules	320
8 VPC Peering Connection	324
8.1 VPC Peering Connection Overview	324
8.2 VPC Peering Connection Usage	326
8.2.1 VPC Peering Connection Usage Examples	
8.2.2 Using a VPC Peering Connection to Connect Two VPCs	327
8.2.3 Using a VPC Peering Connection to Connect Subnets in Two VPCs	367
8.2.4 Using a VPC Peering Connection to Connect ECSs in Two VPCs	381
8.2.5 Unsupported VPC Peering Configurations	
8.3 Creating a VPC Peering Connection to Connect Two VPCs in the Same Account	393
8.4 Creating a VPC Peering Connection to Connect Two VPCs in Different Accounts	400
8.5 Obtaining the Peer Project ID of a VPC Peering Connection	408
8.6 Modifying a VPC Peering Connection	409
8.7 Viewing VPC Peering Connections	410
8.8 Deleting a VPC Peering Connection	410
8.9 Modifying Routes Configured for a VPC Peering Connection	411
8.10 Viewing Routes Configured for a VPC Peering Connection	412
8.11 Deleting Routes Configured for a VPC Peering Connection	414
9 IPv4/IPv6 Dual-Stack Network	416
10 VPC Flow Log	422
10.1 VPC Flow Log.	422
10.2 Creating a VPC Flow Log	424
10.3 Viewing a VPC Flow Log	426
10.4 Enabling or Disabling a VPC Flow Log	427
10.5 Deleting a VPC Flow Log	428

10.6 VPC Flow Log Configuration Examples	428
10.6.1 Viewing the Traffic of ECSs from the Same VPC	428
10.6.2 Viewing the Traffic Between VPCs Connected by a VPC Peering Connection	435
10.6.3 Viewing the Traffic Between ECSs in Different VPCs Connected by an Enterprise Router	442
11 Elastic IP	449
11.1 EIP Overview	449
11.2 Assigning an EIP and Binding It to an ECS	450
11.3 Unbinding an EIP from an ECS and Releasing the EIP	455
11.4 Modifying an EIP Bandwidth	456
11.5 Exporting EIP Information	457
11.6 Managing EIP Tags	457
11.7 IPv6 EIP	459
12 Shared Bandwidth	460
12.1 Shared Bandwidth Overview	460
12.2 Assigning a Shared Bandwidth	461
12.3 Adding EIPs to a Shared Bandwidth	463
12.4 Removing EIPs from a Shared Bandwidth	464
12.5 Modify a Shared Bandwidth	464
12.6 Deleting or Unsubscribing from a Shared Bandwidth	465
13 Monitoring and Auditing	467
13.1 Cloud Eye Monitoring	467
13.1.1 Supported Metrics	467
13.1.2 VPC Events That Can Be Monitored	470
13.1.3 Viewing Metrics	470
13.1.4 Creating an Alarm Rule	471
13.2 CTS Auditing	472
13.2.1 Key Operations Recorded by CTS	472
13.2.2 Viewing VPC Traces	476
14 Managing Quotas	479

Permissions Management

1.1 Creating an IAM 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

- 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

• 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"
    }
]
```

• Example 4: Allowing users to view associated resources

To allow users to view resources associated with a specific resource, you need to assign them permissions to query that resource and its associated resources. The following is an example policy containing actions for allowing users to view the servers, extended network interfaces, and supplementary network interfaces associated with a security group:

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 Know 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).





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



D 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?

Once created, the CIDR block of a VPC or subnet cannot be modified. To ensure smooth service expansion and O&M, plan your VPC and subnet CIDR blocks carefully based on your service size and communication requirements.

D 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 enough 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.

The IPv4 CIDR block you specify when you create a VPC is the primary one. The primary CIDR block cannot be changed after the VPC is created. If IP addresses in the primary CIDR block are insufficient, you can add a secondary IPv4 CIDR block to the VPC. For details, see **Extending a VPC CIDR Block Using a Secondary IPv4 CIDR Block**.

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**.

VPC CIDR Block	IP Address Range	Netmask	Example CIDR Block
10.0.0/8-24	10.0.0.0- 10.255.255.255	8-24	10.0.0/8
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

Table 2-1 VPC CIDR blocks (RFC 1918)

In addition to these 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:

Fable 2-2 Reserved system	n and public CIDR blocks
---------------------------	--------------------------

Reserved System CIDR Blocks	Reserved Public CIDR Blocks	
• 100.64.0.0/10	• 0.0.0.0/8	
• 214.0.0.0/7	• 127.0.0.0/8	
• 198.18.0.0/15	• 240.0.0/4	
• 169.254.0.0/16		

Planning Subnet CIDR Blocks

• **Subnet mask planning:** A subnet CIDR block must be within its VPC CIDR block. Each subnet CIDR block in a VPC must be unique. A subnet mask can be between the netmask of its VPC CIDR block and the /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 10.0.0/16, 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 the CIDR block size:** After a subnet is created, the CIDR block cannot be changed. You need to plan the CIDR block in advance based on the number of IP addresses required by your service.
 - The subnet CIDR block cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements.
 Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address.
 - The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks from the VPC available for new subnets, which can be a problem when you want to scale out services.
- Avoiding subnet CIDR block conflicts: If you need to connect two VPCs or connect a VPC to an on-premises data center, there cannot be any CIDR block conflicts.

If the subnet CIDR blocks at both ends of the network conflict, **create a subnet**.

How Do I Know 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 information presented here.

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 (VPC-A, VPC-B, and VPC-X) in region A. 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.

on A		
VPC-A 172.16.0.0/16		VPC-B 172.17.0.0/16
Subnet-A01 172.16.0.0/24	VPC peering connection Peer-AB	Subnet-B01 172.17.0.0/24 ECS CCE RDS
		VPC-X 172.17.0.0/16
		Subnet-X01 172.17.9.0/24 ECS CCE RDS

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 networking services described in **Table 2-3**, you can flexibly connect VPCs in the same region, in different regions, or in different accounts.

Networking 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.	 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.	 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
Central Network	Cloud Connect can connect VPCs under the same account or different accounts across regions. You can attach VPCs in the same region to an enterprise router, and then add enterprise routers in different regions to a central network as attachments, so the VPCs can communicate with each other. This solution features higher scalability and is suitable for complex networking with many VPCs from different regions.	 VPC in different regions can be connected in minutes. Routes can be automatically added. Low latency and high speed
VPN	You can use VPN to connect VPCs in different regions. This will enable them to communicate with each other over the Internet.	 Low costs Simple configuration Immediate use Unstable networks dependent on the Internet quality
Direct Connect	You can use Direct Connect to connect VPCs in different regions.	 Dedicated connections with high security Low latency and high speed

Table 2-3 Networking services that can connect VPCs

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

Networking 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.	 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 • SNAT • DNAT	 NAT Gateway supports both source NAT (SNAT) and destination NAT (DNAT). SNAT enables multiple instances to share one or more EIPs to access the public network. 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. 	 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.

Networking Service	Function	Highlights
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.	 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.

Networking Service	Function	Highlights
VPN	VPN provides an encrypted, Internet- based channel that connects an on- premises data center and the cloud.	 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.	 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.	 VPC Peering is free. Routes can be configured on the console easily.

Table 2-5 Networking services that can connect VPCs to an on-premises data center

Networking Service	Function	Highlights
Central Network	You can use Direct Connect or VPN to connect on-premises data centers to VPCs in multiple regions and use a central network to connect the VPCs, so that the on-premises data centers can communicate with all the VPCs. You can attach VPCs and Direct Connect global DC gateways in the same region to an enterprise router, and then add the enterprise router, and then add the enterprise routers in different regions to a central network . In this way, VPCs in different regions can communicate with on-premises data centers in multiple cities. Compared with a cloud connection, a central network features a simpler network architecture and higher	 Route learning is supported. There is no need to configure routes manually. Network connection policies can be defined flexibly.

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.

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:

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

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

Central Network

You can attach VPCs in the same region to an enterprise router, and then add enterprise routers in different regions to a central network as attachments, so the VPCs can communicate with each other. This solution features higher scalability and is suitable for complex networking if there are multiple VPCs in different regions.

For details, see Connecting VPCs Across Regions Using Enterprise Router and Central Network.

In **Figure 2-10**, there are four VPCs in three regions: VPC-A in region A, VPC-B in region B, and VPC-C and VPC-D in region C. There is an enterprise router in each region: ER-A for VPC-A, ER-B for VPC-B, and ER-C for VPC-C and VPC-D. The VPCs are attached to the enterprise router in each region, and the enterprise routers in the three regions are added to a central network for cross-region network connectivity. If there will be more VPCs in the future, you only need to attach the VPCs to the enterprise router in the same region. Compared with a cloud connection, this solution features simpler network topology.



Figure 2-10 Connecting VPCs in different regions using a central network

VPN

You can use **VPN** to connect VPCs in different regions. This will enable them to communicate with each other over the Internet.

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 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.





Direct Connect

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

In **Figure 2-12**, 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 2-12 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 2-13**, 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.





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 you want ECSs in a VPC to share an EIP, see Using a Public NAT Gateway to Enable Servers to Share One or More EIPs to Access the Internet
- If you want ECSs in different VPCs to share an EIP, see Allowing VPCs to Share an EIP to Access the Internet Using Enterprise Router and SNAT.

ECSs in a VPC Sharing an EIP

In **Figure 2-14**, ECSs deployed on two subnets (Subnet-A01 and Subnet-A02) of VPC-A in region A need to access the Internet. 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 2-14 Enabling ECSs in a VPC to access the public network using a NAT gateway



ECSs in Different VPCs Sharing an EIP

In **Figure 2-15**, 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 2-15 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 a Public NAT Gateway to Enable Servers to Be Accessed by the Internet.

In **Figure 2-16**, 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 2-16 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 2-17**, 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 2-17 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.

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 Onpremises Data Center to a VPC**.

In **Figure 2-18**, 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.





Direct Connect

Direct Connect establishes a dedicated network connection between an onpremises data center and the cloud.

For details, see Accessing a VPC over a Direct Connect Connection and Using BGP to Route Traffic.

In **Figure 2-19**, 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 2-19 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.

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 2-20**, 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 2-20 Connecting an on-premises data center to VPCs using Direct Connect and VPC Peering



Enterprise Router

You can use VPN or Direct Connect to connect an on-premises data center to a VPC, and then use an enterprise router to connect multiple VPCs if there are in the same region.

- Setting Up a Hybrid Cloud Network Using Enterprise Router and Direct Connect (Global DC Gateway)
- Setting Up a Hybrid Cloud Network Using Enterprise Router and a Pair of Direct Connect Connections (Global DC Gateway)
- Setting Up a Hybrid Cloud Network Using Enterprise Router and a Pair of Active/Standby Direct Connect Connections (Global DC Gateway)
- Setting Up a Hybrid Cloud Network Using Enterprise Router, VPN, and Direct Connect (Global DC Gateway)

In **Figure 2-21**, 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 two VPCs and global DC gateways are attached to an enterprise router in the same region, so that the two VPCs can communicate with each other and also with the on-premises data center.

In this example, two Direct Connect connections are deployed to balance loads, improving the network performance and reliability. When both connections work normally, the network transmission capability is greatly improved. If one connection becomes faulty, the other connection can take over services, and your on-premises data center can still access the VPCs.

Figure 2-21 Connecting an on-premises data center to VPCs in the same region using Direct Connect and Enterprise Router



Connecting Multiple VPCs in Different Regions to On-Premises Data Centers

To connect multiple VPCs in different regions to on-premises data centers, you can use Direct Connect or VPN to connect each on-premises data center to a VPC, and then use a cloud connection or central network to connect all 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.

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.

Central Network

You can attach VPCs and Direct Connect global DC gateways in the same region to an enterprise router, and then add the enterprise routers in different regions to

a **central network**. In this way, VPCs in different regions can communicate with on-premises data centers in multiple cities. Compared with a cloud connection, using a central network features a simpler architecture and higher scalability.

In **Figure 2-22**, VPCs and global DC gateways in each region are attached to different enterprise routers, so the on-premises data center in each city can access the VPCs in the corresponding region. Then the two enterprise routers (ER-A and ER-C) are connected over a central network. In this way, the three VPCs (VPC-A, VPC-B, and VPC-C) and two on-premises data centers (IDC-A and IDC-C) are on the same cloud network and can communicate with each other. In this solution, only the enterprise router in each region is added to the central network, simplifying the network architecture. Also, with global DC gateways attached to enterprise routers, VPCs can share Direct Connect connections to communicate with the on-premises data centers. Route learning of enterprise routers eliminates complex configurations and simplifies maintenance.





2.3 VPC

2.3.1 Creating a VPC with a 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 page for **creating a VPC**.
2. On the **Create VPC** page, set parameters for the VPC and subnets as prompted.

You can click $^{\scriptsize \textcircled{}}$ to create more subnets. A maximum of three subnets can be created at a time.

Basic Information	
Region	EU-Dublin Pegions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.
Name	vpc-19e3
IPv4 CIDR Block	192 · 168 · 0 / 16 • Recommended: 10.0.0.0%-24 (Select) 172.16.0.0/12-24 (Select) 192.168.0.0/16-24 (Select) 192.168.0.0/16-24 (Select)
	A 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. Very VPC CIDR blocks in current region
Enterprise Project	-Select- • C Create Enterprise Project ?
Enterprise Project	Select C Create Enterprise Project (2)
Enterprise Project Advanced Settings	Seleci C Create Enterprise Project (2)
Advanced Settings	-Select- C Create Enterprise Project Tag Description AZ1
Advanced Settings	Select C Create Enterprise Project C
Enterprise Project Advanced Settings Advanced Settings	-Select • C Create Enterprise Project (*) Tag Description AZ1 • (*) Subnet-ta16 I122 • 168 • (*) • (*) / (24 • *) Available IP Addresser :251 The CIDR Book cannot be modified after the subnet has been created.
Content of the second	C Create Enterprise Project (*) Tag Description AZ1

Figure 2-23 Creating a VPC and subnet

Table 2-6 VPC parameter	descriptions
---------------------------------	--------------

Parameter	Description	Example Value
Region	Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the region nearest to where your services will be accessed.	EU-Dublin
Name	The VPC name. The name:Can contain 1 to 64 characters.	vpc-test
	 Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	

Parameter	Description	Example Value
Enterprise Project	The enterprise project to which the VPC belongs.	default
	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.	
IPv4 CIDR Block	The CIDR block of the VPC. The CIDR block of a subnet can be the same as the CIDR block for the VPC (for a single subnet in the VPC) or a subset of the CIDR block for the VPC (for multiple subnets in the VPC).	10.0.0/8
	The following CIDR blocks are supported:	
	• 10.0.0/8-24	
	• 172.16.0.0/12-24	
	• 192.168.0.0/16-24	
	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.	
Advanced Settings (Optional) > Tag	The VPC tag. Click \checkmark to expand the configuration area and set this parameter.	Key: vpc_key1Value: vpc-01
	Add tags to help you quickly identify, classify, and search for your VPCs.	
	For details, see Managing VPC Tags.	

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

Table 2-7 Subnet parameter descriptions

Parameter	Description	Example Value
Name	The subnet name. The name:Can contain 1 to 64 characters.	subnet-01
	 Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	

Description	Example Value
An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network. Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.	AZ1
 By default, all instances within a given VPC can communicate with each other, even if they are in different subnets that are 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 and its subnet can be in different AZs. 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. For details, see Region and 	
	 Description An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network. Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services. By default, all instances within a given VPC can communicate with each other, even if they are in different subnets that are 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 and its subnet can be in different AZs. 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. For details, see Region and AZ.

Parameter	Description	Example Value
Parameter IPv4 CIDR Block	 Description This parameter is displayed only in regions where IPv4/ IPv6 dual stack is supported. The IPv4 CIDR block of the subnet. A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets: Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to plan the CIDR block in advance based on the number of IP addresses required by your service. The subnet CIDR block cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. The first and last three addresses in a subnet are reserved for system use. For example, in subnet 10.0.0/24, 10.0.0.1 is the gateway address, 10.0.253 is the system interface address. The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks from the VPC available for new subnets, which can be a problem when you 	Example Value 10.0.0/24
	services.	

Parameter	Description	Example Value
	 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. For details about subnet planning, see VPC and Subnet 	
	Planning.	
IPv6 CIDR Block	This parameter is displayed only in regions where IPv4/ IPv6 dual stack is supported.	-
	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.	
	For details, see IPv4 and IPv6 Dual-Stack Network.	

Parameter	Description	Example Value
Associated Route Table	The default route table with which the subnet will be associated. A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. A default route table automatically comes with a VPC. Subnets in the VPC are automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other. 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.	-
Advanced Settings (Optional) > Gateway	The gateway address of the subnet. Click \checkmark to expand the configuration area and set this parameter. Retain the default value unless there are special requirements.	10.0.0.1

Parameter	Description	Example Value
Advanced Settings (Optional) > DNS Server Address	The DNS server addresses. Click ✓ to expand the configuration area and set this parameter. 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.	100.125.x.x
	You can change the default DNS server addresses if needed. This may interrupt your access to cloud services.	
	You can also click Reset on the right to restore the DNS server addresses to the default value.	
	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	

Parameter	Description	Example Value
Advanced Settings (Optional) > Domain Name	The domain name. Click \checkmark to expand the configuration area and set this parameter. Enter domain names separated by spaces, up to 254 characters total. Each label in a domain name can contain a maximum of 63 characters. (For example, test and com are two labels in test.com.)	test.com
	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.	
	If the domain names are changed, ECSs newly added to this subnet will use the new domain names.	
	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.	

Parameter	Description	Example Value
Advanced Settings (Optional) > IPv4 DHCP Lease Time	The period during which a client can use an IP address automatically assigned by the DHCP server. Click \checkmark to expand the configuration area and set this parameter.	-
	The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client.	
	• Limited: Set the DHCP lease time. The unit can be day or hour.	
	• Unlimited : The DHCP lease time does not expire.	
	After you change the DHCP lease time on the console, the change is applied automatically when the DHCP lease of an instance (such as ECS) is renewed. You can wait for the system to renew the lease or manually renew the lease. Renewing lease will not change the IP address used by the instance. If you want the new lease time to take effect immediately, manually renew the lease or restart the ECS.	
Advanced Settings (Optional) > Tag	The subnet tag. Click \checkmark to expand the configuration area and set this parameter. Add tags to help you quickly identify, classify, and search for your subnets.	Key: subnet_key1Value: subnet-01
	For details, see Managing Subnet Tags.	

Parameter	Description	Example Value
Advanced Settings (Optional) > Description	Supplementary information about the subnet. Click \checkmark to expand the configuration area and set this parameter.	N/A
	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 >).	

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 Network in a VPC** and **Setting Up an IPv4/IPv6 Dual-Stack Network In a VPC**.

2.3.2 Extending a VPC CIDR Block Using a Secondary IPv4 CIDR Block

Scenarios

Generally, the number of IP addresses in a VPC CIDR block determines how many cloud resources that can be deployed in the VPC. If there are no sufficient IP addresses in the VPC CIDR block, you can add a secondary IPv4 CIDR block to expand the VPC CIDR block and increase the number of IP addresses.

The IPv4 CIDR block you specify when you create a VPC is the primary one. The primary CIDR block cannot be changed after the VPC is created. If IP addresses in the primary CIDR block are insufficient, you can add a secondary CIDR block to the VPC. The secondary CIDR block can be used in the same way as the primary CIDR block.

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 a VPC" in the *Virtual Private Cloud API Reference*.

Configuration Example

As services develop, the IP addresses in VPC-A (primary CIDR block: 192.168.10.0/24) were insufficient. To solve this problem, you can add a secondary CIDR block (172.17.10.0/24) to this VPC and create two subnets (Subnet-A03 and

Subnet-A04) in the secondary CIDR block to support future resource deployment and service expansion.



Figure 2-24 Configuration example with a secondary CIDR block

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 you create a subnet (100.20.0.0/16) from a secondary CIDR block of a VPC, the system will automatically generate a **Local** route with the destination of 100.20.0.0/16. If the VPC route table already has a route with the VPC peering connection as the next hop and 100.20.0.0/24 as the destination, the two destinations (100.20.0.0/16 and 100.20.0.0/24) overlap 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** provides you with IP address ranges that cannot be used as secondary IPv4 CIDR blocks. For example, the CIDR block 192.168.0.0/16 has IP addresses from 192.168.0.0 to 192.168.255.255, indicating that none of the IP addresses can be included in a secondary IPv4 CIDR block, for example, 192.168.0.0/16, 192.168.31.0/24, 192.168.100.0/24, and 192.168.255.255/32.

Туре	CIDR Block	IP Address Range	
Primary CIDR blocks	10.0.0/8	10.0.0-10.255.255.255	
and existing CIDR blocks	172.16.0.0/12	172.16.0.0-172.31.255.255	
	192.168.0.0/16	192.168.0.0-192.168.255.255	
Reserved system	100.64.0.0/10	100.64.0.0-100.127.255.255	
CIDR blocks	214.0.0.0/7	214.0.0.0-215.255.255.255	
	198.18.0.0/15	198.18.0.0-198.19.255.255	
	169.254.0.0/16	169.254.0.0-169.254.255.255	
Reserved public	0.0.0/8	0.0.0.0-0.255.255.255	
CIDR blocks	127.0.0.0/8	127.0.0.0-127.255.255.255	
	240.0.0/4	240.0.0.0-255.255.255.255	

 Table 2-8 IP address ranges that cannot be used as secondary IPv4 CIDR

 blocks

Adding a Secondary IPv4 CIDR Block

- 1. Log in to the management console.
- 2. Click \equiv 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 target VPC and click **Edit CIDR Block** in the **Operation** column.

The Edit CIDR Block dialog box is displayed.

- 4. Click Add Secondary IPv4 CIDR Block.
- 5. Enter a secondary IPv4 CIDR block in the text box and click **OK**.

Deleting a Secondary IPv4 CIDR Block

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

The VPC list page is displayed.

3. In the VPC list, locate the target 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.
 - 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.
- 5. 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 \equiv 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 \square next to ID to copy the VPC ID.

Figure 2-25 VPC ID

< vpc-8					
Summary Topology	Summary Topology Tags				
VPC Information	VPC Information				
Name	vpc-В 🖉		ID	17cd7278-	
Status	Available		CIDR Block	172.17.0.0/16 Edit CIDR Block	
Enterprise Project	default		Description	🖉	

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

Constraints

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 a VPC" 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 \checkmark on the right of the VPC name.
 - ii. Enter a VPC name and click **OK**.
 - Method 2:
 - In the VPC list, locate the target VPC and click its name.
 The **Summary** page is displayed.
 - ii. Click \checkmark on the right of the VPC name or description, enter the information, and click \checkmark .

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 target 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.

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.



Figure 2-26 Modifying network address and subnet mask

6. Click OK.

2.3.5 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.

2.3.6 Exporting VPCs

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 \equiv 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.

2.3.7 Managing VPC Tags

Scenarios

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

- Add tags to a VPC.
- Modify a VPC tag.
- Delete a VPC tag.

For details about VPC tag requirements, see Table 2-9.

Paramete r	Requirements	Example Value
Кеу	 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: Uppercase letters Lowercase letters Digits Special characters, including hyphens (-) and underscores (_) 	vpc_key1
Value	 Can contain a maximum of 43 characters. Can contain only the following character types: Uppercase letters Lowercase letters Digits Special characters, including periods (.), hyphens (-) and underscores (_) 	vpс-01

Table 2-9 VPC	tag key and valu	le requirements
---------------	------------------	-----------------

Notes and Constraints

Each cloud resource can have a maximum of 20 tags.

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 tag you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

2.3.8 Deleting a VPC

Scenarios

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

The VPC service has multiple resources. Some are free, while some are not. For details about VPC resource pricing, see Pricing Details.

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 the VPC.

4. Confirm the information and click Yes.

If a VPC cannot be deleted, a message will be displayed on the console. Delete the resources 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 an Existing 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 \bigcirc in the upper left corner and select the desired region and project.
- 3. Click \equiv 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 subnet parameters as needed.

You can click $^{\bigoplus}$ to create more subnets. A maximum of three subnets can be created at a time.

Parameter	Description	Example Value
Region	The region where the 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: Can contain 1 to 64 characters. Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	subnet-01

 Table 2-10
 Subnet parameter descriptions

Parameter	Description	Example Value
AZ	An AZ is a geographic location with independent power supply and network facilities in a region. AZs are physically isolated, and AZs in the same VPC are interconnected through an internal network.	AZ1
	Each region contains multiple AZs. If one AZ is unavailable, other AZs in the same region continue to provide services.	
	• By default, all instances within a given VPC can communicate with each other, even if they are in different subnets that are 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 and its subnet can be in different AZs. 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.	
	For details, see Region and AZ .	

Parameter	Description	Example Value
Parameter IPv4 CIDR Block	 Description This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported. The IPv4 CIDR block of the subnet. A subnet is a unique CIDR block with a range of IP addresses in a VPC. Comply with the following principles when planning subnets: Planning CIDR block size: After a subnet is created, the CIDR block cannot be changed. You need to plan the CIDR block in advance based on the number of IP addresses required by your service. The subnet CIDR block cannot be 	Example Value 10.0.0/24
	 too small. Ensure that the number of available IP addresses in the subnet meets service requirements. The first and last three addresses in a subnet are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address. The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you 	
	 may not have enough CIDR blocks from the VPC available for new subnets, which can be a problem when you want to scale out services. Avoiding subnet CIDR block 	
	 conflicts: Avoid CIDR block conflicts if you need to connect two VPCs or connect a VPC to an on-premises data center. If the subnet CIDR blocks at both ends of a network conflict, create a subnet. 	
	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.	

Parameter	Description	Example Value
IPv6 CIDR Block (Optional)	This parameter is displayed only in regions where IPv4/IPv6 dual stack is supported.	N/A
	If you select this option, the system automatically assigns an IPv6 CIDR block to the created subnet. Currently, the IPv6 CIDR block cannot be customized. IPv6 CIDR block cannot be disabled after the subnet is created. For details, see IPv4 and IPv6 Dual- Stack Network.	
Associated Route Table	A route table contains a set of routes that are used to control the traffic routing for your subnets in a VPC. A default route table automatically comes with a VPC. Subnets in the VPC are automatically associated with the default route table. The default route table ensures that subnets in a VPC can communicate with each other. 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.	N/A
Advanced Settings (Optional) > Gateway	Click ✓ to expand the configuration area and set this parameter. The gateway address of the subnet. Retain the default value unless there are special requirements.	10.0.0.1

Parameter	Description	Example Value
Advanced Settings (Optional) > DNS Server Address	Click ✓ to expand the configuration area and set this parameter. 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. You can change the default DNS server addresses if needed. This may interrupt your access to cloud services. You can also click Reset on the right to restore the DNS server addresses to the default value. A maximum of two DNS server IP addresses must be separated using commas (.).	100.125.x.x
Advanced Settings (Optional) > IPv4 DHCP Lease Time	 Click ✓ to expand the configuration area and set this parameter. The period during which a client can use an IP address automatically assigned by the DHCP server. After the lease time expires, a new IP address will be assigned to the client. Limited: Set the DHCP lease time. The unit can be day or hour. Unlimited: The DHCP lease time does not expire. 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. 	N/A
Advanced Settings (Optional) > Tag	Click V 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.	 Key: subnet_key1 Value: subnet-01

Parameter	Description	Example Value
Advanced Settings >	Click \checkmark to expand the configuration area and set this parameter.	N/A
Description	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 >).	

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.

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.
- 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 $\overset{\frown}{=}$ on the right of the parameter to be modified and modify the parameter as prompted.

Table 2-11 Parameter descrip	otions
------------------------------	--------

Parameter	Description	Example Value
Name	The subnet name. The name:Can contain 1 to 64 characters.	Subnet
	 Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	

Parameter	Description	Example Value
DNS Server Address	By default, two DNS server addresses are configured. You can change them as required. A maximum of two DNS server addresses are supported. Use commas (,) to separate every two addresses.	100.125.x.x
	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.	
	You can change the default DNS server addresses if needed. This may interrupt your access to cloud services.	
	You can also click Reset on the right to restore the DNS server addresses to the default value.	
	A maximum of two DNS server IP addresses can be configured. Multiple IP addresses must be separated using commas (,).	

Parameter	Description	Example Value
DHCP Lease Time	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.	-
	• Limited: Set the DHCP lease time. The unit can be day or hour.	
	• Unlimited : The DHCP lease time does not expire.	
	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.	
Description	Supplementary information about the subnet. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

2.4.3 Exporting Subnets

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 \equiv 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.
- 5. In the upper left corner of the subnet list, click **Export**.
 - **Export selected data to an XLSX file**: Select one or more subnets and export information about the selected subnets.
 - Export all data to an XLSX file: Export information about all the subnets in the current region.

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

2.4.4 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.

NOTE

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

Procedure

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

The Virtual Private Cloud page is displayed.

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

The subnet details page is displayed.

- 6. On the **Summary** page, view the resources in the subnet.
 - a. In the **Resources** area in the lower part of the page, view the number of resources in the subnet. Click the number to the right of each resource to view the resources in the subnet.
 - b. In the **Networking Components** area on the right of the page, view the NAT gateways in the subnet.

Figure 2-28 Viewing resources in a subnet

< subnet Summary IP Addre	! Isses Tags		C
Subnet Information Name Az Status VPC IPv4 CIDR Block Description	иллин — С Az1 Avaladar ecc — (192 168 0.016) 122 168 0.024 — С	Natura ID 21100000000000000000000000000000000000	Networking Components Image: Components Roads Tables 10-yes (2013)/387 (2014un); (2) Image: Components Image: Compon
Gateway and DNS DHCP DNS Server Address NTP Server Address	Information Examts 100 125 1.256 100 125 64,250 🖉 ① - 🖉 ②	Cutreary 192.168.8.1 DI-CP Lasse Time Unimited <i>L</i>	Interface of the second sec
VPC Resources ECSs Add Network Interfaces	1 Mas Add	0 Last Balances 1	0

7. Delete resources from the subnet.

Table 2-12	Viewing and	l deleting r	esources in	a subnet
------------	-------------	--------------	-------------	----------

Resource	Reference
ECS	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.
	 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	You cannot jump to the target BMS from the current page. To delete a BMS from the subnet, you need to go to the BMS console, search for the target BMS in the BMS list, and delete it.
	 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	You can directly jump to the target load balancer page.
	 Click the number to the right of Load Balancers. The load balancer list is displayed.
	 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	You can directly jump to the target network interface page.
	 Click the number to the right of Network Interfaces. The Network Interfaces page is displayed.
	 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	You can directly jump to the target NAT gateway page. 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.
	Deleting a Public NAT Gateway
	Deleting a Private NAT Gateway

2.4.5 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, RDS instances, and load balancers.

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 \equiv 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.
- 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 2-29 Viewing IP addresses in a subnet

< subnet-fi					5
ummary IP Addresses Tags					
Assign Virtual IP Address Unbind EIP Learn	nmore about virtual IP address configuration.				Virtual IP Address
Virtual IP Address	Bound EIP		Bound Server (NIC)		Operation
192.168.0.25	-				Bind to EIP Bind to Server More -
					IP Address V Enter an IP address. Q C
IP Address		Used By		Operation	
IPv4: 192,165.0.1 IPv6:		Gateway		Reference	
192.168.0.23		Dedicated Load Balancer		Release	
192.168.0.67		Dedicated Load Balancer		Release	
192.168.0.114		Server		Release	
192.168.0.160		Supplementary network interface		Release	
192.168.0.222		NAT Gateway		Release	
IPv4 192.168.0.253 IPv6		System interface		Release	
192.168.0.254		DHCP		Release	

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?

2.4.6 Managing Subnet Tags

Scenarios

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

- Add a tag to a subnet.
- Modify a subnet tag.
- Delete a subnet tag.

For details about subnet tag requirements, see Table 2-13.

Parame ter	Requirements	Example Value
Key	 Cannot be left blank. Must be unique for each subnet. Can contain a maximum of 36 characters. Can contain only the following character types: Uppercase letters Lowercase letters Digits Special characters, including hyphens (-) and underscores (_) 	subnet_key1
Value	 Can contain a maximum of 43 characters. Can contain only the following character types: Uppercase letters Lowercase letters Digits Special characters, including hyphens (-) and underscores (_) 	subnet-01

Table 2-13 Subnet tag k	ey and value requirements
-------------------------	---------------------------

Notes and Constraints

Each cloud resource can have a maximum of 20 tags.

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.

- 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 \equiv 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.
- 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 tag you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **OK**.

2.4.7 Deleting a Subnet

Scenarios

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

D NOTE

The VPC service has multiple resources. Subnets can be used for free. For details about VPC resource pricing, see **Pricing Details**.

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 \bigcirc 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.
- 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 Table and Route

3.1 Route Table and Route Overview

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 a route table can be associated with multiple subnets.

There are **default route tables** and **custom route tables**. You can add IPv4 and IPv6 routes to them.



Figure 3-1 Route tables

• **Default route table:** Each VPC comes with a default route table. If you create a subnet in a VPC, the subnet is automatically associated 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 of a subnet only controls outbound traffic. The default route table of the VPC that the subnet belongs to controls inbound traffic, for example, traffic from VPC peering connections, Direct Connect, and VPN connections to the VPC.

NOTE

By default, there is no quota for custom route tables. 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 route: A system route is automatically added by the VPC service or other services (such as VPN and Direct Connect) and cannot be deleted or modified.

Each route table comes with routes whose next hops are Local. Generally, a route table contains the following local routes:

- Routes whose destination is 100.64.0.0/10, which is used to deploy public services, for example, the DNS servers. The route directs 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 that enables 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 route: After a route table is created, you can add custom routes and configure information such as the destination and next hop in the route to determine where network traffic is directed. In addition to manually added

custom routes, there are custom routes added by other cloud services, such as Cloud Container Engine (CCE) or NAT Gateway.

Route tables include default route tables and custom route tables. They support the next hop types described in **Table 3-1** and **Table 3-2**.

Next Hop Type	Description	
Server	Traffic intended for the destination is forwarded to an ECS in the VPC.	
Extension NIC	Traffic intended for the destination is forwarded to the extended network interface of an ECS in the VPC.	
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.	
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.	
VPC endpoint	Traffic intended for the destination is forwarded to a VPC endpoint.	
Cloud container	Traffic intended for the destination is forwarded to a cloud container.	
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.	
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.	

Table 3-1 Next hop types supported by the default route table

Table 3-2 Next hop types supported by a custom route table

Next Hop Type	Description	
Server	Traffic intended for the destination is forwarded to an ECS in the VPC.	
Extension NIC	Traffic intended for the destination is forwarded to the extended network interface of an ECS in the VPC.	
BMS user-defined network	Traffic intended for the destination is forwarded to a BMS user-defined network.	

Next Hop Type	Description	
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.	
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.	
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.	
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.	
VPC endpoint	Traffic intended for the destination is forwarded to a VPC endpoint.	
Cloud container	Traffic intended for the destination is forwarded to a cloud container.	
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.	
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.	

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.

You cannot add a route whose next hop type is **VPC endpoint** or **Cloud container** to a route table. These routes are automatically added by the VPC Endpoint or CCE service.

How Route Tables Work

Each subnet in a VPC must have a route table associated. A subnet can be associated with the default route table or a custom route table. For details about the association between a subnet and a route table, see **Changing the Route Table Associated with a Subnet**.

Figure 3-2 shows two subnets in VPC-A. They are associated with different route tables to meet different network access requirements.

- Subnet 1 is associated with the default route table that contains a route pointing to the VPN gateway. This route allows ECSs in subnet 1 to access the on-premises data center through VPN.
- Subnet 2 is associated with a custom route table that contains a route pointing to the NAT gateway. This route allows ECSs in subnet 2 to access the Internet through the NAT gateway and EIP.



Figure 3-2 Default and custom route tables of subnets

Route Priority

In a VPC route table, **routes are matched in descending order of priority: local routes > specific routes > EIP routes > default route**. For details about each type of routes, see **Table 3-3**.

Table 3-3 Route priorities

No.	Route Type	Description
1	Local routes	Local routes are system routes used for communications within a VPC and have the highest priority. Table 3-4 provides examples of local routes.

No.	Route Type	Description			
2	Specific routes	Excepting local routes, if there are multiple routes that match the request destination, the longest prefix match is used. This means that the route with the longest subnet mask is preferentially used to determine the next hop.			
		For example, if the destination of traffic entering a VPC is 192.168.1.12/32, the VPC route table has the following routes:			
		 The destination of route A is 192.168.0.0/16, with ECS-A as the next hop. 			
		• The destination of route B is 192.168.1.0/24, with a VPC peering connection as the next hop.			
		• The destination of route C is 0.0.0/0, with an NAT gateway as the next hop.			
		According to the longest prefix match, the request preferentially matches route B and will be forwarded to the VPC peering connection.			
3	EIP routes	If an ECS in a subnet has an EIP bound, the EIP route takes precedence over the default route (destination: 0.0.0.0/0) in the route table. In this case, the EIP is used to access the Internet. Example:			
		 The destination of route A is 0.0.0/0, with an NAT gateway as the next hop. 			
		• ECS-A in a VPC subnet has an EIP bound.			
		In this case, ECS-A will use the EIP to access the Internet instead of the NAT gateway.			
4	Default route	The route with the destination 0.0.0.0/0 is the default route, which can match any traffic. According to the longest prefix match, 0.0.0.0/0 has the lowest priority.			

Route Table Configuration

You can configure routes with different next hop types in a VPC route table to meet specific network access requirements. For example, you can set the next hop to a VPC peering connection to enable communications between VPCs, or set the next hop to a NAT gateway to access the Internet.

Connecting VPCs Through a VPC Peering Connection and Routes

As shown in **Figure 3-3**, VPC-A and VPC-B in region A are connected by VPC peering connection peering-AB and the routes that point to the peer VPC in the route tables of the two VPCs.



Figure 3-3 Connecting VPCs in the same region over a VPC peering connection

Connecting VPCs Through an Enterprise Router and Routes

As shown in **Figure 3-4**, there is an enterprise router in region A with VPCs attached. The system automatically adds routes to point to the enterprise router to each VPC route table, and adds routes to point to each VPC to the enterprise router route table. In this way, the enterprise router can forward traffic across the four VPCs.



Figure 3-4 Connecting VPCs in the same region using an enterprise router

Connecting VPCs to the Internet Using a NAT Gateway and Routes

In **Figure 2-14**, ECSs in Subnet-A01 of VPC-A in region A need to access the Internet. You need to create a public NAT gateway in Subnet-NAT and configure an SNAT rule for Subnet-A01. The system automatically adds a route pointing to the NAT gateway to the route table of VPC-A. This route forwards ECS traffic to the NAT gateway and then the ECSs can access the Internet using the EIP.



Figure 3-5 Enabling ECSs in a VPC to access the Internet using a NAT gateway

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 each VPC route table, there are local routes and custom routes.

 Generally, the destination of a custom route cannot overlap with that of a local route. The destination of a local route can be a subnet CIDR block or CIDR blocks that are used for internal communications.

For example, if VPC-A has a subnet that supports IPv4/IPv6 dual stack. Its IPv4 CIDR block is 192.168.2.0/24 and IPv6 CIDR block is 2407:c080:802:be7::/64. The system automatically adds the local routes in **Table 3-4** to the route table of VPC-A. In this case, the destinations of the custom routes you set cannot overlap with those of the existing local routes.

Table 3-	4 VPC-A	local	routes
----------	---------	-------	--------

Local Route Destination	Description
192.168.2.0/24	IPv4 CIDR block of the subnet
2407:c080:802:be7::/64	IPv6 CIDR block of the subnet
100.64.0.0/10	Network used by public services on the cloud, such as DNS
198.19.128.0/20	Network used by internal services, such as VPC Endpoint
127.0.0.0/8	Loopback address

• You cannot add two routes with the same destination to a VPC route table even if their next hop types are different.

Custom Route Table Configuration Process

Figure 3-6 Process for configuring a custom route table



Table 3-5 Process for configuring a custom route table

N o.	Step	Description	Reference
1	Create a custom route table.	If the default route table cannot meet your service requirements, you can create a custom route table.	Creating a Custom Route Table
2	Add a custom 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 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 the default route table cannot meet your service requirements, you can create a custom route table and associate subnets with it to control traffic in and out of the subnets.

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.

Parameter	Description	Example Value
Name	Name (Mandatory) The name of the route table. The name:	
	Can contain 1 to 64 characters.	
	 Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	
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.	vpс-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 \oplus to add more routes.	

Table 3-6 Parameter descriptions

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.

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 \equiv 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.
- 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 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.

If a route table is associated with a subnet, adding rules to the route table may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

Notes and Constraints

• Generally, the destination of a custom route cannot overlap with that of a local route. The destination of a local route can be a subnet CIDR block or CIDR blocks that are used for internal communications.

For example, if VPC-A has a subnet that supports IPv4/IPv6 dual stack. Its IPv4 CIDR block is 192.168.2.0/24 and IPv6 CIDR block is 2407:c080:802:be7::/64. The system automatically adds the local routes in **Table 3-7** to the route table of VPC-A. In this case, the destinations of the custom routes you set cannot overlap with those of the existing local routes.

Local Route Destination	Description
192.168.2.0/24	IPv4 CIDR block of the subnet
2407:c080:802:be7::/64	IPv6 CIDR block of the subnet

Table 3-7 VPC-A local routes

Local Route Destination	Description
100.64.0.0/10	Network used by public services on the cloud, such as DNS
198.19.128.0/20	Network used by internal services, such as VPC Endpoint
127.0.0.0/8	Loopback address

• You cannot add two routes with the same destination to a VPC route table even if their next hop types are different.

Procedure

- 1. Log in to the management console.
- 2. Click 🔍 in the upper left corner and select the desired region and project.
- 3. Click \equiv 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 \oplus to add more routes.

Table 3-8 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

Parameter	Description	Example Value
Destination	Mandatory Enter the destination of the route. You can enter a single IP address or an IP address range in CIDR notation.	IPv4: 192.168.0.0/16
	The destination of the default route is 0.0.0.0/0, indicating that any traffic is matched.	
	NOTE If an IP address group contains an IP address range in the format of <i>Start IP address-End IP</i> <i>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.	
Next Hop Type	Mandatory Set the type of the next hop. NOTE When you add or modify a custom route in a	VPC peering connection
	default route table, the next hop type of the route cannot be set to VPN gateway , Direct Connect gateway , or Cloud connection .	
Next Hop	Mandatory	peer-AB
	Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.	
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.

If a route table is associated with a subnet, modifying rules in the route table may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

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 \bigcirc in the upper left corner and select the desired region and project.
- 3. Click \equiv 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 3-9 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.	IPv4: 192.168.0.0/16
	The destination of the default route is 0.0.0.0/0, indicating that any traffic is matched.	
	NOTE If an IP address group contains an IP address range in the format of <i>Start IP address-End IP</i> <i>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.	

Parameter	Description	Example Value
Next Hop	Mandatory	VPC peering
Туре	Set the type of the next hop.	connection
	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.	
Next Hop	Mandatory	peer-AB
	Set the next hop. The resources in the drop-down list box are displayed based on the selected next hop type.	
Description	Optional	-
	Enter the description of the route in the text box as required.	

8. Click **OK**.

3.3.3 Replicating a Route

Scenarios

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

Notes and Constraints

Table 3-10 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.

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

Table 3-10 Route replication

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 \equiv 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.

If a route table is associated with a subnet, deleting rules from the route table may affect how and where traffic is directed. Be careful with this operation as it may interrupt services.

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.

Differences Between SNAT ECSs and NAT Gateways

NAT Gateway provides network address translation (NAT) for servers, such as ECSs and BMSs, in a VPC or servers in local data centers that connect to a VPC through Direct Connect or VPN, allowing these servers to access the Internet using EIPs or to provide services for the Internet.

NAT Gateway 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 \equiv . In the service list, choose **Compute** > **Elastic Cloud Server**.
- 4. On the displayed page, locate the target ECS in the ECS list and click its name to go to the page showing ECS details.
- 5. On the displayed page, click the **Network Interfaces** tab.
- 6. Click the network interface IP address to view details and disable **Source**/ **Destination Check**.

By default, the source/destination check option is enabled to check whether the source IP address contained in the packets sent by the ECS is correct. If the IP address is incorrect, the system does not allow the ECS 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** or **Unbinding an EIP**.
 - 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, remotely 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) 56(84) bytes of data. 64 bytes from xxx.xxx.xxx: icmp_seq=1 ttl=51 time=9.34 ms 64 bytes from xxx.xxx.xxx: icmp_seq=2 ttl=51 time=9.11 ms 64 bytes from 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: Example command:

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

Figure 3-7 Configuring SNAT

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

Check whether the configuration is successful. If information similar to Figure 3-8 (for example, 192.168.1.0/24) is displayed, the configuration was successful.

iptables -t nat --list

Figure 3-8 Verifying configuration



16. Add a route. For details, see section Adding Routes to a Route Table.

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

Generally, cloud servers 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 to the on-premises servers using VPN connections or 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 cloud server goes down, the standby server becomes the active server 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's network interface 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

4 Virtual IP Address

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.

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 cloud servers and network interfaces) that a virtual IP address can be bound to	10	No

 Table 4-1
 Virtual IP address quotas

Notes and Constraints

- If a cloud server has multiple network interfaces that are in the same subnet, you are not advised to bind virtual IP addresses to the network interfaces. There may be route conflicts on the ECS, which would cause communication failure using the virtual IP address.
- 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.

- Switch to the IP Addresses tab and click Assign Virtual IP Address.
 The Assign Virtual IP Address page is displayed.
- 7. Set the parameters as required based on the below table.

Table 4-2 Virtual IP address parameters

Parameter	Description	Example Value
Subnet	Subnet from which a virtual IP address will be assigned. The current subnet is selected by default.	Subnet-A01

Parameter	Description	Example Value
IP Address Version	 IP address version. This parameter is only shown when IPv6 is enabled for the subnet. There are two options: IPv4 IPv6 	IPv4
Assignment Mode	 How virtual IP addresses are assigned. There are two options: Automatic: The system assigns a virtual IP address from the 	Manual
	subnet.Manual: You can specify a virtual	
	IP address.	
IP Address	Virtual IP addresses. This parameter is required if Assignment Mode is set to Manual .	192.168.0.15
	Specify an available IP address from the subnet CIDR block.	

8. After setting the parameters, click **OK**.

You can then check the assigned virtual IP address in the virtual 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.

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 \bigcirc 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.
- 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 the bound EIP.

- 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 instance and click **OK**.

In the virtual IP address list, you can view the bound instance.

- After you bind one or more virtual IP addresses to an ECS on the console, you need to manually configure the virtual IP addresses on the 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. The configurations for ECSs will not be lost after restart.

- 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 network interface that the virtual IP address is to be bound and the connection of the network interface:

nmcli connection

Information similar to the following is displayed:



The command output in this example is described as follows:

- eth0 in the DEVICE column indicates the network interface that the virtual IP address is to be bound.
- Wired connection 1 in the NAME column indicates the connection of the network interface.
- 2. Add the virtual IP address for the connection:

nmcli connection modify "<connection-name-of-the-network-interface>"
+ipv4.addresses <virtual-IP-address>

Configure the parameters as follows:

- connection-name-of-the-network-interface. The connection name of the network interface 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-network-interface>"

In this example, run the following command:

nmcli connection up "Wired connection 1"

Information similar to the following is displayed:

J#nmcli connection up "Wired connection 1" onnection 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, virtual IP address 172.16.0.125 is bound to network interface eth0.

[1]	22.16.0.247_subnet0-ecs-pod6-gaea-dpdk-ipv6 ~]#ip_a
L :	lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN group default glen 1000</loopback,up,lower_up>
	link/loopback 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</broadcast,multicast,up,lower_up>
	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:

 Delete the virtual IP address from the connection of the network interface: nmcli connection modify "<connection-name-of-the-network-interface>" -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 network interface that the virtual IP address is to be bound:

ifconfig

Information similar to the following is displayed. In this example, the network interface bound to the virtual IP address is **eth0**.

root@ecs-X-ubantu:~# ifconfig

eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 172.16.0.210 netmask 255.255.255.0 broadcast 172.16.0.255 inet6 fe80::f816:3eff:fe01:f1c3 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:01:f1:c3 txqueuelen 1000 (Ethernet) RX packets 43915 bytes 63606486 (63.6 MB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 3364 bytes 455617 (455.6 KB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

2. Switch to the /etc/netplan directory:

cd /etc/netplan

- 3. Add a virtual IP address to the network interface.
 - a. Open the configuration file **01-netcfg.yaml**:

vim 01-netcfg.yaml

- b. Press i to enter the editing mode.
- c. In the network interface configuration area, add a virtual IP address.

In this example, add a virtual IP address for **eth0**:

addresses:

- 172.16.0.26/32

```
The file content is as follows:

network:

version: 2

renderer: NetworkManager

ethernets:

eth0:

dhcp4: true

addresses:

- 172.16.0.26/32

eth1:

dhcp4: true
```

```
eth2:
dhcp4: true
eth3:
dhcp4: true
eth4:
dhcp4: true
```

- d. Press **Esc**, enter :**wq**!, save the configuration, and exit.
- 4. Make the configuration in **3** take effect:

netplan apply

5. Check whether the virtual IP address has been bound:

ip a

Information similar to the following is displayed. In the command output, virtual IP address 172.16.0.26 is bound to network interface eth0. root@ecs-X-ubantu:/etc/netplan# ip a

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

```
link/ether fa:16:3e:01:f1:c3 brd ff:ff:ff:ff:ff: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 network interface by referring to **3**.
- 2. Make the deletion take effect by referring to 4.

