

Virtual Private Cloud

FAQs

Issue 01
Date 2024-10-25



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Contents

1 Billing and Payments	1
1.1 Will I Be Billed for Using the VPC Service?	1
1.2 Why Is My VPC Still Being Billed After It Was Deleted?	2
1.3 How Do I View My VPC Bills?	3
1.4 How Is an EIP Charged?	4
1.5 How Do I Change My EIP Billing Mode Between Pay-per-Use and Yearly/Monthly?	5
1.6 How Do I Change the Billing Option of a Pay-per-Use EIP Between By Bandwidth and By Traffic?	7
2 VPCs and Subnets	9
2.1 What Is Virtual Private Cloud?	9
2.2 Which CIDR Blocks Are Available for the VPC Service?	12
2.3 How Many VPCs Can I Create?	13
2.4 Can Subnets Communicate with Each Other?	13
2.5 What Subnet CIDR Blocks Are Available?	16
2.6 Can I Change the CIDR Block of a Subnet?	16
2.7 How Many Subnets Can I Create?	16
2.8 How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?	16
2.9 How Can I Make a Domain Name in a Subnet Take Effect Immediately After Being Changed?	21
2.10 Why Can't I Delete My VPCs and Subnets?	21
2.11 Can I Change the VPC of an ECS?	27
2.12 Why Is the ECS IP Address Released After the System Time Is Changed?	27
2.13 How Do I Change the DNS Server Address of an ECS?	27
3 EIPs	30
3.1 How Do I Assign or Retrieve a Specific EIP?	30
3.2 What Are the Differences Between EIP, Private IP Address, and Virtual IP Address?	30
3.3 Can I Change the Dedicated Bandwidth Used by an EIP to a Shared Bandwidth?	33
3.4 How Many ECSs Can I Bind an EIP To?	33
3.5 How Do I Access an ECS with an EIP Bound from the Internet?	33
3.6 What Is the EIP Assignment Policy?	33
3.7 Can I Bind an EIP of an ECS to Another ECS?	34
3.8 Can I Buy a Specific EIP?	34
3.9 How Do I Query the Region of My EIPs?	34
3.10 How Can I Unbind an Existing EIP from an Instance and Bind Another EIP to the Instance?	35

3.11 Can I Bind an EIP to a Cloud Resource in Another Region?.....	36
3.12 Can I Change the Region of an EIP?.....	37
4 VPC Peering Connections.....	38
4.1 How Many VPC Peering Connections Can I Create in an Account?.....	38
4.2 Can a VPC Peering Connection Connect VPCs in Different Regions?.....	38
4.3 Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?....	39
5 Virtual IP Addresses.....	48
5.1 Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?.....	48
5.2 How Do I Bind a Virtual IP Address on Huawei Cloud to a Server in an On-premises Data Center?....	53
5.3 Why Is the Network Disconnected Between Servers Using a Virtual IP Address After an Active/ Standby Switchover?.....	53
6 Bandwidth.....	54
6.1 What Are Inbound Bandwidth and Outbound Bandwidth?.....	54
6.2 What Are the Differences Among Static BGP, Dynamic BGP, and Premium BGP?.....	55
6.3 How Do I Know If My EIP Bandwidth Limit Has Been Exceeded?.....	57
6.4 What Are the Differences Between Public Bandwidth and Private Bandwidth?.....	60
6.5 What Bandwidth Types Are Available?.....	61
6.6 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth?.....	61
6.7 How Many EIPs Can I Add to Each Shared Bandwidth?.....	61
6.8 Can I Increase a Yearly/Monthly Bandwidth and Decrease It Later?.....	62
6.9 What Is the Relationship Between Bandwidth and Upload/Download Rate?.....	62
7 Connectivity.....	63
7.1 Does a VPN Allow Communication Between Two VPCs?.....	63
7.2 Why Cannot I Access Public Websites Through Domain Names or Access Internal Domain Names on the Cloud When My ECS Has Multiple NICs?.....	63
7.3 What Are the Priorities of the Custom Route and EIP If Both Are Configured for an ECS to Enable the ECS to Access the Internet?.....	64
7.4 Why Are There Intermittent Interruptions When a Local Host Accesses a Website Built on an ECS?....	64
7.5 Why Do ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communication?	65
7.6 Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur When They Communicate?.....	66
7.7 Why Can't My ECS Use Cloud-init?.....	68
7.8 Why Can't My ECS Access the Internet Even After an EIP Is Bound?.....	73
7.9 Why Is My ECS Unable to Communicate at a Layer 2 or Layer 3 Network?.....	77
7.10 How Do I Handle a BMS Network Failure?.....	79
7.11 Why Does My ECS Fail to Obtain an IP Address?.....	81
7.12 How Do I Handle a VPN or Direct Connect Connection Network Failure?.....	83
7.13 Why Can My Server Be Accessed from the Internet But Cannot Access the Internet?.....	85
7.14 Why Can't I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?.....	87
7.15 Why Does My ECS Fail to Communicate with Other After It Has Firewall Installed?.....	88
8 Routing.....	91