Windows: Windows Server

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 II	P address
------------	-------------	------------	-----------

Internet Protocol Version 4 ((TCP/IPv4) Properties
-------------------------------	-----------------------

 \times

You can get IP settings assigned aut this capability. Otherwise, you need for the appropriate IP settings.	omatically if your network support to ask your network administrator
◯ Obtain an IP address automatic	ally
• Use the following IP address: -	
IP address:	10 . 0 . 0 . 101
Subnet mask:	255.255.255.0
Default gateway:	10 . 0 . 0 . 1
Obtain DNS server address aut	omatically
Ose the following DNS server a	ddresses:
Preferred DNS server:	100 . 125 . 1 . 250
Alternate DNS server:	114 . 114 . 114 . 114
Validate settings upon exit	Advanced

- 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.

Advanced TCP/IP Settings ×	
IP Settings DNS WINS	
IP addresses	
IP address	Subnet mask
10.0.0.101	255.255.255.0
Add	Edit Remove
Default gateways:	
Gateway	Metric
10.0.0.1	Automatic
TCP/IP Address	×
IP address: 10	. 0 . 0 . 154
Subnet mask: 255	. 255 . 255 . 0
In	Add Cancel
	OK Cancel

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's network interface has been correctly configured.

Helpful Links

• Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS Network Interface?

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

4.4 Unbinding a Virtual IP Address from an Instance or EIP

Scenarios

You can unbind a virtual IP address from a cloud server 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.

- 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.

- 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 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.
- 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.

Constraints

If you want to release a virtual IP address that is being used by a resource, refer to **Table 4-3**.

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.
	RDS instance
	CCE instance
	API gateway

Table 4-3 Releasing a virtual IP address that is being used by a resource
Figure 4-6 Scenario 1: Virtual IP address cannot be released

< subnet A-01			С
Summary IP Addresses Tags			
Assign Virtual IP Address Unbind EIP Lee	im more about virtual IP address configuration.		Virtual IP Address v Enter a knyword. Q
Virtual IP Address	Bound EIP	Bound Instance	Operation
172.16.0.2	(3) 120	ecs-(172.16.0.118) View All (2)	Unbind from EIP Bind to Instance More +
			This operation cannot be performed because the IP address is bound to an instance or an EIP Untion the IP address and try again.

Figure 4-7 Scenario 2: Virtual IP address cannot be released

Summary IP Addresses Tags			C
Assign Virtual IP Address Unbind EIP	Learn more about virtual IP address configuration.		Virtual IP Address • Enter a keyword Q C
Virtual IP Address	Bound EIP	Bound Instance	Operation
172.16.1.76	-	Relational Database Service (172.16.1.149)	Bind to EIP Bind to Instance More -
		- I	This operation cannot be performed because the virtual # address is being used by a system component.

Procedure

- 1. Log in to the management console.
- 2. Click 🔍 in the upper left corner and select the desired region and project.
- 3. Click \equiv 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.
- 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 bound to the standby ECS to provide services. This section describes how to use a virtual IP address and Keepalived to set up a high-availability web cluster.

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.

D NOTE

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

Table 4-4 Resource planning

Resource Type	Quan tity	Description
VPC and subnet	1	• 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	In this example, two ECSs are required for active/standby switchover. Configure the two ECSs as follows:
		• 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
		 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	Quan tity	Description
Virtual IP	1	Assign a virtual IP address from Subnet-A01 .
address		• 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	• 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

- :	4 0	D	c			- 1-	:		. .	-1
Figure 4	4-9	Process 1	ror	setting	up	a n	ign-a	ivailabili	ty web	cluster

Create cloud resources.	Configure Keepalived on the active and standby ECSs	Bind the virtual IP address to the active and standby ECSs and EIP.	Disable IP forwarding on the standby ECS.	Verify the automatic switchover between the active and standby ECSs.
Create a VPC.		Bind the virtual IP address to the active and standby ECSs.		
Create ECSs.		Disable the source/ destination check option for the NICs of the active and standby ECSs.		
Assign a virtual IP address.		Bind the virtual IP address to an EIP.		
Assign an EIP.				

Step 1: Create Cloud Resources

- Create a VPC and subnet.
 For details, see Creating a VPC and Subnet.
- 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 4-5** to **Sg-A** and associate **Sg-A** with **ECS-HA1** and **ECS-HA2**.

Direc tion	Act ion	Туре	Protoc ol & Port	Source/ Destination	Description
Inbou nd	Allo w	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows remote logins to Linux ECSs over SSH port 22.
Inbou nd	Allo w	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows remote logins to Windows ECSs over RDP port 3389.
Inbou nd	Allo w	IPv4	TCP: 80	Source: 0.0.0.0/0	Allows external access to the website deployed on the ECSs over HTTP port 80.
Inbou nd	Allo w	IPv4	All	Source: current security group (Sg-A)	Allows the ECSs in Sg-A to communicate with each other using IPv4 addresses.
Inbou nd	Allo w	IPv6	All	Source: current security group (Sg-A)	Allows the ECSs in sg-A to communicate with each other using IPv6 addresses.
Outb ound	Allo w	IPv4	All	Destination: 0.0.0.0/0	Allows ECSs in Sg-A to access the external networks using IPv4 addresses.
Outb ound	Allo w	IPv6	All	Destination: : :/0	Allows ECSs in Sg-A to access the external networks using IPv6 addresses.

Table 4-5 Sg-A rules

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 4-6** to allow the ECSs in the two security groups to communicate with each other.

Se cur ity Gr ou p	Dire ction	A ct io n	Ту pe	Prot ocol & Port	Source/ Destinati on	Description
Sg- A	Inbo und	Al lo w	IP v4	All	Source: Sg-B	Allows ECSs in Sg-B to access those in Sg-A over any IPv4 protocol and port.
Sg- B	Inbo und	Al lo w	IP v4	All	Source: Sg-A	Allows ECSs in Sg-A to access those in Sg-B over any IPv4 protocol and port.

Table 4-6 Rules of security groups Sg-A and Sg-B

- **EIP**: Select **Not required**.
- 3. Assign a virtual IP address from **Subnet-A01**. For details, see **Assigning a Virtual IP Address**.
- 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**.
 - 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 lm_sensors-libs.x86_64 0:3.4.0-8.20160601gitf9185e5.el7_9.1 0:2.6.1-1.el7 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-filesystem.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-HA1
- 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



Welcome to ECS-HA1

- 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 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 }

```
vrrp_script chk_http_port {
       script "/etc/keepalived/chk_nginx.sh"
       interval 2
       weight -5
       fall 2
       rise 1
     }
vrrp_instance VI_1 {
   state MASTER
   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
```

- 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: gperftools-libs.x86_64 centos-indexhtml.noarch 0:7-9.el7.centos lm_sensors-libs.x86_64 0:3.4.0-8.20160601gitf9185e5.el7_9.1 0:2.6.1-1.el7 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-filesystem.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:
                     }
       }
3
```

- 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 90
   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
   }
3
```

- 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

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

- 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 \checkmark to expand the details area and check whether **Source/Destination Check** is disabled.

nmary	Disks N	etwork Interfaces	Security Groups	EIPs	Monitoring	Tags	Cloud Backup and Recove
	-						
40		0 E	d # 500 kr				0 1 1/0
After you add a	In extension NI	C, configure policy-base	d routing on the ECS to e	nable network	communication bet	ween the EC	S and NIC.
After you attac	h or detach a ne	etwork interface or chan	ge a VPC, enable NIC mu	Iti-queue to in	nprove network perfe	ormance.	
Attach Net	work Interface	You can attach 1 m	ore network interfaces.				
∧ 192.10	38.0.195						
Name							
NIC ID		Obca4a90-	d45907ea3	5a			
Status		Activated					
EID.		Mullivaleu					
EIP		-					
Security Gr	oup	Sg-A	_				
Source/Des	tination Check	0					
IPv4 Subne	t ID	2b0b04ca	9d13b4d651	1			

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



Welcome to ECS-HA1

- 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 4-7**. In this example, the ECS runs the Linux OS.

OS	Operations
Linux	 Run the following command to switch to user root: su root
	 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 .
	 If 0 is displayed, no further action is required.
	 If 1 is displayed, go to the next step.
	 Use either of the following methods to modify the configuration file: Method 1
	 a. Run the following command to open the /etc/ sysctl.conf file: vim /etc/sysctl.conf
	b. Press i to enter the editing mode.
	c. Set net.ipv4.ip_forward to 0 .
	 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.conf
	 Run the following command to apply the modification: sysctl -p /etc/sysctl.conf
Windows	 In the search box, enter cmd to open the command prompt window, and run the following command: ipconfig/all
	 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.
	• If the value is set to 0 , IP forwarding will be disabled.
	• If the value is set to 1 , IP forwarding will be enabled.

Table 4-7 Disabling IP forwarding

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

- 1. Restart the active and standby ECSs.
 - a. Remotely log in to ECS-HA1.For details, see Logging In to an ECS.
 - b. Run the following command to restart **ECS-HA1**: **reboot**
 - c. Repeat **1.a** to **1.b** to restart **ECS-HA2**.
- 2. 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

b. Remotely log in to **ECS-HA1** and run the following command to check whether the virtual IP address is bound to the network interface (eth0) 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 network interface (eth0) 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:fefe:5619/64 scope link
- valid_lft forever preferred_lft forever
- c. Run the following command to disable Keepalived on ECS-HA1:

systemctl stop keepalived.service

- 3. Check whether **ECS-HA2** becomes the active ECS.
 - a. Remotely log in to **ECS-HA2** and run the following command to check whether the virtual IP address is bound to the network interface (eth0) 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 network interface (eth0) 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 glen 1000

link/ether fa:16:3e:fe:56:3f brd 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



Welcome to ECS-HA2

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 cloud servers (such as 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 can be created on the Network Interfaces tab, and can be attached to or detached from an instance.

Application Scenarios

Elastic network interfaces help you flexibly migrate and separate services.

- Flexible migration: You can detach an **extended** network interface from a cloud server and attach it to another one. The private IP address, EIP, and security group rules of the original cloud server can be migrated together, so you do not need to reconfigure them. This allows the service traffic on the faulty cloud server to be quickly switched to the standby one, achieving quick service recovery.
- Service separation: You can configure multiple network interfaces for a cloud server. These network interfaces can be in different subnets of the same VPC and process the internal, external, and management traffic of the cloud server respectively. You can configure access control policies and routes for each subnet, and define security group rules for each network interface to isolate networks and service traffic.

In **Figure 5-1**, the cloud server has one primary network interface and four extended network interfaces. These network interfaces can be in different subnets. In this example, extended network interface 01 and the primary network interface are in Subnet-A01, and extended network interface 04 is in Subnet-A03.



Figure 5-1 Cloud server network interfaces

Each cloud server can have a limited number of elastic network interfaces attached. If the cloud server specifications support supplementary network interfaces, you can attach supplementary network interfaces to the elastic network interfaces. For details, see Application Scenarios.

Constraints

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

5.1.2 Creating a Network Interface

Scenarios

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

Notes and Constraints

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

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	description	าร
-------	-----	-----------	-------------	----

Paramet er	Parameter Description	Example Value		
Region	Region where the network interface is created. Select the region nearest to you to ensure the lowest latency possible.	EU-Dublin		
Name	Name of the network interface. The name:	networkInterface-891e		
	 Can contain 1 to 64 characters. Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 			
VPC	VPC where the network interface is created.	vpс-001		
Subnet	Subnet where the network interface is created.	subnet-001		
Private IP Address	Whether to automatically assign a private IP address.	-		
Security Group	Security group that will be associated with the network interface.	sg-001		

4. Click **OK**.

5.1.3 Managing Network Interfaces

You can perform the following operations on network interfaces:

- Attaching a Network Interface to a Cloud Server
- Binding an EIP to a Network Interface
- Binding a Virtual IP Address to a Network Interface
- Enabling or Disabling the Instance-dependent Deletion Function for a Network Interface
- Detaching a Network Interface from an Instance or Unbinding an EIP from a Network Interface
- Changing Security Groups That Are Associated with a Network Interface
- Viewing the Basic Information About a Network Interface

Attaching a Network Interface to a Cloud Server

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

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

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

- 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**.

Binding an EIP to a Network Interface

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

Each network interface has a private IP address. After an EIP is bound to a network interface, 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 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 more flexible access services.

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

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

- 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.

Binding a Virtual IP Address to a Network Interface

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

A virtual IP address can only be bound to a network interface attached to an instance.

For more information about virtual IP addresses, see Virtual IP Address Overview.

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

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

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 is displayed.

- 5. Locate the row that contains the target virtual IP address and click **Bind to Instance** in the **Operation** column.
- 6. Select the instance that the virtual IP address to be bound and click **OK**.

Enabling or Disabling the Instance-dependent Deletion Function for a Network Interface

This function is available only in certain regions, which is subject to that displayed on the console.

- **Instance-dependent Deletion** is disabled by default. With this function disabled, the network interface will not be deleted if it is detached from an instance or if the instance is deleted. You can attach the network interface to another instance.
- If **Instance-dependent Deletion** is enabled, the network interface will be deleted after it is detached from an instance.
- 1. Log in to the management console.
- 2. Click = in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

- 4. Click the private IP address of the target network interface.
- 5. On the **Summary** tab page, enable or disable **Instance-dependent Deletion**.

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

- If **Instance-dependent Deletion** is enabled for a network interface, detaching the network interface from its instance will also delete the network interface.
 - Deleting a network interface will also delete 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, detaching the network interface from its instance will not delete the network interface.

If the network interface has an EIP bound, detaching the network interface from its instance will also unbind the EIP from the network interface.

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

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

- 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 choose to release the EIP when unbinding it.

Changing Security Groups That Are Associated with a Network Interface

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.

- 1. Log in to the management console.

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

- 4. You can change the security groups associated with a network interface:
 - On the network interface list page:

- i. In the network interface list, locate the row that contains the target network interface, and choose **More** > **Change Security Group** in the **Operation** column.
- ii. On the **Change Security Group** page, select the security groups to be associated and click **OK**.
- On the network interface details page:
 - i. Click the private IP address of the target network interface.
 - ii. Choose the Associated Security Groups tab and click Change Security Group.
 - iii. On the **Change Security Group** page, select the security groups to be associated and click **OK**.

Viewing the Basic Information About a Network Interface

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.

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

The VPC list is displayed.

3. In the navigation pane on the left, choose **Virtual Private Cloud** > **Network Interfaces**.

The network interface list is displayed.

4. Click the private IP address of the target network interface.

5.1.4 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 network 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 \equiv 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 cloud server cannot meet your requirements, you can use supplementary network interfaces, which can be attached to VLAN subinterfaces of elastic network interfaces.

Application Scenarios

Each cloud server can have a limited number of elastic network interfaces attached. If the cloud server specifications support supplementary network interfaces, you can attach supplementary network interfaces to the elastic network interfaces.

You can attach supplementary network interfaces to a cloud server. Each supplementary network interface belongs to a different subnet in the same VPC. They have dedicated private IP addresses and EIPs and process the internal, external, and management traffic of the cloud server. You can configure access control policies and routes for each subnet, and define security group rules for each supplementary network interface to isolate networks and service traffic.

Supplementary network interfaces are attached to VLAN sub-interfaces of network interfaces. **Figure 5-2** shows the attachment relationship. **Both the primary and extended network interfaces of an ECS can have supplementary network interfaces attached.** In this example, the cloud server has one primary network interface and four extended network interfaces. If the number of elastic network interfaces that can be attached to a cloud server reaches the upper limit, you can attach supplementary network interface 01 to the primary network interface 01, and supplementary network interfaces 03 and 04 to extended network interface 04. So there are four supplementary network interfaces. **Elastic network**

interfaces and supplementary network interfaces can be in different subnets. In this example, supplementary network interface 04 and extended network interface 04 are in Subnet-A03, but supplementary network interface 03 is in Subnet-A01.



Figure 5-2 Attached supplementary network interfaces

Constraints

- A maximum of 256 supplementary network interfaces can be attached to a cloud server of certain specifications. The number of supplementary network interfaces that can be attached to a cloud server varies by server specification.
- A cloud server 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

If the number of network interfaces attached to an instance exceeds the upper limit, you can attach supplementary network interfaces to the network interfaces, including the primary and extended network interfaces, of the instance. This helps you set up flexible and highly available networks.

Notes and Constraints

- Supplementary network interfaces must be in the same VPC as the network interface they are attached to, but they can be in different subnets and security groups.
- After supplementary network interfaces are created, you need to create VLAN subinterfaces on the network interface of the instance and configure

corresponding rules by referring to **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

Paramet er	Description	Example Value	
Region	Region where the supplementary network interface will be created. Select the region nearest to you to ensure the lowest latency possible.	EU-Dublin	
Network Interface	Network interface that you want the supplementary network interface to attach to. Select an elastic network interface from the drop-down list.	(172.16.0.145)	
VPC	VPC where the supplementary network interface will be created. The VPC of the network interface that the supplementary network interface is attached to is selected by default.	vрс-А	
Subnet	Subnet where the supplementary network interface will be created. The supplementary network interface and its network interface can be in different subnets.	subnet-A01	
Quantity	Number of supplementary network interfaces to be created.	1	

Paramet er	Description	Example Value
Private IP Address	Whether to assign a private IPv4 address or IPv6 address to the supplementary network interface. There are two options:	IPv4
	• Private IPv4 network : a private IPv4 address will be assigned. This option is selected by default and cannot be deselected.	
	 IPv6 network (Public and private network traffic): a private IPv6 address will be assigned. Both private and public IPv6 networks are supported. IPv6 is shown only when IPv6 is enabled for the subnet of the supplementary network interface. 	
IPv4 Address	How a private IPv4 address will be assigned to the supplementary network interface. There are two options:	Automatically assign IP address
	• Automatically assign IP address: The system assigns an IP address from the subnet you have selected.	
	 Manually specify IP address: You can specify an IP address. If you select Manually specify IP address, enter a private IPv4 address. 	
IPv6 Address	How an IPv6 address will be assigned to the supplementary network interface if IPv6 network (Public and private network traffic) is selected for Private IP Address.	Automatically assign IP address
	There are two options:	
	• Automatically assign IP address: The system assigns an IP address from the subnet.	
	• Manually specify IP address: You can specify an IP address. If you select Manually specify IP address, enter an IPv6 address.	
Security Group	Security group that the supplementary network interface will be associated with.	sg-001

Paramet er	Description	Example Value
Descripti on	Description of the supplementary network interface.	-
(Optiona l)	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

4. Click **Create Now**.

To use a supplementary network interface, you need to create a VLAN subinterface. 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 subinterface for the network interface of the instance and configure a private IP address and default routes for the supplementary network interface.

Before doing so, you need to obtain:

- The information described in **Table 5-3** when you configure a supplementary network interface for a Linux ECS.
- The information described in **Table 5-3** and **Table 5-4** when you configure a supplementary network interface for a Windows ECS.

Table 5-3 Information about the supplementary network interface and subn	net
--	-----

Item	How to Obtain			
VLAN ID	 In the supplementary network interface list, click the private IP address of the target supplementary network interface. 			
MAC address				
Private IP address	The Summary page is displayed.			
	2. On the displayed page, check and record the following information:			
	VLAN ID			
	MAC address			
	Private IP address			

Item	How to Obtain			
Subnet mask	1. In the supplementary network interface list,			
Gateway address	locate the target supplementary network interface and click the subnet name in the Network column. The Summary page of the subnet is displayed.			
	On the displayed page, check and record the following information:			
	 Subnet mask: subnet mask of the IPv4 CIDR block. For example, if the IPv4 CIDR block is 192.168.0.0/24, the mask is 24. 			
	 Subnet gateway: In the Gateway and DNS Information area, check the gateway address. 			

Table 5-4 Information about the network interface and subnet to which the supplementary network interface belongs

ltem	How to Obtain			
MAC address	1. In the ECS list, click the name of the ECS with			
Private IP address	The Summary page is displayed.			
	 2. Switch to the Network Interface tab and click to check and record the following information: MAC address Drivets ID address 			
Subnet mask	1. In the network interface list, locate the target			
Gateway address	the Network column. The Summary page of the subnet is displayed.			
	On the displayed page, check and record the following information:			
	 Subnet mask: subnet mask of the IPv4 CIDR block. For example, if the IPv4 CIDR block is 192.168.0.0/24, the mask is 24. 			
	 Subnet gateway: In the Gateway and DNS Information area, check the gateway address. 			

Configuring a Supplementary Network Interface for a Linux ECS

The following describes how to create a VLAN subinterface for a supplementary network interface of an ECS network interface. Huawei Cloud EulerOS 2.0

Standard 64 bits is used as an example. In this example, the information about the supplementary network interface and subnet is as follows:

- VLAN ID: 1937
- MAC address: fa:16:3e:6d:c5:5a
- Private IP address: 192.168.0.149
- Subnet mask: 24
- Subnet gateway address: 192.168.0.1

NOTE

This example describes how to configure the supplementary network interface for the primary network interface of an ECS. If you want to do the same thing for the extended network interface of the ECS, follow the similar steps.

- 1. Log in to the ECS.
- 2. Run the following command to check and record the network interface name of the ECS:

ifconfig

Information similar to the following is displayed. In this example, the network interface name is **eth0**.

```
[root@ecs-subeni-linux ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 192.168.0.125 netmask 255.255.0 broadcast 192.168.0.255
inet6 fe80::f816:3eff:fe6d:c542 prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:6d:c5:42 txqueuelen 1000 (Ethernet)
RX packets 78131 bytes 111604802 (106.4 MiB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 8686 bytes 1422159 (1.3 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Run the following command to create a VLAN subinterface on the network interface:

ip link add link *<network-interface-name>* **name** *<VLAN-subinterface-name>* **type vlan id** *<VLAN-ID-of-the-supplementary-network-interface>*

Variables in the preceding command are as follows:

- *network-interface-name*: the network interface name queried in **2**. In this example, the name is **eth0**.
- VLAN-subinterface-name. Name the subinterface in the format of <network-interface-name>.<VLAN-ID-of-the-supplementary-networkinterface>. In this example, the VLAN subinterface name is eth0.1937.
- VLAN-ID-of-the-supplementary-network-interface: In this example, the ID is 1937.

Example command:

ip link add link eth0 name eth0.1937 type vlan id 1937

4. Run the following command to create a namespace:

ip netns add <namespace-name>

namespace-name: Name it in the format of **ns***<supplementary-network-interface-VLAN-ID>*. In this example, the name is **ns1937**.

Example command:

ip netns add ns1937

 Run the following command to add the VLAN subinterface to the namespace: ip link set </LAN-subinterface-name> netns </namespace-name> Example command:

ip link set eth0.1937 netns ns1937

6. Run the following command to change the MAC address of the VLAN subinterface to that of the supplementary network interface:

ip netns exec *<namespace-name>* **ifconfig** *<VLAN-subinterface-name>* **hw ether** *<MAC-address-of-the-supplementary-network-interface>* Example command:

ip netns exec ns1937 ifconfig eth0.1937 hw ether fa:16:3e:6d:c5:5a

7. Run the following command to enable the VLAN subinterface:

ip netns exec *<namespace-name>* **ifconfig** *<VLAN-subinterface-name>* **up** Example command:

ip netns exec ns1937 ifconfig eth0.1937 up

8. Run the following command to configure a private IP address for the VLAN subinterface:

ip netns exec <namespace-name> ip addr add <private-IP-address> dev <VLAN-subinterface-name>

private-IP-address: private IP address of the supplementary network interface/ subnet mask. In this example, the value is **192.168.0.149/24**.

Example command:

ip netns exec ns1937 ip addr add 192.168.0.149/24 dev eth0.1937

9. Run the following command to configure the default route for the VLAN subinterface:

ip netns exec *<namespace-name>* **ip route add default via** *<gateway-address-of-the-subnet-where-the-supplementary-network-interface-is-created>*

Example command:

ip netns exec ns1937 ip route add default via 192.168.0.1

- 10. Check whether the supplementary network interface has worked.
 - a. Run the following command to verify the connectivity between network interface and the test ECS:

ping <private-IP-address-of-the-test-ECS>

Plan the same VPC and security group for the test ECS and the ECS with network interface attached, so that the two ECSs can communicate with each other by default.

Example command:

ping 192.168.0.133

If information similar to the following is displayed, the two ECSs can communicate with each other. If the communication is normal, proceed with **10.b**.

[root@ecs-subeni-linux ~]# ping 192.168.0.133 PING 192.168.0.133 (192.168.0.133) 56(84) bytes of data. 64 bytes from 192.168.0.133: icmp_seq=1 ttl=64 time=0.302 ms 64 bytes from 192.168.0.133: icmp_seq=2 ttl=64 time=0.262 ms --- 192.168.0.133 ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 999ms rtt min/avg/max/mdev = 0.262/0.282/0.302/0.020 ms

b. Run the following command to verify the connectivity between the supplementary network interface and the test ECS:

ip netns exec <namespace-name> ping <private-IP-address-of-the-test-ECS>

Plan the same VPC and security group for the test ECS and the ECS with the supplementary network interface attached. This allows the two ECSs to communicate with each other by default.

Example command:

ip netns exec ns1937 ping 192.168.0.133

If information similar to the following is displayed, the two ECSs can communicate with each other. This means the supplementary network interface has worked.

[root@ecs-subeni-linux ~]# ip netns exec ns1937 ping 192.168.0.133 PING 192.168.0.133 (192.168.0.133) 56(84) bytes of data. 64 bytes from 192.168.0.133: icmp_seq=1 ttl=64 time=0.420 ms 64 bytes from 192.168.0.133: icmp_seq=2 ttl=64 time=0.233 ms

--- 192.168.0.133 ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 999ms rtt min/avg/max/mdev = 0.233/0.326/0.420/0.095 ms

▲ CAUTION

- The route configured above is a temporary route that is applied once configured, and will be lost after ECS restarts. To avoid network disruptions, take step **11** to configure permanent routes instead.
- If the ECS needs to access a public domain name through the supplementary network interface, you need to take step **11** to configure DNS for that supplementary network interface and then restart the ECS.
- 11. Configure a permanent route and DNS for the supplementary network interface. The configuration will work after the ECS is restarted.
 - a. Configure a permanent route for the supplementary network interface.
 - i. Run the following command to open the **/etc/rc.local** file: **vi /etc/rc.local**
 - ii. Press i to enter the editing mode.
 - iii. Add the following content to the end of the file.

The parameters and values must be the same as those in steps **3** to **9**.

ip link add link eth0 name eth0.1937 type vlan id 1937 ip netns add ns1937 ip link set eth0.1937 netns ns1937 ip netns exec ns1937 ifconfig eth0.1937 hw ether fa:16:3e:6d:c5:5a ip netns exec ns1937 ifconfig eth0.1937 up ip netns exec ns1937 ip addr add 192.168.0.149/24 dev eth0.1937 ip netns exec ns1937 ip route add default via 192.168.0.1

iv. Press **Esc** to exit and enter :wq! to save the configuration.

v. Run the following command to assign execute permissions to the **/etc/rc.local** file:

chmod +x /etc/rc.local

If your operating system is Red Hat or EulerOS, run the following command after you perform **11.a.v**:

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

b. (Optional) Configure DNS for the supplementary network interface if the ECS needs to access the public domain name through the supplementary network interface.

If DNS resolution is not required, take step **11.c** to restart the ECS.

i. Run the following command to go to the **/etc/sysconfig/network-scripts/** directory that stores the network interface configuration file:

cd /etc/sysconfig/network-scripts/

ii. Run the following command to modify the network interface configuration file:

vi ifcfg-<network-interface-name>

network-interface-name: the name queried in **2**. In this example, the name is **vi ifcfg-eth0**.

- iii. Press i to enter the editing mode.
- iv. Add the following content to the end of the file.

114.114.114.114 is the public recursive DNS address. DNS1=114.114.114.114

- v. Press **Esc** to exit and enter :wq! to save the configuration.
- c. Run the following command to restart the ECS:

reboot

- d. Check whether the permanent route has worked by referring to **10**.
- e. (Optional) If DNS is configured, check whether the DNS configuration has worked.
 - i. Bind an EIP to the supplementary network interface by referring to Binding or Unbinding an EIP to or from a Supplementary Network Interface.
 - ii. Run the following command to check whether the supplementary network interface can access the public domain name:

ip netns exec <namespace-name> ping <public-domain-name>

Example command:

ip netns exec ns1937 ping support.huaweicloud.com

If information similar to the following is displayed, the DNS configuration has worked.

[root@ecs-subeni-linux ~]# ip netns exec ns1937 ping support.huaweicloud.com PING support.huaweicloud.com (36.150.72.70) 56(84) bytes of data. 64 bytes from 36.150.72.70 (36.150.72.70): icmp_seq=1 ttl=54 time=2.68 ms 64 bytes from 36.150.72.70 (36.150.72.70): icmp_seq=2 ttl=54 time=2.61 ms 64 bytes from 36.150.72.70 (36.150.72.70): icmp_seq=3 ttl=54 time=2.60 ms ^C

⁻⁻⁻ support.huaweicloud.com ping statistics ---

3 packets transmitted, 3 received, 0% packet loss, time 3015ms rtt min/avg/max/mdev = 2.604/2.633/2.681/0.068 ms

- 12. (Optional) Remotely log in to the ECS using the private IP address of the supplementary network interface.
 - a. Add an inbound rule to allow traffic over SSH port 22 to the security group associated with the supplementary network interface.

For details, see Adding a Security Group Rule.

Direc tion	Prior ity	Action	Туре	Protocol & Port	Source
Inbo und	1	Allow	IPv4	TCP: 22	Set the IP address based on service requirements. For example, to remotely log in to the ECS from a local PC, set the source to the IP address of the local PC.

 Table 5-5 A security group rule that allows traffic over SSH port 22

b. Run the following command to check whether port 22 in the namespace is listened on:

ip netns exec <namespace-name> netstat -antp | grep 22

Example command:

ip netns exec ns1937 netstat -antp | grep 22

- If the command output is empty, port 22 in the namespace is not listened on. Go to 12.c.
- If information similar to the following is displayed, port 22 is listened on. No further action is required. [root@ecs-subeni-linux ~]# ip netns exec ns1937 netstat -antp | grep 22 tcp 0 0 0.0.0.0:22 0.0.0.0:* LISTEN 2797/sshd tcp6 0 0 :::22 :::* LISTEN 2979/sshd
- c. Run the following command to start the SSH service and enable listening port 22:

ip netns exec <namespace-name> /sbin/sshd

Example command:

ip netns exec ns1937 /sbin/sshd

d. Run the following command to check whether port 22 in the namespace is listened on:

ip netns exec <namespace-name> netstat -antp | grep 22

Example command:

ip netns exec ns1937 netstat -antp | grep 22

If information similar to the following is displayed, port 22 is listened on:[root@ecs-subeni-linux ~]# ip netns exec ns1937 netstat -antp | grep 22tcp000.0.0.0:220.0.0.0:*LISTEN2797/sshdtcp6000:::22:::*LISTEN2979/sshd

- 13. (Optional) Allow traffic over HTTP port 80 of the supplementary network interface if the ECS needs to provide the web service through the supplementary network interface.
 - a. Add an inbound rule to allow traffic over HTTP port 80 to the security group associated with the supplementary network interface.

For details, see Adding a Security Group Rule.

Table 5-6 A security	/ group rule	that allows	traffic	over H1	TTP port 80
----------------------	--------------	-------------	---------	---------	-------------

Direc tion	Prior ity	Action	Туре	Protocol & Port	Source
Inbo und	1	Allow	IPv4	TCP: 80	0.0.0/0 Allows any IP address to access the supplementary network interface over port 80.

b. Run the following command to check whether port 80 in the namespace is listened on:

ip netns exec <namespace-name> netstat -antp | grep 80
Example command:

ip netns exec ns1937 netstat -antp | grep 80

- If the command output is empty, port 80 in the namespace is not listened on. In this case, enable port 80 for the web service.
- If information similar to the following is displayed, port 80 is listened on. No further action is required. [root@ecs-subeni-linux ~]# ip netns exec ns1937 netstat -antp | grep 80

tcp6 0 0 :::80 :::* LISTEN ...

Configuring a Supplementary Network Interface for a Windows ECS

The following describes how to create a VLAN subinterface on the network interface of a Windows ECS. Windows Server 2019 Standard 64bit is used as an example. In this example, the information about the supplementary network interface, primary network interface, and subnet is as follows:

- Supplementary network interface
 - VLAN ID: 2242
 - MAC address: fa:16:3e:6d:c5:db
 - Private IP address: 192.168.0.22
 - Subnet mask: 24 (255.255.255.0)
 - Subnet gateway address: 192.168.0.1
- Network interface
 - MAC address: fa:16:3e:6d:c5:d5
 - Private IP address: 192.168.0.16
 - Subnet mask: 24 (255.255.255.0)
Subnet gateway address: 192.168.0.1

This example describes how to configure the supplementary network interface for the primary network interface of an ECS. If you want to do the same thing for the extended network interface of the ECS, follow the similar steps.

- 1. Log in to the ECS.
- 2. Enter **Windows PowerShell** in the search box in the lower left corner of the desktop and press **Enter**.
- 3. On the displayed window, run the following command to query the Ethernet adapter information of the network interface:

ipconfig

Information similar to the following is displayed. In this example, the Ethernet adapter name is **tap7888b905-ee**.

Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.
PS C:\Users\Administrator> ipconfig
Windows IP Configuration
Ethernet adapter tap7888b905-ee:
Connection-specific DNS Suffix . : openstacklocal
Link-local IPv6 Address : fe80::1e55:468d:da2a:e16%3
IPv4 Address
Subnet Mask
Default Gateway 192.168.0.1

- 4. Create a bond group.
 - a. Run the following command to create a bond group for the custom VLAN:

New-NetLbfoTeam -Name *<bond-group-name>* -TeamMembers "*<Ethernet-adapter-name-of-the-network-interface>*" -TeamingMode SwitchIndependent -LoadBalancingAlgorithm IPAddresses -Confirm:\$false

Variables in the preceding command are as follows:

- bond-group-name: the bond group name of the custom VLAN. In this example, the bond group name is Team1.
- *Ethernet-adapter-name-of-the-network-interface*: information queried in 3. In this example, the name is tap7888b905-ee.

Example command:

New-NetLbfoTeam -Name Team1 -TeamMembers "tap7888b905-ee" -TeamingMode SwitchIndependent -LoadBalancingAlgorithm IPAddresses -Confirm:\$false

Information similar to the following is displayed.

PS C:\Users\Administra dBalancingAlgorithm IP	ator> New-NetLbfoTeam -Name PAddresses -Confirm:\$false	e Team1 -TeamMembers		SwitchIndependent -L	
Name Members TeamingMode LoadBalancingAlgorithm Status	: Team1 : tap7888b905-ee : Team1 : SwitchIndependent n : IPAddresses : Up				

b. Run the following commands to query the bond group you have created: Get-NetLbfoTeamMember

Information similar to the following is displayed.

PS C:\Users\Administrato	or> Get-NetLbfoTeamMember
Name	: tap7888b905-ee
InterfaceDescription	: Red Hat VirtIO Ethernet Adapter
Team	: Team1
AdministrativeMode	: Active
OperationalStatus	: Active
TransmitLinkSpeed(Gbps)	: 100
ReceiveLinkSpeed(Gbps)	: 100
FailureReason	: NoFailure

Get-NetAdapter

Information similar to the following is displayed:

PS C:\Users\Administ	rator> Get-NetAdapter				
Name	InterfaceDescription	ifIndex	Status	MacAddress	LinkSpeed
Team1	Microsoft Network Adapter Multiplexo		Up	FA-16-3E-6D-C5-D5	100 Gbps
tap7888b905-ee	Red Hat VirtIO Ethernet Adapter		Up	FA-16-3E-6D-C5-D5	100 Gbps

- 5. Configure a custom VLAN network.
 - a. Run the following command to create a VLAN subinterface:

Add-NetLbfoTeamNIC -Team "<bond-group-name>" -VlanID <VLAN-ID-of-the-supplementary-network-interface> -Confirm:\$false

Example command:

Add-NetLbfoTeamNIC -Team "Team1" -VlanID 2242 -Confirm:\$false Information similar to the following is displayed:

PS C:\Users\Administrato	<pre>r> Add-NetLbfoTeamNIC -Team "Team1" -VlanID 2242 -Confirm:\$false</pre>
Name	: Team1 - VLAN 2242
InterfaceDescription	: Microsoft Network Adapter Multiplexor Driver #2
Team	: Team1
VlanID	: 2242
Primary	: False
Default	: False
TransmitLinkSpeed(Gbps)	: 100
ReceiveLinkSpeed(Gbps)	: 100

b. Run the following command to open the **Network Connections** page:

ncpa.cpl

On the displayed page, **Team1** is the bond group created in **4.a**, and **Team1 – VLAN 2242** is the VLAN subinterface created in **5.a**.



6. Configure the network for the network interface.

a. On the Network Connections page, double-click Team1.
 The Team1 Status page is displayed.

🛋 Team1 Status			×
General			
Connection			
IPv4 Connectivi	ity:		Internet
IPv6 Connectivi	ity:	No net	work access
Media State:			Enabled
Duration:			00:13:42
Speed:			100.0 Gbps
D <u>e</u> tails			
Activity			
	Sent —		Received
Bytes:	85,328	Ĭ	39,356
Properties	€Disable	Diagnose	
			<u>C</u> lose

b. On the **Team1 Status** page, click **Properties**. The **Team1 Properties** page is displayed.

Team1 Properties	Х
Networking Sharing	
Connect using:	
Microsoft Network Adapter Multiplexor Driver	
Configure This connection uses the following items:]
 Client for Microsoft Networks File and Printer Sharing for Microsoft Networks Microsoft MAC Bridge QoS Packet Scheduler Microsoft Load Balancing/Failover Provider Internet Protocol Version 4 (TCP/IPv4) Microsoft Network Adapter Multiplexor Protocol × 	
Install Uninstall Properties	
Description Allows your computer to access resources on a Microsoft network.	
OK Cance	

c. On the Team1 Properties page, click Configure....
 The Microsoft Network Adapter Multiplexor Driver Properties page is displayed.

×

General Advanced Driver Details Events	
Microsoft Network Adapter Multiplexor Driver	
Device type: Network adapters	
Manufacturer: Microsoft	
Location: {E477CEEE-C3B2-4DF6-B923- 8FCCF627463C}	
Device status	
This device is working properly.	
\sim	
OK Cancel	

Microsoft Network Adapter Multiplexor Driver Properties

d. On the **Microsoft Network Adapter Multiplexor Driver Properties** page, choose the **Advanced** tab, click **MAC Address**, enter the MAC address of the network interface, and click **OK**.

When entering the MAC address, remove the colons (:) and use the uppercase letters. For example, if the MAC address of the network interface is **fa:16:3e:6d:c5:d5**, enter **FA163E6DC5D5**.

 \times

Microsoft Network Ada	pter Multiplexor	Driver Properties
Triffer 03011 Treet of R Add	prer manuprexer	onver i ropercies

General Advanced Driver Details Events The following properties are available for this network adapter. Click the property you want to change on the left, and then select its value on the right. Property: Value: Encapsulated Task Offload FA163EDC5D5 ۸ Header Data Split IPsec Offload O Not Present IPv4 Checksum Offload Large Send Offload Version 2 (IPv4 Large Send Offload Version 2 (IPv! MAC Address Receive Side Scaling Recv Segment Coalescing (IPv4) Recv Segment Coalescing (IPv6) TCP Checksum Offload (IPv4) TCP Checksum Offload (IPv6) UDP Checksum Offload (IPv4) UDP Checksum Offload (IPv6) OK Cancel

e. Return to the **Team1 Properties** page, double-click **Internet Protocol Version 4 (TCP/IPv4)**.

The Internet Protocol Version 4 (TCP/IPv4) Properties page is displayed.

Microsoft Netwo	ork Adapter Multiplexor	Driver
This connection uses	the following items:	Configure
Client for Mic Client for Mic File and Print Microsoft MA QoS Packet	rosoft Networks er Sharing for Microsoft C Bridge Scheduler	Networks
Microsoft Loa	ad Balancing/Failover F ocol Version 4 (TCP/IP) twork Adapter Multiplex	Provider (4) or Protocol
Microsoft Loa	ad Balancing/Failover F ocol Version 4 (TCP/IP) twork Adapter Multiplex	Provider (4) or Protocol >
Microsoft Loa	ad Balancing/Failover F cool Version 4 (TCP/IP) twork Adapter Multiplex Uninstall	Provider (4) or Protocol
Microsoft Loa Internet Proto Microsoft Net Install Description Transmission Contro wide area network across diverse inter	ad Balancing/Failover F cool Version 4 (TCP/IP) twork Adapter Multiplex Uninstall of Protocol/Internet Prot protocol that provides of connected networks.	Provider (4) or Protocol Properties tocol. The default communication

Internet Protocol Version 4 (TCP/IPv4	4) Properties	~
General		
You can get IP settings assigned auto this capability. Otherwise, you need t for the appropriate IP settings.	omatically if your network suppo to ask your network administrati	orts or
Obtain an IP address automatica	ally	
• Use the following IP address:		
IP address:		
Subnet mask:		
Default gateway:		
Obtain DNS server address auto	matically	
• Use the following DNS server ad	dresses:	
Preferred DNS server:		
Alternate DNS server:		
Validate settings upon exit	Advanced	I
	OK Ca	ncel

- f. On the **Internet Protocol Version 4 (TCP/IPv4) Properties** page, configure the network information of the network interface and click **OK**.
 - Select Use the following IP address:.
 - IP address: Enter the private IP address of the network interface. In this example, the private IP address is 192.168.0.16.
 - **Subnet mask**: Enter the mask of the subnet where the network interface is created. In this example, the mask is **255.255.255.0**.
 - Default gateway: Enter the gateway of the subnet where the network interface is created. In this example, the gateway is 192.168.0.1.

You can get IP settings assigne	ed automatically if your network supports
this capability. Otherwise, you for the appropriate IP settings.	need to ask your network administrator
O Obtain an IP address auto	omatically
Use the following IP addre	ess:
IP address:	192.168.0.16
Subnet mask:	255 . 255 . 255 . 0
Default gateway:	192.168.0.1
Obtain DNS server addres	ss automatically
• Use the following DNS ser	ver addresses:
Preferred DNS server:	
Alternate DNS server:	• • •

g. On the **Team1 Properties** page, click **OK** to save the settings.

Microsoft Netwo	ork Adapter Multiplexor	Driver	
his connection uses t	the following items:	Configur	e
Client for Microsoft MA Control of Microsoft MA Control of Microsoft MA Control of Microsoft Loa Control of Microsoft Loa Control of Microsoft Loa Control of Microsoft Net	rosoft Networks er Sharing for Microsoft C Bridge Scheduler ad Balancing/Failover P ocol Version 4 (TCP/IP) work Adapter Multiplex	Networks Provider (4) or Protocol	
<			>
Install	Uninstall	Propertie	es
Description Transmission Contro	ol Protocol/Internet Prot	tocol. The defa	ult

h. Return to the Team1 Status page and click Details....

On the **Network Connection Details** page, check whether the following information is correctly configured:

- Physical Address: MAC address of the network interface.
- IPv4 Address: the private IP address of the network interface.
- IPv4 Subnet Mask: the mask of the subnet where the network interface is created.
- **IPv4 Default Gateway**: the gateway of the subnet where the network interface is created.

eneral		
Connection		
IPv4 Connectiv	ity:	Internet
IPv6 Connectiv	ity:	No network access
Media State:		Enabled
Duration:		00:13:42
Speed:		100.0 Gbps
<u>Internet Theorem (1997)</u>	9	
Activity	Sent —	Received
Activity	Sent — 85,328	Received 39,356

Network Connection Details

X

Property	Value
Connection-specific DN	
Description	Microsoft Network Adapter Multiplexor Dr
Physical Address	FA-16-3E-6D-C5-D5
DHCP Enabled	No
IPv4 Address	192.168.0.16
IPv4 Subnet Mask	255.255.255.0
IPv4 Default Gateway	192.168.0.1
IPv4 DNS Server	
IPv4 WINS Server	
NetBIOS over Tcpip En	Yes
Link-local IPv6 Address	fe80::801a:e3:2406:1ad6%9
IPv6 Default Gateway	
IPv6 DNS Servers	fec0:0:0.ffff::1%1
	fec0:0:0.ffff::2%1
	fec0:0:0.ffff::3%1
<	>

i. Check the settings and click **Close**.

The Network Connections page is displayed.



- 7. Configure the network for the supplementary network interface.
 - a. On the Network Connections page, double-click Team1 VLAN 2242.
 The Team1 VLAN 2242 Status page is displayed.

Connection		
IPv4 Connectivi	ty: No network	access
IPv6 Connectivi	ty: No network	access
Media State:		Enabled
Duration:	0	0:58:22
Creat	100	0 Ghos
Details	100	
Activity	100	
Activity	Sent — R	eceived
Activity	Sent — R 447	eceived 0

b. On the **Team1 - VLAN 2242 Status** page, click **Properties**. The **Team1 - VLAN 2242 Properties** page is displayed.

Microsoft Ne	twork Adapter Multiplexor	Driver #2	
This connection us	es the following items:	<u>C</u> onfigure	
Client for I File and P Grosoft I Grosoft I Grosoft I Microsoft I Microsoft I Microsoft I Microsoft I	Microsoft Networks rinter Sharing for Microsoft MAC Bridge set Scheduler Load Balancing/Failover P	Networks rovider	^
Internet Pr Microsoft I	rotocol Version 4 (TCP/IPv Network Adapter Multiplex	or Protocol	v
 ✓ Microsoft I ✓ Internet Pr ✓ Microsoft I 	rotocol Version 4 (TCP/IPv Network Adapter Multiplex <u>U</u> ninstall	Properties	~

 c. On the Team1 - VLAN 2242 Properties page, click Configure....
 The Microsoft Network Adapter Multiplexor Driver #2 Properties page is displayed.

ICIAI	Advanced Dr	ver Details Eve	nts
	Microsoft Netw	ork Adapter Multiple	xor Driver #2
	Device type:	Network adapt	ers
	Manufacturer:	Microsoft	
	Location:	(BEB59F3A-10 5A850DB1BE0	04C-47E9-8B28- 01}
This	device is workin) properly.	^
			~
			~

d. On the **Microsoft Network Adapter Multiplexor Driver #2 Properties** page, choose the **Advanced** tab, click **MAC Address**, enter the MAC address of the supplementary network interface, and click **OK**.

When entering the MAC address, remove the colons (:) and use the uppercase letters. For example, if the MAC address of the supplementary network interface is **fa:16:3e:6d:c5:db**, enter **FA163E6DC5DB**.

The foll the prop on the r	owing proper perty you war right.	ties are av It to chang	ailable fo e on the	or this n e left, ar	etwork adapter. nd then select its	Click value
Propert	y:				Value:	
Encap Heade IPsec IPv4 C Large Large MAC A Receiv Recv Recv	sulated Task offload Checksum Offload Send Offload Send Offload Address ve Side Scalir Segment Coa Segment Coa becksum Offl	Offload load Version 2 Version 2 lescing (IP lescing (IP lescing (IP	(IPv: (IPv:	•	FA163E6DC5D)B
TCP C UDP C UDP C	hecksum Offl Thecksum Off Thecksum Off	oad (IPv6) load (IPv4 load (IPv6)) ~			

e. Return to the **Team1 - VLAN 2242 Properties** page, double-click **Internet Protocol Version 4 (TCP/IPv4)**.

The Internet Protocol Version 4 (TCP/IPv4) Properties page is displayed.

7	Microsoft Netw	vork Adapter Multiplex	or Driver #2	
his co	onnection uses	s the following items:	Configur	e
	Client for Mi File and Prin Microsoft M QoS Packet Microsoft Lo Internet Pro	crosoft Networks hter Sharing for Micros AC Bridge t Scheduler bad Balancing/Failove tocol Version 4 (TCP/ etwork Adapter Multip	oft Networks er Provider <mark>IPv4)</mark> lexor Protocol	^
<				>
	Install	Uninstall	Propertie	s
Deso Tran wide	cription nsmission Cont e area network	rol Protocol/Internet F protocol that provide	Protocol. The defau s communication	ult

nternet Protocol Version 4 (TCP/IPv4) Properties	\times
General		
You can get IP settings assigned auto this capability. Otherwise, you need t for the appropriate IP settings.	matically if your network supports o ask your network administrator	
Obtain an IP address automatica	ally	
• Use the following IP address:		
IP address:		
Subnet mask:		
Default gateway:		
Obtain DNS server address auto	matically	
• Use the following DNS server add	dresses:	
Preferred DNS server:		
Alternate DNS server:		
Validate settings upon exit	Advanced	
	OK Cance	el

- f. On the **Internet Protocol Version 4 (TCP/IPv4) Properties** page, configure the network information of the supplementary network interface and click **OK**.
 - Select Use the following IP address:
 - IP address: Enter the private IP address of the supplementary network interface. In this example, the private IP address is 192.168.0.22.
 - Subnet mask: Enter the mask of the subnet where the supplementary network interface is created. In this example, the mask is 255.255.255.0.
 - Default gateway: Enter the gateway of the subnet where the supplementary network interface is created. In this example, the gateway is 192.168.0.1.

. .

his capability. Otherwise, you need or the appropriate IP settings.	to ask your network administrate	or
Obtain an IP address automatic	ally	
Use the following IP address: -		
IP address:	192.168.0.22	
S <u>u</u> bnet mask:	255.255.255.0	
Default gateway:	192 . 168 . 0 . 1	
Obtain DNS server address aut	omatically	
Use the following DNS server as	ddresses:	
Preferred DNS server:	× • •	
Alternate DNS server:		
Validate settings upon exit	Ad <u>v</u> anced	

Warning - Multiple default gateways are intended to provide redundancy to a single network (such as an intranet or the Internet). They will not function properly when the gateways are on two separate, disjoint networks (such as one on your intranet and one on the Internet). Do you want to save this configuration?

Yes

g. On the **Team1 - VLAN 2242 Properties** page, click **OK** to save the settings.

No

Microsoft Netwo	ork Adapter Multiplexor	Driver #2
his connection uses	the following items:	<u>C</u> onfigure
File and Print Hicrosoft MA Gos Packet Microsoft Loa Microsoft Loa Internet Proto	er Sharing for Microsoft C Bridge Scheduler ad Balancing/Failover P ocol Version 4 (TCP/IPv work Adapter Multiplex	Networks rovider (4) or Protocol
Microsoft Net		>
<		1.22
<	<u>U</u> ninstall	Properties

- h. Return to the Team1 VLAN 2242 Status page and click Details....
 On the Network Connection Details page, check whether the following information is correctly configured:
 - **Physical Address**: MAC address of the supplementary network interface.
 - **IPv4 Address**: the private IP address of the supplementary network interface.
 - **IPv4 Subnet Mask**: the mask of the subnet where the supplementary network interface is created.
 - IPv4 Default Gateway: the gateway of the subnet where the supplementary network interface is created.

Connection		
IPv4 Connectivi	ty:	No network access
IPv6 Connectivi	ty:	No network access
Media State:		Enabled
Duration:		00:05:35
Speed:		100.0 Gbps
Activity		4 📷
Activity	Sent —	Received
Activity Bytes:	Sent — 346	Received

Network Connection Details

×

Property	Value
Connection-specific DN	
Description	Microsoft Network Adapter Multiplexor Dr
Physical Address	FA-16-3E-6D-C5-DB
DHCP Enabled	No
IPv4 Address	192.168.0.22
IPv4 Subnet Mask	255.255.255.0
IPv4 Default Gateway	192.168.0.1
IPv4 DNS Server	
IPv4 WINS Server	
NetBIOS over Tcpip En	Yes
Link-local IPv6 Address IPv6 Default Gateway	fe80::9756:7398:cfe0:6149%14
IPv6 DNS Servers	fec0:0:0:ffff::1%1
	fec0:0:0:ffff::2%1
	fec0:0:0:ffff::3%1
<	>

- i. Check the settings and click **Close**.
- 8. On the Windows PowerShell CLI page, check whether the network interface and supplementary network interface are connected to the test ECS.
 - a. Run the following command to verify the connectivity between network interface and the test ECS:

Ping <*private-IP-address-of-the-test-ECS>* **-S** <*private-IP-address-of-the-network-interface>*

Plan the same VPC and security group for the test ECS and the ECS with network interface attached, so that the two ECSs can communicate with each other by default.

Example command:

Ping 192.168.0.133 -S 192.168.0.16

If information similar to the following is displayed, the two ECSs can communicate with each other.

PS C:\Users\Administrator> Ping 192.168.0.133 -5 192.168.0.16
Pinging 192.168.0.133 from 192.168.0.16 with 32 bytes of data:
Reply from 192.168.0.133: bytes=32 time=1ms TTL=64
Reply from 192.168.0.133: bytes=32 time<1ms TTL=64
Reply from 192.168.0.133: bytes=32 time<1ms TTL=64
Ping statistics for 192.168.0.133:
 Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
 Minimum = 0ms, Maximum = 1ms, Average = 0ms</pre>

b. Run the following command to verify the connectivity between the supplementary network interface of **eth0** and the test ECS:

Ping *<private-IP-address-of-the-test-ECS>* **-S** *<private-IP-address-of-the-supplementary-network-interface>*

Plan the same VPC and security group for the test ECS and the ECS with the supplementary network interface attached. This allows the two ECSs to communicate with each other by default.

Example command:

Ping 192.168.0.133 -S 192.168.0.22

If information similar to the following is displayed, the two ECSs can communicate with each other.

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, bound 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 \equiv 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, deleting a supplementary network interface that is used as the next hop for a custom route in a VPC route table will also delete the associated 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 subinterfaces configured on the ECS.

5.3 Network Interface Configuration Examples

5.3.1 Binding an EIP to the Extended Network Interface of an ECS to Enable Internet Access

Scenarios

As shown in **Figure 5-3**, the ECS has two network interfaces, one primary network interface and one extended network interface. You can bind an EIP to the extended network interface of the ECS and configure policy-based routes to ensure that the ECS can access the Internet through the EIP.

This section uses a Linux ECS as an example.

- After the configuration, the ECS will access the Internet through the **EIP bound** to the extended network interface instead of the **EIP bound to the primary** network interface. The ECS will not be able to communicate with the Internet through the primary network interface, and the original network connection will be interrupted. Exercise caution when performing this operation.
- If you need to communicate with the Internet through both the primary and extended network interfaces, see Configuring Policy-based Routes for an ECS with Multiple Network Interfaces.



Figure 5-3 Accessing the Internet through the EIP bound to the extended network interface

Step 1: Create Cloud Resources and Attach an Extended Network Interface

 Create a VPC and two subnets in the VPC.
 In this example, the primary and extended network interfaces 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, see **Purchasing an ECS**.

3. Create a network interface and attach it to the ECS as an extended network interface.

When creating a network interface, select a different subnet from where the primary network interface is created. For details, see **Creating a Network Interface**.

Attach the network interface to the ECS. For details, see **Attaching a Network Interface to a Cloud Server**.

 Assign an EIP and bind it to the extended network interface of the ECS. For details, see Assigning an EIP.
 Bind the EIP to the extended network interface 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 extended network interface, you need to obtain the network information in **Table 5-7**.

ltem	Primary Network Interface	Extended Network Interface
Private IP address of the network interface	192.168.11.42	192.168.17.191
Subnet gateway address	192.168.11.1	192.168.17.1

Table 5-7 Required ECS network information

- 1. Obtain the private IP addresses of the ECS's network interfaces.
 - 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 extended network interfaces 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 5-4 Viewing the gateway address of the subnet

Gateway and DNS Infor			
DHCP	Enabled	Gateway	192.168.0.1
DNS Server Address	100.125.1.250, 100.125.129.250 🖉	Domain Name	- 2 3
IPv4 DHCP Lease Time	1250 days 🖉 💮	NTP Server Address	- 2 3

Step 3: Configure Policy-based Routes for the Extended Network Interface

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 network interface:

route -n

The following figure is displayed. In this figure:

- The destination of the route for the primary network interface is 192.168.11.0/24.
- The destination of the route for the extended network interface is 192.168.17.0/24.

[root@ecs-b926 ~]# route -n							
Kernel IP routin	ng 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	eth8
169.254.0.0	0.0.0.0	255.255.0.0	U	1002	0	0	eth8
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 network interface names of the ECS:

ifconfig

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

- 192.168.11.42 is the IP address of the primary network interface, and the network interface name is eth0.
- 192.168.17.191 is the IP address of the extended network interface, and the network interface name is eth1.

[root@ecs-b926~]# if config
eth0: flags=4163 <up,broadcast,running,multicast> mtu 1500</up,broadcast,running,multicast>
inet 192.168.11.42 netmask 255.255.255.0 broadcast 192.168.11.255
inet6 fe80::f816:3eff:fef7:1c44
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
-414. 61
etni: Hags-HD3CUF, BRUHUCHSI, RUMPING, RULFICHSI> mtu 1500
inet 192.100.17.191 metmask 253.253.253.0 Droadcast 192.100.17.253
The to feed and the start field burning the first of the start of the
ether fa:16:3e:16:55:7f tqueuelen 1998 (Ethernet)
KX packets 11 bytes 1283 (1.2 K18)
KX errors 8 aropped 8 overruns 8 frame 8
TX packets 12 bytes 1388 (1.3 KiB)
TX errors & dropped & overruns & carrier & collisions &
lo: flags=73(UP_LOOPBACK_RUNNING>mtu_65536
inet 122 A A 1 network 255 A A A
inet6 ::1 metivlen 128 scaneid 8×18
loon tyrueuelen 1 (Local Loonback)
RY nackets 51 hites 12018 (11 2 KiR)
RY property of drouged a constraints of frame of
TY notes 51 Nutes 12018 (11 2 ViR)
TX propers d dropped & compress & contrier & collicions &
in citors o aroppea o overtains o carrier o corristons o

- 4. Configure the default route for the ECS so that it can access the Internet through the extended network interface.
 - a. Run the following command to delete the default route of the primary network interface:

route del -net 0.0.0.0 gw <subnet-gateway-IP-address> dev <network
interface-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 network interface collected in section Table 5-7.
- Network interface name: Enter the name of the primary network interface obtained in **3**.

Example command:

route del -net 0.0.0.0 gw 192.168.11.1 dev eth0

This operation will interrupt the ECS traffic. Ensure that services will not be affected before deleting the default route of the primary network interface.

b. Run the following command to configure the default route for the extended network interface:

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 extended network interface collected in section Table 5-7.

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 Network Interfaces

5.3.2.1 Overview

Background

If a cloud server has multiple network interfaces, the primary network interface can communicate with external networks by default, but the extended network interfaces cannot. To enable extended network interfaces to communicate with external networks, you need to configure policy-based routes for these network interfaces.

Scenarios

This example describes how to configure policy-based routes for an ECS with two network interfaces. **Figure 5-5** shows the networking.

- The primary and extended network interfaces 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, so the two ECSs can communicate with each other through primary network interfaces without configuring policy-based routes.
- After policy-based routes are configured for the two network interfaces of the source ECS, both the primary and extended network interfaces can be used to communicate with the destination ECS.

NOTE

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 network interface to communicate with the destination ECS.



Figure 5-5 Networking of an ECS with two network interfaces

5.3.2.2 Collecting ECS Network Information

Scenarios

Before configuring policy-based routes for an ECS with multiple network interfaces, you need to collect network information about the ECS.

• **Table 5-8** lists the information to be collected for a Linux ECS using IPv4.

Ту pe	Primary Network Interface	Extended Network Interface	How to Obtain
So ur ce	 IPv4 address: 10.0.0.115 Subnet IPv4 CIDR block: 10.0.0.0/24 Subnet IPv4 gateway: 10.0.0.1 	 IPv4 address: 10.0.1.183 Subnet IPv4 CIDR block: 10.0.1.0/24 Subnet IPv4 gateway: 10.0.1.1 	 Obtaining ECS Network Interface Addresses Obtaining Subnet
De sti na tio n	IPv4 address: 10.0.2.12	N/A	CIDR Blocks and Gateway Addresses

 Table 5-8 Linux ECS using IPv4 addresses

• **Table 5-9** lists the information to be collected for a Windows ECS using IPv4.

Table 5-9 Windows ECS using IPv4

Ту pe	Primary Network Interface	Extended Network Interface	How to Obtain
So ur ce	 IPv4 address: 10.0.0.59 Subnet IPv4 gateway: 10.0.0.1 	 IPv4 address: 10.0.1.104 Subnet IPv4 gateway: 10.0.1.1 	Obtaining ECS Network Interface Addresses
De sti na tio n	IPv4 address: 10.0.2.12	N/A	 Obtaining Subnet CIDR Blocks and Gateway Addresses

The above information is only for your reference.

Obtaining ECS Network Interface 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**.
- 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 extended network interfaces.

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.
- 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 name of its VPC. 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.
- 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 Automatically Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Network Interfaces (Huawei Cloud EulerOS 2.0/ CentOS 8.0 or Later)

Scenarios

This section describes how to use an automation script to configure policy-based routes for a Linux ECS with two network interfaces. **The automation script supports Huawei Cloud EulerOS 2.0, CentOS 8.0 and later versions.**

- IPv4: If IPv4 communication between cloud servers with multiple network interfaces is required, you need to configure IPv4 routes by referring to Configuring Policy-based Routes for a Linux ECS with IPv4 Addresses.
- IPv6: If IPv6 communication between cloud servers with multiple network interfaces is required, you need to configure IPv6 routes by referring to Configuring Policy-based Routes for a Linux ECS with IPv6 Addresses.
- IPv4/IPv6 dual stack: If both IPv4 and IPv6 communications between cloud servers with multiple network interfaces are required, you need to configure both IPv4 and IPv6 routes by referring to Configuring Policy-based Routes for a Linux ECS with Both IPv4 and IPv6 Addresses.

For details about the background knowledge and networking of an ECS with two network interfaces, see **Overview**.

Configuring Policy-based Routes for a Linux ECS with IPv4 Addresses

1. Collect the information, such as IPv4 addresses, about ECS network interfaces for configuring policy-based routes.

For details, see Collecting ECS Network Information.

In this example, the network information of the ECS is shown in **Table 5-10**.

Table 5-10 Linux ECSs using IPv4 addresses

Тур e	Primary Network Interface	Extended Network Interface
Sour ce	 IPv4 address: 10.0.0.115 Subnet IPv4 CIDR block: 10.0.0.0/24 Subnet IPv4 gateway: 10.0.0.1 	 IPv4 address: 10.0.1.183 Subnet IPv4 CIDR block: 10.0.1.0/24 Subnet IPv4 gateway: 10.0.1.1
Dest inati on	IPv4 address: 10.0.2.12	N/A

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the source ECS can use its primary network interface to communicate with the destination ECS using IPv4 addresses:

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

ping -I <*IPv4-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv4-address-of-the-network-interface-on-the-destination-ECS>*

Example 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 network interface 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 ---

4. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The IPv4 address of the primary network interface is 10.0.0.115, and its name is eth0.
- The IPv4 address of the extended network interface is 10.0.1.183, and its name is eth1.

```
[root@ecs-resource ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
```

5.

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 Create an automation script to configure IPv4 routes: a. Create the automation configuration script file **15-policy-route.sh**: vi /etc/NetworkManager/dispatcher.d/15-policy-route.sh b. Press i to enter the editing mode. Add the automation script to the file. C. #!/bin/bash if ["\$2" != "up"]; then exit 0 fi interface=\$1 check_route_table() { local route_table_id=\$1 local ip_version=\$2 local max_attempts=100 local attempts=0 while [\$attempts -lt \$max_attempts]; do output=\$(ip -\$ip_version route show table \$route_table_id &>/dev/null) if [\$? -ne 0] || [-z "\$output"]; then break else route_table_id=\$((route_table_id + 1)) attempts=\$((attempts + 1)) fi done echo \$route_table_id } add_ipv4_route_table() { local DEFAULT_IPV4_ROUTE_TABLE_ID=1000 v4_route_table_id=\$(ip route show table all | grep -F "default via \$gateway dev \$interface table" | awk '{print \$NF}') if [-z \$v4_route_table_id]; then interface_number=\$(echo \$interface | grep -o '[0-9]*\$') if [-z "\$interface_number"]; then v4_route_table_id="\$DEFAULT_IPV4_ROUTE_TABLE_ID" else v4_route_table_id=\$((DEFAULT_IPV4_ROUTE_TABLE_ID + interface_number)) fi v4_route_table_id=\$(check_route_table \$v4_route_table_id 4) echo "add policy route for dev: \$interface table: \$v4_route_table_id subnet: \$subnet gateway: \$gateway" ip route add default via \$gateway dev \$interface table \$v4_route_table_id
ip route add \$subnet dev \$interface table \$v4_route_table_id fi } generate_ipv4_policy_route() { subnet=\$(nmcli device show \$interface | grep -F 'IP4.ROUTE' | grep -F 'nh = 0.0.0.0' | cut -d'=' f2 | cut -d',' -f1 | tr -d ' ' | head -n 1) gateway=\$(ipcalc -i \$subnet | awk '/HostMin/ {print \$2}') add_ipv4_route_table nmcli device show \$interface | grep -F 'IP4.ADDRESS' | while read -r line; do IP=\$(echo \$line | awk '{print \$2}' | cut -d'/' -f1) if ip rule list | grep -F "\$v4_route_table_id" | grep -q "\$IP"; then echo "ip rule already exists for \$IP with table \$v4_route_table_id" continue fi echo "Adding rule for \$IP on \$interface with table \$v4_route_table_id" ip rule add from \$IP table \$v4_route_table_id done } generate_ipv4_policy_route

- d. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Check the permissions on the automatic configuration script file **15**-**policy-route.sh**:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15policy-route.sh** is **-rw-r--r-**, which is different from that on other system files, for example, permission **-rwxr-xr-x** on the file **10-ifcfg-rh-routes.sh**. [root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient -rw-r--r-- 1 root root 2019 Jun 18 12:21 15-policy-route.sh -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-dhcp -rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-onoffline

- -rwxrwxr-x 1 root root 719 May 10 2019 hook-network-manager
- drwxr-xr-x. 2 root root 4096 Apr 18 17:12 no-wait.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d
- drwxr-xr-x. 2 root root 4096 Apr 18 17:12 pre-up.d
- f. Grant permissions on the automation configuration script file **15-policyroute.sh**:

sudo chown root:root /etc/NetworkManager/dispatcher.d/15-policyroute.sh

sudo chmod 755 /etc/NetworkManager/dispatcher.d/15-policyroute.sh

g. Check the permissions on the automatic configuration script file **15**-**policy-route.sh** again:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15**-**policy-route.sh** is **-rwxr-xr-x**.

[root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh

- -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient
- -rwxr-xr-x 1 root root 2019 Jun 18 12:21 15-policy-route.sh
- -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-dhcp

-rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-onoffline

-rwxrwxr-x 1 root root 719 May 10 2019 hook-network-manager drwxr-xr-x. 2 root root 4096 Apr 18 17:12 no-wait.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d drwxr-xr-x. 2 root root 4096 Apr 18 17:12 pre-up.d

h. Restart the network service to execute the **15-policy-route.sh** script:

systemctl restart NetworkManager

Before restarting the network service, ensure that services are not affected.

6. Check whether the policy-based routes are added for the network interfaces with IPv4 addresses.

ip rule

ip route show table 1000

ip route show table 1001

table 1000 is the route table name of the primary network interface, and **table 1001** is the route table name of the extended network interface.

If information similar to the following is displayed, the policy-based routes have been added for the network interfaces with IPv4 addresses.

- [root@ecs-resource ~]# ip rule 0: from all lookup local 32764: from 10.0.1.183 lookup 1001 32765: from 10.0.0.115 lookup 1000 32766: from all lookup main 32767: from all lookup default [root@ecs-resource ~]# ip route show table 1000 default via 10.0.0.1 dev eth0 10.0.0/24 dev eth0 scope link [root@ecs-resource ~]# ip route show table 1001 default via 10.0.1.1 dev eth1 10.0.1.0/24 dev eth1 scope link
- 7. Check whether the source and destination ECSs can communicate with each other using IPv4 addresses.

ping -l <*IPv4-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv4-address-of-the-network-interface-on-the-destination-ECS>*

ping -I <*IPv4-address-of-the-extended-network-interface-on-the-source-ECS>* <*IPv4-address-of-the-network-interface-on-the-destination-ECS>*

Example 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 IPv4 addresses of the source ECS can communicate with the destination ECS, indicating that the IPv4 policy-based routes are successfully configured.

[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=0.231 ms

64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.195 ms ^C

--- 10.0.2.12 ping statistics ---

2 packets transmitted, 2 received, 0% packet loss, time 1043ms rtt min/avg/max/mdev = 0.195/0.213/0.231/0.018 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=0.274 ms 64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.196 ms 64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.125 ms ^C --- 10.0.2.12 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2042ms rtt min/avg/max/mdev = 0.125/0.198/0.274/0.060 ms

Configuring Policy-based Routes for a Linux ECS with IPv6 Addresses

1. Collect the information, such as IPv4 and IPv6 addresses, about ECS network interfaces for configuring policy-based routes.

For details, see **Collecting ECS Network Information**.

In this example, the network information of the ECS is shown in Table 5-11.

Тур e	Primary Network Interface	Extended Network Interface
Sour ce	 IPv4 address: 10.0.0.133 IPv6 address: 2407:c080:802:1f85:918b:9039 :41b2:24a8 	 IPv4 address: 10.0.1.120 IPv6 address: 2407:c080:802:2107:ab85:4a2 7:d20:2119
	 Subnet IPv6 CIDR block: 2407:c080:802:1f85::/64 Subnet IPv6 gateway: 2407, 000 000 1/05, 1 	 Subnet IPv6 CIDR block: 2407:c080:802:2107::/64 Subnet IPv6 gateway: 2407.000.002.2107.1
Dest inati on	 IPv4 address: 10.0.2.3 IPv6 address: 2407:c080:802:1185::1 IPv6 address: 2407:c080:802:2108:96ec:3c49: 391a:5ebc 	N/A

Table 5-11 Linux ECSs using IPv4 and IPv6 addresses

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the ECSs have IPv6 enabled and have IPv6 addresses.

Perform this step for both the source and destination ECSs to ensure that the ECSs have IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

ECSs in this example run **CentOS 8.0 64bit** or **Huawei Cloud EulerOS 2.0 Standard 64 bit**. For details about how to obtain IPv6 addresses for ECSs running other OSs, see Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses).

a. Check whether the ECS has IPv6 addresses: ip addr In the following command output, eth0 and eth1 are the network interfaces of the ECS. Each network interface has one **inet6** entry starting with **fe80**. This indicates that the ECS has IPv6 enabled but does not have IPv6 addresses assigned. In this case, perform **3.b** to **3.g** to obtain IPv6 addresses.

[root@ecs-resource ~]# ip addr

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:72:72:f7 brd ff:ff:ff:ff:ff

inet 10.0.0.133/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0

valid_lft 315359994sec preferred_lft 315359994sec

inet6 fe80::f816:3eff:fe72:72f7/64 scope link noprefixroute

valid_lft forever preferred_lft forever

3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:72:73:ea brd ff:ff:ff:ff:ff:ff

inet 10.0.1.120/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid_ft 315359994sec preferred_ft 315359994sec

inet6 fe80::f816:3eff:fe72:73ea/64 scope link noprefixroute

valid_lft forever preferred_lft forever

b. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The IPv4 address of the primary network interface is 10.0.0.133, and its name is eth0.
- The IPv4 address of the extended network interface is 10.0.1.120, and its name is eth1.

[root@ecs-resource ~]# ifconfig eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.0.0.133 netmask 255.255.255.0 broadcast 10.0.0.255 inet6 fe80::f816:3eff:fe72:72f7 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:72:72:f7 txqueuelen 1000 (Ethernet) RX packets 50917 bytes 71068144 (67.7 MiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 4969 bytes 1123356 (1.0 MiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 eth1: flags=4163<UPBROADCAST PLINNING MULTICAST> mtu 1500

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet **10.0.1.120** netmask 255.255.255.0 broadcast 10.0.1.255 inet6 fe80::f816:3eff:fe72:73ea prefixlen 64 scopeid 0x20<link> ether fa:16:3e:72:73:ea txqueuelen 1000 (Ethernet) RX packets 21 bytes 3190 (3.1 KiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 26 bytes 2934 (2.8 KiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

- c. Configure the **ifcfg** file of the primary network interface.
 - i. Run the following command to open the **ifcfg** file of the primary network interface:

vi /etc/sysconfig/network-scripts/ifcfg-<Primary-network-interfacename>

The name of the primary network interface is obtained in **3.b**. Example command:

vi /etc/sysconfig/network-scripts/ifcfg-eth0

ii. Press i to enter the editing mode.

- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- d. Configure the **ifcfg** file of the extended network interface.
 - i. Run the following command to open the **ifcfg** file of the extended network interface:

vi /etc/sysconfig/network-scripts/ifcfg-<Extended-networkinterface-name>

The name of the extended network interface is obtained in **3.b**. Example command:

vi /etc/sysconfig/network-scripts/ifcfg-eth1

- ii. Press **i** to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Edit the **/etc/sysconfig/network** file.
 - i. Run the following command to open the **/etc/sysconfig/network** file:

vi /etc/sysconfig/network

- ii. Press **i** to enter the editing mode.
- iii. Add the following content to the end of the file: NETWORKING_IPV6="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- f. Run the following command to restart the network service for the configuration to take effect:

systemctl restart NetworkManager

g. Check whether the ECS has IPv6 addresses:

ip addr

In the following command output, each network interface has an additional **inet6** entry starting with **2407**, followed by an entry starting with **fe80**. This indicates that the ECS has IPv6 addresses. [root@ecs-resource ~]# ip addr

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:72:72:f7 brd ff:ff:ff:ff:ff:ff

inet 10.0.0.133/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0

valid_lft 315359994sec preferred_lft 315359994sec

inet6 2407:c080:802:1f85:918b:9039:41b2:24a8/128 scope global dynamic noprefixroute valid_lft 7194sec preferred_lft 7194sec

inet6 fe80::f816:3eff:fe72:72f7/64 scope link noprefixroute

valid_lft forever preferred_lft forever

3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:72:73:ea brd ff:ff:ff:ff:ff:ff

inet 10.0.1.120/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid_lft 315359994sec preferred_lft 315359994sec

inet6 2407:c080:802:2107:ab85:4a27:d20:2119/128 scope global dynamic noprefixroute valid [ft 7195sec preferred [ft 7195sec

inet6 fe80::f816:3eff:fe72:73ea/64 scope link noprefixroute

valid_lft forever preferred_lft forever

- h. Log in to the destination ECS and obtain an IPv6 address by performing operations from **3.a** to **3.g**.
- 4. Log in to the source ECS and check whether it can use its primary network interface to communicate with the destination ECS:

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

ping6 -I <*IPv6-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv6-address-of-the-network-interface-on-the-destination-ECS>*

Example command:

ping6 -I 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc

If information similar to the following is displayed, the source ECS can use its primary network interface to communicate with the destination ECS. [root@ecs-resource ~]# ping6 -I 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:1f85:918b:9039:41b2:24a8 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.283 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.212 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc; icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc; icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc; icmp_seq=3 ttl=64 time=0.122 ms 7C

s packets transmitted, 3 received, 0% packet loss, time 2 rtt min/avg/max/mdev = 0.122/0.205/0.283/0.065 ms

- 5. Create an automation script to configure IPv6 routes:
 - a. Create the automation configuration script file **15-policy-route.sh**:

vi /etc/NetworkManager/dispatcher.d/15-policy-route.sh

- b. Press **i** to enter the editing mode.
- c. Add the automation script to the file. #!/bin/bash

```
if [ "$2" != "up" ]; then
  exit 0
fi
interface=$1
check_route_table() {
  local route_table_id=$1
  local ip_version=$2
  local max_attempts=100
  local attempts=0
  while [ $attempts -lt $max_attempts ]; do
     output=$(ip -$ip_version route show table $route_table_id &>/dev/null)
     if [ $? -ne 0 ] || [ -z "$output" ]; then
        break
     else
        route_table_id=$((route_table_id + 1))
        attempts=$((attempts + 1))
     fi
  done
  echo $route_table_id
}
add_ipv6_route_table() {
  local DEFAULT_IPV6_ROUTE_TABLE_ID=1100
```