8.1 How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?.....	91
8.2 Can a Route Table Span Multiple VPCs?.....	92
8.3 Are There Any Restrictions on Using a Route Table?.....	93
9 Security.....	94
9.1 Does a Modified Security Group Rule or a Network ACL Rule Take Effect Immediately for Existing Connections?.....	94
9.2 Why Is Outbound Access on TCP Port 25 Blocked?.....	95
9.3 How Do I Know the Instances Associated with a Security Group?.....	95
9.4 Why Can't I Delete a Security Group?.....	97
9.5 Can I Change the Security Group of an ECS?.....	97
9.6 How Do I Configure a Security Group for Multi-Channel Protocols?.....	97
9.7 Why Are Some Ports of ECSs Inaccessible?.....	98
9.8 Why Is Access from a Specific IP Address Still Allowed After a Network ACL Rule That Denies the Access from the IP Address Has Been Added?.....	98
9.9 Why Are My Security Group Rules Not Working?.....	99

1 Billing and Payments

1.1 Will I Be Billed for Using the VPC Service?

VPC provides a wide range of cloud resources. Some are free, while some are not. [Table 1-1](#) describes how these resources are billed.

Table 1-1 VPC resource billing

Resource	Billing Description
VPC	Free
Subnet	Free
Route table	Free
VPC peering connection	Free
Elastic network interface	Free
Supplementary network interface	Free
IP address group	Free
Security group	Free
Network ACL	Free
VPC flow log	Free
Traffic Mirroring	Free

Resource	Billing Description
EIP and bandwidth	<p>If you use EIPs and bandwidths, you need to pay for their prices.</p> <ul style="list-style-type: none">• EIP reservation price If your pay-per-use EIP has no instance bound, you will be billed for the EIP reservation price.• Fixed bandwidth<ul style="list-style-type: none">- EIP bandwidth prices: bandwidth prices of yearly/monthly EIPs and pay-per-use EIPs (by bandwidth); traffic price of pay-per-use EIPs (by traffic)- Shared bandwidth price- Shared data package price <p>For details, see EIP Billing.</p>
VPC endpoint	<p>If you use VPC endpoints, you need to pay for them.</p> <p>For details, see VPC Endpoint Billing.</p>

NOTE


Currently, free resources are not billed. You will be notified in advance if the billing starts.

1.2 Why Is My VPC Still Being Billed After It Was Deleted?

Symptom

As shown in [Figure 1-1](#), you deleted all VPCs under your account. However, there are still bills for VPC.

Figure 1-1 VPC list



Reasons

- In the Billing Center, on the **By service type** tab, click **Virtual Private Cloud (VPC)**, and check any of the resources listed in [Table 1-2](#) are billed. For details, see section [How Do I View My VPC Bills?](#)

Table 1-2 VPC billing

Service	Resource
VPC	EIP
	Fixed bandwidth
	Add-on package
	VPC endpoint

- Due to billing delay, the consumption amount will not be deducted immediately after a pay-per-use resource is deleted. Bills are generated and the consumption amount is deducted only after the settlement period ends.

1.3 How Do I View My VPC Bills?

VPC provides a wide range of cloud resources. Some are free, while some are not. [Table 1-3](#) describes how these resources are billed.

Table 1-3 VPC billing

Service	Resource
VPC	EIP
	Fixed bandwidth
	Add-on package
	VPC endpoint

You can perform the following operations to view VPC bills in the Billing Center.

Procedure

1. In the upper right corner of the console, choose **Billing & Costs > Bills**.
The **Dashboard** page is displayed.
2. In the navigation pane on the left, choose **Expenditure Items**.
The **Expenditure Items** page is displayed.
3. On the top of the expenditure item list, choose **Service Type > Virtual Private Cloud VPC**. The VPC billed items are displayed.
 - EIP reservation price
If your pay-per-use EIP has no instance bound, you will be billed for the EIP reservation price.
 - Fixed bandwidth
 - EIP bandwidth price
 - Yearly/Monthly EIPs: billed based the bandwidth size and required duration

- Pay-per-use EIPs (by bandwidth): billed based the bandwidth size and usage duration
- Pay-per-use EIPs (by traffic): billed based on how much traffic you use
 - Shared bandwidth price
 - Shared data package price
- VPC endpoint: billed based on how long the instance is retained in your account

Figure 1-2 The expenditure item list

The screenshot shows the 'Expenditure Items' page in the Billing Center. It features a table with columns: Billing Cycle, Enterprise, Account No., Service Type, Resource Type, Billing Mode, Expenditure Time, Order No./Transacti..., Bill Type, Transaction Time, Region, Specificatio..., and Usage. The table lists several items for March 2024, including Fixed Bandwidth and VPC Endpoint services, all billed on a Pay-per-Use basis.