```
v6_route_table_id=$(ip -6 route show table all | grep -F "default via $v6_gateway dev
$interface table" | grep -oP "table \K\d+")
   if [ -z $v6_route_table_id ]; then
      interface_number=$(echo $interface | grep -o '[0-9]*$')
      if [ -z "$interface_number" ]; then
        v6_route_table_id=$DEFAULT_IPV6_ROUTE_TABLE_ID
      else
        v6_route_table_id=$((DEFAULT_IPV6_ROUTE_TABLE_ID + interface_number))
      fi
      v6_route_table_id=$(check_route_table $v6_route_table_id 6)
      echo "add policy route for dev: $interface table: $v6_route_table_id v6_subnet: $v6_subnet
v6_gateway: $v6_gateway"
      ip -6 route add default via $v6_gateway dev $interface table $v6_route_table_id
     ip -6 route add $v6_subnet dev $interface table $v6_route_table_id
   fi
}
generate_ipv6_policy_route() {
IPV6INIT=$(awk -F= '/^IPV6INIT=/{gsub(/"/, ""); print $2}' "/etc/sysconfig/network-scripts/
ifcfg-$interface" | tr -d ' ' | tail -n 1)
   if [ "$IPV6INIT" != "yes" ]; then
      echo "$interface IPV6INIT is not set to yes, skip IPv6 policy route configurate."
      return 0
   fi
   for ((x=0; x<10; x++)); do
      if (ip address show $interface | grep -F "inet6" | grep -F "dynamic"); then
        v6_subnet=$(nmcli device show $interface | grep -F 'IP6.ROUTE' | grep -F 'nh = ::'| grep -
v 'dst = fe80' | grep -v '/128' | cut -d'=' -f2 | cut -d',' -f1 | tr -d ' ' | head -n 1)
        v6_gateway=$(ipcalc -6 -i $v6_subnet | awk '/HostMin/ {print $2}')1
         add_ipv6_route_table
         nmcli device show $interface | grep -F 'IP6.ADDRESS' | grep -v "/64" | while read -r line;
do
           IP=$(echo $line | awk '{print $2}' | cut -d'/' -f1)
           if ip -6 rule list | grep -F "$v6_route_table_id" | grep -q "$IP"; then
              echo "ip rule already exists for $IP with table $v6_route_table_id"
              continue
           fi
           echo "Adding rule for $IP on $interface with table $v6 route table id"
           ip -6 rule add from $IP table $v6_route_table_id
         done
        break
      fi
       sleep 1
   done
}
```

generate_ipv6_policy_route

- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Check the permissions on the automatic configuration script file **15**-**policy-route.sh**:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15policy-route.sh** is **-rw-r--r-**, which is different from that on other system files, for example, permission **-rwxr-xr-x** on the file **10-ifcfg-rh-routes.sh**. [root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient -rw-r-r-- 1 root root 2515 Jun 18 15:13 15-policy-route.sh -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-dhcp

Issue 01 (2025-07-22) Copyright © Huawei Cloud Computing Technologies Co., Ltd.

-rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-onoffline -rwxrwxr-x 1 root root 719 May 10 2019 hook-network-manager drwxr-xr-x. 2 root root 4096 Apr 18 17:12 no-wait.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d drwxr-xr-x. 2 root root 4096 Apr 18 17:12 pre-up.d

f. Grant permissions on the automation configuration script file **15-policyroute.sh**:

sudo chown root:root /etc/NetworkManager/dispatcher.d/15-policyroute.sh

sudo chmod 755 /etc/NetworkManager/dispatcher.d/15-policyroute.sh

g. Check the permissions on the automatic configuration script file **15**-**policy-route.sh** again:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15**-**policy-route.sh** is **-rwxr-xr-x**.

[root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient -rwxr-xr-x 1 root root 2019 Jun 18 12:21 15-policy-route.sh -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-dhcp -rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-onoffline -rwxrwxr-x 1 root root 719 May 10 2019 hook-network-manager drwxr-xr-x. 2 root root 4096 Apr 18 17:12 no-wait.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d drwxr-xr-x. 2 root root 4096 Apr 18 17:12 pre-up.d

h. Restart the network service to execute the **15-policy-route.sh** script:

systemctl restart NetworkManager

Before restarting the network service, ensure that services are not affected.

6. Check whether the policy-based routes are added for the network interfaces with IPv6 addresses.

ip -6 rule

ip -6 route show table 1100

ip -6 route show table 1101

table 1100 is the route table name of the primary network interface, and **table 1101** is the route table name of the extended network interface.

If information similar to the following is displayed, the policy-based routes have been added for the network interfaces with IPv6 addresses.

[root@ecs-resource ~]# ip -6 rule
0: from all lookup local
32764: from 2407:c080:802:2107:ab85:4a27:d20:2119 lookup 1101
32765: from 2407:c080:802:1f85:918b:9039:41b2:24a8 lookup 1100
32766: from all lookup main
[root@ecs-resource ~]# ip -6 route show table 1100
2407:c080:802:1f85::/64 dev eth0 metric 1024 pref medium
default via 2407:c080:802:1f85::1 dev eth0 metric 1024 pref medium
[root@ecs-resource ~]# ip -6 route show table 1101
2407:c080:802:1f85::1/64 dev eth0 metric 1024 pref medium
[root@ecs-resource ~]# ip -6 route show table 1101
2407:c080:802:2107::/64 dev eth1 metric 1024 pref medium
[root@ecs-resource ~]# ip -6 route show table 1101
2407:c080:802:2107::/64 dev eth1 metric 1024 pref medium
default via 2407:c080:802:2107::1 dev eth1 metric 1024 pref medium

7. Check whether the source and destination ECSs can communicate with each other using IPv6 addresses.

ping -6 -1 <*IPv6-address-of-the-primary-network-interface-on-the-source-ECS> <IPv6-address-of-the-network-interface-on-the-destination-ECS>*

ping -6 -1 *<IPv6-address-of-the-extended-network-interface-on-the-source-ECS> <IPv6-address-of-the-network-interface-on-the-destination-ECS>*

Example commands:

ping -6 -l 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc

ping -6 -l 2407:c080:802:2107:ab85:4a27:d20:2119 2407:c080:802:2108:96ec:3c49:391a:5ebc

If information similar to the following is displayed, both the IPv6 addresses of the source ECS can communicate with the destination ECS, indicating that the IPv6 policy-based routes are successfully configured. [root@ecs-resource ~]# ping -6 -I 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:1f85:918b:9039:41b2:24a8 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.328 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.209 ms ^C --- 2407:c080:802:2108:96ec:3c49:391a:5ebc ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1016ms rtt min/avg/max/mdev = 0.209/0.268/0.328/0.059 ms [root@ecs-resource ~]# ping -6 -I 2407:c080:802:2107:ab85:4a27:d20:2119 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:2107:ab85:4a27:d20:2119 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.345 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.203 ms ^C --- 2407:c080:802:2108:96ec:3c49:391a:5ebc ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1051ms rtt min/avg/max/mdev = 0.203/0.274/0.345/0.071 ms

Configuring Policy-based Routes for a Linux ECS with Both IPv4 and IPv6 Addresses

1. Collect the information, such as IPv4 and IPv6 addresses, about ECS network interfaces for configuring policy-based routes.

For details, see Collecting ECS Network Information.

In this example, the network information of the ECS is shown in Table 5-12.

Тур е	Primary Network Interface	Extended Network Interface
Sour ce	 IPv4 address: 10.0.0.133 IPv6 address: 2407:c080:802:1f85:918b:9039 :41b2:24a8 Subnet IPv6 CIDR block: 2407:c080:802:1f85::/64 Subnet IPv6 gateway: 2407:c080:802:1f85::1 	 IPv4 address: 10.0.1.120 IPv6 address: 2407:c080:802:2107:ab85:4a2 7:d20:2119 Subnet IPv6 CIDR block: 2407:c080:802:2107::/64 Subnet IPv6 gateway: 2407:c080:802:2107::1
Dest inati on	 IPv4 address: 10.0.2.3 IPv6 address: 2407:c080:802:2108:96ec:3c49: 391a:5ebc 	N/A

Table 5-12 Linux ECSs using IPv4 and IPv6 addresses

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the source ECS can use its primary network interface to communicate with the destination ECS using IPv4 addresses:

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

ping -l <*IPv4-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv4-address-of-destination-ECS>*

Example 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 network interface to communicate with the destination ECS. [root@ecs-resource ~]# ping -1 10.0.0.133 10.0.2.3 PING 10.0.2.3 (10.0.2.3) from 10.0.0.133 : 56(84) bytes of data. 64 bytes from 10.0.2.3: icmp_seq=1 ttl=64 time=0.261 ms 64 bytes from 10.0.2.3: icmp_seq=2 ttl=64 time=0.219 ms ^C --- 10.0.2.3 ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1032ms rtt min/avg/max/mdev = 0.219/0.240/0.261/0.021 ms

4. Check whether the ECSs have IPv6 enabled and have IPv6 addresses.

Perform this step for both the source and destination ECSs to ensure that the ECSs have IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

ECSs in this example run **CentOS 8.0 64bit** or **Huawei Cloud EulerOS 2.0 Standard 64 bit**. For details about how to obtain IPv6 addresses for ECSs running other OSs, see Linux (Manually Enabling Dynamic Assignment of IPv6 Addresses).

a. Check whether the ECS has IPv6 addresses:

ip addr

In the following command output, eth0 and eth1 are the network interfaces of the ECS. Each network interface has one **inet6** entry starting with **fe80**. This indicates that the ECS has IPv6 enabled but does not have IPv6 addresses assigned. In this case, perform **4.b** to **4.g** to obtain IPv6 addresses.

[root@ecs-resource ~]# ip addr

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

- link/ether fa:16:3e:72:72:f7 brd ff:ff:ff:ff:ff:ff
- inet 10.0.0.133/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0 valid lft 315359994sec preferred lft 315359994sec
- inet6 fe80::f816:3eff:fe72:72f7/64 scope link noprefixroute
- valid_lft forever preferred_lft forever

3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

- link/ether fa:16:3e:72:73:ea brd ff:ff:ff:ff:ff:ff
- inet 10.0.1.120/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid_lft 315359994sec preferred_lft 315359994sec
- inet6 fe80::f816:3eff:fe72:73ea/64 scope link noprefixroute valid_lft forever preferred_lft forever
- b. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The IPv4 address of the primary network interface is 10.0.0.133, and its name is eth0.
- The IPv4 address of the extended network interface is 10.0.1.120, and its name is eth1.

[root@ecs-resource ~]# ifconfig eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.0.0.133 netmask 255.255.255.0 broadcast 10.0.0.255 inet6 fe80::f816:3eff:fe72:72f7 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:72:72:f7 txqueuelen 1000 (Ethernet) RX packets 50917 bytes 71068144 (67.7 MiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 4969 bytes 1123356 (1.0 MiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet **10.0.1.120** netmask 255.255.255.0 broadcast 10.0.1.255 inet6 fe80::f816:3eff:fe72:73ea prefixlen 64 scopeid 0x20<link> ether fa:16:3e:72:73:ea txqueuelen 1000 (Ethernet) RX packets 21 bytes 3190 (3.1 KiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 26 bytes 2934 (2.8 KiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

- c. Configure the **ifcfg** file of the primary network interface.
 - i. Run the following command to open the **ifcfg** file of the primary network interface:

vi /etc/sysconfig/network-scripts/ifcfg-<Primary-network-interfacename>

The name of the primary network interface is obtained in **4.b**. Example command:

vi /etc/sysconfig/network-scripts/ifcfg-eth0

- ii. Press **i** to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- d. Configure the **ifcfg** file of the extended network interface.
 - i. Run the following command to open the **ifcfg** file of the extended network interface:

vi /etc/sysconfig/network-scripts/ifcfg-<Extended-networkinterface-name>

The name of the extended network interface is obtained in **4.b**. Example command:

vi /etc/sysconfig/network-scripts/ifcfg-eth1

- ii. Press i to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Edit the /etc/sysconfig/network file.
 - i. Run the following command to open the **/etc/sysconfig/network** file:

vi /etc/sysconfig/network

- ii. Press i to enter the editing mode.
- iii. Add the following content to the end of the file: NETWORKING_IPV6="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- f. Run the following command to restart the network service for the configuration to take effect:

systemctl restart NetworkManager

g. Check whether the ECS has IPv6 addresses:

ip addr

In the following command output, each network interface has an additional **inet6** entry starting with **2407**, followed by an entry starting with **fe80**. This indicates that the ECS has IPv6 addresses. [root@ecs-resource ~]# ip addr 2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

- link/ether fa:16:3e:72:72:f7 brd ff:ff:ff:ff:ff:ff inet 10.0.0.133/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0 valid_lft 315359994sec preferred_lft 315359994sec inet6 2407:c080:802:1f85:918b:9039:41b2:24a8/128 scope global dynamic noprefixroute valid_lft 7194sec preferred_lft 7194sec inet6 fe80::f816:3eff:fe72:72f7/64 scope link noprefixroute valid_lft forever preferred_lft forever 3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000 link/ether fa:16:3e:72:73:ea brd ff:ff:ff:fff inet 10.0.1.120/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid_lft 315359994sec preferred_lft 315359994sec inet6 2407:c080:802:2107:ab85:/427:d20:2119/128 scope global dynamic noprefixroute
- inet6 2407:c080:802:2107:ab85:4a27:d20:2119/128 scope global dynamic noprefixroute valid_lft 7195sec preferred_lft 7195sec inet6 fe80::f816:3eff:fe72:73ea/64 scope link noprefixroute
- valid_lft forever preferred_lft forever
- h. Log in to the destination ECS and obtain an IPv6 address by performing operations from **4.a** to **4.g**.
- 5. Log in to the source ECS and check whether it can use its primary network interface to communicate with the destination ECS:

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

ping6 -I <*IPv6-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv6-address-of-the-network-interface-on-the-destination-ECS>*

Example command:

ping6 -I 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc

If information similar to the following is displayed, the source ECS can use its primary network interface to communicate with the destination ECS. [root@ecs-resource ~]# ping6 -1 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:1f85:918b:9039:41b2:24a8 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.283 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.212 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=3 ttl=64 time=0.122 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=3 ttl=64 time=0.122 ms 7C --- 2407:c080:802:2108:96ec:3c49:391a:5ebc ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2050ms

rtt min/avg/max/mdev = 0.122/0.205/0.283/0.065 ms

- 6. Create an automation script to configure both IPv4 and IPv6 routes:
 - a. Create the automation configuration script file **15-policy-route.sh**:

vi /etc/NetworkManager/dispatcher.d/15-policy-route.sh

- b. Press i to enter the editing mode.
- c. Add the automation script to the file. #!/bin/bash

```
if [ "$2" != "up" ]; then
    exit 0
fi
interface=$1
check_route_table() {
    local_route_table.
```

```
local route_table_id=$1
local ip_version=$2
```

```
local max_attempts=100
   local attempts=0
   while [ $attempts -lt $max_attempts ]; do
     output=$(ip -$ip_version route show table $route_table_id &>/dev/null)
     if [ $? -ne 0 ] || [ -z "$output" ]; then
        break
     else
        route_table_id=$((route_table_id + 1))
        attempts=$((attempts + 1))
     fi
   done
   echo $route_table_id
}
add_ipv4_route_table() {
   local DEFAULT_IPV4_ROUTE_TABLE_ID=1000
   v4_route_table_id=$(ip route show table all | grep -F "default via $gateway dev $interface
table" | awk '{print $NF}')
  if [ -z $v4_route_table_id ]; then
     interface_number=$(echo $interface | grep -o '[0-9]*$')
     if [ -z "$interface_number" ]; then
        v4_route_table_id="$DEFAULT_IPV4_ROUTE_TABLE_ID"
     else
        v4_route_table_id=$((DEFAULT_IPV4_ROUTE_TABLE_ID + interface_number))
     fi
     v4_route_table_id=$(check_route_table $v4_route_table_id 4)
     echo "add policy route for dev: $interface table: $v4_route_table_id subnet: $subnet
gateway: $gateway"
     ip route add default via $gateway dev $interface table $v4_route_table_id
     ip route add $subnet dev $interface table $v4_route_table_id
   fi
}
generate_ipv4_policy_route() {
   subnet=$(nmcli device show $interface | grep -F 'IP4.ROUTE' | grep -F 'nh = 0.0.0.0' | cut -d'=' -
f2 | cut -d',' -f1 | tr -d ' ' | head -n 1)
  gateway=$(ipcalc -i $subnet | awk '/HostMin/ {print $2}')
   add_ipv4_route_table
   nmcli device show $interface | grep -F 'IP4.ADDRESS' | while read -r line; do
     IP=$(echo $line | awk '{print $2}' | cut -d'/' -f1)
     if ip rule list | grep -F "$v4_route_table_id" | grep -q "$IP"; then
        echo "ip rule already exists for $IP with table $v4_route_table_id"
        continue
     fi
     echo "Adding rule for $IP on $interface with table $v4_route_table_id"
     ip rule add from $IP table $v4_route_table_id
   done
}
generate_ipv4_policy_route
add_ipv6_route_table() {
   local DEFAULT_IPV6_ROUTE_TABLE_ID=1100
   v6_route_table_id=$(ip -6 route show table all | grep -F "default via $v6_gateway dev
$interface table" | grep -oP "table \K\d+")
   if [ -z $v6_route_table_id ]; then
     interface_number=$(echo $interface | grep -o '[0-9]*$')
     if [ -z "$interface_number" ]; then
        v6_route_table_id=$DEFAULT_IPV6_ROUTE_TABLE_ID
     else
        v6_route_table_id=$((DEFAULT_IPV6_ROUTE_TABLE_ID + interface_number))
```

```
fi
      v6_route_table_id=$(check_route_table $v6_route_table_id 6)
      echo "add policy route for dev: $interface table: $v6 route table id v6 subnet: $v6 subnet
v6 gateway: $v6 gateway"
      ip -6 route add default via $v6_gateway dev $interface table $v6_route_table_id
     ip -6 route add $v6_subnet dev $interface table $v6_route_table_id
   fi
}
generate_ipv6_policy_route() {
   IPV6INIT=$(awk -F= '/^IPV6INIT=/{gsub(/"/, ""); print $2}' "/etc/sysconfig/network-scripts/
ifcfg-$interface" | tr -d ' ' | tail -n 1)
   if [ "$IPV6INIT" != "yes" ]; then
      echo "$interface IPV6INIT is not set to yes, skip IPv6 policy route configurate."
      return 0
   fi
   for ((x=0; x<10; x++)); do
      if (ip address show $interface | grep -F "inet6" | grep -F "dynamic"); then
        v6_subnet=$(nmcli device show $interface | grep -F 'IP6.ROUTE' | grep -F 'nh = ::'| grep -
v 'dst = fe80' | grep -v '/128' | cut -d'=' -f2 | cut -d',' -f1 | tr -d ' ' | head -n 1)
        v6_gateway=$(ipcalc -6 -i $v6_subnet | awk '/HostMin/ {print $2}')1
         add ipv6 route table
         nmcli device show $interface | grep -F 'IP6.ADDRESS' | grep -v "/64" | while read -r line;
do
           IP=$(echo $line | awk '{print $2}' | cut -d'/' -f1)
           if ip -6 rule list | grep -F "$v6_route_table_id" | grep -q "$IP"; then
              echo "ip rule already exists for $IP with table $v6_route_table_id"
              continue
           fi
           echo "Adding rule for $IP on $interface with table $v6_route_table_id"
           ip -6 rule add from $IP table $v6_route_table_id
         done
        break
      fi
       sleep 1
   done
}
```

generate_ipv6_policy_route

- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Check the permissions on the automatic configuration script file **15**-**policy-route.sh**:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15policy-route.sh** is **-rw-r--r-**, which is different from that on other system files, for example, permission **-rwxr-xr-x** on the file **10-ifcfg-rh-routes.sh**. [root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient -rw-r-r-- 1 root root 2515 Jun 18 15:13 15-policy-route.sh -rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-dhcp -rwxr-xr-x. 1 root root 455 Feb 26 17:45 20-chrony-onoffline -rwxrwxr-x 1 root root 455 Feb 26 17:45 20-chrony-onoffline -rwxrwxr-x 1 root root 4096 Apr 18 17:12 no-wait.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d drwxr-xr-x. 2 root root 4096 Apr 18 17:12 pre-up.d

f. Grant permissions on the automation configuration script file 15-policyroute.sh:

sudo chown root:root /etc/NetworkManager/dispatcher.d/15-policyroute.sh

sudo chmod 755 /etc/NetworkManager/dispatcher.d/15-policy-route.sh

g. Check the permissions on the automatic configuration script file **15**-**policy-route.sh** again:

ll /etc/NetworkManager/dispatcher.d/

Information similar to the following is displayed. The permission on **15**policy-route.sh is -rwxr-xr-x. [root@ecs-resource ~]# ll /etc/NetworkManager/dispatcher.d/ total 36 -rwxr-xr-x. 1 root root 3840 Feb 26 17:45 10-ifcfg-rh-routes.sh -rwxr-xr-x. 1 root root 1062 Feb 26 17:45 11-dhclient -rwxr-xr-x. 1 root root 2019 Jun 18 12:21 15-policy-route.sh -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-dhcp -rwxr-xr-x. 1 root root 1412 Feb 26 17:45 20-chrony-onoffline -rwxrwxr-xr x. 1 root root 719 May 10 2019 hook-network-manager drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d drwxr-xr-x. 2 root root 4096 Feb 26 17:45 pre-down.d

h. Restart the network service to execute the **15-policy-route.sh** script:

systemctl restart NetworkManager

Before restarting the network service, ensure that services are not affected.

7. Check whether the policy-based IPv4 routes are added.

ip rule

ip route show table 1000

ip route show table 1001

table 1000 is the route table name of the primary network interface, and **table 1001** is the route table name of the extended network interface.

If information similar to the following is displayed, the policy-based routes have been added for the network interfaces with IPv4 addresses. [root@ecs-resource ~]# ip rule 0: from all lookup local 32764: from 10.0.1.120 lookup 1001 32765: from 10.0.0.133 lookup 1000 32766: from all lookup main 32767: from all lookup default [root@ecs-resource ~]# ip route show table 1000 default via 10.0.0.1 dev eth0 10.0.0.0/24 dev eth0 scope link [root@ecs-resource ~]# ip route show table 1001 default via 10.0.1.1 dev eth1 10.0.1.0/24 dev eth1 scope link

8. Check whether the source and destination ECSs can communicate with each other using IPv4 addresses.

ping -l <*IPv4-address-of-the-primary-network-interface-on-the-source-ECS>* <*IPv4-address-of-the-network-interface-on-the-destination-ECS>*

ping -l <*IPv4-address-of-the-extended-network-interface-on-the-source-ECS>* <*IPv4-address-of-the-network-interface-on-the-destination-ECS>*

Example commands:

ping -I 10.0.0.133 10.0.2.3

ping -I 10.0.1.120 10.0.2.3

If information similar to the following is displayed, both the IPv4 addresses of the source ECS can communicate with the destination ECS, indicating that the

IPv4 policy-based routes are successfully configured. [root@ecs-resource ~]# ping -I 10.0.0.133 10.0.2.3 PING 10.0.2.3 (10.0.2.3) from 10.0.0.133 : 56(84) bytes of data. 64 bytes from 10.0.2.3: icmp_seq=1 ttl=64 time=0.241 ms 64 bytes from 10.0.2.3: icmp_seq=2 ttl=64 time=0.198 ms ^C ---- 10.0.2.3 ping statistics ----

2 packets transmitted, 2 received, 0% packet loss, time 1064ms rtt min/avg/max/mdev = 0.198/0.219/0.241/0.021 ms [root@ecs-resource ~]# ping -I 10.0.1.120 10.0.2.3 PING 10.0.2.3 (10.0.2.3) from 10.0.1.120 : 56(84) bytes of data. 64 bytes from 10.0.2.3: icmp_seq=1 ttl=64 time=0.242 ms 64 bytes from 10.0.2.3: icmp_seq=2 ttl=64 time=0.184 ms ^C

--- 10.0.2.3 ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1062ms rtt min/avg/max/mdev = 0.184/0.213/0.242/0.029 ms

9. Check whether the policy-based IPv6 routes are added.

ip -6 rule

ip -6 route show table 1100

ip -6 route show table 1101

table 1100 is the route table name of the primary network interface, and **table 1101** is the route table name of the extended network interface.

If information similar to the following is displayed, the policy-based routes have been added for the network interfaces with IPv6 addresses. [root@ecs-resource ~]# ip -6 rule 0: from all lookup local 32764: from 2407:c080:802:2107:ab85:4a27:d20:2119 lookup 1101 32765: from 2407:c080:802:1f85:918b:9039:41b2:24a8 lookup 1100 32766: from all lookup main [root@ecs-resource ~]# ip -6 route show table 1100 2407:c080:802:1f85::/64 dev eth0 metric 1024 pref medium

default via 2407:c080:802:1f85::1 dev eth0 metric 1024 pref medium [root@ecs-resource ~]# ip -6 route show table 1101

2407:c080:802:2107::/64 dev eth1 metric 1024 pref medium

default via 2407:c080:802:2107::1 dev eth1 metric 1024 pref medium

10. Check whether the source and destination ECSs can communicate with each other using IPv6 addresses.

ping -6 -1 *<IPv6-address-of-the-primary-network-interface-on-the-source-ECS> <IPv6-address-of-the-network-interface-on-the-destination-ECS>*

ping -6 -1 *<IPv6-address-of-the-extended-network-interface-on-the-source-ECS> <IPv6-address-of-the-network-interface-on-the-destination-ECS>*

Example commands:

ping -6 -l 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc

ping -6 -l 2407:c080:802:2107:ab85:4a27:d20:2119 2407:c080:802:2108:96ec:3c49:391a:5ebc

If information similar to the following is displayed, both the IPv6 addresses of the source ECS can communicate with the destination ECS, indicating that the IPv6 policy-based routes are successfully configured.

[root@ecs-resource ~]# ping -6 -I 2407:c080:802:1f85:918b:9039:41b2:24a8 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:1f85:918b:9039:41b2:24a8 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.328 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.209 ms ^C --- 2407:c080:802:2108:96ec:3c49:391a:5ebc ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1016ms rtt min/avg/max/mdev = 0.209/0.268/0.328/0.059 ms [root@ecs-resource ~]# ping -6 -I 2407:c080:802:2107:ab85:4a27:d20:2119 2407:c080:802:2108:96ec:3c49:391a:5ebc PING 2407:c080:802:2108:96ec:3c49:391a:5ebc(2407:c080:802:2108:96ec:3c49:391a:5ebc) from 2407:c080:802:2107:ab85:4a27:d20:2119 : 56 data bytes 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=1 ttl=64 time=0.345 ms 64 bytes from 2407:c080:802:2108:96ec:3c49:391a:5ebc: icmp_seq=2 ttl=64 time=0.203 ms ^C --- 2407:c080:802:2108:96ec:3c49:391a:5ebc ping statistics ---2 packets transmitted, 2 received, 0% packet loss, time 1051ms rtt min/avg/max/mdev = 0.203/0.274/0.345/0.071 ms

5.3.2.4 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Network Interfaces (CentOS)

Scenarios

This section describes how to configure policy-based routes for a CentOS 8.0 64bit ECS with two network interfaces.

- IPv4: If IPv4 communication between cloud servers with multiple network interfaces is required, you need to configure IPv4 routes by referring to **Configuring IPv4 Policy-based Routes for a CentOS ECS**.
- IPv6: If IPv6 communication between cloud servers with multiple network interfaces is required, you need to configure IPv6 routes by referring to **Configuring IPv6 Policy-based Routes for a CentOS ECS**.
- IPv4/IPv6 dual stack: If both IPv4 and IPv6 communications between cloud servers with multiple network interfaces are required, you need to configure both IPv4 and IPv6 routes by referring to Configuring IPv4 Policy-based Routes for a CentOS ECS and Configuring IPv6 Policy-based Routes for a CentOS ECS.

For details about the background knowledge and networking of an ECS with two network interfaces, see **Overview**.

Configuring IPv4 Policy-based Routes for a CentOS ECS

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see Collecting ECS Network Information.

In this example, the network information of the ECS is shown in Table 5-13.

Тур е	Primary Network Interface	Extended Network Interface	
Sour ce	IP address: 10.0.0.115Subnet: 10.0.0/24Subnet gateway: 10.0.0.1	IP address: 10.0.1.183Subnet: 10.0.1.0/24Subnet gateway: 10.0.1.1	
Dest inati on	IP address: 10.0.2.12	N/A	

Table 5-13 Linux ECSs using IPv4 addresses (CentOS)

2. Log in to the source ECS.

For details, see **Logging In to an ECS**.

3. Check whether the source ECS can use its primary network interface to communicate with the destination ECS:

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

ping -l <*IP-address-of-the-primary-network-interface-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 network interface 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 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. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The primary network interface address is 10.0.0.115, and its name is eth0.
- The extended network interface address is 10.0.1.183, and its name is eth1.

[root@ecs-resource ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.0.0.115 netmask 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 255.255.0 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 source ECS.

Temporary routes are applied immediately but are lost after ECS restart. To avoid network interruptions, perform **6** to configure permanent routes instead.

- a. Configure policy-based routes for both the primary and extended network interfaces.
 - Primary network interface

ip route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip rule add from <network-interface-address> table <route-tablename>

Extended network interface

ip route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip rule add from <network-interface-address> table <route-tablename>

Configure the parameters as follows:

- Network interface name: Enter the name obtained in 4.
- Route table name: Name the route table with a number.
- Other network information: Enter the IP addresses collected in 1.

In this example, run the following commands:

- Primary network interface
 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
- Extended network interface
 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

D NOTE

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

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

ip rule

ip route show table *<route-table-name-of-the-primary-network-interface>*

ip route show table *<route-table-name-of-the-extended-network-interface>*

The route table name is the one configured 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 and destination ECSs can communicate with each other.

ping -l <*IP-address-of-the-primary-network-interface-on-the-source-ECS>* <*IP-address-of-the-destination-ECS>*

ping -l <*IP-address-of-the-extended-network-interface-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 network interfaces 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 permanent routes for the source 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:

```
# check eth0
for ((x=0; x<10; x++)); do
 if (ip addr show eth0 | grep -w 10.0.0.115 >/dev/null 2>&1); then
  break
 fi
 sleep 1
done
# 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
# check eth1
for ((x=0; x<10; x++)); do
 if (ip addr show eth1 | grep -w 10.0.1.183 >/dev/null 2>&1); then
  break
 fi
 sleep 1
done
# 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

The parameters are as follows:

- check eth0: Check for the presence of the IP address 10.0.0.115 on primary network interface eth0 every second up to 10 times.
- Add v4 routes for eth0: Add policy-based routes for the primary network interface. Set the value to be the same as that configured in 5.a.
- check eth1: Check for the presence of the IP address 10.0.1.183 on extended network interface eth1 every second up to 10 times.
- Add v4 routes for eth1: Add policy-based routes for the extended network interface. Set the value to be the same as that configured in 5.a.
- Add v4 routes for cloud-init: Set the Cloud-Init address to the same value as that in this example.
- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Run the following command to add execute permissions to the **/etc/ rc.local** file:

chmod +x /etc/rc.local

D 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

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.

Configuring IPv6 Policy-based Routes for a CentOS ECS

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see **Collecting ECS Network Information**.

Тур е	Primary Network Interface	Extended Network Interface
Sour ce	 IPv4 address: 10.0.0.102 IPv6 address: 2407:c080:1200:1dd8:859c:e5d 5:8b3d:a2d9 Subnet IPv6 CIDR block: 2407:c080:1200:1dd8::/64 Subnet IPv6 gateway: 2407:c080:1200:1dd8::1 	 IPv4 address: 10.0.1.191 IPv6 address: 2407:c080:1200:1a9c:7cc0:63b 5:8e65:4dd8 Subnet IPv6 CIDR block: 2407:c080:1200:1a9c::/64 Subnet IPv6 gateway: 2407:c080:1200:1a9c::1
Dest inati on	 IPv4 address: 10.0.2.3 IPv6 address: 2407:c080:1200:1dd9:16a7:fe7 a:8f71:7044 	N/A

Table 5-14 Linux ECSs	using IPv6	addresses	(CentOS)
-----------------------	------------	-----------	----------

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the ECSs have IPv6 enabled and have IPv6 addresses.

Perform this step for both the source and destination ECSs to ensure that the ECSs have IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

a. Run the following command to check whether the ECS has IPv6 addresses:

ip addr

In the following command output, eth0 and eth1 are the network interfaces of the ECS. Each network interface has one **inet6** entry starting with **fe80**. This indicates that the ECS has IPv6 enabled but does not have IPv6 addresses assigned. In this case, perform **3.b** to **3.g** to obtain IPv6 addresses.

[root@ecs-resource ~]# ip addr

2: **eth0**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

link/ether fa:16:3e:22:22:88 brd ff:ff:ff:ff:ff
inet 10.0.0.102/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0
valid_lft 107943256sec preferred_lft 107943256sec
inet6 fe80::f816:3eff:fe22:2288/64 scope link
valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group
default qlen 1000
link/ether fa:16:3e:22:23:e1 brd ff:ff:ff:ff:ff:ff
inet 10.0.1.191/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1
valid_lft 107943256sec preferred_lft 107943256sec
inet6 fe80::f816:3eff:fe22:23e1/64 scope link

valid_lft forever preferred_lft forever

b. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The primary network interface address is 10.0.0.102, and its name is eth0.
- The extended network interface address is 10.0.1.191, and its name is eth1.

[root@ecs-resource ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.0.0.102 netmask 255.255.255.0 broadcast 10.0.0.255
inet6 fe80::f816:3eff:fe22:2288 prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:22:22:88 txqueuelen 1000 (Ethernet)
RX packets 135116 bytes 132321802 (126.1 MiB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 60963 bytes 23201005 (22.1 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.0.1.191 netmask 255.255.255.0 broadcast 10.0.1.255 inet6 fe80::f816:3eff:fe22:23e1 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:22:23:e1 txqueuelen 1000 (Ethernet) RX packets 885 bytes 97676 (95.3 KiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 47 bytes 4478 (4.3 KiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

c. Configure the **ifcfg** file of the primary network interface.

i. Run the following command to open the **ifcfg** file of the primary network interface:

vi /etc/sysconfig/network-scripts/ifcfg-Primary network interface name

The name of the primary network interface is obtained in **3.b**.

In this example, run the following command:

vi /etc/sysconfig/network-scripts/ifcfg-eth0

- ii. Press **i** to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- d. Configure the **ifcfg** file of the extended network interface.
 - i. Run the following command to open the **ifcfg** file of the extended network interface:

vi /etc/sysconfig/network-scripts/ifcfg-Extended network interface name

The name of the extended network interface is obtained in **3.b**.

In this example, run the following command:

vi /etc/sysconfig/network-scripts/ifcfg-eth1

- ii. Press i to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Edit the **/etc/sysconfig/network** file.
 - i. Run the following command to open the **/etc/sysconfig/network** file:

vi /etc/sysconfig/network

- ii. Press **i** to enter the editing mode.
- iii. Add the following content to the end of the file: NETWORKING_IPV6="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- f. Run the following command to restart the network service for the configuration to take effect:

systemctl restart NetworkManager

g. Run the following command to check whether the ECS has IPv6 addresses:

ip addr

In the following command output, each network interface has an additional **inet6** entry starting with **2407**, followed by an entry starting with **fe80**. This indicates that the ECS has IPv6 addresses. [root@ecs-resource ~]# ip addr

2: **eth0**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

link/ether fa:16:3e:22:22:88 brd ff:ff:ff:ff:ff:ff

inet 10.0.0.102/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0
valid_lft 107999994sec preferred_lft 107999994sec
inet6 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9/128 scope global dynamic noprefixroute
valid_lft 7195sec preferred_lft 7195sec

inet6 fe80::f816:3eff:fe22:2288/64 scope link noprefixroute

valid_lft forever preferred_lft forever 3: **eth1**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

- link/ether fa:16:3e:22:23:e1 brd ff:ff:ff:ff:ff:ff
- inet 10.0.1.191/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid lft 107999994sec preferred lft 107999994sec
- inet6 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8/128 scope global dynamic noprefixroute valid_lft 7198sec preferred_lft 7198sec
- inet6 fe80::f816:3eff:fe22:23e1/64 scope link noprefixroute
- valid_lft forever preferred_lft forever
- h. Log in to the destination ECS and obtain an IPv6 address by performing operations from **3.a** to **3.g**.
- 4. Log in to the source ECS and check whether it can use its primary network interface to communicate with the destination ECS:

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

ping6 -I <*IP-address-of-the-primary-network-interface-on-the-source-ECS* <*IP-address-of-the-destination-ECS*>

In this example, run the following command:

ping6 -I 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

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

[root@ecs-resource ~]# ping6 -l 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 : 56 data bytes 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.635 ms

64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.320 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.287 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=4 ttl=64 time=0.193 ms ^C

--- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics ---4 packets transmitted, 4 received, 0% packet loss, time 3074ms rtt min/avg/max/mdev = 0.193/0.358/0.635/0.167 ms

5. Log in to the source ECS and configure temporary routes for the ECS.

Temporary routes are applied immediately but are lost after ECS restart. To avoid network interruptions, perform **6** to configure permanent routes instead.

- a. Configure policy-based routes for both the primary and extended network interfaces.
 - Primary network interface

ip -6 route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip -6 route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip -6 rule add from <network-interface-address> table <route-tablename>

Extended network interface

ip -6 route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip -6 route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip -6 rule add from <network-interface-address> table <route-tablename>

Configure the parameters as follows:

- Network interface name: Enter the name obtained in 3.b.
- Route table name: Name the route table with a number.
- Other network information: Enter the IP addresses collected in 1.

In this example, run the following commands:

Primary network interface

ip -6 route add default via 2407:c080:1200:1dd8::1 dev eth0 table 10

ip -6 route add 2407:c080:1200:1dd8::/64 dev eth0 table 10 ip -6 rule add from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 table 10

Extended network interface

ip -6 route add default via 2407:c080:1200:1a9c::1 dev eth1 table 20

ip -6 route add 2407:c080:1200:1a9c::/64 dev eth1 table 20 ip -6 rule add from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 table 20

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

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

ip -6 rule

ip -6 route show table *<route-table-name-of-the-primary-network-interface>*

ip -6 route show table *<route-table-name-of-the-extended-network-interface>*

The route table name is the one configured in **5.a**.

In this example, run the following commands:

ip -6 rule

ip -6 route show table 10

ip -6 route show table 20

If information similar to the following is displayed, the policy-based routes have been added. [root@ecs-resource ~]# ip -6 rule 0: from all lookup local 32764: from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 lookup 20 32765: from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 lookup 10 32766: from all lookup main [root@ecs-resource ~]# ip -6 route show table 10 2407:c080:1200:1dd8::/64 dev eth0 metric 1024 pref medium default via 2407:c080:1200:1dd8::1 dev eth0 metric 1024 pref medium [root@ecs-resource ~]# ip -6 route show table 20 2407:c080:1200:1a9c::/64 dev eth1 metric 1024 pref medium default via 2407:c080:1200:1a9c::1 dev eth1 metric 1024 pref medium

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

ping -6 -l <*IP-address-of-the-primary-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

ping -6 -l *<IP-address-of-the-extended-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

In this example, run the following commands:

ping -6 -l 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

ping -6 -l 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

If information similar to the following is displayed, both the network interfaces of the source ECS can communicate with the destination ECS. [root@ecs-resource ~]# ping -6 -I 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 : 56 data bytes

64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.770 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.295 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.245 ms ^C

--- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics ---

3 packets transmitted, 3 received, 0% packet loss, time 2080ms

rtt min/avg/max/mdev = 0.245/0.436/0.770/0.237 ms

[root@ecs-resource ~]# ping -6 -I 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 : 56 data bytes

64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.922 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.307 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.174 ms ^C

--- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics --- 3 packets transmitted, 3 received, 0% packet loss, time 2059ms

- rtt min/avg/max/mdev = 0.174/0.467/0.922/0.326 ms
- 6. Configure permanent routes for the source 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: # check eth0 for ((x=0; x<10; x++)); do if (ip addr show eth0 | grep -w 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 >/dev/null 2>&1); then

break fi sleep 1 done # Add v6 routes for eth0 ip -6 route flush table 10 ip -6 route add default via 2407:c080:1200:1dd8::1 dev eth0 table 10 ip -6 route add 2407:c080:1200:1dd8::/64 dev eth0 table 10 ip -6 rule add from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 table 10 # check eth1 for ((x=0; x<10; x++)); do if (ip addr show eth1 | grep -w 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 >/dev/null 2>&1); then break fi sleep 1 done # Add v6 routes for eth1 ip -6 route flush table 20 ip -6 route add default via 2407:c080:1200:1a9c::1 dev eth1 table 20 ip -6 route add 2407:c080:1200:1a9c::/64 dev eth1 table 20 ip -6 rule add from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 table 20

The parameters are as follows:

- check eth0: Check whether primary network interface eth0 has obtained an IPv6 address (2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9) every second up to 10 times.
- Add v6 routes for eth0: Add policy-based routes for the primary network interface. Set the value to be the same as that configured in 5.a.
- check eth1: Check whether extended network interface eth1 has obtained an IPv6 address (2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8) every second up to 10 times.
- Add v6 routes for eth1: Add policy-based routes for the extended network interface. Set the value to be the same as that configured in 5.a.
- 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

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.

5.3.2.5 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Linux ECS with Multiple Network Interfaces (Ubuntu)

Scenarios

This section describes how to configure policy-based routes for an Ubuntu 22.04 server 64-bit ECS with two network interfaces.

- IPv4: If IPv4 communication between cloud servers with multiple network interfaces is required, you need to configure IPv4 routes by referring to Configuring Policy-based Routes for Ubuntu ECSs Using IPv4 Addresses.
- IPv6: If IPv6 communication between cloud servers with multiple network interfaces is required, you need to configure IPv6 routes by referring to Configuring IPv6 Policy-based Routes for an Ubuntu ECS.
- IPv4/IPv6 dual stack: If both IPv4 and IPv6 communications between cloud servers with multiple network interfaces are required, you need to configure both IPv4 and IPv6 routes by referring to Configuring Policy-based Routes for Ubuntu ECSs Using IPv4 Addresses and Configuring IPv6 Policy-based Routes for an Ubuntu ECS.

For details about the background knowledge and networking of an ECS with two network interfaces, see **Overview**.

Configuring Policy-based Routes for Ubuntu ECSs Using IPv4 Addresses

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see Collecting ECS Network Information.

In this example, the network information of the ECS is shown in Table 5-15.

Туре	Primary Network Interface	Extended Network Interface
Sour	• IPv4 address: 10.0.0.138	• IPv4 address: 10.0.1.25
ce	 Subnet IPv4 CIDR block: 10.0.0.0/24 	 Subnet IPv4 CIDR block: 10.0.1.0/24
	• Subnet IPv4 gateway: 10.0.0.1	• Subnet IPv4 gateway: 10.0.1.1

Table 5-15 Linux ECSs using	IPv4 addresses	(Ubuntu)
-----------------------------	----------------	----------

Туре	Primary Network Interface	Extended Network Interface
Dest inati on	IPv4 address: 10.0.2.146	N/A

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the source ECS can use its primary network interface to communicate with the destination ECS:

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

ping -I <*IP-address-of-the-primary-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

In this example, run the following command:

ping -I 10.0.0.138 10.0.2.146

If information similar to the following is displayed, the source ECS can use its primary network interface to communicate with the destination ECS. root@ecs-s:~# ping -I 10.0.0.138 10.0.2.146 PING 10.0.2.146 (10.0.2.146) from 10.0.0.138 : 56(84) bytes of data. 64 bytes from 10.0.2.146: icmp_seq=1 ttl=64 time=0.247 ms 64 bytes from 10.0.2.146: icmp_seq=2 ttl=64 time=0.194 ms 64 bytes from 10.0.2.146: icmp_seq=3 ttl=64 time=0.190 ms ^C --- 10.0.2.146 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2049ms rtt min/avg/max/mdev = 0.190/0.210/0.247/0.025 ms

4. Query the network interface names of the source ECS:

ip addr

Search for the network interface names based on IP addresses.

- The primary network interface address is 10.0.0.138, and its name is eth0.
- The extended network interface address is 10.0.1.25, and its name is eth1.

root@ecs-s:~# ip addr

```
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000
```

```
link/ether fa:16:3e:22:22:ac brd ff:ff:ff:ff:ff
altname enp0s3
altname ens3
inet 10.0.0.138/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0
valid_lft 107999167sec preferred_lft 107999167sec
inet6 fe80::f816:3eff:fe22:22ac/64 scope link
valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen
1000
link/ether fa:16:3e:22:23:3b brd ff:ff:ff:ff:ff
altname enp4s1
inet 10.0.1.25/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1
valid_lft 107999167sec preferred_lft 107999167sec
inet6 fe80::f816:3eff:fe22:233b/64 scope link
valid_lft forever preferred_lft forever
```

5. Configure temporary routes for the source ECS.

Temporary routes are applied immediately but are lost after ECS restart. To avoid network interruptions, perform **6** to configure permanent routes instead.

- a. Configure policy-based routes for both the primary and extended network interfaces.
 - Primary network interface

ip route add default via <subnet-gateway> dev <network-interfacename> table <route-table-name>

ip route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip rule add from <network-interface-address> table <route-tablename>

• Extended network interface

ip route add default via <subnet-gateway> dev <network-interfacename> table <route-table-name>

ip route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip rule add from <network-interface-address> table <route-tablename>

Configure the parameters as follows:

- Network interface name: Enter the name obtained in 4.
- Route table name: Name the route table with a number.
- Other network information: Enter the IP addresses collected in **1**.