Billing Cycle	Enterprise	Account No.	Service Type	Resource Type	Billing Mode	Expenditure Time	Order No./Transacti...	Bill Type	Transaction Time	Region	Specificatio...	Usage
Mar 20...	default		Virtual Private Cloud VPC	Fixed Bandwidth	Pay-per-Use	Mar 25, 2024 15:00:0... Mar 25, 2024 16:00:0...	b57e6b0b-8bee-45...	Expenditure...	Mar 25, 2024 16:28:0...			
Mar 20...	default		Virtual Private Cloud VPC	Fixed Bandwidth	Pay-per-Use	Mar 25, 2024 15:00:0... Mar 25, 2024 16:00:0...	e30850fa-e0f5-40c3...	Expenditure...	Mar 25, 2024 16:28:0...			
Mar 20...	default		Virtual Private Cloud VPC	VPC Endpoint	Pay-per-Use	Mar 25, 2024 15:00:0... Mar 25, 2024 16:00:0...	67d21187-0d6f-46b3...	Expenditure...	Mar 25, 2024 16:23:1...			
Mar 20...	default		Virtual Private Cloud VPC	Fixed Bandwidth	Pay-per-Use	Mar 25, 2024 15:00:0... Mar 25, 2024 16:00:0...	b3372766-8014-41b...	Expenditure...	Mar 25, 2024 16:19:4...			
Mar 20...	default		Virtual Private Cloud VPC	Fixed Bandwidth	Pay-per-Use	Mar 25, 2024 15:00:0... Mar 25, 2024 16:00:0...	b7912d64-95fa-403...	Expenditure...	Mar 25, 2024 16:18:0...			
Mar 20...	default		Virtual Private Cloud VPC	VPC Endpoint	Pay-per-Use	Mar 25, 2024 14:00:0... Mar 25, 2024 15:00:0...	cb634862-7c1c-488...	Expenditure...	Mar 25, 2024 15:23:3...			

1.4 How Is an EIP Charged?

The EIP and bandwidth bills are included in VPC bills. The billed items are as follows:

- EIP reservation price
 - If your pay-per-use EIP has no instance bound, you will be billed for the EIP reservation price.
- Fixed bandwidth
 - EIP bandwidth price
 - Yearly/Monthly EIPs: billed based the bandwidth size and required duration
 - Pay-per-use EIPs (by bandwidth): billed based the bandwidth size and usage duration
 - Pay-per-use EIPs (by traffic): billed based on how much traffic you use
 - Shared bandwidth price
 - Shared data package price

For details, see [EIP Billing](#).

Figure 1-3 The expenditure item list

The screenshot shows the 'Expenditure Items' page in the Billing Center. It features a table with columns: Billing..., Enterprise..., Account No..., Service Type, Resource Type, Billing Mode, Expenditure Time, Order No./Transact..., Bill Type, Transaction Time, Region, Specificatio..., and Usage. The table lists several items for 'Virtual Private Cloud VPC' services, including 'Fixed Bandwidth' and 'VPC Endpoint', with various billing modes like 'Pay-per-Use' and 'Pay-per-Use'.

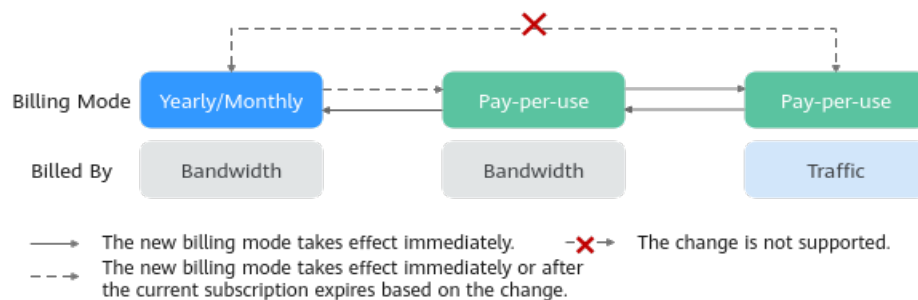
1.5 How Do I Change My EIP Billing Mode Between Pay-per-Use and Yearly/Monthly?