In this example, run the following commands:

Primary network interface

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 10.0.0.138 table 10

Extended network interface
 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 10.0.1.25 table 20

NOTE

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

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

ip rule

ip route show table *<route-table-name-of-the-primary-network-interface>*

ip route show table *<route-table-name-of-the-extended-network-interface>*

The route table name is the one configured 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-s:~# ip rule 0: from all lookup local 32764: from 10.0.1.25 lookup 20 32765: from 10.0.0.138 lookup 10 32766: from all lookup main 32767: from all lookup default root@ecs-s:~# ip route show table 10 default via 10.0.0.1 dev eth0 10.0.0./24 dev eth0 scope link root@ecs-s:~# 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 and destination ECSs can communicate with each other.

ping -l *<IP-address-of-the-primary-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

ping -l <*IP-address-of-the-extended-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

In this example, run the following commands:

ping -I 10.0.0.138 10.0.2.146

ping -I 10.0.1.25 10.0.2.146

If information similar to the following is displayed, both the network interfaces of the source ECS can communicate with the destination ECS. root@ecs-s:~# ping -I 10.0.0.138 10.0.2.146 PING 10.0.2.146 (10.0.2.146) from 10.0.0.138 : 56(84) bytes of data. 64 bytes from 10.0.2.146: icmp_seq=1 ttl=64 time=0.258 ms 64 bytes from 10.0.2.146: icmp_seq=2 ttl=64 time=0.242 ms 64 bytes from 10.0.2.146: icmp_seq=3 ttl=64 time=0.165 ms ^C --- 10.0.2.146 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2039ms rtt min/avg/max/mdev = 0.165/0.221/0.258/0.040 ms root@ecs-s:~# ping -I 10.0.1.25 10.0.2.146 PING 10.0.2.146 (10.0.2.146) from 10.0.1.25 : 56(84) bytes of data. 64 bytes from 10.0.2.146: icmp_seq=1 ttl=64 time=0.498 ms 64 bytes from 10.0.2.146: icmp_seq=2 ttl=64 time=0.427 ms 64 bytes from 10.0.2.146: icmp_seq=3 ttl=64 time=0.185 ms ^C --- 10.0.2.146 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2031ms rtt min/avg/max/mdev = 0.185/0.370/0.498/0.133 ms

- 6. Configure permanent routes for the source ECS.
 - a. Run the following command to add **network-routes.service** to the systemd service:

vi /etc/systemd/system/network-routes.service

- b. Press i to enter the editing mode.
- c. Add the following content to the end of the file:

Description=Network Routes Configuration After=network.target

[Service]

Type=oneshot

RemainAfterExit=yes ExecStart=/bin/bash -c 'for((x=0; x<10; x++)); do [[(ip addr show eth0 | grep -w 10.0.138 >/dev/null 2>&1 && echo 1)]] && break; sleep 1; done; 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 10.0.0.138 table 10; for((x=0; x<10; x++)); do [[<math>(ip addr show eth1 | grep -w 10.0.1.25 >/dev/null 2>&1 && echo 1)]] && break; sleep 1; done; 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 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 10.0.1.25 table 20; ip rule add to 169.254.169.254 table main'

[Install]

WantedBy=multi-user.target

The parameters are as follows:

- for loop: Check whether eth0 or eth1 has obtained an IPv4 address (eth0: 10.0.0.138; eth1: 10.0.1.25) every second up to 10 times.
- **ip route flush table** *route-table-name*: Running this command will delete existing routes in the specified route table. This prevents new routes from being affected.
- Policy-based routes of the primary network interface: Set it to the same value as that in 5.a.
- Policy-based routes of the extended network interface: Set it to the same value as that in 5.a.
- **ip rule add to 169.254.169.254 table main**: Set the Cloud-Init address to the same value as that in this example.
- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Run the following commands to reload the systemd configuration and start the service:

systemctl daemon-reload

systemctl enable network-routes.service

If information similar to the following is displayed, the service is started: root@ecs-s:~# systemctl daemon-reload root@ecs-s:~# systemctl enable network-routes.service Created symlink /etc/systemd/system/multi-user.target.wants/network-routes.service \rightarrow /etc/ systemd/system/network-routes.service.

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

reboot

▲ CAUTION

Policy-based routes added to the **network-routes.service** file only work after the source 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.

Configuring IPv6 Policy-based Routes for an Ubuntu ECS

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see **Collecting ECS Network Information**.

In this example, the network information of the ECS is shown in Table 5-16.

Тур е	Primary Network Interface	Extended Network Interface
Sour ce	 IPv4 address: 10.0.0.138 IPv6 address: 2407:c080:1200:1dd8:1473:49d b:22d7:13c7 Subnet IPv6 CIDR block: 2407:c080:1200:1dd8::/64 Subnet IPv6 gateway: 2407:c080:1200:1dd8::1 	 IPv4 address: 10.0.1.25 IPv6 address: 2407:c080:1200:1a9c:691e:fffe :7e22:12c4 Subnet IPv6 CIDR block: 2407:c080:1200:1a9c::/64 Subnet IPv6 gateway: 2407:c080:1200:1a9c::1
Dest inati on	 IPv4 address: 10.0.2.146 IPv6 address: 2407:c080:1200:1dd9:f5e1:94d 1:2822:dede 	N/A

 Table 5-16 Linux ECSs using IPv6 addresses (Ubuntu)

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Check whether the ECSs have IPv6 enabled and have IPv6 addresses.

Perform this step for both the source and destination ECSs to ensure that the ECSs have IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

a. Run the following command to check whether the source ECS has IPv6 addresses:

ip addr

In the following command output, eth0 and eth1 are the network interfaces of the ECS. Each network interface has one **inet6** entry starting with **fe80**. This indicates that the ECS has IPv6 enabled but does not have IPv6 addresses assigned. In this case, perform **3.b** to **3.h** to obtain IPv6 addresses.

root@ecs-s:~# ip addr

2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group defaul</broadcast,multicast,up,lower_up>	t
qlen 1000	
link/ether fa:16:3e:22:22:ac brd ff:ff:ff:ff:ff:ff	
altname enp0s3	
altname ens3	
inet 10.0.0.138/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0	
valid_lft 107999781sec preferred_lft 107999781sec	
inet6 fe80::f816:3eff:fe22:22ac/64 scope link	
valid_lft forever preferred_lft forever	
3: eth1: <broadcast,multicast,up,lower_up> mtu 1500 qdisc mq state UP group defaul</broadcast,multicast,up,lower_up>	t
qlen 1000	
link/ether fa:16:3e:22:23:3b brd ff:ff:ff:ff:ff:ff	
altname enp4s1	
inet 10.0.1.25/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1	
valid_lft 107999781sec preferred_lft 107999781sec	
inet6 fe80::f816:3eff:fe22:233b/64 scope link	
valid_lft forever preferred_lft forever	

b. Query the network interface names of the source ECS:

ifconfig

Search for the network interface names based on IP addresses.

- The primary network interface address is 10.0.0.138, and its name is eth0.
- The extended network interface address is 10.0.1.25, and its name is eth1.

```
root@ecs-s:~# ifconfig
```

```
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.0.0.138 netmask 255.255.255.0 broadcast 10.0.0.255
inet6 fe80::f816:3eff:fe22:22ac prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:22:22:ac txqueuelen 1000 (Ethernet)
RX packets 863 bytes 269089 (269.0 KB)
RX errors 0 dropped 0 overruns 0 frame 0
TX packets 1117 bytes 359807 (359.8 KB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.0.1.25 netmask 255.255.255.0 broadcast 10.0.1.255
```

```
inet6 fe80::f816:3eff:fe22:233b prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:22:23:3b txqueuelen 1000 (Ethernet)
RX packets 10 bytes 1358 (1.3 KB)
RX errors 0 dropped 0 overruns 0 frame 0
```

TX packets 10 bytes 973 (973.0 B)

TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

- ...
- c. Configure the **01-netcfg.yaml** file.
 - Run the following command to access /etc/netplan/:
 cd /etc/netplan
 - ii. Run the following command to open the **01-netcfg.yaml** file:vi **01-netcfg.yaml**
 - iii. Press i to enter the editing mode.
 - iv. Add **dhcp6: true** to the network interfaces for which you want to assign IPv6 addresses as follows:

In this example, the primary network interface name queried in **3.b** is eth0, and the extended network interface name is eth1. network: version: 2
```
renderer: NetworkManager
ethernets:
eth0:
dhcp4: true
eth1:
dhcp6: true
eth2:
dhcp6: true
eth2:
dhcp4: true
eth3:
dhcp4: true
eth4:
dhcp4: true
```

- v. Press **ESC** to exit and enter **:wq!** to save the configuration.
- d. Run the following commands to change the permissions on the **01**-**netcfg.yaml** file and ensure that only the file owner has the read and write permissions:

chmod 600 /etc/netplan/01-netcfg.yaml

chown root:root /etc/netplan/01-netcfg.yaml

- e. Run the following command to apply the modification: **netplan apply**
- f. Configure the NetworkManager.conf file.
 - Run the following command to open the NetworkManager.conf file: vi /etc/NetworkManager/NetworkManager.conf
 - ii. Press i to enter the editing mode.
 - iii. Add **dhcp=dhclient** to the file as follows:

[main] plugins=ifupdown,keyfile **dhcp=dhclient**

[ifupdown] managed=true

[device] wifi.scan-rand-mac-address=no

- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- g. Run the following command to restart the network service for the configuration to take effect:

systemctl restart NetworkManager

h. Run the following command to check whether the source ECS has IPv6 addresses:

ip addr

In the following command output, each network interface has an additional **inet6** entry starting with **2407**, followed by an entry starting with **fe80**. This indicates that the ECS has IPv6 addresses. root@ecs-s:/etc/netplan# ip addr

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000 link/ether fa:16:3e:22:22:ac brd ff:ff:ff:ff:ff altname enp0s3 altname ens3 inet 10.0.0.138/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0 valid_lft 107999982sec preferred_lft 107999982sec inet6 2407:c080:1200:1dd8:1473:49db:22d7:13c7/128 scope global dynamic noprefixroute

- valid_lft 7182sec preferred_lft 7182sec
- inet6 fe80::f816:3eff:fe22:22ac/64 scope link noprefixroute
- valid_lft forever preferred_lft forever
- 3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:22:23:3b brd ff:ff:ff:ff:ff:ff

altname enp4s1

inet 10.0.1.25/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1

valid_lft 107999982sec preferred_lft 107999982sec

inet6 2407:c080:1200:1a9c:691e:fffe:7e22:12c4/128 scope global dynamic noprefixroute valid_lft 7182sec preferred_lft 7182sec

inet6 fe80::f816:3eff:fe22:233b/64 scope link noprefixroute

valid_lft forever preferred_lft forever

i. Log in to the destination ECS and obtain an IPv6 address by performing operations from **3.a** to **3.h**.

In the following command output, **eth0** has an additional **inet6** entry starting with **2407**, followed by an entry starting with **fe80**. This indicates that the ECS has been assigned an IPv6 address. root@ecs-d:/etc/netplan# ip addr

2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP group default qlen 1000

link/ether fa:16:3e:22:24:b4 brd ff:ff:ff:ff:ff:ff
altname enp0s3
altname ens3
inet 10.0.2.146/24 brd 10.0.2.255 scope global dynamic noprefixroute eth0
valid_lft 107999994sec preferred_lft 107999994sec
inet6 2407:c080:1200:1dd9:f5e1:94d1:2822:dede/128 scope global dynamic noprefixroute
valid_lft 7195sec preferred_lft 7195sec
inet6 fe80::f816:3eff:fe22:24b4/64 scope link noprefixroute
valid_lft forever preferred lft forever

4. Log in to the source ECS and check whether it can use its primary network interface to communicate with the destination ECS:

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

ping6 -I <*IP-address-of-the-primary-network-interface-on-the-source-ECS>* <*IP-address-of-the-destination-ECS>*

In this example, run the following command:

ping6 -I 2407:c080:1200:1dd8:1473:49db:22d7:13c7 2407:c080:1200:1dd9:f5e1:94d1:2822:dede

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

root@ecs-s:/etc/netplan# ping6 -I 2407:c080:1200:1dd8:1473:49db:22d7:13c7 2407:c080:1200:1dd9:f5e1:94d1:2822:dede PING 2407:c080:1200:1dd9:f5e1:94d1:2822:dede

PING 2407:c080:1200:1dd9:f5e1:94d1:2822:dede(2407:c080:1200:1dd9:f5e1:94d1:2822:dede) from 2407:c080:1200:1dd8:1473:49db:22d7:13c7 : 56 data bytes

64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=1 ttl=64 time=0.244 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=2 ttl=64 time=0.212 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=3 ttl=64 time=0.169 ms ^C

--- 2407:c080:1200:1dd9:f5e1:94d1:2822:dede ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2033ms rtt min/avg/max/mdev = 0.169/0.208/0.244/0.030 ms

5. Log in to the source ECS and configure temporary routes for the ECS.

Temporary routes are applied immediately but are lost after ECS restart. To avoid network interruptions, perform **6** to configure permanent routes instead.

- a. Configure policy-based routes for both the primary and extended network interfaces.
 - Primary network interface

ip -6 route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip -6 route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip -6 rule add from *<network-interface-address>* **table** *<route-table-name>*

Extended network interface

ip -6 route add default via *<subnet-gateway>* **dev** *<network-interface-name>* **table** *<route-table-name>*

ip -6 route add <subnet-CIDR-block> dev <network-interface-name>
table <route-table-name>

ip -6 rule add from <network-interface-address> table <route-tablename>

Configure the parameters as follows:

- Network interface name: Enter the name obtained in 3.b.
- Route table name: Name the route table with a number.
- Other network information: Enter the IP addresses collected in 1.

In this example, run the following commands:

Primary network interface

ip -6 route add default via 2407:c080:1200:1dd8::1 dev eth0 table 10

ip -6 route add 2407:c080:1200:1dd8::/64 dev eth0 table 10 ip -6 rule add from 2407:c080:1200:1dd8:1473:49db:22d7:13c7 table 10

Extended network interface

ip -6 route add default via 2407:c080:1200:1a9c::1 dev eth1 table 20

ip -6 route add 2407:c080:1200:1a9c::/64 dev eth1 table 20 ip -6 rule add from 2407:c080:1200:1a9c:691e:fffe:7e22:12c4 table 20

NOTE

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

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

ip -6 rule

ip -6 route show table *<route-table-name-of-the-primary-network-interface>*

ip -6 route show table *<route-table-name-of-the-extended-network-interface>*

The route table name is the one configured in **5.a**.

In this example, run the following commands:

ip -6 rule

ip -6 route show table 10

ip -6 route show table 20

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

root@ecs-s:/etc/netplan# ip -6 rule 0: from all lookup local 32764: from 2407:c080:1200:1a9c:691e:fffe:7e22:12c4 lookup 20 32765: from 2407:c080:1200:1d8:1473:49db:22d7:13c7 lookup 10 32766: from all lookup main root@ecs-s:/etc/netplan# ip -6 route show table 10 2407:c080:1200:1d8::/64 dev eth0 metric 1024 pref medium default via 2407:c080:1200:1d8::1 dev eth0 metric 1024 pref medium root@ecs-s:/etc/netplan# ip -6 route show table 20 2407:c080:1200:1a9c::/64 dev eth1 metric 1024 pref medium default via 2407:c080:1200:1a9c::1 dev eth1 metric 1024 pref medium

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

ping -6 -l <*IP-address-of-the-primary-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

ping -6 -1 *<IP-address-of-the-extended-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

In this example, run the following commands:

ping6 -I 2407:c080:1200:1dd8:1473:49db:22d7:13c7 2407:c080:1200:1dd9:f5e1:94d1:2822:dede

ping6 -I 2407:c080:1200:1a9c:691e:fffe:7e22:12c4 2407:c080:1200:1dd9:f5e1:94d1:2822:dede

If information similar to the following is displayed, both the network interfaces of the source ECS can communicate with the destination ECS. root@ecs-s:/etc/netplan# ping6 -I 2407:c080:1200:1dd8:1473:49db:22d7:13c7 2407:c080:1200:1dd9:f5e1:94d1:2822:dede

PING 2407:c080:1200:1dd9:f5e1:94d1:2822:dede(2407:c080:1200:1dd9:f5e1:94d1:2822:dede) from 2407:c080:1200:1dd8:1473:49db:22d7:13c7 : 56 data bytes

64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=1 ttl=64 time=0.260 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=2 ttl=64 time=0.248 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=3 ttl=64 time=0.165 ms ^C

--- 2407:c080:1200:1dd9:f5e1:94d1:2822:dede ping statistics ---

3 packets transmitted, 3 received, 0% packet loss, time 2043ms

rtt min/avg/max/mdev = 0.165/0.224/0.260/0.042 ms

root@ecs-s:/etc/netplan# ping6 -I 2407:c080:1200:1a9c:691e:fffe:7e22:12c4

2407:c080:1200:1dd9:f5e1:94d1:2822:dede

PING 2407:c080:1200:1dd9:f5e1:94d1:2822:dede(2407:c080:1200:1dd9:f5e1:94d1:2822:dede) from 2407:c080:1200:1a9c:691e:fffe:7e22:12c4 : 56 data bytes

64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=1 ttl=64 time=0.592 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=2 ttl=64 time=0.208 ms 64 bytes from 2407:c080:1200:1dd9:f5e1:94d1:2822:dede: icmp_seq=3 ttl=64 time=0.162 ms ^C

--- 2407:c080:1200:1dd9:f5e1:94d1:2822:dede ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2031ms rtt min/avg/max/mdev = 0.162/0.320/0.592/0.192 ms

- 6. Configure permanent routes for the source ECS.
 - a. Run the following command to create the **network-routes6.service** file for the systemd service:

vi /etc/systemd/system/network-routes6.service

- b. Press **i** to enter the editing mode.
- c. Add the following content to the end of the file:

[Unit] Description=Network Routes Configuration After=network.target

[Service] Type=oneshot RemainAfterExit=yes

[Install] WantedBy=multi-user.target

The parameters are as follows:

- for loop: Check whether eth0 or eth1 has obtained an IPv6 address (eth0: 2407:c080:1200:1dd8:1473:49db:22d7:13c7; eth1:2407:c080:1200:1a9c:691e:fffe:7e22:12c4) every second up to 10 times.
- **ip route flush table** *route-table-name*: Running this command will delete existing routes in the specified route table. This prevents new routes from being affected.
- Policy-based routes of the primary network interface: Set it to the same value as that in 5.a.
- Policy-based routes of the extended network interface: Set it to the same value as that in 5.a.
- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Run the following commands to reload the systemd configuration and start the service:

systemctl daemon-reload

systemctl enable network-routes6.service

If information similar to the following is displayed, the service is started: root@ecs-s:/etc/netplan# systemctl daemon-reload root@ecs-s:/etc/netplan# systemctl enable network-routes6.service Created symlink /etc/systemd/system/multi-user.target.wants/network-routes6.service \rightarrow /etc/ systemd/system/network-routes6.service.

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

reboot

Policy-based routes added to the **network-routes6.service** file only work after the source 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.

5.3.2.6 Manually Configuring IPv4 and IPv6 Policy-based Routes for a Windows ECS with Multiple Network Interfaces

Scenarios

This section describes how to configure policy-based routes for a Windows Server 2012 64-bit ECS with two network interfaces.

- IPv4: If IPv4 communication between cloud servers with multiple network interfaces is required, you need to configure IPv4 routes by referring to **Configuring IPv4 Policy-based Routes for a Windows ECS**.
- IPv6: If IPv6 communication between cloud servers with multiple network interfaces is required, you need to configure IPv6 routes by referring to Configuring IPv6 Policy-based Routes for a Windows ECS.
- IPv4/IPv6 dual stack: If both IPv4 and IPv6 communications between cloud servers with multiple network interfaces are required, you need to configure both IPv4 and IPv6 routes by referring to Configuring IPv4 Policy-based Routes for a Windows ECS and Configuring IPv6 Policy-based Routes for a Windows ECS.

For details about the background knowledge and networking of an ECS with two network interfaces, see **Overview**.

Configuring IPv4 Policy-based Routes for a Windows ECS

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see **Collecting ECS Network Information**.

In this example, the network information of the ECS is shown in Table 5-17.

Тур е	Primary Network Interface	Extended Network Interface
Sour	• IPv4 address: 10.0.0.59	• IPv4 address: 10.0.1.104
ce	• Subnet IPv4 gateway: 10.0.0.1	• Subnet IPv4 gateway: 10.0.1.1

Table 5-17 Windows ECSs using IPv4 addresses

Тур е	Primary Network Interface	Extended Network Interface
Dest inati on	IPv4 address: 10.0.2.12	N/A

2. Log in to the source ECS.

For details, see **Logging In to an ECS**.

3. Check whether the source ECS can use its primary network interface to communicate with the destination ECS:

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

ping -S <*IP-address-of-the-primary-network-interface-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 network interface to communicate with the destination ECS.

C:\Users\A	dministrator	>ping -S	10.0.0.59	10.0.2.12
Pinging 10 Reply from Reply from Reply from Reply from Reply from	.0.2.12 from 10.0.2.12: 10.0.2.12: 10.0.2.12: 10.0.2.12: 10.0.2.12:	10.0.0.9 bytes=32 bytes=32 bytes=32 bytes=32 bytes=32	59 with 32 time<1ms time<1ms time<1ms time<1ms	bytes of data: TIL=64 TIL=64 TIL=64 TIL=64 TIL=64

4. Configure a policy-based route for the extended network interface.

route add -p 0.0.0.0 mask 0.0.0.0 <subnet-gateway-of-the-extendednetwork-interface> metric <route-priority>

Configure the parameters as follows:

- **0.0.0.0/0**: Default route. Do not change it.
- Subnet gateway of the extended network interface: Enter the IP address collected in 1.
- Route priority: Set its value to 261. The priority of the extended network interface must be lower than that of the primary network interface. 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 network interface already has policy-based routes and you do not need to configure again.
- If the ECS has multiple extended network interfaces, configure policy-based routes for all extended network interfaces one by one.
- 5. Check whether the policy-based route is added.

route print

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

C:\Users\Administrator>route print				
Interface List 19fa 16 3e fc 7b 76Red Hat VirtIO Ethernet Adapter #3 14fa 16 3e 5d 3e b6Red Hat VirtIO Ethernet Adapter 1Software Loopback Interface 1 1600 00 00 00 00 00 00 e0 Microsoft ISATAP Adapter #2				
IPv4 Route Table				
Active Routes:				
Network Destinatio	on Netmask	Gateway	Interface	Metric
0.0.0.0	0.0.0.0	10.0.1.1	10.0.1.104	1 266
0.0.0.0	0.0.0.0	10.0.0.1	10.0.0.59	7 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	1 261
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	1 261
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	0n-link	10.0.1.104	1 261
Persistent Routes: Network Address 0.0.0.0	Netmask 0.0.0.0	Gateway Address 10.0.1.1	Metric 261	
IPv6 Route Table				
Active Routes:				
If Metwic Network	Destination	Gateway		
1 306 ::1/128	Beschlacion	0n-link		
14 261 fe80::/	64	0n-link		
19 261 fe80::/	64	0n-link		
19 261 fe80::1	97h:3504:e05:5a4d	/128		
11 101 1000-11		0n-link		
14 261 fe80::e	115:8e6a:5dcc:671	5/128		
		On-link		
1 306 ff00::/	/8	On-link		
14 261 ff00::/	/8	On-link		
19 261 ff00::/	′ 8	On-link		
Persistent Routes:				

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

ping -S <IP-address-of-the-primary-network-interface-on-the-source-ECS>
<IP-address-of-the-destination-ECS>

ping -S <IP-address-of-the-extended-network-interface-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 network interfaces 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 TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=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: bytes=32 time<1ms TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=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

Configuring IPv6 Policy-based Routes for a Windows ECS

1. Collect the ECS network interface information required for configuring policybased routes.

For details, see Collecting ECS Network Information.

In this example, the network information of the ECS is shown in Table 5-18.

Тур е	Primary Network Interface	Extended Network Interface
Sour ce	IPv6 address: 2407:c080:802:aba:6788:fb94:d71f :8deb	IPv6 address: 2407:c080:802:be6:71c8:42e0:d44 e:eeb4
Dest inati on	IPv6 address: 2407:c080:802:be7:c2e6:d99c:b68 5:c6c8	N/A

Table 5-18 Windows ECSs using IPv6 addresses

2. Log in to the source ECS.

For details, see Logging In to an ECS.

3. Run the following command to check whether the ECS has IPv6 enabled and has IPv6 addresses:

ipconfig

If information similar to the following is displayed, each network interface has an IPv6 address starting with 2407, which indicates that the ECS has IPv6 addresses.

C:\Users\Administrator>ipconfig	
Windows IP Configuration	
Ethernet adapter Ethernet 4:	
Connection-specific DNS Suffix .	: openstacklocal
IPv6 Address	: 2407:c080:802:be6:ec23:ec4:c886:cc1
LINK-local IPvb Hadress	: fe80::883b:ab73:1003:a170%17
IPv4 Address	: 192.168.1.12
Subnet Mask	: 255.255.255.0
Default Gateway	: fe80::f816:3eff:fe3e:1e1e%19
Sthernet adapter Ethernet 2: Connection-specific DNS Suffix . IPv6 Address	: openstacklocal : 2407:c080:802:aba:8999:5e61:e19:cf7e : te80::180d:t3b5:27ac:2acb%14 : 192.168.0.57 : 255.255.255.0
Default Gateway	: fe80::f816:3eff:fede:c837%14 192.168.0.1
Default Gateway	: fe80::f816:3eff:fede:c837%14 192.168.0.1 :
Default Gateway	: fe80::f816:3eff:fede:c837%14 192.168.0.1 : : Media disconnected : openstacklocal

Perform this step for both the source and destination ECSs to ensure that the ECSs have IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

4. Check whether the source and destination ECSs can communicate with each other.

ping -6 -S *<IP-address-of-the-primary-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

ping -6 -S *<IP-address-of-the-extended-network-interface-on-the-source-ECS> <IP-address-of-the-destination-ECS>*

In this example, run the following commands:

ping -6 -S 2407:c080:802:aba:8999:5e61:e19:cf7e 2407:c080:802:be7:c2e6:d99c:b685:c6c8

ping -6 -S 2407:c080:802:be6:ec23:ec4:c886:cc1 2407:c080:802:be7:c2e6:d99c:b685:c6c8

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

C:\Users\Administrator>ping -6 -8 2407:c080:802:aba:8999:5e61:e19:cf7e 2407:c08 :802:be7:c2e6:d99c:b685:c6c8
Pinging 2407:c080:802:be7:c2e6:d99c:b685:c6c8 from 2407:c080:802:aba:8999:5e61: 19:cf7e with 32 bytes of data: Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms
Ping statistics for 2407:c080:802:be7:c2e6:d99c:b685:c6c8: Packets: Sent = 4, Received = 4, Lost = 0 (0% loss), Approximate round trip times in milli-seconds: Minimum = Oms, Maximum = Oms, Average = Oms
C:\Users\Administrator>ping -6 -8 2407:c080:802:be6:ec23:ec4:c886:cc1 2407:c080 802:be7:c2e6:d99c:b685:c6c8
Pinging 2407:c080:802:be7:c2e6:d99c:b685:c6c8 from 2407:c080:802:be6:ec23:ec4:c 86:cc1 with 32 bytes of data: Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time=3ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms
Ping statistics for 2407:c080:802:be7:c2e6:d99c:b685:c6c8: Packets: Sent = 4, Received = 4, Lost = 0 (0% loss), Approximate round trip times in milli-seconds: Minimum = 0ms, Maximum = 3ms, Average = 0ms

NOTE

ECSs in this example run Windows Server 2012 (64-bit). You do not need to configure policy-based routes for these ECSs because both the network interfaces of such an ECS can communicate with others using IPv6.

6 Access Control

6.1 Access Control Overview

A VPC is your private network on the cloud. You can configure access control 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 Access Control Options

Table 6-1 provides differences between access control options. You can select one or more as needed.

	Table 6-1	Differences	between	access	control	options
--	-----------	-------------	---------	--------	---------	---------

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 or Optional	Mandatory . Instances must be added to at least one security group.	Optional . You can determine whether to associate a subnet with a network ACL based on service requirements.
Stateful	Stateful . The response traffic of inbound and outbound requests is allowed to flow to and out of an instance.	Stateful . The response traffic of inbound and outbound requests is allowed to flow to and out of a subnet.
Action	Does not support Allow or Deny rules.	Supports both Allow and Deny rules.
Rule Packets	Packet filtering based on 3-tuple (protocol, port, and source/destination)	Packet filtering based on 5-tuple (protocol, source port, destination port, source, and destination)
Matching Rule	 If an instance is associated with multiple security groups that have multiple rules: 1. Rules are first matched based on how early a security group is associated with an instance. The security group associated earlier takes precedence over those associated later. 2. Rules are then matched by priority in that security group. A rule with a smaller value has a higher priority. 3. Deny rules take precedence over allow rules if the rules have the same priority. 	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.

ltem	Security Group	Network ACL
Usage	 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: 	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 the network ACL, and enable the network ACL. The network ACL then protects the associated subnets and instances in the subnets.
	 Add or remove the instance to or from a security group on the security group console. 	
	 Add or remove the instance to or from a security group on the instance console. 	

6.2 Security Group

6.2.1 Security Group and Security Group Rule Overview

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. For details, see **Default Security Groups**. 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

Security groups are 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 first 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.
 - Protocol: the protocol that is used to match traffic. The protocol can be TCP, UDP, ICMP, or GRE.
 - **TCP** is ideal for applications that require reliable connections and high data integrity, such as remote login, web browsing, email, and file transfers.
 - **UDP** is ideal for applications demanding high speed and low latency, such as online gaming and video meetings.
 - ICMP is used to communicate data transmission problems. For example, the ping command can be used to check the connectivity between network devices, error reports can be generated for O&M, and diagnosis information can be transmitted for network analysis and optimization.
 - **GRE** is widely used and enables communications between networks using different protocols by encapsulating one protocol within another, for example, encapsulating IP packets.
 - **Port**: the destination port range that is used to match traffic. 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. Changes to inbound rules take effect immediately for existing connections. Changes to outbound security group rules do not affect existing persistent connections and take effect only for new connections.

If you add, delete, or update rules in a security group, or add or remove instances in a security group, the details are as follows:

- For connections established by inbound traffic, the system automatically clears the connection tracking entries corresponding to the existing persistent connections based on Table 6-2. That is, the connection tracking entries are expired in advance. Then, the system re-establishes connections to match the new inbound rules of the security group.
 - If the security group rules allow the traffic of the connections, the connections can be established and network communication is not affected.
 - If the security group rules deny the traffic of the connections, the connections cannot be established again and the network communication will be interrupted.

Scenario	Clearing Policy	
Adding instances to security group A	 Clear inbound connection tracking entries of the instances newly added to security group A. If an inbound rule of another security group (for example, security group B) denies access from security group A, clear inbound connection tracking entries of all instances in security group B. 	
Removing instances from security group A	 Clear inbound connection tracking entries of all instances in security group A. If an inbound rule of another security group (for example, security group B) allows access from security group A, clear inbound connection tracking entries of all instances in security group B. 	
Adding rules to security group A	If a Deny rule is added in the inbound or outbound direction, clear inbound connection tracking entries of all instances in security group A.	
Deleting rules from security group A	If an Allow rule is deleted in the inbound direction, clear inbound connection tracking entries of all instances in security group A.	
Modifying rules in security group AIf the priority, action, protocol, port, or source address of a rule is modified in the inbound direction, clear inbound connection tracking ent of all instances in security group A.		

Table 6-2 Scenarios and policies for clearing connection tracking entries

Scenario	Clearing Policy
Changing IP address entries in an IP address group	If an IP address group is associated with an inbound rule in security group A, deleting or adding an IP address entry from or to the IP address group will clear inbound connection tracking entries of all instances in security group A.

 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.

D NOTE

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 6-3** 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 6-3 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.

Direction	Actio n	Туре	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

 Table 6-3 Security group rules

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 first 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

Security Group Examples

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 the ECSs in Sg-B SSH port 22. That is, ECSs in Subnet-A can remotely log 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. That is, ECSs in Subnet-B can communicate with each other.
- 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

If you use an intermediate network instance to forward traffic between instances in different subnets, setting the source of the inbound rule to the security group associated with the peer instance does not allow the instances to communicate with each other. To enable communications, set the source to the private IP address or subnet CIDR block of the intermediate network instance. For example, to connect ECSs in **Subnet-A** and **Subnet-B** in **Figure 6-5**, set the source of the inbound rule to the 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 pair, 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 network, 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-A01** and **Subnet-B01** 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 6-4 Security group configuration process description

N o.	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
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

Constraints on Using Security Groups

- 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**, of the rules with **Source** set to **IP address group**, and of the rules with inconsecutive ports, cannot exceed 120. If there are both IPv4 and IPv6 security group rules, up to 120 rules can be added for each type.

The limits on outbound security group rules are the same as those on inbound rules.

For example, to add inbound IPv4 rules to a security group (Sg-A), you can refer to **Table 6-5** for rules that meet the restrictions. Of these rules, rule A02 uses inconsecutive ports (TCP: 22,25,27) and security group Sg-B as the source. In this case, only one quota is occupied.

Rule No.	Action	Туре	Protocol & Port	Source
Rule A01	Allow	IPv4	All	Current security group: Sg-A
Rule A02	Allow	IPv4	TCP: 22,25,27	Another security group: Sg-B
Rule A03	Allow	IPv4	TCP: 80-82	IP address group: ipGroup-A
Rule A04	Allow	IPv4	TCP: 22-24,25	IP address: 192.168.0.0/16

 Table 6-5 Inbound security group rules

• Traffic from load balancers is not restricted by network ACL and security group rules if:

Transfer Client IP Address is enabled for the listeners 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 Working?**

6.2.2 Default Security Groups

When creating an instance, you must associate it with a security group. If there is no security group available, a default security group will be created and associated with the instance. Note the following when using default security group:

- The default security group name is **default**. It is recommended that you do not change the name of the default security group in order to distinguish it from custom security groups.
- A default security group cannot be deleted, but you can modify its rules or add rules to it.
- The default security group allows instances in the security group to communicate with each other and 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 accordingly.

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

Directi on	Ac tio n	Тур е	Proto col & Port	Source/ Destination	Description
Inboun d	All ow	IPv 4	All	Source: default security group (default)	Allows IPv4 instances in the security group to communicate with each other using any protocol over any port.
Inboun d	All ow	IPv 6	All	Source: default security group (default)	Allows IPv6 instances in the security group to communicate with each other using any protocol over any port.
Outbo und	All ow	IPv 4	All	Destination: 0.0.0.0/0	Allows all traffic from the instances in the security group to any IPv4 address over any port.
Outbo und	All ow	IPv 6	All	Destination: : :/0	Allows all traffic from the instances in the security group to any IPv6 address over any port.

Table 6	-6 Ru	les in	the	default	security	aroun
	- u i\u	10.5 111	unc	ucraute	Security	gioup

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 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 creating instances, such as cloud servers, containers, and databases, in a VPC subnet, you can use the default security group or create a custom security group, and then add inbound and outbound rules 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 Files over FTP
- Setting Up a Website on an ECS to Provide Internet-Accessible Services
- 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 Only Specific External Websites

NOTE

If a configured security group rule does not take effect, locate the cause by referring to **Why Are My Security Group Rules Not Working?** or **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, while allowing instances in it to communicate with each other.

If required, you can add inbound rules to allow specific traffic to access 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 6-7**.

Direc tion	Pri ori ty	Ac ti on	Ty pe	Prot ocol & Port	Destinatio n	Description
Outb ound	1	All o w	IPv 4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outb ound	1	All o w	IPv 6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

Table 6-7 Default outbound rules in a security group

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 6-8**.
- To remotely log in to a Windows ECS using RDP, enable port 3389. For details, see **Table 6-9**.

Direction	Priori ty	Action	Туре	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 22	IP address: 0.0.0.0/0

Table 6-8 Remotely logging in to a Linux ECS using SSH

Table 6-9 Remotely logging in to a Windows ECS using RDP

Direction	Priori ty	Action	Туре	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3389	IP address: 0.0.0.0/0

If the source is set to 0.0.0.0/0, all external IP addresses are allowed to remotely log in to the ECS. To ensure network security and prevent service interruptions caused by network intrusions, set the source to a trusted IP address. For details, see **Table 6-10**.

Table 6-10 Remote	y logging i	n to an ECS using	g a trusted IP address
-------------------	-------------	-------------------	------------------------

ECS Type	Direc tion	Pri ori ty	Actio n	Туре	Protocol & Port	Source
Linux ECS	Inbou nd	1	Allow	IPv4	TCP: 22	IP address: 192.168.0.0/24
Window s ECS	Inbou nd	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 Files over FTP

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 over FTP, you need to enable FTP ports 20 and 21.

Table 6-11 Remotely connecting to an ECS from any server to upload ordownload files over FTP

Direction	Priori ty	Action	Туре	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 0.0.0.0/0

- If the source is set to 0.0.0/0, all external IP addresses are allowed to remotely log in to the ECS to upload or download files. To ensure network security and prevent service interruptions caused by network intrusions, set the source to a trusted IP address. For details, see Table 6-12.
- You must first install the FTP server program on the ECSs and then check whether ports 20 and 21 are working properly.

Direction	Priori ty	Action	Туре	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 20-21	IP address: 192.168.0.0/24

Table 6-12 Remotely connecting to an ECS from a trusted server to upload or download files

Setting Up a Website on an ECS to Provide Internet-Accessible Services

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

Table 6-13 Setting up a website on an ECS to provide internet-accessible services

Direction	Priori ty	Action	Туре	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.

Direction	Priori ty	Action	Туре	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

 Table 6-14 Using ping command to verify network connectivity

Enabling Communications Between Instances in Different Security Groups

Instances in the same VPC but in different security groups cannot communicate with each other. If you want ECSs in security group **sg-A** to access MySQL databases in security group **sg-B**, you need to add an inbound rule to security group **sg-B** to allow access from ECSs in security group **sg-A**.

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

Direction	Priori ty	Action	Туре	Protocol & Port	Source
Inbound	1	Allow	IPv4	TCP: 3306	Security group: sg-A

NOTE

If you use an intermediate network instance to forward traffic between instances in different subnets, setting the source of the inbound rule to the security group associated with the peer instance does not allow the instances to communicate with each other. To enable communications, set the source to the private IP address or subnet CIDR block of the intermediate network instance. For example, to connect ECSs in **Subnet-A** and **Subnet-B** as described in the second security group example in **Security Group Examples**, set the source of the inbound rule to the virtual IP address.

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

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

Directio n	Prio rity	Acti on	Туре	Protocol & Port	Source	Description
Inbound	1	Allo w	IPv4	TCP: 3306	Security group: sg- A	Allows the ECSs in security group sg-A to access the MySQL database.
Inbound	1	Allo w	IPv4	TCP: 1521	Security group: sg- B	Allows the ECSs in security group sg-B to access the Oracle database.

Table 6-16 Allowing	external	instances t	o access	the database	e deploved	on an ECS
Tuble o To Autowing	externat	motunees t	.0 466655	the databas	. acpioyea	

Directio n	Prio rity	Acti on	Туре	Protocol & Port	Source	Description
Inbound	1	Allo w	IPv4	TCP: 1433	IP address: 172.16.3.2 1/32	Allows the ECS whose private IP address is 172.16.3.21 to access the MS SQL database.
Inbound	1	Allo w	IPv4	TCP: 5432	IP address: 192.168.0. 0/24	Allows ECSs whose private IP addresses are in the 192.168.0.0/24 network to access the PostgreSQL database.
Inbound	1	Allo w	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.

Allowing ECSs to Access Only Specific External Websites

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

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

Dire ctio n	Prio rity	Ac tio n	Ty pe	Protoc ol & Port	Destinatio n	Description
Out bou nd	1	All ow	IP v4	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.
Out bou nd	1	All ow	IP v4	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.

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

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

Direc tion	Pri ori ty	Ac ti on	Ty pe	Prot ocol & Port	Destinatio n	Description
Outb ound	1	All o w	IPv 4	All	0.0.0.0/0	Allows the instances in the security group to access any IPv4 address over any port.
Outb ound	1	All o w	IPv 6	All	::/0	Allows the instances in the security group to access any IPv6 address over any port.

 Table 6-18 Default outbound rules in a security group

6.2.4 Common ECS Ports

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

Table 6-19 lists some high-risk ports that are blocked by default. Even if you have added a security group rule to allow access over these ports, traffic over these ports in restricted regions is still denied. In this case, do not use these high-risk ports for your services.

Table 6	5-19 Hig	h-risk	ports
---------	----------	--------	-------

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

Common Ports

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

Table 6-20 Common	ports use	d 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 Files over FTP .		
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	НТТР	Used to access websites over HTTP. For configuration examples, see Setting Up a Website on an ECS to Provide Internet-Accessible Services .		
110	POP3	Used to receive emails using Post Office Protocol version 3 (POP3).		
143	ΙΜΑΡ	Used to receive emails using Internet Message Access Protocol (IMAP).		
443	HTTPS	Used to access websites over HTTPS. For configuration examples, see Setting Up a Website on an ECS to Provide Internet-Accessible Services.		
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.		

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

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

Security Group Templates

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

Templa te	Direc tion	Protocol/ Port/ Type	Source/ Destina tion	Description	Scenario
General - purpose web server	Inbou nd	TCP: 22 (IPv4)	0.0.0.0/0	Allows all IPv4 addresses to access instances in the security group over port 22 (SSH) for remotely logging in to Linux instances.	• Remotely log in to an instance (such as an ECS) in a security group from
		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.	 an external network. Enable external servers to ping the instances in a security group to
		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.	 group to verify network connectivity. Use instances in a security group as web servers to provide website services accessible from the Internet.
		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.	

 Table 6-21
 Security group rules
Templa te	Direc tion	Protocol/ Port/ Type	Source/ Destina tion	Description	Scenario
	Outb ound	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	Inbou nd	All (IPv4) All (IPv6)	Current security group	Allows the instances in a security group to communicate with each other over a private network over any protocol and port.	Allowing any traffic to enter and leave a security group over any port may be risky.
		All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows any IP address to access the instances in a security group over any protocol and port.	
	Outb ound	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	Inbou nd	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.	

Templa te	Direc tion	Protocol/ Port/ Type	Source/ Destina tion	Description	Scenario
	Outb ound	All (IPv4) All (IPv6)	0.0.0.0/0 ::/0	Allows all traffic from the instances in the security group to external resources using any protocol.	

Procedure

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

Table 6-22 Parameter description

Parameter	Description	Example Value
Name	Mandatory	sg-AB
	The name of the security group. The name:	
	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.	
Enterprise	Mandatory	default
Project	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</i> <i>Management User Guide</i> .	

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 6-21 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 with source or destination set to IP addresses or the current security group. Rules with source or destination set to an IP address group or another security group will not be cloned.
- Only security group rules are cloned, but not the instances associated with the security group. After the clone is complete, you need to add the instances to the new security group. For details, see Adding an Instance to or Removing an Instance from a 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 \bigcirc in the upper left corner and select the desired region and project.
- 3. Click = in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The Virtual Private Cloud page is displayed.

 In the navigation pane on the left, choose Access Control > 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.

Constraints

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

Figure 6-10 Default security group

Network Console	Security Groups ③				Feedback Guick Links Create Security Group	,
Dashboard	Dalata				C	
Virtual Private Cloud 👻	Specify filter criteria.				Q	
Access Control 🔺	Name1D	Security Group Rules	Associated Instances Description	Enterprise Project	Operation	
Security Groups Network ACLs	59-60 e4ft	6	1	default	Manage Rule Manage Instance More +	
IP Address Groups	c3d5a383	4	0 Default security gr	oup default	Manage Rule Manage Instance Clone	

- If you want to delete a security group that is associated with instances, such as cloud servers, containers, and databases, you need to remove the instances from the security group first. For details, see Adding an Instance to or Removing an Instance from a Security Group.
- 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.

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

You can add a security group rule using any of the following methods:

- Adding Rules to a Security Group: You need to specify the action, priority, type, protocol, port, and source or destination of the security group rule as prompted.
- Fast-Adding Multiple Security Group Rules: You can quickly add rules with common ports and protocols for remote logins, ping tests, common web services, and database services.
- Allowing Common Ports with a Few Clicks: You can allow common ports with just a few clicks. This function can be used in the following scenarios:
 - Remote login to ECSs
 - Using the ping command to test ECS connectivity
 - ECSs functioning as web servers to provide website access services

Precautions

Before adding a security group rule, note the following:

- 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. For details about the rules, see What Are Security Group Rules?
- If an instance is associated with multiple security groups, the traffic matches security group rules by priority. For details about the matching sequence, see How Traffic Matches Security Group Rules.
- The number of rules in a security group is limited. Keep only the rules you need. For details, see **Constraints on Using Security Groups**.
- 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 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 6-23 shows the inbound rules.

Directi on	Prior ity	Actio n	Туре	Protocol & Port	Source/Destination
Inboun d	1	Allow	IPv4	All	Source: current security group (Sg-A)
Inboun d	1	Allow	IPv6	All	Source: current security group (Sg-A)

Table 6-23 Inbound rules for communication between instances

- 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 different VPCs.

Configuration Example

Before configuring security group rules, you need to plan access policies for instances in the security group.

- If an instance needs to provide services for external systems, add an inbound rule to allow external requests to the instance.
- If there are attacks to an instance from external networks, add an inbound rule to deny external requests that have security risks.
- If an instance needs to access the Internet, add an outbound rule to allow requests from the instance to the Internet.
- If you no longer need to control certain inbound or outbound traffic, you can delete the corresponding security group rules to simplify the rule configuration.

Security Group Examples shows more security group rule configuration examples.

Adding Security Group Rules

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 \equiv 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.

On the Inbound Rules tab, click Add Rule.
 The Add Inbound Rule dialog box is displayed.

7. Configure inbound rule parameters as prompted.

You can click \oplus to add more inbound rules.

Figure 6-11 Add Inbound Rule

Add Inbound	d Rule Learn	more about sec	urity group configuration.			
Some securi If you select	ity group rules will not IP address for Source	take effect for ECS e, you can enter mu	s with certain specifications. Learn mo tiple IP addresses in the same IP add	ress box. Each IP address represe	nts a different security <u>c</u>	jroup rule.
Security Group s	g-test-					
1	Action (?)	Iype IPv4 •	Protocols/TCP (Custo • 22	Source (?) IP address • 0.0.0.0/0 (2)	Description	Replicate Delete
			 Add Rule OK 	ncel		

	Table 6-24	Inbound	rule	parameter	descriptio
--	------------	---------	------	-----------	------------

Param eter	Description	Example Value
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default value is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	 The value can be Allow or Deny. If the Action is set to Allow, traffic is allowed to access the cloud servers in the security group over specified ports. If the Action is set to Deny, traffic is denied to access the cloud servers 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
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Protoc ol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	ТСР

Param eter	Description	Example Value
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.	22, 22-30
	Inbound rules control incoming traffic over specific ports to instances in the security group.	
	Enter ports in any of the following formats:	
	• 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 .	
Source	Source of the security group rule. The value can be IP address or Security group , to allow access from the IP addresses or the instances in the security group.	192.168.0.0 /24
	IP address	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 Any IP address: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	If the source is a security group, this rule will apply to all instances associated with the selected security group.	
Descrip tion	(Optional) Supplementary information about the security group rule.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The inbound rule list is displayed.

- On the Outbound Rules tab, click Add Rule.
 The Add Outbound Rule dialog box is displayed.
- 10. Configure outbound rule parameters as prompted.

You can click $\textcircled{\oplus}$ to add more outbound rules.

Figure 6-12 Add Outbound Rule

Add Outbound Rule Learn more abo	ut security group configuration.	×				
Some security group rules will not take effect for ECSs with certain specifications. Learn more If you select IP address for Destination, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.						
Security Group sg-test-zcy You can import multiple rules in a batch.						
Priority ? Action ? Type	Protocol & Port ? Destination ?	Description Operation				
1 Allow v IPv4	Protocols/All IP address I-65535 0.0.0.0/0	Replicate Delete				
Add Rule Cancel						

Table 6-25 Outbound rule parameter description

Param eter	Description	Example Value
Priority	The security group rule priority.	1
	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.	
Action	The value can be Allow or Deny .	Allow
	• 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 .	
Туре	Destination IP address version. You can select:	IPv4
	• IPv4	
	• IPv6	
Protoc ol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	ТСР

Param eter	Description	Example Value
	Destination port used to match traffic in a security group rule. The value can be from 1 to 65535.	22, 22-30
	Outbound rules control outgoing traffic over specific ports from instances in the security group.	
	Enter ports in any of the following formats:	
	• 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 .	
Destina tion	Destination of the security group rule. The value can be IP address or Security group , to allow access to the IP addresses or the instances in the security group.	0.0.0/0
	IP address	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 Any IP address: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Descrip tion	(Optional) Supplementary information about the security group rule.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The outbound rule list is displayed.

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 details, see Verifying Security Group Rules.

Fast-Adding Multiple Security Group Rules

- 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 VPC list 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.

- On the Inbound Rules tab, click Fast-Add Rule.
 The Fast-Add Inbound Rule dialog box is displayed.
- 7. Configure inbound rule parameters as prompted.

Figure 6-13 Fast-Add Inbound Rule

Fast-Add Inbound	d Rule Learn more about s	ecurity group configuration.			×
Some security group If you select IP addr different security group	p rules will not take effect for EC ess for Source, you can enter m pup rule.	Ss with certain specifications. I nultiple IP addresses in the sam	Learn more le IP address box. Each IF	9 address represents a	×
Security Group sg-tes	st				
* 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	4 🔻]			
		OK Cancel			

Table 6-26 Inbound rule parameter description

Param eter	Description	Example Value
Protoco ls and Ports	Common protocols and ports are provided for:Remote login and pingWeb serviceDatabase	SSH (22)
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4

Param eter	Description	Example Value
Source	Source of the security group rule. The value can be IP address or Security group to allow access from the IP addresses or the instances in the security group. You can specify:	192.168.0.0 /24
	 A single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6) 	
	 An IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	 Any IP address: 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.	
Action	The value can be Allow or Deny .	Allow
	• If the Action is set to Allow , traffic is allowed to access the cloud servers in the security group over specified ports.	
	• If the Action is set to Deny , traffic is denied to access the cloud servers 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 .	
Priority	The security group rule priority.	1
	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.	
Descrip tion	(Optional) Supplementary information about the security group rule.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

The inbound rule list is displayed 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 outbound rule parameters as prompted.

 \times

Figure 6-14 Fast-Add Outbound Rule

Fast-Add Outbound Rule Learn more about security group configuration.					;
Some security group rules will not take effect for ECSs with certain specifications. Learn more If you select IP address for Destination, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.				×	
Security Group sg-te	est				
* Protocols and Ports					
Remote Login an	d 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	√4 ▼				
Г	1				
		OK Cancel			

Table 6-27 Outbound rule parameter description

Param eter	Description	Example Value
Protoc ols and Ports	Common protocols and ports are provided for:Remote login and pingWeb serviceDatabase	SSH (22)
Туре	Destination IP address version. You can select:IPv4IPv6	IPv4

Param eter	Description	Example Value
Destin ation	Destination of the security group rule. The value can be IP address or Security group to allow access to the IP addresses or the instances in the security group. You can specify:	0.0.0.0/0
	• xxx.xxx.xxx.xxx/32 (IPv4 address)	
	• xxx.xxx.xxx.0/24 (IPv4 address range)	
	• 0.0.0.0/0 (any IPv4 address)	
	• sg-abc (security group)	
	 A single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/127 (IPv6) 	
	 An IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
	• Any IP address: 0.0.0.0/0 (IPv4); ::/0 (IPv6)	
	Security group: sg-abc	
Priority	The security group rule priority.	1
	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.	
Action	The value can be Allow or Deny .	Allow
	• 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 .	
Descrip tion	(Optional) Supplementary information about the security group rule.	N/A
	The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

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 details, see Verifying Security Group Rules.

Allowing Common Ports with a Few Clicks

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

The VPC list page is displayed.

- In the navigation pane on the left, choose Access Control > Security Groups. The security group list is displayed.
- 5. Locate the target security group and click its name.

The security group details page is displayed.

6. Click the **Inbound Rules** or **Outbound Rules** tab as required, and then click **Allow Common Ports**.

The Allow Common Ports page is displayed.

Table 6-28 describes the common ports that can be opened with a few clicks.

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.

 Table 6-28
 Common ports

Direction	Protocol & Port & Type	Source/ Destination	Description
	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.
Outboun d	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.

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

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 details, see Verifying Security Group Rules.

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 6-29** shows the rule.

Table 6-29	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

tcv

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

Figure 6-15 Command output for the Linux ECS

0 0 0.0.0:<mark>80</mark> 0.0.0.0:*

- Checking the port of a Windows server

LISTEN

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

netstat -an | findstr 80

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

Figure 6-16 Command output for the Windows ECS

TCP	0.0.0.0:80	0.0.0:0	LISTENING

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 ECS Ports**.

Procedure

- 1. Log in to the management console.
- 2. Click 🔍 in the upper left corner and select the desired region and project.
- 3. Click \equiv 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.

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

Figure 6-17 Fast-Add Inbound Rule

Fast-Add Inbound Rule Learn more about security group configuration.								×
e	Some security group rules will not take effect for ECSs with certain specifications. Learn more If you select IP address for Source, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.							
S	ecurity Group	sg-test						
★ Pi	rotocols and Po	rts						
	Remote Lo	gin and Pin	g:					
	SSH (22)		RDP (3389)	FTP (20-21)		Telnet (23)	ICMP (All)	
	Web Servic	ce:						
	HTTP (80)		HTTPS (443)	HTTP_ALT	(8080)			
	Database:							
	MySQL (33	06)	MS SQL (1433)	PostgreSQL	. (5432)	Oracle (1521)	Redis (6379)	
* Tj	/pe	IPv4		•				
		-						•
				ОК	Cancel			

Table 6-30 Inbound rule parameter description

Param eter	Description	Example Value		
Protoco ls and Ports	Common protocols and ports are provided for:Remote login and pingWeb servicesDatabases	SSH (22)		
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4		
Source	Source of the security group rule. The value can be a single IP address or a security group to allow access from the IP addresses or the instances in the security group. You can specify:	192.168.0.0 /24		
	 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.			

Param eter	Description	Example Value
Action	The value can be Allow or Deny .	Allow
	• If the Action is set to Allow , traffic is allowed to access the cloud servers in the security group over specified ports.	
	• If the Action is set to Deny , traffic is denied to access the cloud servers 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 .	
Priority	Security group rule priority.	1
	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.	
Descrip tion	(Optional) Supplementary information about the security group rule.	-
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

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

 \times

Figure 6-18 Fast-Add Outbound Rule

Fast-Add Outbound Rule Learn more about security group configuration.							
Son If yo diffe	Some security group rules will not take effect for ECSs with certain specifications. Learn more If you select IP address for Destination, you can enter multiple IP addresses in the same IP address box. Each IP address represents a different security group rule.						
Security	Group sg-test						
* Protocol:	Is and Ports						
Re	emote Login and Pi	ing:					
ss ss	SH (22)	RDP (3389)	FTP (20-21)	Telnet (23)	ICMP (All)		
We	eb Service:						
HT	TTP (80)	HTTPS (443)	HTTP_ALT (8080)				
Da	atabase:						
My	ySQL (3306)	MS SQL (1433)	PostgreSQL (5432)	Oracle (1521)	Redis (6379)		
★ Туре	IPv4	•					
	r	ı					
			OK Cancel				

Table 6-31 Outbound rule parameter description

Param eter	Description	Example Value
Protoc ols and Ports	Common protocols and ports are provided for:Remote login and pingWeb servicesDatabases	SSH (22)
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4

Param eter	Description	Example Value
Destin ation	Destination of the security group rule. The value can be a single IP address or a security group to allow access to the IP addresses or the instances in the security group. You can specify:	0.0.0.0/0
	 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	
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 .	Allow
	• 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 .	
Descrip tion	(Optional) Supplementary information about the security group rule.	-
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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

6.2.6.3 Modifying a Security Group Rule

Scenarios

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

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

Notes and Constraints

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

- The inbound rules in **Table 6-32** ensure that instances in the security group can communicate with each other. Do not modify these rules.
- The outbound rules in **Table 6-32** 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.

Direction	Action	Туре	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

 Table 6-32
 Security group rules

Procedure

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

The Virtual Private Cloud page is displayed.

- 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**.

6.2.6.4 Replicating a Security Group Rule

Scenarios

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

Procedure

- 1. Log in to the management console.
- 2. Click 🔍 in the upper left corner and select the desired region and project.
- 3. Click \equiv 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.
- 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**.

6.2.6.5 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 6-33 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.

Rule	Directi on	Priority	Action	Туре	Protocol & Port	Destination
Rule A	Inboun d	1	Allow	IPv4	TCP: 22	0.0.0.0/0
Rule B	Inboun d	5	Deny	IPv4	TCP: 22	0.0.0/0

Table 6-33 Rules with different actions

- If you want to import rules of the security group in one region to another under the same account, rules with **Source** or **Destination** set to an IP address group or another security group cannot be imported.
- If you want to import rules of the security group in one account to another account, rules with **Source** or **Destination** set to an IP address group or another security group cannot be imported.

Procedure

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

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > 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 6-34 describes the parameters in the template for importing rules.

Table 6-34	Template	parameters
------------	----------	------------

Param eter	Description	Example Value
Directi on	 The direction in which the security group rule takes effect. 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
Priorit y	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. If the Action is set to Allow, traffic is allowed to access the cloud servers in the security group over specified ports. If the Action is set to Deny, traffic is denied to access the cloud servers 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
Protoc ol & Port	The network protocol used to match traffic in a security group rule. The protocol can be All , TCP , UDP , GRE , or ICMP .	ТСР
	 Destination port used to match traffic in a security group rule. The value can be from 1 to 65535. Inbound rules control incoming traffic over specific ports to instances in the security group. Outbound rules control outgoing traffic over specific ports from instances in the security group. Enter ports in any of the following formats: 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, 22-30

Param eter	Description	Example Value
Туре	Source IP address version. You can select: • IPv4 • IPv6	IPv4
Source	 Source of the security group rule. The value can be IP address or Security group, to allow access from the IP addresses or the instances in the security group. IP address Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) Any IP address: 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[96a8a 93f-XXX- d7872990c 314]
Destin ation	Destination of the security group rule. The value can be IP address or Security group , to allow access to the IP addresses or the instances in the security group.	sg- test[96a8a 93f-XXX- d7872990c 314]
Descri ption	(Optional) Supplementary information about the security group rule. The security group rule description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-
Last Modifi ed	The time when the security group was modified.	-

6.2.6.6 Deleting One or More Security Group Rules

Scenarios

If you no longer need one or more security group rules to control the traffic to and from the instances in a security group, you can delete them.

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 specific traffic, the security group denies all traffic to or from the instances in it.

• The inbound rules in **Table 6-35** ensure that instances in the security group can communicate with each other. Do not delete these rules.

• The outbound rules in **Table 6-35** allow instances in the security group to access external networks. If you delete these rules, the instances in the security group cannot access external networks.

Direction	Action	Туре	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

 Table 6-35
 Security group rules

Deleting a Security Group Rule

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

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Security Groups. The security group list is displayed.
- 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.
- Locate the target rule and click **Delete** in the **Operation** column. A confirmation dialog box is displayed.
- 8. Click **OK**.

Deleting Multiple Security Group Rules

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

The Virtual Private Cloud page is displayed.

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

- 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, select the target security group rules and click **Delete** up above the upper left corner of the list.
 - A confirmation dialog box is displayed.
- 8. Click OK.

6.2.6.7 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 6-36 lists the operations on security group rules supported by CTS.

	Table 6-36 Ope	rations on s	ecurity grou	p rules supp	orted 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 6-37** that is added to security group **Sg-A**.

 Table 6-37
 The new security group rule

Dire ctio n	Act ion	Туре	Protoc ol & Port	Source	Last Modified
Inbo und	Allo w	IPv4	TCP: 23	10.0.0.0/1 6	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 6-19 The trace list for new security group rules

Tra	ce List 💿							Procedure for Using CTS	🕞 Go to Old Editio
	Export ~								
	Last 1 day	V Q Traci	Name: createSecurity-group-rule \times	Add filter				×	00
	Trace Name	Trace Source	Resource Type	Resource Name	Resource ID	Operator	Trace Status	Operation Time	
	createSecurity-group-rule	VPC	security-group-rules			NOCET-BERT	o normal	Jun 19, 2024 15:49:54	GMT+08:00
	createSecurity-group-rule	VPC	security-group-rules		-	NCOTABLE .	o normal	Jun 19, 2024 15:48:31	GMT+08:00
	createSecurity-group-rule	VPC	security-group-rules	-	-	100073887	O normal	Jun 19, 2024 10:46:07	GMT+08:00

2. 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 6-38** provides the detailed information about the operation, including operator ID and details about the security group rules.

NOTE

The trace details in **Table 6-38** are only for your reference. The actual information may vary.

Table 0-30 The trace details for the new security group rule	Table 6-38	The trace details	for the new	security group	rule
---	------------	-------------------	-------------	----------------	------

Example Command Output	Description
"source_ip": "124.71.XX.146",	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 .

Example Command Output	Description		
"user": { "access_key_id": "HSTA205XXXXXC4MHAE", "account_id": "3c24f6f885294XXXX93ce075fbd", "user_name": "cts-test-01", "domain": { "name": "cts_test"	Account of the operator who performs the operation. Key parameters are described as follows:		
"name": "cts-test", "id": "3c24f6f885294XXXX93ce075fbd" }, "name": "cts-test-01", "principal_is_root_user": "false", "id": "a26ee7e7224XXXXXe4a28a9ce503",	 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 3c24f6f885294XXXX93ce 075fbd. 		
	 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 a26ee7e7224XXXXe4a28 a9ce503. 		
	For details about more parameters of CTS traces, see the response parameter description in Trace Structure .		

Example Command Output	Description
<pre>"response": "{\"request_id \":\"8d2d1111cafaXXX9b49d53e2da38f \",\"security_group_rules\":[{\"id\":\"b6acda6e-0976- XXXX-82bc-a8093cbd591d\",\"project_id \":\"15289aca74eXXXa37dea0315d99\",\"security_group _id\":\"3730d371-3111-4ace-XXXX- 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_jp_refix \":\"\"\"10.0.0/16\",\"multiport \":\"23\",\"remote_address_group_id\":null,\"action \":\"allow\",\"priority\":1}]}",</pre>	 Details about the security group rule in response. Key parameters are described as follows: 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 10.0.0.0/16 .
	• 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, extended network interfaces, 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

Notes and Constraints

If you see a message saying you lack the required permissions when viewing a security group's resources on the management console, you need to request the permissions for viewing the security group and its associated resources, such as servers, extended network interfaces, and supplementary network interfaces. For details, see **Example 4: Allowing users to view associated resources**.

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.

- 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 \equiv in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The Virtual Private Cloud page is displayed.

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

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, choose **More** > **Manage Network** > **Change Security Group** in the **Operation** column.

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 protection 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 6-20. 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 6-20 Network ACL rules

NOTE

The above figure 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 has inbound and outbound rules that are used 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 have rule numbers. 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 network ACL rule 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.
 - TCP is ideal for applications that require reliable connections and high data integrity, such as remote login, web browsing, email, and file transfers.
 - **UDP** is ideal for applications demanding high speed and low latency, such as online gaming and video meetings.
 - ICMP is used to communicate data transmission problems. For example, the ping command can be used to check the connectivity between network devices, error reports can be generated for O&M, and diagnosis information can be transmitted for network analysis and optimization.
- Source/Destination: The source or destination of the traffic.
- **Source Port Range/Destination Port Range**: The source or destination ports or port ranges. The value 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 rules. Similarly, if inbound traffic is allowed, responses to such inbound traffic are allowed to flow out, regardless of the outbound rules.
- Network ACLs use connection tracking to track traffic to and from instances. Changes to inbound and outbound rules do not take effect immediately for the existing traffic.

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

D NOTE

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.
- There are default inbound and outbound rules in network ACLs, as shown in Table 6-39. If a network ACL has no custom rules, the default inbound and outbound rules are applied, denying all traffic in and out of a subnet. You can use the default rules only when there is no need for traffic to go in and out of a subnet. If the traffic needs to go in and out of the subnet, you need to add custom rules to control traffic as required.

Direc tion	Prio rity	Action	Proto col	Sourc e	Source Port Range	Destinat ion	Destinati on Port Range
Inbo und	*	Deny	All	0.0.0. 0/0	All	0.0.0.0/0	All
Outb ound	*	Deny	All	0.0.0. 0/0	All	0.0.0.0/0	All

Table 6-39 Default network ACL rules

• The default and custom rules of a network ACL does not block the traffic described in **Table 6-40**.

Directio n	Description					
Inbound	Traffic between the source and destination in the same subnet					
	Broadcast traffic to 255.255.255/32					
	Multicast traffic to 224.0.0.0/24					
Outbou	Traffic between the source and destination in the same subnet					
nd	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					

 Table 6-40 Traffic not blocked by network ACL rules

Directio n	Description
	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 associated with one network ACL. If 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 matching sequence of inbound traffic is the same as that of outbound traffic. The following takes inbound traffic as an example to describe how the rules are applied.

- If a custom rule is matched:
 - If **Action** is set to **Deny**, traffic is denied to flow into the subnet.
 - If **Action** is set to **Allow**, traffic is allowed to flow into the subnet.
- If no custom rule is matched, the default rule is applied, denying traffic to flow into the subnet.



Figure 6-21 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

In **Figure 6-22**, **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 6-22 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 6-41**, 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.

Dire ctio n	Priori ty	Action	Туре	Protoc ol & Port	Source	Description
Inbo und	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.

Table 6-41 Security group rules

NOTE

For more examples of network ACL rule configurations, see **Network ACL Configuration Examples**.

Controlling Communications Between Instances in Different Subnets

In Figure 6-23, 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 6-23 Controlling communications between instances in different subnets

NOTE

For more examples of network ACL rule configurations, see **Network ACL Configuration Examples**.

Network ACL Configuration Procedure

Figure 6-24 Procedure for configuring a network ACL



N o.	Step	Description	Reference
1	Create a network ACL.	A network ACL comes with default inbound and outbound rules that 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
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.	Associating Subnets with a Network ACL
	שטוופנא.	A subnet can be associated with only one network ACL.	

Table 6-42 Procedure for configuring a network ACL

Constraints on Using Network ACLs

- 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.
- For inbound network ACL rules, the sum of the rules with **Source** set to **IP** address group, of the rules with **Destination** set to **IP** address group, of the rules with **Source Port Range** set to inconsecutive ports, and of the rules **Destination Port Range** set to inconsecutive ports, cannot exceed 120. If there are both IPv4 and IPv6 network ACL rules, up to 120 rules can be added for each type.

The limits on outbound network ACL rules are the same as those on inbound network ACL rules.

For example, to add inbound IPv4 rules to a network ACL, you can refer to **Table 6-43** for rules that meet the restrictions. Of these rules, rule A02 uses inconsecutive ports (22-24,25) as the source port range and IP address group ipGroup-A as the source. In this case, only one quota is occupied.

Rule No.	Prio rity	Туре	Acti on	Prot ocol	Source	Sourc e Port Rang e	Destina tion	Destinati on Port Range
Rule A01	1	IPv4	Den y	ТСР	0.0.0.0/ 0	22,25, 27	0.0.0.0/0	1-65535
Rule A02	2	IPv4	Allo W	ТСР	IP addres s group: ipGrou p-A	22-24, 25	0.0.0.0/0	1-65535
Rule A03	3	IPv4	Allo w	All	0.0.0.0/ 0	All	IP address group: ipGroup -B	All
Rule A04	4	IPv4	Allo w	UDP	0.0.0.0/ 0	1-655 35	0.0.0.0/0	80-83,87

 Table 6-43 Inbound network ACL rules

• Traffic from load balancers is not restricted by network ACL and security group rules if:

Transfer Client IP Address is enabled for the listeners 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

NOTE

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 6-44. If a network ACL has no custom rules, the default inbound and outbound rules are applied, denying all traffic in and out of a subnet.

Direc tion	Prior ity	Action	Proto col	Sourc e	Source Port Range	Destinat ion	Destinati on Port Range
Inbo und	*	Deny	All	0.0.0. 0/0	All	0.0.0.0/0	All
Outb ound	*	Deny	All	0.0.0. 0/0	All	0.0.0.0/0	All

 Table 6-44 Default network ACL rules

• You do not need to add rules to allow response traffic to requests. This is because the network ACLs are stateful and allow the responses to flow in or out of 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 6-45** to protect the instances in 10.0.0.0/24.

- 1. The default rule denies any traffic to the subnet. You need to add custom rule 02 to allow inbound traffic.
- 2. 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 Rule Numbers).

Dir ecti on	Pri ori ty	Typ e	Acti on	Prot ocol	Sourc e	Sourc e Port Rang e	Desti natio n	Desti natio n Port Range	Description
Inb oun d	1	IPv 4	Den y	ТСР	0.0.0.0 /0	All	10.0.0 .0/24	445	Custom rule 01
Inb oun d	2	IPv 4	Allo w	All	0.0.0.0 /0	All	10.0.0 .0/24	All	Custom rule 02
Inb oun d	*		Den y	All	0.0.0.0 /0	All	0.0.0. 0/0	All	Default rule

Table 6-45 Inbound rules for denying external access to a specific port in a subnet

Denying Access from a Specific IP Address

You can add inbound rules as described in **Table 6-46** 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.

- 1. The default rule denies any traffic to the subnet. You need to add custom rule 02 to allow inbound traffic.
- 2. 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 Rule Numbers).

Dir ecti on	Pri ori ty	Тур e	Acti on	Prot ocol	Sourc e	Sour ce Port Rang e	Desti natio n	Desti natio n Port Range	Description
Inb oun d	1	IPv 4	Den y	ТСР	10.1.1. 12/32	All	10.5.0. 0/24	All	Custom rule 01
Inb oun d	2	IPv 4	Allo w	All	0.0.0.0 /0	All	10.5.0. 0/24	All	Custom rule 02
Inb oun d	*		Den y	All	0.0.0.0 /0	All	0.0.0.0 /0	All	Default rule

Table 6-46 Inbound rules for denying access from a specific IP address

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 6-47.
 - 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 6-47 No	etwork ACL	rules for	allowing	access t	to specific	ports	on	an
instance in a s	subnet							

Dir ect ion	Pri ori ty	Ty pe	Act ion	Pro toc ol	Sourc e	Sour ce Port Rang e	Desti natio n	Desti natio n Port Rang e	Description
Inb ou nd	1	IPv 4	Allo w	ТСР	0.0.0. 0/0	All	10.8. 0.0/2 4	80	Custom rule 01
Inb ou nd	2	IPv 4	Allo w	ТСР	0.0.0. 0/0	All	10.8. 0.0/2 4	443	Custom rule 02
Inb ou nd	*		De ny	All	0.0.0. 0/0	All	0.0.0. 0/0	All	Default rule
Ou tbo un d	*		De ny	All	0.0.0. 0/0	All	0.0.0. 0/0	All	Default rule

- 2. Add security group rules listed in **Table 6-48**.
 - 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 6-48 Security group rules for allowing access to specific ports

Direc tion	Priorit y	Action	Туре	Protocol & Port	Source/ Destinatio n	Descriptio n
Inbou nd	1	Allow	IPv4	TCP: 80	IP address: 0.0.0.0/0	Rule 01
Inbou nd	1	Allow	IPv4	TCP: 443	IP address: 0.0.0.0/0	Rule 02
Outb ound	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.

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

Table 6-49 Parameter descriptions

4. Click Create Now.

Follow-up Operations

- 1. A 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. 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.

- 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.

- 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

Virtual Private Cloud

User Guide

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.

- 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 number, 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 \equiv in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Network ACLs. The network ACL list is displayed.
- In the network ACL list, locate the target network ACL and 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

Scenarios

You can add inbound and outbound rules to a network ACL to control the traffic in and out of a subnet. Network ACL rules are matched in an ascending order, either by the system-generated rule numbers or those you define.

 Adding a Network ACL Rule (Default Rule Numbers): Rules are matched in order of their number, starting with the lowest. The rule number is automatically assigned.

As shown in **Table 6-50**, 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.

Priority (Rules A and B)		Priority (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	*

Table 6-50 Default priorities

• Adding a Network ACL Rule (Custom Rule Numbers): If you want a rule to be matched earlier or later than a specific rule, you can insert the rule above or below the specific rule.

In **Table 6-51**, 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 matched 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.

Tuble of ST Custom phondes	Table	6-51	Custom	priorities
----------------------------	-------	------	--------	------------

Priority (Rules A and B)		Priority (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	*

Constraints

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

Adding a Network ACL Rule (Default Rule Numbers)

- 1. Log in to the management console.
- 2. Click \bigcirc 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.
- 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.

Parameter	Description	Example Value
Туре	Network ACL type. There are two options: • IPv4 • IPv6	IPv4
Action	 The action for the network ACL rule. There are two options: 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 .	ТСР
Source	 The source from which the traffic is allowed. The source can be an IP address or IP address range. IP address: Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	192.168.0.0/24
	- IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6)	
Source Port Range	The source ports or port ranges 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. IP address: Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	0.0.0.0/0
Destination Port Range	The destination ports or port ranges used to match traffic. The value ranges from 1 to 65535.	22-30

Table 6-52 Parameter descriptions	s
-----------------------------------	---

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

8. Click **OK**.

Return to the rule list to check the new rule.

- Rules are assigned a number based on the order they are added, with lower-numbered rule matched earlier.
- If the status of the new rule is **Enabled**, the rule is applied.

Adding a Network ACL Rule (Custom Rule Numbers)

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

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Network ACLs. The network ACL list is displayed.
- 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 will be matched earlier than the current rule.
 - Locate the target rule and choose More > Insert Rule Below in the Operation column. The new rule will be matched later than the current rule.

6.3.4.2 Modifying a Network ACL Rule

Scenarios

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

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

Notes and Constraints

Default network ACL rules cannot be modified or deleted.

Procedure

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

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Network ACLs. The network ACL list is displayed.
- 5. In the network ACL list, locate the target network ACL and click its name. The network ACL summary page is displayed.
- Click the Inbound Rules or Outbound Rules tab, locate the target rule, click Modify in the Operation column, and modify parameters based on Table 6-53.

Parameter	Description	Example Value
Туре	Network ACL type. There are two options: • IPv4 • IPv6	IPv4
Action	 The action for the network ACL rule. There are two options: Allow: allows matched traffic in and out of a subnot. 	Allow
	 Deny: denies matched traffic in and out of a subnet. 	
Protocol	The protocol supported by the network ACL to match traffic. The value can be TCP , UDP , or ICMP .	ТСР

 Table 6-53 Parameter descriptions

Parameter	Description	Example Value
Source	The source from which the traffic is allowed. The source can be an IP address or IP address range.	192.168.0.0/24
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Source Port Range	The source ports or port ranges 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.	0.0.0.0/0
	IP address:	
	 Single IP address: 192.168.10.10/32 (IPv4); 2002:50::44/128 (IPv6) 	
	 All IP addresses: 0.0.0.0/0 (IPv4); ::/0 (IPv6) 	
	 IP address range: 192.168.1.0/24 (IPv4); 2407:c080:802:469::/64 (IPv6) 	
Destination Port Range	The destination ports or port ranges used to match traffic. The value ranges from 1 to 65535.	22-30
Description	Supplementary information about the network ACL rule. This parameter is optional.	N/A
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

7. Click **OK**.

6.3.4.3 Enabling or Disabling One or More Network ACL Rules

Scenarios

After a rule is added, it is in the **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 is 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.

Enabling or Disabling a Network ACL Rule

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

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Network ACLs. The network ACL list is displayed.
- 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**.

6.3.4.4 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.

- 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.

6.3.4.5 Deleting One or More Network ACL Rules

Scenarios

You can delete network ACL rules if you no longer need them.

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.

A maximum of 50 network ACL rules can be deleted at a time.

Deleting a Network ACL Rule

1. Log in to the management console.

- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- 3. Click \equiv in the upper left corner and choose **Networking** > **Virtual Private Cloud**.

The Virtual Private Cloud page is displayed.

- In the navigation pane on the left, choose Access Control > Network ACLs. The network ACL list is displayed.
- 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 network ACL rule list, locate the target rule and click **Delete** in the **Operation** column.

A confirmation dialog box is displayed.

8. Confirm the information and click **OK**.

Deleting Network ACL Rules in Batches

- 1. Log in to the management console.
- 2. Click 🔍 in the upper left corner and select the desired region and project.
- 3. Click \equiv 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 network ACL rule list, select multiple rules and click **Delete** in the upper left corner above the list.

A confirmation dialog box is displayed.

8. 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.

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.

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 \oplus 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 network ACL list, locate the target 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 a 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 from 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 cannot see the disassociated subnets 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. In the network ACL list, 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 cannot see the disassociated subnets in the subnet list.

7 IP Address Group

7.1 IP Address Group Overview

What Is an IP Address Group?

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

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

Table 7-1 Resources that can be associated with an IP address group	Table 7-1 Resources that can	be associated with an IP addres	s 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 7-2**.

Paramet er	Description	Example Value
Region	Mandatory	Region A
	Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the region nearest to where your services will be accessed.	
	An IP address group can be associated only with resources in the same region.	
Name	Mandatory	ipGroup-A
	Enter the name of the IP address group. The name:	
	• 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.	

Table 7-2 Parameters for creating an IP address group

Paramet er	Description	Example Value
Max. Entries	Mandatory Set the number of IP address entries that can be added to an IP address group. By default, the system displays the maximum number of IP address entries that can be added to an IP address group. You can change the number as required. If you want to increase the maximum number of IP address entries in a group, submit a service ticket .	20
IP Address Version	Mandatory Select the type of IP addresses that can be added to an IP address group. • IPv4 • IPv6	IPv4
IP Address Entries	 Optional 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: 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 	 Without description: 192.168.0.0/16 With description: 192.168.0.0/16 ECS01
Descripti on (Optiona l)	Optional Enter the description of the IP address group in the text box as required.	-

3. Click Create Now.

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

An IP address group takes effect only after it is associated with 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.
- There are IP address entries in the IP address group. For details, see Adding IP Address Entries to an IP Address Group.

Procedure

You need to associate an IP address group with resources. For details, see **Table 7-3**.

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 Inbound rule: Set Source to an IP address group. Outbound rule: Set Destination to an IP address group.
Network ACL	The Source or Destination of a network ACL is set to IP address group .	 Adding a Network ACL Rule 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.

Tahle	7-3	Associating	an	IP	address	aroun	with	resources
lavie	1-2	Associating	an	IP	auuress	group	VVILII	resources

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. Entries
- Description

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. 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 7-4**.

Paramet er	Description	Example Value
Name	Mandatory	ipGroup-A
	Enter the name of the IP address group. The name:	
	• 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.	
Max.	Mandatory	20
Entries	Set the number of IP address entries that can be added to an IP address group. By default, the system displays the maximum number of IP address entries that can be added to an IP address group. You can change the number as required.	
	If you want to increase the maximum number of IP address entries in a group, submit a service ticket .	

 Table 7-4 IP address group parameters

Paramet er	Description	Example Value
Descripti on	Optional Enter the description of the IP address group in the text box as required.	-

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 address entries
- Associated resources

Procedure

- 1. Go to the **IP address group list page**.
- 2. In the upper left corner above the IP address group list, click **Export**.
 - **Export selected data to an XLSX file**: Select one or more IP address groups and export information about the selected IP address groups.
 - Export all data to an XLSX file: Export information about all the IP address groups in the current region.

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

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 address entries
- 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 address entries 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.

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. Perform the following operations to delete IP address groups.
 - Delete a single IP address group:
 - In the IP address group 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 a batch.
 - i. In the IP address group 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 Address Entries in an IP Address Group

7.3.1 Adding IP Address Entries to an IP Address Group

Scenarios

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

Constraints

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

If an IP address group has security groups and network ACLs associated, the rules associated with the IP address group 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 entry list, click **Add**. The dialog box for adding IP address entries is displayed.
- 4. Add IP address entries to the IP address group as prompted.

Table 7-5 IP address group parameters

Paramet er	Description	Example Value	
Name	The name of the IP address group.	ipGroup-A	
Max. Entries	Mandatory Set the number of IP address entries that can be added to an IP address group. By default, the system displays the maximum number of IP address entries that can be added to an IP address group. You can change the number as required. If you want to increase the maximum	20	
	number of IP address entries in a group, submit a service ticket .		
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: IPv4	IPv4	
Paramet er	Description	Example Value	
--------------------------	--	------------------------------------	
IP Address Entries	Mandatory Enter an IP address or IP address range on each line, and press Enter .	192.168.0.0/16 192.168.10.10/32	
	 You can enter: 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		

7.3.2 Modifying IP Address Entries 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.

Constraints

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

If an IP address group has security groups and network ACLs associated, the rules associated with the IP address group 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 entry list, click **Modify**.
- The dialog box for modifying an IP address entry is displayed.
- Modify the information as prompted.
 For details, see Table 7-6.

Paramet er	Description	Example Value
Name	The name of the IP address group.	ipGroup-A

Table 7-6 Parameters for modifying IP address entries

Paramet er	Description	Example Value
Max. Entries	Mandatory Set the number of IP address entries that can be added to an IP address group. By default, the system displays the maximum number of IP address entries that can be added to an IP address group. You can change the number as required. If you want to increase the maximum number of IP address entries in a group, submit a service ticket .	20
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: IPv4 IPv6 	IPv4
IP Address Entries	 You can modify existing IP addresses, IP address ranges, and their descriptions in an IP address group. The format is "IP address Description". Description is optional, can contain 0 to 255 characters, and cannot contain angle brackets (< or >). You can enter: 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 	 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 entry was modified.

7.3.3 Deleting IP Address Entries from an IP Address Group

Scenarios

This section describes how to delete IP address entries from an IP address group.

Constraints

If an IP address group has resources associated, deleting IP address entries from the group may affect your network communications.

If an IP address group has security groups and network ACLs associated, the rules associated with the IP address group 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 address entries:
 - Delete a single IP address entry.
 - i. In the IP address entry list, locate the target IP address entry and click **Delete** in the **Operation** column.
 - A confirmation dialog box is displayed.
 - ii. Confirm the information and click **OK**.
 - Delete IP address entries in batches.
 - i. In the IP address entry list, select the IP address entries to be deleted.
 - ii. Click the **Delete** button above the IP address entry 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.

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

 Table 7-7 Inbound rules

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 7-8**. The following resource details are only examples. You can modify them as required.

Resource	Quantity	Description		
IP address group	1	Create an IP address group and add IP addresses that need to access the instances.		
		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:		
		– 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 7-9 .		

 Table 7-8 Resource planning

Table 7-9 Inbound rules

Direction	Action	Туре	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 Overview

What Is a VPC Peering Connection?

A VPC peering connection connects two VPCs to enable them to communicate using private IP addresses. The VPCs to be peered can be under 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

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 to 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 a few ECSs in different regions to communicate with each other, you can **assign and bind EIPs to the ECSs**.

• If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. For example, if the CIDR block of a VPC is 192.168.0.0/16 and that of another VPC is 192.168.0.0/16 or 192.168.0.0/18, the CIDR blocks overlap.

In this case, you can configure the network by referring to VPC Peering Connection Usage Examples.

If there are CCE clusters, you need to avoid CIDR block overlapping between the subnets and container subnets in addition to the VPC CIDR blocks at both ends. Otherwise, communications will fail. For details, see Cross-VPC Cluster Interconnection.

• 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.

Locati on	CIDR Block	Description	Example
VPCs in the same region	 VPC CIDR blocks do not overlap. Subnet CIDR blocks of VPCs do not overlap. 	You can create VPC peering connections to connect entire CIDR blocks of VPCs. Then, all resources in the VPCs can communicate with each other.	Using a VPC Peering Connection to Connect Two VPCs

Table o-1 VPC beening connection usage example	Table 8-1 VPC	peerina	connection	usade	exam	ples
--	---------------	---------	------------	-------	------	------

Locati on	CIDR Block	Description	Example
VPCs in the same	Cs • VPC CIDR You can create VPC blocks peering connections to connect specific		Using a VPC Peering Connection to Connect Subnets in Two VPCs
region	 Some subnet CIDR blocks overlap. 	 subnets or ECSs from different VPCs. To connect specific subnets from two VPCs, the subnet CIDR blocks cannot overlap. To connect specific ECSs from two VPCs, each ECS must have a unique private IP address. 	Using a VPC Peering Connection to Connect ECSs in Two VPCs
VPCs in the same region	 VPC CIDR blocks overlap. All subnet CIDR blocks overlap. 	VPC peering connections are not usable.	Unsupported VPC Peering Configurations

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.

If there are CCE clusters, you need to avoid CIDR block overlapping between the subnets and container subnets in addition to the VPC CIDR blocks at both ends. Otherwise, communications will fail. For details, see Cross-VPC Cluster Interconnection.

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 Addr ess Versi on	Example
Two VPCs peered	You have two VPCs that require full access to each other's	IPv4	Two VPCs Peered Together (IPv4)
togetner	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.	IPv6	Two VPCs Peered Together (IPv6)
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	IPv4	Multiple VPCs Peered Together Through Transitive Peering Connections (IPv4)
	access to each other's resources.	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	IPv4	One Central VPC Peered with Two VPCs (IPv4)
	central VPC. However, the two peer VPCs need to be isolated from each other.	IPv6	One Central VPC Peered with Two VPCs (IPv6)
	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.		
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 Addr ess Versi on	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	IPv4	One Central VPC Peered with Multiple VPCs (IPv4)
	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		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)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02		172.16.1.91
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
	6	Subne t-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)
-----------	------------------	--------

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
A	172.16.1.0/24	Local	Syste m		
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.	
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
В	10.0.1.0/24	Local	Syste m		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.	

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**.



Table 8-6 Resource planning details (IPv6)

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	I72.1 Subne IPv4: rtb- 5.0.0/ t-A01 172.1 VPC-A 16 IPv6: 2407: c080: 802:c 34::/ 64 14000000000000000000000000000000000000	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb 		
		Subne t-A02	 IPv4: 172.1 6.1.0 /24 IPv6: 2407: c080: 802:c 37::/ 64 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -B	VPC 10.00 -B .0/16	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407: c080: 802:c 35::/ 64 	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
		Subne t-B02	 IPv4: 10.0. 1.0/2 4 IPv6: 2407: c080: 802:c 38::/ 64 	rtb- VPC-B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700: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)
-----------	-----	-------	--------	--------

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC- A	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	2407:c080:802:c 34::/64	Local	Syst em			
	172.16.1.0/24	Local	Syst em			
	2407:c080:802:c 37::/64	Local	Syst em			
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.		
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the		
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-AB	Cust om	next nop for IPV6 communication.		
rtb- VPC-	10.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.		
В	2407:c080:802:c 35::/64	Local	Syst em			
	10.0.1.0/24	Local	Syst em			
	2407:c080:802:c 38::/64	Local	Syst em			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.		
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	the next hop for IPv6 communication.		

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 Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
		Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	192.1 68.0.	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01	-	192.168.0.194
	0/16	Subne t-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)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC	rtb- 172.16.0.0/24 Local Sys VPC m	Syste m	Local routes are automatically added for communications within a VPC.	
-A	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.

Rou te Tabl e	Destination	Next Hop	Rout e 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.		
rtb- VPC	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-В	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.		
	192.168.0.0/16 (VPC-C)	Peerin g-BC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-BC as the next hop.		
rtb- VPC	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-C	192.168.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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-В)	Peerin g-BC	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop.		

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 VP-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 Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02		172.16.1.91
VPC -B	VPC 10.0. -B 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
		Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	VPC 192.1 -C 68.0. 0/16	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194
		Subne t-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:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC -A	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
	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.
rtb- VPC	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-В	10.0.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
	192.168.0.0/16 (VPC-C)	Peerin g-AB	Cust om	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	Syste m	Local routes are automatically added for communications within a VPC.

 Table 8-14 VPC route tables (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	192.168.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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-В)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AC as the next hop.

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)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC 172.1 -A 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c080 :802: c34::/ 64 	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb 	
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c080 :802: c37::/ 64 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407 :c080 :802: c35::/ 64 	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
		Subne t-B02	 IPv4: 10.0. 1.0/2 4 IPv6: 2407 :c080 :802: c38::/ 64 	rtb- VPC-B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700:c1cf

Table 8-15 Resource	planning	details	(IPv6)
---------------------	----------	---------	--------

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -C	192.1 68.0. 0/16	Subne t-C01	 IPv4: 192. 168. 0.0/2 4 IPv6: 2407 :c080 :802: c3c::/ 64 	rtb- VPC-C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af
		Subne t-C02	 IPv4: 192. 168. 1.0/2 4 IPv6: 2407 :c080 :802: c3d:: /64 	rtb- VPC-C	ECS- C02		 IPv4: 192.168.1.200 IPv6: 2407:c080:80 2:c3d:e9ca:16 9a: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:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	2407:c080:802:c 34::/64	Local	Syste m	
	172.16.1.0/24	Local	Syste m	
	2407:c080:802:c 37::/64	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Custo m	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:c 35::/64 (Subnet- B01)	Peerin g-AB	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-AB	Custo m	next hop for IPv6 communication.
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Custo m	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:c 3c::/64 (Subnet- C01)	Peerin g-AC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-AC	Custo m	next hop for IPv6 communication.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	2407:c080:802:c 35::/64	Local	Syste m	
	10.0.1.0/24	Local	Syste m	
	2407:c080:802:c 38::/64	Local	Syste m	

 Table 8-17
 VPC route tables (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Custo m	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:c 34::/64 (Subnet- A01)	Peerin g-AB	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Custo m	the next hop for IPV6 communication.		
	192.168.0.0/16 (VPC-C)	Peerin g-BC	Custo m	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:c 3c::/64 (Subnet- C01)	Peerin g-BC	Custo m	Add routes with the IPv6 CIDR bloc of Subnet-C01 and Subnet-C02 as t destinations and Peering-BC as the		
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-BC	Custo m	next hop for IPV6 communication.		
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
С	2407:c080:802:c 3c::/64	Local	Syste m			
	192.168.1.0/24	Local	Syste m			
	2407:c080:802:c 3d::/64	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Custo m	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:c 34::/64 (Subnet- A01)	Peerin g-AC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Custo m	the next hop for IPv6 communication.		

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	10.0.0.0/16 (VPC-В)	Peerin g-BC	Custo m	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:c 35::/64 (Subnet- B01)	Peerin g-BC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-BC as the
	2407:c080:802:c Peerin Cust 38::/64 (Subnet- g-BC m B02)		Custo m	next hop for IPv6 communication.

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)

VP C Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02		172.16.1.91
VPC -B	10.0.0 .0/16	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
		Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	VPC 192.1 -C 68.0.0	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194
/16	Subne t-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:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
А	172.16.1.0/24	Local	Syste m			
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.		
	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.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
В	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.		
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
C	192.168.1.0/24	Local	Syste m			

Table 8-20 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

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)



VPC Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC 17 -A 6. 16	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c08 0:80 2:c3 4::/6 4 	rtb-VPC- A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c08 0:80 2:c3 7::/6 4 	rtb-VPC- A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/ 24 IPv6: 2407 :c08 0:80 2:c3 5::/6 4 	rtb-VPC- B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
		Subne t-B02	 IPv4: 10.0. 1.0/ 24 IPv6: 2407 :c08 0:80 2:c3 8::/6 4 	rtb-VPC- B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700:c1cf

 Table 8-21 Resource planning details (IPv6)

VPC Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -C	VPC 192.1 Subne t-C01 19 -C 68.0. 0/16 t-C01 19 0/16 4 10 16 0.0 24 0.1 24 10 24 1 24 1		 IPv4: 192. 168. 0.0/ 24 IPv6: 2407 :c08 0:80 2:c3c ::/64 	rtb-VPC- C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af
		Subne t-C02	 IPv4: 192. 168. 1.0/ 24 IPv6: 2407 :c08 0:80 2:c3 d::/6 4 	rtb-VPC- C	ECS- C02		 IPv4: 192.168.1.200 IPv6: 2407:c080:80 2:c3d:e9ca:16 9a: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:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description			
rtb- VPC	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
-A	2407:c080:802:c 34::/64	Local	Syste m				
	172.16.1.0/24	Local	Syste m				
	2407:c080:802:c 37::/64	Local	Syste m				
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering- AB as the next hop for IPv4 communication.			
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the			
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-AB	Cust om	next hop for IPv6 communication.			
	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:c 3c::/64 (Subnet- C01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the			
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-AC	Cust om	next hop for IPv6 communication.			
rtb- VPC	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
-В	2407:c080:802:c 35::/64	Local	Syste m				
	10.0.1.0/24	Local	Syste m				
	2407:c080:802:c 38::/64	Local	Syste m				

Table 8-23 VPC route table details (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description				
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.				
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the				
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	next hop for IPv6 communication.				
rtb- VPC -C	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.				
	2407:c080:802:c 3c::/64	Local	Syste m					
	192.168.1.0/24	Local	Syste m					
	2407:c080:802:c 3d::/64	Local	Syste m					
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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:c 34::/64 (Subnet- A01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the				
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Cust om	next hop for IPv6 communication.				

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)
VP C Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	Prima ry CIDR block: 172.1 6.0.0/ 16	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	Secon dary CIDR block: 192.1 67.0.0 /16	Subne t-A- Exten d01	192.167. 0.0/24	rtb- VPC-A	ECS- A- Exte nd0 1		192.167.0.100
VPC -B	10.0.0 .0/16	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
VPC -C	192.1 68.0.0 /16	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194

 Table 8-24 Resource planning details

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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC-	- 172.16.0.0/24 Local Syste L - m fe		Syste m	Local routes are automatically added for communications within a VPC.		
~	192.167.0.0/24	Local	Syste m			
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.		
	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.		
rtb- VPC-	rtb- 10.0.0/24 Local VPC-		Syste m	Local routes are automatically added for communications within a VPC.		
В	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peerin g-AB	Cust om	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AB as		
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peerin g-AB	Cust om	the next hop.		
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
C	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peerin g-AC	Cust om	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AC as		
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peerin g-AC	Cust om	the next hop.		

Table 8-26 VPC route table details (IPv4)

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 Res	ource planning	details (IPv4)
----------------	----------------	----------------

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02	purpose web server	172.16.1.91

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	192.1 68.0. 0/16	Subne t-C01	192.168 .0.0/24	rtb-VPC- C	ECS- C01		192.168.0.194
VPC -D	10.2. 0.0/1 6	Subne t-D01	10.2.0.0 /24	rtb-VPC- D	ECS- D01		10.2.0.237
VPC -E	10.3. 0.0/1 6	Subne t-E01	10.3.0.0 /24	rtb-VPC- E	ECS- E01		10.3.0.87
VPC -F	172.1 7.0.0/ 16	Subne t-F01	172.17. 0.0/24	rtb-VPC- F	ECS- F01		172.17.0.103
VPC -G	10.4. 0.0/1 6	Subne t-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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
	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.
	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.
	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.
	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.
	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.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
C	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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	Syste m	Local routes are automatically added for communications within a VPC.

Table 8-29 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	172.16.0.0/16 (VPC-A)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop.		
rtb- VPC-	10.3.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
E	172.16.0.0/16 (VPC-A)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.		
rtb- VPC-	172.17.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
F	172.16.0.0/16 (VPC-A)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.		
rtb- VPC-	10.4.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
G	172.16.0.0/16 (VPC-A)	Peerin g-AG	Cust om	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)

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c08 0:80 2:c3 4::/6 4 	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c08 0:80 2:c3 7::/6 4 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c3 5::/6 4 	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
VPC -C	192.1 68.0. 0/16	Subne t-C01	 IPv4: 192. 168. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c3c ::/64 	rtb- VPC-C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af
VPC -D	10.2. 0.0/1 6	Subne t-D01	 IPv4: 10.2. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 5::/6 4 	rtb- VPC-D	ECS- D01		 IPv4: 10.2.0.237 IPv6: 2407:c080:80 2:c45:6bb7:f1 61:3596:6e4c

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -E	10.3. 0.0/1 6	Subne t-E01	 IPv4: 10.3. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 6::/6 4 	rtb- VPC-E	ECS- E01		 IPv4: 10.3.0.87 IPv6: 2407:c080:80 2:c46:2a2f:55 8a:85da:4c70
VPC -F	172.1 7.0.0/ 16	Subne t-F01	 IPv4: 172. 17.0. 0/24 IPv6: 2407 :c08 0:80 2:c4 7::/6 4 	rtb- VPC-F	ECS- F01		 IPv4: 172.17.0.103 IPv6: 2407:c080:80 2:c47:b5e2:e6f 0:c42b:44fd
VPC -G	10.4. 0.0/1 6	Subne t-G01	 IPv4: 10.4. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 8::/6 4 	rtb- VPC-G	ECS- G01		 IPv4: 10.4.0.10 IPv6: 2407:c080:80 2:c48:3020:f4 8c: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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
-A	2407:c080:802:c 34::/64	Local	Syste m		
	172.16.1.0/24	Local	Syste m		
	2407:c080:802:c 37::/64	Local	Syste m		
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering- AB as the next hop for IPv4 communication.	
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add a route with the IPv6 CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop for IPv6 communication.	

Table 8-32 VPC route table details (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e 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:c 3c::/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:c 45::/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:c 46::/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:c 47::/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:c 48::/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.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-В	2407:c080:802:c 35::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.
	2407:c080:802:c 34::/64 (Subnet- A01)	0:802:c Peerin Cus Subnet- g-AB om		Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	next hop for IPv6 communication.
rtb- VPC	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-C	2407:c080:802:c 3c::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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:c 34::/64 (Subnet- A01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Cust om	next hop for IPv6 communication.
rtb- VPC	10.2.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-D	2407:c080:802:c 45::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop for IPv4 communication.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AD	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AD as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AD	Cust om	next hop for IPv6 communication.
rtb- VPC	10.3.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-E	2407:c080:802:c 46::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AE	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AE as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AE	Cust om	next hop for IPv6 communication.
rtb- VPC	172.17.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-F	2407:c080:802:c 47::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AF	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AF as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AF	Cust om	next nop for IPv6 communication.
rtb- VPC	10.4.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-G	2407:c080:802:c 48::/64	Local	Syste m	

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AG	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AG as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AG	Cust om	next hop for IPv6 communication.

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.

Scenario	Scenario Description	IP Addr ess Versi on	Example
Two VPCs peered to two subnets in a	You have a central VPC that requires access to the multiple other VPCs. The other VPCs	IPv4	Two VPCs Peered to Two Subnets in a Central VPC (IPv4)
central VPC	 need to be isolated from each other. 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. 	IPv6/ IPv4	Two VPCs Peered to Two Subnets in a Central VPC (IPv6/ IPv4)

Table 8-33 Scenario description

Scenario	Scenario Description	IP Addr ess Versi on	Example
One central VPC peered to specific subnets in two VPCs	 You have a central VPC that requires access to two other VPCs. The other VPCs need to be isolated from each other. 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	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.	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 Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A01	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A02	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-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

Table 8-36 VPC	route table	details	(IPv4)
----------------	-------------	---------	--------

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A01	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A02	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (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.
rtb- VPC- B	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop.
rtb- 10.0.0/ VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24 (Subnet-A02)	Peerin g-AC	Cust om	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 0-37 Resource planning details (IFV0/IFV4)	Table 8-37	Resource	planning	details	(IPv6/IPv4)
---	------------	----------	----------	---------	-------------

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c080 :802: c34:: /64 	rtb-VPC- A01	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	172.16. 1.0/24	rtb-VPC- A02	ECS- A02		172.16.1.91

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407 :c080 :802: c35:: /64 	rtb-VPC- B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
VPC -C	10.0. 0.0/1 6	Subne t-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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description			
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
A01	2407:c080:802:c 34::/64	Local	Syste m				
	172.16.1.0/24	Local	Syste m				
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.			
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	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-	p- 172.16.0.0/24 Local Syst pC- m		Syste m	Local routes are automatically added for communications within a VPC.			
A02	2407:c080:802:c 34::/64	Local	Syste m				
	172.16.1.0/24	Local	Syste m				
	10.0.0.0/16 (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.			
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
В	2407:c080:802:c 35::/64	Local	Syste m				
	172.16.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	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:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add a route with the IPv6 CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop for IPv6 communication.			

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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
C	172.16.1.0/24 (Subnet-A02)	Peerin g-AC	Cust om	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)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general- purpose	172.16.0.111
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01	web server	10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb-VPC- B	ECS- B02		10.0.1.167
VPC -C	10.0. 0.0/1	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71
	6	Subne t-C02	10.0.1.0 /24	rtb-VPC- C	ECS- C02		10.0.1.116

 Table 8-40 Resource planning details (IPv4)

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

Table 8-42 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC- A	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	10.0.0.0/24 (Subnet-B01)	Peerin g-AB	Cust om	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)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	10.0.1.0/24 Local Sy m	Syste m		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
C	10.0.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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 Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-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. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb-VPC- B	ECS- B02		10.0.1.167
VPC -C	10.0. 0.0/1	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71
	6	Subne t-C02	10.0.1.0 /24	rtb-VPC- C	ECS- C02		10.0.1.167

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

 Table 8-44
 Peering relationships (IPv4)

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.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	10.0.1.0/24 (Subnet-C02)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	10.0.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

Table 8-45 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
С	10.0.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	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.

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description
rtb- VPC -A	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.
	10.0.1.167/32 (ECS-B02)	Peerin g-AB	Cust om	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)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.

 Table 8-46
 VPC route table details

• 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.

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description
rtb- VPC -A	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24 (Subnet-B02)	Peerin g-AB	Cust om	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)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C01 as the destination and Peering-AC as the next hop.

Table 8-47 VPC route table details

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 ECSs 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.

Scenario	Scenario Description	IP Addr ess Versi on	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)

Table 8-48 Scenario description

Scenario	Scenario Description	IP Addr ess Versi on	Example
A central VPC peered with two other VPCs using longest prefix match	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:	IPv4	A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)
	• 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		
	This configuration expands the communication scope.		

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.





Fable 8-49 Res	ource planning	details (IPv4)
-----------------------	----------------	----------------

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01- 1	sg-web: general- purpose	172.16.0.111
					ECS- A01- 2	web server	172.16.0.218
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-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

Table 8-51 VPC route table details (IP)	v4)
---	-----

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	10.0.0.139/32 (ECS-B01)	Peerin g-AB	Cust om	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)	Peerin g-AC	Cust om	Add a route with the private IP address of ECS-C01 as the destination and Peering-AC as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	172.16.0.111/32 (ECS-A01-1)	Peerin g-AB	Cust om	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	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.218/32 (ECS-A01-2)	Peerin g-AC	Cust om	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 Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-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

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	172.16.1.0/24	Local	Syste m	
	10.0.0.139/32 (ECS-B01)	Peerin g-AB	Cust om	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)	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.
rtb- VPC- B	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
rtb- VPC- C	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

Table 8-54 VPC route table details (IPv4)

8.2.5 Unsupported VPC Peering Configurations

Scenarios

The VPC peering connection configurations are not supported in Table 8-55.

Scenario	Example
 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 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-	Invalid VPC Peering for Sharing an EIP
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 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 **Figure 8-20**, 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

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description
rtb- VPC	10.0.0/24	Local	Syst em	Local routes are automatically added for communications within
-A	10.0.1.0/24	Local	Syst em	a VPC.
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC	10.0.0.0/24	Local	Syst em	Local routes are automatically added for communications within
-В	10.0.1.0/24	Local	Syst em	a VPC.
	10.0.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

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

Ro ute Tab le	Destination	Next Hop	Rou te Typ e	Description	
rtb- VPC	10.0.0/24LocalSystLocal routes are automadded for communicat	Local routes are automatically added for communications			
-A	10.0.1.0/24	Local	Syst em	within a VPC.	
	10.0.2.0/24 (Subnet-B02)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.	
rtb- VPC	10.0.0.0/24	Local	Syst em	Local routes are automatically added for communications	
-B	10.0.2.0/24	Local	Syst em	within a VPC.	
	10.0.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	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 (the service server **ECS-A01** and database server **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

Currently, VPC peering connections are free.

Constraints

• If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. For example, if the CIDR block of a VPC is 192.168.0.0/16 and that of another VPC is 192.168.0.0/16 or 192.168.0.0/18, the CIDR blocks overlap.

In this case, you can configure the network by referring to VPC Peering Connection Usage Examples.

If there are CCE clusters, you need to avoid CIDR block overlapping between the subnets and container subnets in addition to the VPC CIDR blocks at both ends. Otherwise, communications will fail. For details, see Cross-VPC Cluster Interconnection.

- 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 a few ECSs in different regions to communicate with each other, you can **assign and bind EIPs to the ECSs**.

Prerequisites

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

 \times

Figure 8-27 Creating a VPC peering connection

Create VPC Peering Connection

 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 N	ame peering		
Local VPC Settings			
* Local VPC	vpc-1	C	
Local VPC CIDR Block	192.1		
Peer VPC Settings			
* Account	My account Another account	0	
★ Peer Project	eu-w 1 If you select My account, the project is filled in by	r default.	
* Peer VPC	vpc-8	ring connection may not	
	OK Cancel		

Table 8-58 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
Name	Mandatory	peering-AB
	peering connection.	
	The name can contain a maximum of 64 characters, including letters, digits, hyphens (-), and underscores (_).	

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.	vрс-А
Local VPC CIDR Block	CIDR block of the selected local VPC	172.16.0.0/16
Account	 Mandatory My account: The local and peer VPCs are from the same account. Another account: The local and peer VPCs are from different accounts. 	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.	vрс-В
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 Create Now.

A dialog box for adding routes is displayed.

 In the displayed dialog box, click Add Now. On the displayed page about the VPC peering connection details, go to Step 2: Add Routes for the VPC Peering Connection to add a route.

Step 2: Add Routes for the VPC Peering Connection

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.

Parameter	Description	Example Value
VPC	Select a VPC that is connected by the VPC peering connection.	vpc-А
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.	rtb-vpc-A (Default)
	• 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.	

Table 8-59 Parameter description

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.	vрс-В

Parameter	Description	Example Value
Route Table	Select the route table of the VPC. The route will be added to this route table.	rtb-vpc-B (Default)
	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.	
	• 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.	
Destination	An IP address or address range in the other VPC connected by the VPC peering connection. The value can be a VPC CIDR block, subnet CIDR block, or ECS IP address. For details about the route configuration example, see 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

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.
- Check whether ECS-A01 can communicate with RDS-B01.
 ping <peer-server-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 ---

NOTE

In this example, ECS-A01 and RDS-B01 are associated with the same security group, so they can communicate with each other once a VPC peering connection is created between VPC-A and VPC-B. If the instances are associated with different security groups, you need to add inbound rules to allow access from instances of 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 to 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 (the service server **ECS-A01** and database server **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

Subnet-B01-172.17.0.0/24

Destination

172.17.0.21

vpc-B route table

172.16.0.0/16 peering-AB

Next Hop

Account A Vpc-A-Region A 172.16.0.0/16 Security group Security group

VPC peering

connection

peering-AB



ECS-A01

172.16.0.8

Currently, VPC peering connections are free.

Next Hop

peering-AB

Subnet-A01-172.16.0.0/24

Destination

172.17.0.0/16

vpc-A route table

Notes and Constraints

• If the local and peer VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect. For example, if the CIDR block of a VPC is 192.168.0.0/16 and that of another VPC is 192.168.0.0/16 or 192.168.0.0/18, the CIDR blocks overlap.

In this case, you can configure the network by referring to VPC Peering Connection Usage Examples.

If there are CCE clusters, you need to avoid CIDR block overlapping between the subnets and container subnets in addition to the VPC CIDR blocks at both ends. Otherwise, communications will fail. For details, see Cross-VPC Cluster Interconnection.

- 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 a few ECSs in different regions to communicate with each other, you can **assign and bind EIPs to the ECSs**.
- 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 untrusted 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 with a Subnet**.

Step 1: Create a VPC Peering Connection

1. Go to the VPC peering connection list page.

 \times

- 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
--------	-----	---------	------------

 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 peering-4				
Local VPC Settings				
* Local VPC	vpc-1		С	
Local VPC CIDR Block	192.			
Peer VPC Settings				
* Account	My account	Another account	0	
	The VPC peering connection accepts the connection	ection will be activated only request.	after the peer account	
* Peer Project ID				
	If you select Another ac VPC of the peer accourt	ccount, enter the project ID on the isin. Learn more	of the region that the	
* Peer VPC ID				
	ОК	Cancel		

Table 8-60 Parameters for creating a VPC peering connection

Parameter	Description	Example Value
Name	Mandatory	peering-AB
	Enter a name for the VPC peering connection.	
	The name can contain a maximum of 64 characters, including letters, digits, hyphens (-), and underscores (_).	

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.	vpс-А
Local VPC CIDR Block	CIDR block of the selected local VPC	172.16.0.0/16
Account	 Mandatory Options: My account and Another account Select Another account. 	Another account
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: 067cf8aecf3XXX08322f 13b
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 Create Now.

- 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

 NametD 0
 Status 0
 Local VPC 0
 Local VPC CDR Block 0
 Peer Project ID 0
 Peer VPC 0
 Peer VPC CDR Block 0
 Descrittle
 Operation

 peerng-AB
 0472595-484444692 3 Awating acceptance
 vpcA
 172.16.0.16
 076640177
 vpcB
 172.17.0.015
 Modify Desire

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.	vрс-А
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. If there is only 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.	Route from vpc- A to vpc-B
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

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.	vрс-В

Parameter	Description	Example Value
Route Table	Select the route table of the VPC. The route will be added to this route table.	rtb-vpc-B (Default route
	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.	table)
	 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. 	
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 <peer-server-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 ----

NOTE

In this example, ECS-A01 and RDS-B01 are associated with the same security group, so they can communicate with each other once a VPC peering connection is created between VPC-A and VPC-B. If the instances are associated with different security groups, you need to add inbound rules to allow access from instances of 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



3. In the project list, obtain the project ID.

Locate the region of the peer VPC and obtain the project ID corresponding to the region.

Figure 8-32 Project ID

Projects			
	Project ID ↓Ξ	Project Name ↓=	Region J⊟
	067 13b	4	
	92f3 1d5	9	
	152 5d99	3	n hu hu hu hu
	857 5ad	1	Columnation (Columnation)
	59/5 iba5	-4	A day income where the

8.6 Modifying a VPC Peering Connection

Scenarios

This section describes how to modify the basic information about a VPC peering connection, including its name and description.

Either owner of a VPC in a peering connection can modify the VPC peering connection in any state.

Procedure

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

The Virtual Private Cloud page is displayed.

4. In the navigation pane on the left, choose **Virtual Private Cloud** > **VPC Peering Connections**.

The VPC peering connection list is displayed.

- 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.

- 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 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 \equiv 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 name of the target route table in the **Route Table** column.

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 \equiv 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 \equiv 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 I in the upper left corner and select the desired region and project.
 - b. Click \equiv 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 IPv4/IPv6 Dual-Stack Network

What Is an IPv4/IPv6 Dual-Stack Network?

An IPv4/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/IPv6 dual-stack network works.

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



	Table 9-1	Steps fo	r deploying	a dual-stack	network
--	-----------	----------	-------------	--------------	---------

Ste p	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.

Ste p	e Description
2	Subnets in the same VPC can communicate with each other by default. Network ACLs protect subnets, and security groups protect the instances in it.
	1. 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.
	2. 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.
	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:
	• 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	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.
	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.
4	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.
	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.

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 ECSs with certain flavors support IPv6. You need to select such an ECS to use an IPv4/IPv6 dual-stack network.

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

IPv4/IPv6 Dual-Stack Application Scenarios

If your ECS supports IPv6, you can build an IPv4/IPv6 dual-stack network. **Table 9-2** shows where IPv4/IPv6 dual-stack networks can be used.

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

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

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 A	Application scenarios of IPv6 EIPs

Applica tion Scenari o	Description	Subnet	ECS
Public commu nicatio n using IPv6 address es	Your applications deployed on ECSs need to provide services accessible from the Internet using IPv6 addresses.	IPv4 CIDR block	 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.

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. 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 an Existing 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
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. 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

Table 9-4 Operation guide on IPv6 networks

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

VPC Flow Log

VPC flow logs help you collect traffic information about instances in a specified VPC, including inbound and outbound traffic. After creating a flow log, you can view the flow log records in the log group that you configured.

Flow logs can help you:

- Monitor the traffic of security groups and network ACL and optimize their rules.
- Monitor the traffic of network instances and analyze network attacks.
- Determine the direction of the traffic to and from network interfaces.

The collection of flow log data does not affect the throughput or latency of your network. You can create or delete flow logs as required, which does not affect your network performance.

NOTE

The VPC flow log function itself is free of charge, but you may be charged for other resources used. For example, if data is stored in Log Tank Service (LTS), you will be billed based on the LTS standards. For details, see the *Log Tank Service User Guide*.

VPC Flow Log Data

You can create a flow log for a network interface, subnet, or VPC. If you create a flow log for a subnet or a VPC, each network interface in the subnet or VPC is monitored.

The traffic of a monitored network interface is collected and flow log data is generated, including the network interface ID, source address, destination address, source port, destination port, and packet size of the traffic.

Field	Description	Example
version	VPC flow log version.	1
project-id	ID of the project that the object monitored by flow log belongs to.	5f67944957444bd6bb 4fe3b367de8f3d
interface-id	ID of the network interface that the flow log data is generated for.	1d515d18-1b36-47dc -a983-bd6512aed4bd
srcaddr	Source address.	192.168.0.154
dstaddr	Destination address.	192.168.3.25
srcport	Source port.	38929
dstport	Destination port.	53
protocol	Internet Assigned Numbers Authority (IANA) protocol number. 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 that is associated with the traffic.	ACCEPT
	• ACCEPT: The traffic was allowed by security groups or network ACLs.	
	 REJECT: The traffic was denied by security groups or network ACLs. 	

Table 10-1	VPC flow	log field	description
-------------------	----------	-----------	-------------

Field	Description	Example
log-status	 The logging status of the VPC flow log. OK: Data is logged normally to the chosen destinations. 	ОК
	• NODATA : There was no traffic to or from the network interface during the capture window.	
	• SKIPDATA : Some flow log records were skipped during the capture window. This may be caused by an internal capacity constraint or an internal error.	
	Example:	
	When Filter is set to Accepted traffic , if there is accepted traffic, the value of log-status is OK . If there is no accepted traffic, the value of log-status is NODATA regardless of whether there is rejected traffic. If some accepted traffic is abnormally skipped, the value of log- status is SKIPDATA .	

Constraints

Currently, S2, M2, Hc2, D2, P1, G3, Pi1, S3, C3, M3, H3, D3, Ir3, I3, Sn3, E3, C3ne, M3ne, G5, P2v, Ai1, C6, M6, D6, S6, C6s, S7, C7, M7, E7, D7, Ir7, I7, S7n, C7n, M7n, and I7n ECSs support VPC flow logs.

Currently, ECSs of series 7 support VPC flow logs only in certain regions. This function will be successively launched in each region.

For details about ECS types, see ECS Types.

- Each account can have up to 10 VPC flow logs in a region.
- By default, up to 400,000 flow log records can be generated for a single network interface in a collection period (10 minutes). Excess records will be discarded.

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.

Parameter	Description	Example Value
Name	 The VPC flow log name. The name: Can contain 1 to 64 characters. Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	flowlog-495d
Resource Type	Type of the resource whose traffic is to be logged. The options can be one of the following: • NIC • Subnet • VPC	NIC
Resource	The specific resource whose traffic is to be logged. NOTE We recommend that you select an ECS 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 of the ECS's network interface.	N/A
Filter	 All traffic: Both accepted and rejected traffic of the specified resource will be logged. Accepted traffic: Only accepted traffic of the specified resource will be logged. Accepted traffic refers to the traffic allowed by the security group and network ACL. Rejected traffic: Only rejected traffic of the specified resource will be logged. Rejected traffic refers to the traffic denied by the security group and 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

Table 10-2 Parameter descriptions
Parameter	Description	Example Value
Description	Supplementary information about the VPC flow log. This parameter is optional.	N/A
	The VPC flow log description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

D 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. Click Create Now.

Return to the VPC flow log list and 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 flow log and click **View Log Record** in the **Operation** column.

The Log Management page is displayed.

3. In the log group list, locate the target log group and click the name of the target log stream under it.

The log stream details page is displayed.

4. Enter key information in the search box to quickly find the flow log to be viewed.

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 10-3 provides you with flow log examples.

Scenario	Example Value				
A flow log record in which data was recorded during the capture window	Value 1 indicates the VPC flow log version. Traffic with a size of 96 bytes to the network interface (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 . 1 5f67944957444bd6bb4fe3b367de8f3d 1d515d18-1b36-47dc-a983- bd6512aed4bd 192.168.0.154 192.168.3.25 38929 53 17 1 96 1548752136 1548752736 ACCEPT OK				
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 a 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 network interface 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**.

10.6 VPC Flow Log Configuration Examples

10.6.1 Viewing the Traffic of ECSs from the Same VPC

Solution Architecture

In this example, there are two subnets (**Subnet-A01** and **Subnet-A02**) in a VPC (VPC-A). ECS-01 is running in **Subnet-A01** and ECS-02 and ECS-03 are running in **Subnet-A02**. ECS-01 communicates with both ECS-02 and ECS-03. If something goes wrong to the communication between ECS-01 and ECS-02, the O&M engineer needs to check the traffic between the two ECSs. To locate issues, the O&M engineer needs to create a VPC flow log and collect the flow log of the network interface attached to ECS-01.



Figure 10-1 Viewing the traffic of ECSs in a VPC

Constraints

For details about the restrictions on flow logs, see Constraints.

Resource Planning

In this example, the VPC, subnets, flow log, 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.

Resource	Quan tity	Description
VPC and subnet	VPC: 1	• Name: Set it as needed. In this example, VPC-A is used.
	Subne t: 2	 IPv4 CIDR Block (VPC): 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 and Subnet-A02 are used.
		 IPv4 CIDR Block (Subnet): Set it as needed. In this example, the CIDR block of Subnet-A01 is 192.168.0.0/24 and that of Subnet-A02 is 192.168.1.0/24.

Resource	Quan tity	Description		
ECS	2	Configure the two ECSs as follows:		
		• ECS Name: Set it as needed. In this example, the ECSs are named ECS-01 and ECS-02.		
		• ECS flavor: In this example, flow logs of the network interface attached to ECS-01 are collected. Select the ECS flavor that supports flow logs. For details, see Constraints . There are no such restrictions on selecting the flavor for ECS-02 .		
		• Image : Select an image as needed. In this example, a public image (CentOS 8.0 64bit) is used.		
		• System Disk : In this example, a general-purpose SSD disk of 40 GiB is used.		
		• Data Disk : Set it as needed. In this example, no data disk is used.		
		Network		
		 VPC: Select your required VPC. In this example, VPC-A is used. 		
		 Subnet: Select your required subnet. In this example, select Subnet-A01 for ECS-01 and Subnet-A02 for ECS-02. 		
		• Security Group: In this example, the two ECSs are associated with the same security group (Sg-X). Ensure that all rules in Table 10-5 are added. If the ECSs are associated with different security groups, you also need to add additional rules.		
		For example, if ECS-01 is associated with Sg-X and ECS-02 is associated with Sg-A , add the rules in Table 10-6 to Sg-X and Sg-A to allow the two ECSs to communicate with each other.		
		• EIP: Select Not required.		
		 Private IP address: In this example, use 192.168.0.66 for ECS-01 and 192.168.1.31 for ECS-02. 		

Resource	Quan tity	Description		
VPC flow log	1	 Name: Set it as needed. In this example, name it flowlog-A. 		
		• Resource Type : Select NIC in this example.		
		 Resource: Set it as needed. In this example, select the network interface (IP address: 192.168.0.66) of ECS-01. 		
		• Filter: Select All traffic in this example.		
		• Log Group : Select an existing or create a log group. The log group of this example is as follows:		
		 Log Group Name: Set it as needed. In this example, lts-group-A is used. 		
		 Log Retention (Days): Set it as needed. In this example, 30 is used. 		
		• Log Stream: Select an existing or create a log stream. The log stream of this example is as follows:		
		 Log Group Name: In this example, the log group name is lts-group-A. 		
		 Log Stream Name: Set it as needed. In this example, lts-topic-A is used. 		
		 Log Storage: You are advised to enable this function for log search and analysis. 		
		 Log Retention (Days): Set it as needed. In this example, 30 is used. 		

Table 10-5 Security group Sg-X rules

Direct ion	Acti on	Туре	Protoco l & Port	Source/ Destination	Description
Inbou nd	Allo w	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows remote logins to Linux ECSs over SSH port 22.
Inbou nd	Allo w	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows remote logins to Windows ECSs over RDP port 3389.
Inbou nd	Allo w	IPv4	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv4 addresses.
Inbou nd	Allo w	IPv6	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv6 addresses.

Direct ion	Acti on	Туре	Protoco l & Port	Source/ Destination	Description
Outbo und	Allo w	IPv4	All	Destination: 0.0.0.0/0	Allows ECSs in Sg-X to access the external networks using IPv4 addresses.
Outbo und	Allo w	IPv6	All	Destination: ::/ 0	Allows ECSs in Sg-X to access the external networks using IPv6 addresses.

If the source of an inbound rule is set to 0.0.0.0/0, all external IP addresses are allowed to remotely log in to your instances. Exposing port 22 or 3389 to the public network will leave your instances vulnerable to network risks. To address this issue, set the source to a trusted IP address, for example, the IP address of your local PC.

Table 10-6 Rules of security	groups Sg-X and Sg-A
------------------------------	----------------------

Sec urit y Gro up	Direc tion	Act ion	Typ e	Protocol & Port	Source	Description
Sg- X	Inbou nd	Allo w	IPv4	All	Security group Sg-A	Allows IPv4 traffic from ECSs in Sg-A to reach ECSs in Sg-X .
Sg- A	Inbou nd	Allo w	IPv4	All	Security group Sg-X	Allows IPv4 traffic from ECSs in Sg-X to reach ECSs in Sg-A .

Procedure

Figure 10-2 shows the process for viewing the traffic of ECSs in a VPC.

Figure 10-2 Viewing the traffic of ECSs in a VPC



Step 1: Create Cloud Resources

- 1. Create a VPC with two subnets.
 - For details, see Creating a VPC and Subnet.
- Create two ECSs.
 For details, see Purchasing an ECS.

Step 2: Create a VPC Flow Log

- 1. Create a log group and log stream on the LTS console.
- 2. Create a VPC flow log.

For details, see Creating a VPC Flow Log.

Step 3: View the VPC Flow Log

The flow log collects the information about the traffic flowing through the network interface attached to **ECS-01**.

1. Remotely log in to **ECS-01**.

For details, see Logging In to an ECS.

2. Ping ECS-02 from ECS-01 and collect logs:

ping <private-IP-address-of-ECS-02>

Example command:

ping 192.168.1.31

Information similar to the following is displayed. You can view the flow log records in about 10 minutes. Do not stop the ping command during flow log collection.

[root@ecs-01 ~]# ping 192.168.1.31 PING 192.168.1.31 (192.168.1.31) 56(84) bytes of data. 64 bytes from 192.168.1.31: icmp_seq=1 ttl=64 time=0.292 ms 64 bytes from 192.168.1.31: icmp_seq=2 ttl=64 time=0.186 ms 64 bytes from 192.168.1.31: icmp_seq=3 ttl=64 time=0.162 ms

3. Wait for about 10 minutes and view the VPC flow log information by referring to **Viewing a VPC Flow Log**.

You can enter the IP address (192.168.1.31) of **ECS-02** in the search box to quickly filter the logs of the communication between **ECS-01** and **ECS-02**.

The flow log record is in the following format: <version> <project-id> <interface-id> <srcaddr> <dstaddr> <srcport> <dstport> <protocol> <packets> <bytes> <start> <end> <action> <log-status>

- Example log: 1 f0512a6441dc47189f5e03a428f48267
 ef676eb6-0a0a-4939-85c9-9f8db1d1937c 192.168.0.66 192.168.1.31 8 0 1
 585 57330 1739877133 1739877733 ACCEPT OK
- Log description: The VPC flow log version is 1. The log shows that 585 echo request packets (type=8,code=0) were sent from the source (192.168.0.66) to the destination (192.168.1.31) via the network interface ef676eb6-0a0a-4939-85c9-9f8db1d1937c using ICMP (protocol=1) during 19:12:13 to 19:22:13 (10 minutes), on February 18, 2025. The size of all packets is 57,330 bytes.

For details about flow log data, see VPC Flow Log Data.

Step 4: Configure Cloud Structuring Parsing and Analyze Visualized Logs for the VPC Flow Log

LTS allows you to search for and analyze collected logs and displays log analysis results in a visualized manner.

1. Configure cloud structuring parsing.

For details, see **Cloud Structuring Parsing**.

Table 10-7 Parameters for configuring of	cloud structuring parsing
--	---------------------------

Step	Operation
1	Set the structuring mode to Delimiter .
2	Enter the VPC flow log: 1 f0512a6441dc47189f5e03a428f48267 ef676eb6-0a0a-4939-85c9-9f8db1d1937c 192.168.0.66 192.168.1.31 8 0 1 585 57330 1739877133 1739877733 ACCEPT OK
3	Select Space as the delimiter.
4	Click Intelligent Extraction.
5	In the intelligent extraction field list, change the field name to the flow log parameters:
	version, project-id, interface-id, srcaddr, dstaddr, srcport, dstport, protocol, packets, bytes, start, end, action, and log-status

2. Analyze the logs based on the cloud structuring parsing.

The following shows two visualized log analysis methods:

- Visualize logs in statistical charts. Statistical charts, such as tables, bar charts, and line charts, are rendered by LTS based on SQL query syntax.
 - i. On the **Log Analysis** tab, enter the required statement in the search box by referring to **Using SQL Analysis Syntax** to obtain the required logs.

The following takes the traffic of **ECS-01** every hour as an example. SELECT TIME_FORMAT(TIME_CEIL(__time, 'PT1H'), 'yyyy-MM-dd HH:mm:ss') as "time", count(1) as pv group by "time"

- ii. On the right of the page, configure the time and other information.For more information about statistical charts, see Statistical Charts.
- Visualize logs in dashboards. The dashboard is a real-time data visualization tool provided by LTS.
 - i. Ingest VPC logs to LTS by referring to Ingesting VPC Logs to LTS.
 - ii. After VPC logs are ingested, choose **Dashboards** > **VPC dashboard templates** > **VPC Flow Logs** on the LTS console.

Wait for a few minutes and view the log data. For more information about the dashboard, see **VPC Dashboard Template**.

10.6.2 Viewing the Traffic Between VPCs Connected by a VPC Peering Connection

Solution Architecture

In this example, a VPC peering connection **Peer-AB** is used to enable communications between two VPCs (**VPC-A** and **VPC-B**). If O&M engineers need to view the traffic between the two VPCs, you can create a VPC flow log and collect the flow log of **VPC-A**.



Figure 10-3 Viewing traffic between VPCs connected by a VPC peering connection

Constraints

For details about the restrictions on flow logs, see **Constraints**.

Resource Planning

In this example, the VPCs, subnets, ECSs, VPC peering connection, and VPC flow log 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 10)-8 Resourc	e planning
----------	-------------	------------

Resource	Quan tity	Description	
VPC and subnet	VPC: 2	 Name: Set it as needed. In this example, VPC-A and VPC-B are used. 	
	Subne t: 2	• IPv4 CIDR Block (VPC): Set it as needed. In this example, 192.168.0.0/16 is used for VPC-A, and 172.16.0.0/16 is used for VPC-B.	
		• Subnet Name: Set it as needed. In this example, Subnet-A01 and Subnet-B01 are used.	
		 IPv4 CIDR Block (Subnet): Set it as needed. In this example, the CIDR block of Subnet-A01 is 192.168.0.0/24 and that of Subnet-B01 is 172.16.0.0/24. 	
		 Route table: A VPC comes with a default route table. In this example, the default route table of VPC-A is rtb-VPC-A, and that of VPC-B is rtb-VPC-B. 	

Resource	Quan tity	Description
ECS	2	Configure the two ECSs as follows:
		• ECS Name: Set it as needed. In this example, the ECSs are named ECS-01 and ECS-02.
		• ECS flavor: In this example, flow logs of the network interface attached to ECS-01 are collected. Select the ECS flavor that supports flow logs. For details, see Constraints . There are no such restrictions on selecting the flavor for ECS-02 .
		• Image: Set it as needed. In this example, public image Huawei Cloud EulerOS 2.0 Standard 64 bit is used.
		• System Disk : In this example, a general-purpose SSD disk of 40 GiB is used.
		• Data Disk : Set it as needed. In this example, no data disk is used.
		Network
		 VPC: Select your required VPC. In this example, select VPC-A for ECS-01 and VPC-B for ECS-02.
		 Subnet: Select your required subnet. In this example, select Subnet-A01 for ECS-01 and Subnet-B01 for ECS-02.
		• Security Group : In this example, the two ECSs are associated with the same security group (Sg-X). Ensure that all rules in Table 10-9 are added. If the ECSs are associated with different security groups, you also need to add additional rules.
		For example, if ECS-01 is associated with Sg-X and ECS-02 is associated with Sg-A , add the rules in Table 10-10 to Sg-X and Sg-A to allow the two ECSs to communicate with each other.
		• EIP: Select Not required.
		• Private IP address: In this example, use 192.168.0.66 for ECS-01 and 172.16.0.31 for ECS-02 .

Resource	Quan tity	Description
VPC peering connection	1	• VPC Peering Connection Name: Set it as needed. In this example, Peering-AB is used.
		• Local VPC: Set it as planned. In this example, select VPC-A with the CIDR block of 192.168.0.0/16.
		• Account: Set it as planned. In this example, select My account, indicating that the VPCs connected by the VPC peering connection are in the same account.
		• Peer VPC : Set it as planned. In this example, select VPC-B with the CIDR block of 172.16.0.0/16.
		• Routes: After the VPC peering connection is created, you need to add routes to the route tables of the local and peer VPCs to connect them. For details about the required routes in this example, see Table 10-11.
VPC flow log	1	 Name: Set it as needed. In this example, name it flowlog-A.
		• Resource Type : In this example, set it to VPC .
		 Resource: Select a resource as needed. In this example, select VPC-A with the CIDR block of 192.168.0.0/16.
		• Filter: Select All traffic in this example.
		• Log Group : Select an existing or create a log group. The log group of this example is as follows:
		 Log Group Name: Set it as needed. In this example, lts-group-A is used.
		 Log Retention (Days): Set it as needed. In this example, 30 is used.
		• Log Stream: Select an existing or create a log stream. The log stream of this example is as follows:
		 Log Group Name: In this example, the log group name is lts-group-A.
		 Log Stream Name: Set it as needed. In this example, lts-topic-A is used.
		 Log Storage: You are advised to enable this function for log search and analysis.
		 Log Retention (Days): Set it as needed. In this example, 30 is used.

Direct ion	Acti on	Туре	Protoco l & Port	Source/ Destination	Description
Inbou nd	Allo w	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows remote logins to Linux ECSs over SSH port 22.
Inbou nd	Allo w	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows remote logins to Windows ECSs over RDP port 3389.
Inbou nd	Allo w	IPv4	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv4 addresses.
Inbou nd	Allo w	IPv6	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv6 addresses.
Outbo und	Allo w	IPv4	All	Destination: 0.0.0.0/0	Allows ECSs in Sg-X to access the external networks using IPv4 addresses.
Outbo und	Allo w	IPv6	All	Destination: ::/ 0	Allows ECSs in Sg-X to access the external networks using IPv6 addresses.

Table 10-9 Security group Sg-X rules

If the source of an inbound rule is set to 0.0.0/0, all external IP addresses are allowed to remotely log in to your instances. Exposing port 22 or 3389 to the public network will leave your instances vulnerable to network risks. To address this issue, set the source to a trusted IP address, for example, the IP address of your local PC.

Table 10-10 Rules of security gro	oups Sg-X and Sg-A
-----------------------------------	--------------------

Sec urit y Gro up	Direc tion	Act ion	Тур e	Protocol & Port	Source	Description
Sg- X	Inbou nd	Allo w	IPv4	All	Security group Sg-A	Allows IPv4 traffic from ECSs in Sg-A to reach ECSs in Sg-X .

Sec urit y Gro up	Direc tion	Act ion	Тур e	Protocol & Port	Source	Description
Sg- A	Inbou nd	Allo w	IPv4	All	Security group Sg-X	Allows IPv4 traffic from ECSs in Sg-X to reach ECSs in Sg-A .

Table 10-11 VPC route tables

VPC Name	Route Table	Destination	Next Hop Type	Next Hop	Route Type
VPC-A	Default route table: rtb-VPC-A	VPC-B CIDR block: 172.16.0.0/16	VPC peering connecti on	Peering- AB	Custom
VPC-B	Default route table: rtb-VPC-B	VPC-A CIDR block: 192.168.0.0/16	VPC peering connecti on	Peering- AB	Custom

Procedure

Figure 10-4 shows the process for viewing the traffic between ECSs in different VPCs.

Figure 10-4 Process for viewing traffic between VPCs connected by a VPC peering connection



Step 1: Create Cloud Resources

1. Create two VPCs, each with a subnet.

For details, see Creating a VPC and Subnet.

- Create two ECSs.
 For details, see Purchasing an ECS.
- 3. Create a VPC peering connection and add routes to the route tables of the two VPCs.
 - If the two VPCs belong to the same account, refer to Creating a VPC
 Peering Connection to Connect Two VPCs in the Same Account.

If the two VPCs belong to different accounts, refer to Creating a VPC
 Peering Connection to Connect Two VPCs in Different Accounts.

Step 2: Create a VPC Flow Log

- 1. Create a log group and log stream on the LTS console.
- Create a VPC flow log.
 For details, see Creating a VPC Flow Log.

Step 3: View the VPC Flow Log

The flow log collects the information about the traffic flowing through VPC-A.

- Remotely log in to ECS-01 in VPC-A.
 For details, see Logging In to an ECS.
- 2. Ping ECS-02 in VPC-B from ECS-01 in VPC-A and collect logs:

ping <private-IP-address-of-ECS-02>

Example command:

ping 172.16.0.31

Information similar to the following is displayed. You can view the flow log records in about 10 minutes. Do not stop the ping command during flow log collection.

[root@ecs-01 ~]# ping 172.16.0.31 PING 172.16.0.31 (172.16.0.31) 56(84) bytes of data. 64 bytes from 172.16.0.31: icmp_seq=1 ttl=63 time=0.510 ms 64 bytes from 172.16.0.31: icmp_seq=2 ttl=63 time=0.392 ms 64 bytes from 172.16.0.31: icmp_seq=3 ttl=63 time=0.332 ms

3. Wait for about 10 minutes and view the VPC flow log information by referring to **Viewing a VPC Flow Log**.

You can enter the IP address (172.16.0.31) of **ECS-02** in the search box to quickly filter the logs of the communication between **ECS-01** and **ECS-02**.

The flow log record is in the following format: <version> <project-id> <interface-id> <srcaddr> <dstaddr> <srcport> <dstport> <protocol> <packets> <bytes> <start> <end> <action> <log-status>

- Example log: 1 857dcccea8644ce1abcfc57b6474c5ad
 10b6d5df-8abe-4bc5-85ba-f01ce445dacc 192.168.0.66 172.16.0.31 8 0 1
 258 25284 1740022820 1740023420 ACCEPT OK
- Log description: The VPC flow log version is 1. The log shows that 258 echo request packets (type=8,code=0) were sent from the source (192.168.0.66) to the destination (172.16.0.31) via the network interface 10b6d5df-8abe-4bc5-85ba-f01ce445dacc using ICMP (protocol=1) during 11:40:20 to 11:50:20 (10 minutes), on February 20, 2025. The size of all packets is 25,284 bytes.

For details about flow log data, see VPC Flow Log Data.

Step 4: Configure Cloud Structuring Parsing and Analyze Visualized Logs for the VPC Flow Log

LTS allows you to search for and analyze collected logs and displays log analysis results in a visualized manner.

Configure cloud structuring parsing.
 For details, see Cloud Structuring Parsing.

Table 10-12	Parameters	for	configuring	cloud	structuring	parsing
	rurumeters	101	connigannig	ciouu	Judecuring	pursning

Step	Operation
1	Set the structuring mode to Delimiter .
2	Enter the VPC flow log: 1 f0512a6441dc47189f5e03a428f48267 ef676eb6-0a0a-4939-85c9-9f8db1d1937c 192.168.0.66 192.168.1.31 8 0 1 585 57330 1739877133 1739877733 ACCEPT OK
3	Select Space as the delimiter.
4	Click Intelligent Extraction.
5	In the intelligent extraction field list, change the field name to the flow log parameters: version, project-id, interface-id, srcaddr, dstaddr, srcport, dstport, protocol, packets, bytes, start, end, action, and log-status

2. Analyze the logs based on the cloud structuring parsing.

The following shows two visualized log analysis methods:

- Visualize logs in statistical charts. Statistical charts, such as tables, bar charts, and line charts, are rendered by LTS based on SQL query syntax.
 - i. On the **Log Analysis** tab, enter the required statement in the search box by referring to **Using SQL Analysis Syntax** to obtain the required logs.

The following takes the traffic of **ECS-01** every hour as an example. SELECT TIME_FORMAT(TIME_CEIL(__time, 'PT1H'), 'yyyy-MM-dd HH:mm:ss') as "time", count(1) as pv group by "time"

- ii. On the right of the page, configure the time and other information. For more information about statistical charts, see **Statistical Charts**.
- Visualize logs in dashboards. The dashboard is a real-time data visualization tool provided by LTS.
 - i. Ingest VPC logs to LTS by referring to Ingesting VPC Logs to LTS.
 - ii. After VPC logs are ingested, choose **Dashboards** > **VPC dashboard templates** > **VPC Flow Logs** on the LTS console.

Wait for a few minutes and view the log data. For more information about the dashboard, see **VPC Dashboard Template**.

10.6.3 Viewing the Traffic Between ECSs in Different VPCs Connected by an Enterprise Router

Solution Architecture

In this example, enterprise router **ER-X** is used to connect two VPCs (**VPC-A** and **VPC-B**). To view the traffic between the two VPCs and locate issues, you need to create a flow log for the enterprise router to collect the logs of **VPC-A** attachment.



Figure 10-5 Viewing the traffic between ECSs in different VPCs

Resource Planning

In this example, the VPCs, subnets, ECSs, enterprise router, and flow log 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 10-13 Resource planning

Resource	Quan tity	Description	
VPC and subnet	VPC: 2	• Name: Set it as needed. In this example, VPC-A and VPC-B are used.	
	Subne t: 2	 IPv4 CIDR Block (VPC): Set it as needed. In this example, 192.168.0.0/16 is used for VPC-A, and 172.16.0.0/16 is used for VPC-B. 	
		 Subnet Name: Set it as needed. In this example, Subnet-A01 and Subnet-B01 are used. 	
		 IPv4 CIDR Block (Subnet): Set it as needed. In this example, the CIDR block of Subnet-A01 is 192.168.0.0/24 and that of Subnet-B01 is 172.16.0.0/24. 	
		 Route table: A VPC comes with a default route table. In this example, the default route table of VPC-A is rtb-VPC-A, and that of VPC-B is rtb-VPC-B. 	

Resource	Quan tity	Description
ECS	2	Configure the two ECSs as follows:
		• ECS Name : Set it as needed. In this example, the ECSs are named ECS-01 and ECS-02 .
		• ECS flavor: Set it as need. Ensure that the flavor can meet service requirements.
		• System Disk : In this example, a general-purpose SSD disk of 40 GiB is used.
		• Data Disk : Set it as needed. In this example, no data disk is used.
		Network
		 VPC: Select your required VPC. In this example, select VPC-A for ECS-01 and VPC-B for ECS-02.
		 Subnet: Select your required subnet. In this example, select Subnet-A01 for ECS-01 and Subnet-B01 for ECS-02.
		• Security Group: In this example, the two ECSs are associated with the same security group (Sg-X). Ensure that all rules in Table 10-14 are added. If the ECSs are associated with different security groups, you also need to add additional rules.
		For example, if ECS-01 is associated with Sg-X and ECS-02 is associated with Sg-A , add the rules in Table 10-15 to Sg-X and Sg-A to allow the two ECSs to communicate with each other.
		• EIP: Select Not required.
		• Private IP address: In this example, use 192.168.0.66 for ECS-01 and 172.16.0.31 for ECS-02 .
Enterprise	1	• Name: Set it as needed. In this example, ER-X is used.
router		• ASN : Set it as needed. In this example, 64513 is used.
		• Default Route Table Association : Enable this option.
		 Default Route Table Propagation: Enable this option.
		• Auto Accept Shared Attachments: Set it as needed. In this example, enable this option.
		 In this example, you need to add two VPC attachments to the enterprise router.
		 VPC-A attachment: er-attach-vpc-A
		 VPC-B attachment: er-attach-vpc-B

Resource	Quan tity	Description
Enterprise router flow	1	 Name: Set it as needed. In this example, name it flowlog-ER.
log		• Resource Type: In this example, set it to VPC.
		 Resource: Select a resource as needed. In this example, select the er-attach-vpc-A attachment corresponding to VPC-A.
		• Log Group: Select an existing or create a log group. The log group of this example is as follows:
		 Log Group Name: Set it as needed. In this example, lts-group-ER is used.
		 Log Retention (Days): Set it as needed. In this example, 30 is used.
		• Log Stream: Select an existing or create a log stream. The log stream of this example is as follows:
		 Log Group Name: In this example, the log group name is lts-group-ER.
		 Log Stream Name: Set it as needed. In this example, lts-topic-ER is used.
		 Log Storage: You are advised to enable this function for log search and analysis.
		 Log Retention (Days): Set it as needed. In this example, 30 is used.

Table 10-14 Security group Sg-X rules

Direct ion	Acti on	Туре	Protoco l & Port	Source/ Destination	Description
Inbou nd	Allo w	IPv4	TCP: 22	Source: 0.0.0.0/0	Allows remote logins to Linux ECSs over SSH port 22.
Inbou nd	Allo w	IPv4	TCP: 3389	Source: 0.0.0.0/0	Allows remote logins to Windows ECSs over RDP port 3389.
Inbou nd	Allo w	IPv4	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv4 addresses.
Inbou nd	Allo w	IPv6	All	Source: current security group (Sg-X)	Allows the ECSs in Sg-X to communicate with each other using IPv6 addresses.

Direct ion	Acti on	Туре	Protoco l & Port	Source/ Destination	Description
Outbo und	Allo w	IPv4	All	Destination: 0.0.0.0/0	Allows ECSs in Sg-X to access the external networks using IPv4 addresses.
Outbo und	Allo w	IPv6	All	Destination: ::/ 0	Allows ECSs in Sg-X to access the external networks using IPv6 addresses.

If the source of an inbound rule is set to 0.0.0.0/0, all external IP addresses are allowed to remotely log in to your instances. Exposing port 22 or 3389 to the public network will leave your instances vulnerable to network risks. To address this issue, set the source to a trusted IP address, for example, the IP address of your local PC.

 Table 10-15 Rules of security groups Sg-X and Sg-A

Sec urit y Gro up	Direc tion	Act ion	Typ e	Protocol & Port	Source	Description
Sg- X	Inbou nd	Allo w	IPv4	All	Security group Sg-A	Allows IPv4 traffic from ECSs in Sg-A to reach ECSs in Sg-X .
Sg- A	Inbou nd	Allo w	IPv4	All	Security group Sg-X	Allows IPv4 traffic from ECSs in Sg-X to reach ECSs in Sg-A .

Procedure

Figure 10-6 shows the process for viewing the traffic between ECSs in different VPCs.

Figure 10-6 Process for viewing traffic between VPCs connected by an enterprise router



Step 1: Create Cloud Resources

1. Create two VPCs, each with a subnet.

For details, see Creating a VPC and Subnet.

- Create two ECSs.
 For details, see Purchasing an ECS.
- 3. Create an enterprise router.

For details, see Creating an Enterprise Router.

4. Attach the two VPCs to the enterprise router.

If you enable **Auto Add Routes** when creating a VPC attachment, you do not need to manually add static routes to the VPC route table. After the VPC attachments are added, the two VPCs can communicate with each other.

For details, see Creating VPC Attachments for the Enterprise Router.

For details about network planning using an enterprise router, see .

Step 2: Create an Enterprise Router Flow Log

- 1. Create a log group and log stream on the LTS console.
- 2. Create an enterprise router flow log.

Step 3: View the Flow Log

The flow log collects the information of traffic flowing through attachment **er-attach-vpc-A** corresponding to **VPC-A**.

- Remotely log in to ECS-01 in VPC-A.
 For details, see Logging In to an ECS.
- 2. Ping ECS-02 in VPC-B from ECS-01 in VPC-A and collect logs:

ping <private-IP-address-of-ECS-02>

Example command:

ping 172.16.0.31

Information similar to the following is displayed. You can view the flow log records in about 10 minutes. Do not stop the ping command during flow log collection.

```
[root@ecs-01 ~]# ping 172.16.0.31
PING 172.16.0.31 (172.16.0.31) 56(84) bytes of data.
64 bytes from 172.16.0.31: icmp_seq=1 ttl=63 time=0.510 ms
64 bytes from 172.16.0.31: icmp_seq=2 ttl=63 time=0.392 ms
64 bytes from 172.16.0.31: icmp_seq=3 ttl=63 time=0.332 ms
```

3. Wait for about 10 minutes and view the flow log information by referring to **Viewing Details About a Flow Log**.

You can enter the IP address (172.16.0.31) of **ECS-02** in the search box to quickly filter the logs of the communication between **ECS-01** and **ECS-02**.

Flow log format:

<version> <project_id> <resource_id> <instance_id> <srcaddr> <dstaddr> <srcport> <dstport> <protocol> <packets> <bytes> <start> <end> <direct>

Example log: 1 f0512a6441dc47189f5e03a428f48267
 37befd9d-58a8-4a5f-9cb1-13a3fe563c20 bdc50d41-

a33a-4bf5-9391-4957369d17b6 192.168.0.66 172.16.0.31 8 0 1 586 57428 1742872080 1742872680 ingress

Log description: The enterprise router log version is 1. Within the 10 minutes from 11:08:00 to 11:18:00 on March 25, 2025, the traffic from VPC-A attachment (ID: 37befd9d-58a8-4a5f-9cb1-13a3fe563c20) flowing to (ingress) the enterprise router (bdc50d41-a33a-4bf5-9391-4957369d17b6) was recorded. The log shows that 586 echo request (type=8,code=0) packets were sent from the source 192.168.0.66 to the destination 172.16.0.31 through ICMP (protocol=1). The size of all packets is 57,428 bytes.

Step 4: Configure Cloud Structuring Parsing and Analyze Visualized Logs for the Enterprise Router Flow Log

LTS allows you to search for and analyze collected logs and displays log analysis results in a visualized manner.

1. Configure cloud structuring parsing.

For details, see **Cloud Structuring Parsing**.

fable 10-16 Parameters	s for	configuring	cloud	structuring	parsing
------------------------	-------	-------------	-------	-------------	---------

Step	Operation
1	Select Structuring Template to structure logs.
2	In the system template list, select ER Enterprise Router .

2. Analyze the logs based on the cloud structuring parsing.

The following shows two visualized log analysis methods:

- Visualize logs in statistical charts. Statistical charts, such as tables, bar charts, and line charts, are rendered by LTS based on SQL query syntax.
 - i. On the **Log Analysis** tab, enter the required statement in the search box by referring to **Using SQL Analysis Syntax** to obtain the required logs.

The following takes the traffic of **ECS-01** every hour as an example. SELECT TIME_FORMAT(TIME_CEIL(__time, 'PT1H'), 'yyyy-MM-dd HH:mm:ss') as "time", count(1) as pv group by "time"

ii. On the right of the page, configure the time and other information.

For more information about statistical charts, see **Statistical Charts**.

- Visualize logs in dashboards. The dashboard is a real-time data visualization tool provided by LTS.
 - After an enterprise router flow log is created, choose Dashboards > ER dashboard templates > Enterprise Router Flow Log Center on the LTS console.

On the enterprise router dashboard details page, wait for several minutes and view the flow log data. In this example, select the instance and the attachment to view flow log information. For more dashboard information, see **ER Dashboard Template**.

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.





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 a New EIP

- 1. Go to the **Buy EIP** page.
- 2. Configure parameters as prompted.

lte m	Parameter	Description	Example Value
Bas ic Co nfi gur ati on	Billing Mode	 The following options are available: Yearly/Monthly: You pay upfront for the amount of time you expect to use the instance. You need to make sure you have a valid payment method configured first. Pay-per-use: You can start using the EIP first and then pay as you go. You are billed based on the EIP usage duration (by bandwidth) or used traffic (by traffic). 	Pay-per-use
Bas ic Co nfi gur ati on	Region	The desired region. Resources in different regions cannot communicate with each other over internal networks. For low network latency and quick resource access, select the region nearest to where your services will be accessed. The region selected for the EIP is its geographical location.	-

Table 11-1 Parameter descriptions

lte m	Parameter	Description	Example Value
Ba nd wid th Det ails	EIP Type	 Dynamic BGP: Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails. NOTE Dynamic BGP is suitable for communications in CN- Hong Kong or communications between CN-Hong Kong and regions outside the Chinese mainland. If Dynamic BGP is used to access the regions in the Chinese mainland, data is forwarded through international egress routes, which may result in high latency and packet loss. If you need lower latency and better stability to access to the regions in the Chinese mainland, you are advised to select Premium BGP. 	Dynamic BGP
Ba nd wid th Det ails	Billed By	 How the EIP bandwidth will be billed. This parameter is available only when you set Billing Mode to Pay-per- use. Bandwidth: You specify a maximum bandwidth and pay for the amount of time you use the bandwidth. This is suitable for scenarios with heavy or stable traffic. Traffic: You specify a maximum bandwidth and pay for the total outbound traffic you use. This is suitable for scenarios with light or sharply fluctuating traffic. Shared Bandwidth: The bandwidth can be shared by multiple EIPs and is suitable for scenarios with staggered traffic. 	Bandwidth

lte m	Parameter	Description	Example Value
Ba nd wid th Det ails	Bandwidth (Mbit/s)	The bandwidth size in Mbit/s.	100
Ba nd wid th Det ails	Bandwidth Name	 The name of the bandwidth. The name: Can contain 1 to 64 characters. Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	bandwidth
DD oS Pro tec tio n	DDoS Protection	Cloud Native Anti-DDoS Basic Cloud Native Anti-DDoS Basic provides up to a certain amount (for example, less than 5 Gbit/s) of DDoS mitigation capacity for free. The actual thresholds are displayed on the console. If the attack to an EIP exceeds the threshold, the EIP will be blocked.	-
EIP Det ails	EIP Name	 The name of the EIP. The name: Can contain 1 to 64 characters. Can contain letters, digits, underscores (_), hyphens (-), and periods (.). 	eip-test
EIP Det ails	Enterprise Project	The enterprise project that the EIP belongs to. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is default . For details about creating and managing enterprise projects, see the Enterprise Management User Guide.	default

lte m	Parameter	Description	Example Value
EIP Det ails	Tag	Tags help you quickly identify, organize, and search for your EIPs. For more information about tags, see Managing EIP Tags .	 Key: Ipv4_key1 Value: 3005eip
Mo nit ori ng	Monitoring	Basic monitoring is 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.	-
Pur cha se Det ails	Required Duration	The duration for which the EIP will be used. The duration must be specified if the Billing Mode is set to Yearly/Monthly .	1 month
Pur cha se Det ails	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. Monthly subscription: The subscription is renewed every month. Yearly subscription: The 	-
		subscription is renewed each year.	
Pur cha se Det ails	Quantity	The number of EIPs you want to assign. The quantity must be specified if the Billing Mode is set to Pay-per-use .	1

D 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 is no shared bandwidth, create one 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.
 - 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.
- 3. Click Next.
- 4. On the confirmation page:
 - If you select **Pay-per-use** for **Billing Mode**, click **Submit**.
 - If you select Yearly/Monthly for Billing Mode, click Pay Now.

On the payment page, confirm the order information, and click **Confirm**.

If you click **Buy Shared Bandwidth** when you buy an EIP, you also need to pay for the bandwidth.

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. Go to the **EIP list** page.
- 2. On the displayed page, locate the row that contains the target EIP, and click **Unbind** in the **Operation** column.

A confirmation dialog box is displayed.

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

In the EIP list, the target EIP has no associated instance.

Releasing a single EIP

- 1. Go to the **EIP list** page.
- In the EIP list, locate the row that contains the EIP and choose More > Release in the Operation column.
 A confirmation dialog box is displayed.
- Click Yes in the displayed dialog box.
 You can find that the EIP is not in the EIP list.

Unbinding multiple EIPs at once

- 1. Go to the **EIP list** page.
- 2. On the displayed page, select the EIPs to be unbound.
- 3. In the upper left corner of the EIP list, click **Unbind**. A confirmation dialog box is displayed.
- Click Yes in the displayed dialog box.
 In the EIP list, the target EIPs have no associated instances.

Releasing multiple EIPs at once

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- 3. On the displayed page, select the EIPs to be released.
- 4. Click the **Release** button located above the EIP list.
- 5. Click **Yes** in the displayed dialog box.

11.4 Modifying an EIP Bandwidth

Procedure

- 1. Go to the **EIP list** page.
- 2. Locate the EIP whose bandwidth you want to modify, choose **More** > **Modify Bandwidth** in the **Operation** column.
 - 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
- 3. Modify the bandwidth parameters as prompted.
- 4. Click **Next**.
- 5. 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. Go to the **EIP list** page.
- 2. On the EIP list page, select one or more EIPs and click **Export** in the upper left corner.

The system will automatically export all EIPs to an Excel file and download the file to a local directory.

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 11-2** lists the tag key and value requirements.

Parameter	Requirement	Example Value
Кеу	 Cannot be left blank. Must be unique for each EIP. 	lpv4_key1
	 Can contain a maximum of 36 characters. 	
	 Can contain only the following character types: 	
	 Uppercase letters 	
	 Lowercase letters 	
	– Digits	
	 Special characters, including hyphens (-) and underscores (_) 	
Value	• Can contain a maximum of 43 characters.	eip-01
	 Can contain only the following character types: 	
	 Uppercase letters 	
	 Lowercase letters 	
	– Digits	
	 Special characters, including hyphens (-) and underscores (_) 	

Procedure

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

- 1. Go to the **EIP list** page.
- 2. Click the search box above the EIP list.
- 3. 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.

4. 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. Go to the **EIP list** page.
- 2. On the displayed page, locate the EIP whose tags you want to manage and click the EIP name.
- 3. 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 **Edit Tag** in the upper left corner. In the displayed dialog box, click **Add Tag**, enter the tag key and value, and click **OK**.

- Edit a tag.

Click **Edit Tag** in the upper left corner. In the displayed dialog box, enter the tag key and value, and click **OK**.

The tag key cannot be modified.

- Delete a tag.

Click **Edit Tag** in the upper left corner. In the displayed dialog box, click **Delete** in the row that contains the target tag, and click **OK**.

11.7 IPv6 EIP

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.

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.

Mo dul e	Parameter	Description	Example Value
Bas ic Co nfi gur ati on	Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across different regions. For lower network latency and faster access to your resources, select the region nearest to you.	N/A
Mo dul e	Parameter	Description	Example Value
--	-----------------------	--	----------------
Bas ic Co nfi gur ati on	Billing Mode	 A shared bandwidth can be billed on a yearly/monthly or pay-per-use basis. 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
Bas ic Co nfi gur ati on	Name	The name of the shared bandwidth.	Bandwidth-001
Bas ic Co nfi gur ati on	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
Ba nd wid th Det ails	Bandwidth Type	 Select a type of the shared bandwidth based on your EIP type. Standard: Dynamic BGP and static BGP EIPs can be added to a shared bandwidth of this type. 	Standard
Ba nd wid th Det ails	Billed By	The billing method for the shared bandwidth. You can specify a shared bandwidth to be billed by bandwidth.	Bandwidth
Ba nd wid th Det ails	Bandwidth (Mbit/s)	The bandwidth size in Mbit/s.	10

Mo dul e	Parameter	Description	Example Value
Re qui red Du rati on	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
Re qui red Du rati	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.	N/A
on		 Monthly subscription: The subscription is renewed every month. 	
		 Yearly subscription: The subscription is renewed each year. 	

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 network interfaces to it.

Adding EIPs to a Shared Bandwidth

- 1. Go to the **shared bandwidth list** page.
- 2. 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**.
- 3. On the **Add Public IP Address** page, select the EIPs or IPv6 addresses to be added.

D 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.
- 4. 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.

Removing EIPs from a Shared Bandwidth

- 1. Go to the **shared bandwidth list** page.
- 2. In the shared bandwidth list, locate the row that contains the bandwidth from which EIPs are to be removed, choose **More** > **Remove Public IP Address** in the **Operation** column.
- 3. On the **Remove Public IP Address** page, select the EIPs or IPv6 addresses to be removed.
- 4. Set the EIP bandwidth after the EIP is removed.
- 5. Click OK.

12.5 Modify 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).
- You can perform the following operations on a yearly/monthly shared bandwidth:
 - Increasing a Shared Bandwidth (Yearly/Monthly): The change will be applied immediately and the price difference will be billed accordingly.
 - **Decreasing a Shared Bandwidth (Yearly/Monthly)**: 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. Go to the **shared bandwidth list** page.
- 2. 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.
- 3. Click Next.
- 4. Click Submit.

Increasing a Shared Bandwidth (Yearly/Monthly)

- 1. Go to the **shared bandwidth list** page.
- 2. In the shared bandwidth list, locate the row that contains the target shared bandwidth, and click **Modify Bandwidth** in the **Operation** column.
- 3. Select Increase bandwidth and click Continue.
- 4. In the **New Configuration** area on the **Modify Bandwidth** page, change the bandwidth name and size.
- 5. Click Next.
- 6. 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. Go to the **shared bandwidth list** page.
- 2. In the shared bandwidth list, locate the row that contains the target shared bandwidth, and click **Modify Bandwidth** in the **Operation** column.
- 3. Select **Decrease bandwidth** and click **Continue**.
- 4. In the **New Configuration** area on the **Modify Bandwidth** page, change the bandwidth name and size.
- 5. Click Next.
- 6. 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 or Unsubscribing from 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**.

Deleting a Pay-per-Use Shared Bandwidth

- 1. Go to the **shared bandwidth list** page.
- 2. 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**.
- 3. Confirm the deletion as prompted. Click **OK**.

Unsubscribing from a Yearly/Monthly Shared Bandwidth

- 1. Go to the **shared bandwidth list** page.
- 2. In the shared bandwidth list, locate the row that contains the yearly/monthly shared bandwidth you want to delete, click **More** in the **Operation** column, and then click **Delete**. The unsubscription page is displayed.
- 3. Confirm the information and click **Confirm**. A confirmation dialog box is displayed.
- 4. Confirm the information and click **Yes**.

Return to the shared bandwidth list and check whether the target shared bandwidth is unsubscribed from.

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

ID	Nam e	Description	Value Rang e	Uni t	Co nv ers ion Rul e	Monitored Object (Dimension)	Monitorin g Interval (Raw Data)
upstrea m_band width	Outb ound Band widt h	Network rate of outbound traffic (Previously called "Upstream Bandwidth")	≥ 0	bit/ s	10 00 (SI)	Bandwidth or EIP	1 minute
downstr eam_ba ndwidth	Inbo und Band widt h	Network rate of inbound traffic (Previously called "Downstrea m Bandwidth")	≥ 0	bit/ s	10 00 (SI)	Bandwidth or EIP	1 minute
upstrea m_band width_u sage	Outb ound Band widt h Usag e	Usage of outbound bandwidth in the unit of percent. Outbound bandwidth usage = Outbound bandwidth/ Purchased bandwidth	0-100	%	N/ A	Bandwidth or EIP	1 minute
up_stre am	Outb ound Traffi c	Network traffic going out of the cloud platform in a minute (Previously called	≥ 0	Byt e	10 00 (SI)	Bandwidth or EIP	1 minute

Table 13-1 EIP and bandwidth metrics

"Upstream Traffic")

ID	Nam e	Description	Value Rang e	Uni t	Co nv ers ion Rul e	Monitored Object (Dimension)	Monitorin g Interval (Raw Data)
down_st ream	Inbo und Traffi c	Network traffic going into the cloud platform in a minute (Previously called "Downstrea m Traffic")	≥ 0	Byt e	10 00 (SI)	Bandwidth or EIP	1 minute

Dimensions

Кеу	Value
publicip_id	EIP ID
bandwidth_id	Bandwidth ID

If a monitored object has multiple dimensions, all dimensions are mandatory when you use APIs to query the metrics.

• Query a monitoring metric:

dim.0=bandwidth_id,530cd6b0-86d7-4818-837f-935f6a27414d&dim.1=publici p_id,3773b058-5b4f-4366-9035-9bbd9964714a

• Query monitoring metrics in batches:

```
"dimensions": [
{
    "name": "bandwidth_id",
    "value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    {
        "name": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    /
    value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    /
    //
    value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    //
    value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    //
    //
    value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
    //
    //
    value": "530cd6b0-86d7-4818-837f-935f6a27414d"
    }
}
```

13.1.2 VPC Events That Can Be Monitored

Description

Event monitoring provides data collection, query, and alarm reporting for events. You can create alarm rules for both system events and custom events. When there are specified events, you will receive alarm notifications.

Namespace

SYS.VPC

VPC Events That Can Be Monitored

Event Source	Namespac e	Event Name	Event ID	Event Severity
Virtual	SYS.VPC	VPC deleted	deleteVpc	Major
Private Cloud		VPC modified	modifyVpc	Minor
		Subnet deleted	deleteSubnet	Minor
		Subnet modified	modifySubnet	Minor
		Bandwidth modified	modifyBandwi dth	Minor
		VPN deleted	deleteVpn	Major
		VPN modified	modifyVpn	Minor

Table 13-2 VPC events

13.1.3 Viewing Metrics

Scenarios

You can view VPC monitoring metrics in any of the following ways:

You can view the inbound bandwidth, outbound bandwidth, inbound bandwidth usage, outbound bandwidth usage, inbound traffic, and outbound traffic in a specified period.

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 \equiv to open the service list and choose **Management & Deployment > Cloud Eye**.

- 4. Click **Cloud Service Monitoring** on the left navigation pane, 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.

13.1.4 Creating an Alarm Rule

Scenarios

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

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

Creating an Alarm Rule for Metric Monitoring

- 1. Log in to the **Cloud Eye console**.
- 2. In the navigation pane, choose **Alarm Management** > **Alarm Rules**.
- 3. Click Create Alarm Rule in the upper right corner.
- 4. On the **Create Alarm Rule** page, configure parameters as needed. Configure the key parameters as follows:
 - a. Configure the alarm rule name and description.
 - **Name**: Name of the alarm rule. The name is automatically generated, but you can change it to a custom one.
 - **Description**: (Optional) Describe the alarm rule.
 - b. Configure alarm content parameters.
 - Alarm Type: Select Metric.
 - Cloud Product: Select Virtual Private Cloud-Bandwidths or Virtual Private Cloud-Elastic IPs as required.
 - Resource Level: Select the resource level of the monitored object. This parameter is available only if Alarm Type is set to Metric. The value can be Cloud product or Specific dimension. Cloud product is recommended.
 - Monitoring Scope: Select All resources, Resource groups, or Specific resources that the alarm rule will apply to.
 - Method: Select Associate template or Configure manually.
 - Alarm Policy: Specify the policy for triggering an alarm.
 - c. Configure alarm notification parameters.

To send alarm notifications by email, SMS, HTTP, or HTTPS, enable **Alarm Notifications**.

For details about the related parameters, see **Creating an Alarm Rule**.

5. Click Create.

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

Creating an Alarm Rule for Event Monitoring

- 1. Log in to the **Cloud Eye console**.
- 2. In the navigation pane on the left, choose **Event Monitoring**.
- 3. On the displayed page, click **Create Alarm Rule** in the upper right corner.

NOTE

If you want to create an alarm rule for an existing event, locate the target event in the event list and click **Create Alarm Rule** in the **Operation** column. On the displayed **Create Alarm Rule** page, required parameters have been configured for that event.

- 4. On the **Create Alarm Rule** page, configure parameters as needed.
 - a. Configure the alarm rule name and description.
 - Name: Enter a name for the alarm rule, which can be customized or automatically generated by the system.
 - Description: (Optional) Describe the alarm rule.
 - b. Configure alarm content parameters.
 - Alarm Type: Select Event.
 - Event Type: System event or Custom event. Select System event.
 - **Event Source**: The cloud service for which the event is generated. Select **Virtual Private Cloud**.
 - Monitoring Scope: The resources that the alarm rule applies to. All resources is selected by default.
 - Method: Select Associate template or Configure manually.
 - Alarm Policy: Specify the policy for triggering an alarm.
 - c. Configure alarm notification parameters.

To send alarm notifications by email, SMS, HTTP, or HTTPS, enable **Alarm Notifications**.

For more information, see **Creating an Alarm Rule and Notification for Event Monitoring**.

5. Click **Create**.

For more information about VPC monitoring 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 13-3 lists the 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	createPrivatelp
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
Adding 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 router	routers	createRouter
Updating a router	routers	updateRouter
Adding an interface to a router	routers	addRouterInterface

Table 13-	3 VPC	operations	that can	be rec	orded by CTS
-----------	-------	------------	----------	--------	--------------

Operation	Resource Type	Trace Name
Deleting an interface from a router	routers	removeRouterInterface
Creating a port	ports	createPort
Updating a port	ports	updatePort
Deleting a port	ports	deletePort
Creating a network	networks	createNetwork
Updating a network	networks	updateNetwork
Deleting a network	networks	deleteNetwork
Batch adding or deleting subnet tags	tag	batchUpdateTags
Batch adding 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	firewall-groups	createFirewallGroup
Updating a network ACL	firewall-groups	updateFirewallGroup
Deleting a network ACL	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

Operation	Resource Type	Trace Name
Creating a network ACL rule	firewall-rules	createFirewallRule
Updating a network ACL rule	firewall-rules	updateFirewallRule
Deleting a network ACL rule	firewall-rules	deleteFirewallRule
Creating an IP address group	address_group	createAddress_group
Updating an IP address group	address_group	updateAddress_group
Forcibly deleting an IP address group	address_group	force_deleteAddress_group
Deleting an IP address group	address_group	deleteAddress_group
Creating a VPC flow log	flowlogs	createFlowLog
Updating a VPC flow log	flowlogs	updateFlowLog
Deleting a VPC flow log	flowlogs	deleteFlowLog
Creating a public NAT gateway	natgateways	createNatGateway
Modifying a public NAT gateway	natgateways	updateNatGateway
Deleting a public NAT gateway	natgateways	deleteNatGateway
Creating a DNAT rule	dnatrules	createDnatRule
Modifying a DNAT rule	dnatrules	updateDnatRule
Deleting a DNAT rule	dnatrules	deleteDnatRule
Creating an SNAT rule	snatrules	createSnatRule
Modifying an SNAT rule	snatrules	updateSnatRule
Deleting an SNAT rule	snatrules	deleteSnatRule

13.2.2 Viewing VPC Traces

Scenarios

After you enable Cloud Trace Service (CTS) and the management tracker is created, CTS starts recording operations on cloud resources. After a data tracker is created, CTS starts recording operations on data in Object Storage Service (OBS) buckets. CTS stores operation records (traces) generated in the last seven days.

This section describes how to query or export operation records of the last seven days on the CTS console.

- Viewing Real-Time Traces in the Trace List of the New Edition
- Viewing Real-Time Traces in the Trace List of the Old Edition

Constraints

- You can only query operation records of the last seven days on the CTS console. To store operation records for longer than seven days, configure transfer to OBS or Log Tank Service (LTS) so that you can view them in OBS buckets or LTS log groups.
- After performing operations on the cloud, you can query management traces on the CTS console 1 minute later and query data traces 5 minutes later.
- Data traces are not displayed in the trace list of the new version. To view them, you need to go to the old version.
- These operation records are retained for seven days on the CTS console and are automatically deleted upon expiration. Manual deletion is not supported.

Viewing Real-Time Traces in the Trace List of the New Edition

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Management & Governance** > **Cloud Trace Service**. The CTS console is displayed.
- 3. Choose **Trace List** in the navigation pane on the left.
- 4. On the **Trace List** page, use advanced search to query traces. You can combine one or more filters.
 - **Trace Name**: Enter a trace name.
 - **Trace ID**: Enter a trace ID.
 - Resource Name: Enter a resource name. If the cloud resource involved in the trace does not have a resource name or the corresponding API operation does not involve the resource name parameter, leave this field empty.
 - **Resource ID**: Enter a resource ID. Leave this field empty if the resource has no resource ID or if resource creation failed.
 - **Trace Source**: Select a cloud service name from the drop-down list.
 - **Resource Type**: Select a resource type from the drop-down list.
 - **Operator**: Select one or more operators from the drop-down list.
 - Trace Status: Select normal, warning, or incident.

- **normal**: The operation succeeded.
- warning: The operation failed.
- **incident**: The operation caused a fault that is more serious than the operation failure, for example, causing other faults.
- Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range within the last seven days.
- 5. On the **Trace List** page, you can also export and refresh the trace list, and customize columns to display.
 - Enter any keyword in the search box and press Enter to filter desired traces.
 - Click Export to export all traces in the query result as an .xlsx file. The file can contain up to 5,000 records.
 - Click \bigcirc to view the latest information about traces.
 - Click 🞯 to customize the information to be displayed in the trace list. If

Auto wrapping is enabled (), excess text will move down to the next line; otherwise, the text will be truncated. By default, this function is disabled.

- 6. For details about key fields in the trace structure, see **Trace Structure** and **Example Traces**.
- 7. (Optional) On the **Trace List** page of the new edition, click **Old Edition** in the upper right corner to switch to the **Trace List** page of the old edition.

Viewing Real-Time Traces in the Trace List of the Old Edition

- 1. Log in to the management console.
- 2. Click in the upper left corner and choose **Management & Deployment** > **Cloud Trace Service**. The CTS console is displayed.
- 3. Choose **Trace List** in the navigation pane on the left.
- 4. Each time you log in to the CTS console, the new edition is displayed by default. Click **Old Edition** in the upper right corner to switch to the trace list of the old edition.
- 5. Set filters to search for your desired traces. The following filters are available.
 - **Trace Type, Trace Source, Resource Type**, and **Search By**: Select a filter from the drop-down list.
 - If you select Resource ID for Search By, specify a resource ID.
 - If you select **Trace name** for **Search By**, specify a trace name.
 - If you select **Resource name** for **Search By**, specify a resource name.
 - **Operator**: Select a user.
 - Trace Status: Select All trace statuses, Normal, Warning, or Incident.
 - Time range: Select **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range within the last seven days.

- 6. Click **Query**.
- 7. On the **Trace List** page, you can also export and refresh the trace list.
 - Click **Export** to export all traces in the query result as a CSV file. The file can contain up to 5,000 records.
 - Click \mathbb{C} to view the latest information about traces.
- 8. Click \checkmark on the left of a trace to expand its details.

Trace Name		Resource Type	Trace Source	Resource ID (?)	Resource Name ⑦	Trace Status ⑦	Operator (?)	Operation Time	Operation
createDocker	Config	dockerlogincmd	SWR		dockerlogincmd	📀 normal		Nov 16, 2023 10:54:04 GMT+08:00	View Trace
request									
trace_id									
code	200								
trace_name	createDockerConfig								
resource_type	dockerlogincmd								
trace_rating	normal								
api_version									
message	createDockerConfig,	Method: POST Url=/v2/	/manage/utils/secre	t, Reason:					
source_ip									
domain_id									
trace_type	ApiCall								

9. Click **View Trace** in the **Operation** column. The trace details are displayed.

View Trace

{		*
	"request": "",	
	"trace_id": "",	
	"code": "200",	
	"trace_name": "createDockerConfig",	
	"resource_type": "dockerlogincmd",	
	"trace_rating": "normal",	
	"api_version": "",	
	"message": "createDockerConfig, Method: POST Url=/v2/manage/utils/secret, Reason:",	
	"source_ip": "",	
	"domain_id": "",	
	"trace_type": "ApiCall",	
	"service_type": "SWR",	
	"event_type": "system",	
	"project_id": "",	
	"response": "",	
	"resource_id": "",	
	"tracker_name": "system",	
	"time": "Nov 16, 2023 10:54:04 GMT+08:00",	
	"resource_name": "dockerlogincmd",	
	"user": {	
	"domain": {	
	"name": " ",	
	"id": "	-

- 10. For details about key fields in the trace structure, see **Trace Structure** and **Example Traces** in the *CTS User Guide*.
- 11. (Optional) On the **Trace List** page of the old edition, click **New Edition** in the upper right corner to switch to the **Trace List** page of the new edition.

×

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.
- In the upper right corner of the page, choose Resources > My Quotas. The Quotas page is displayed.
- 4. View the used and total quota of each type of resources on the displayed page.

If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

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