Table 1-4 Billing mode change description

Change	Description
From yearly/monthly to pay-per-use	<ul style="list-style-type: none"> An EIP billed on a yearly/monthly basis can be directly changed to be billed on a pay-per-use basis (billed by bandwidth). An EIP billed on a yearly/monthly basis cannot be directly changed to be billed on a pay-per-use basis (billed by traffic). To change this: <ol style="list-style-type: none"> Change the yearly/monthly EIP to be billed by bandwidth on a pay-per-use basis. Change the EIP billed by bandwidth on a pay-per-use basis to be billed by traffic on a pay-per-use basis. <p>The new billing mode takes effect only after the yearly/monthly subscription expires, if you want to change the EIP to be billed by bandwidth on a pay-per-use basis upon expiration.</p> <p>The new billing mode takes effect immediately, if you want to change the EIP to be billed by bandwidth on a pay-per-use basis immediately.</p>

Change	Description
From pay-per-use to yearly/monthly	<ul style="list-style-type: none"> An EIP that is billed by bandwidth on a pay-per-use basis can be directly changed to be billed on a yearly/monthly basis. An EIP that is billed by traffic on a pay-per-use basis cannot be directly changed to be billed on a yearly/monthly basis. To change this: <ol style="list-style-type: none"> Change the EIP billed by traffic on a pay-per-use basis to be billed by bandwidth on a pay-per-use basis. Change the EIP billed by bandwidth on a pay-per-use basis to be billed on a yearly/monthly basis. <p>After the change is successful, the new billing mode takes effect immediately.</p>

Figure 1-4 EIP billing change



From Yearly/Monthly to Pay-Per-Use upon Expiration (Billed by Bandwidth)

- Go to the [EIP list](#) page.
- In the EIP list, change billing mode of a single EIP or multiple EIPs from yearly/monthly to pay-per-use (billed by bandwidth):
 - Single EIP:
Locate the row that contains the EIP and choose **More > Change to Pay-per-Use upon Expiration** in the **Operation** column.
 - Multiple EIPs:
Select the EIPs in the EIP list and choose **More > Change to Pay-per-Use upon Expiration** in the upper left corner of the list.
- In the displayed dialog box, confirm the information and click **Yes**.
You are switched to a page of the Billing Center.
- Confirm the information and click **Change to Pay-per-Use**.

From Pay-per-Use (Billed by Bandwidth) to Yearly/Monthly

1. Go to the [EIP list](#) page.
2. In the EIP list, change the billing mode of a single EIP or multiple EIPs from pay-per-use (billed by bandwidth) to yearly/monthly.
 - Single EIP:
Locate the row that contains the EIP and choose **More > Change Billing Mode** in the **Operation** column.
 - Multiple EIPs:
Select EIPs and choose **More > Change Billing Mode** in the upper left corner of the EIP list.
3. In the displayed dialog box, confirm the information and click **Yes**.
4. On the **Change Subscriptions** page, set parameters such as **Renewal Duration**.
5. Click **Pay**.

1.6 How Do I Change the Billing Option of a Pay-per-Use EIP Between By Bandwidth and By Traffic?

Table 1-5 EIP billing mode change description

Change	Description
From billing by traffic (pay-per-use) to billing by bandwidth (pay-per-use)	A pay-per-use EIP billed by traffic can be directly changed to be billed by bandwidth. After the change is successful, the new billing mode takes effect immediately.
From billing by bandwidth (pay-per-use) to billing by traffic (pay-per-use)	A pay-per-use EIP billed by bandwidth can be directly changed to be billed by traffic. After the change is successful, the new billing mode takes effect immediately.

Pay-per-Use EIPs: From Billing By Traffic to By Bandwidth

1. Go to the [EIP list](#) page.
2. In the EIP list, locate the row that contains the EIP, click **More** in the **Operation** column, and click **Modify Bandwidth**.
3. On the **Modify Bandwidth** page, change the billing option as prompted.
You can also change the bandwidth name and size.
4. Click **Next**.
5. On the displayed page, confirm the configurations and click **Submit**.

 NOTE

- Changing the billing options does not change EIPs or interrupt their use.
- The preceding change scenarios apply only to **pay-per-use** EIPs.
- **Yearly/monthly** EIPs cannot be directly changed to **pay-per-use EIPs billed by traffic**. If the change is required, refer to [How Do I Change My EIP Billing Mode Between Pay-per-Use and Yearly/Monthly?](#)

2 VPCs and Subnets

2.1 What Is Virtual Private Cloud?

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 subnets, security groups, network ACLs, route tables, and more to manage cloud resources flexibly. You can also use EIPs to connect cloud resources in VPCs to the Internet, and use Direct Connect and VPN to connect on-premises data centers to VPCs to build a hybrid cloud network.

Product Architecture

The following describes the basics, security, connectivity, and O&M of VPCs.

Figure 2-1 VPC architecture

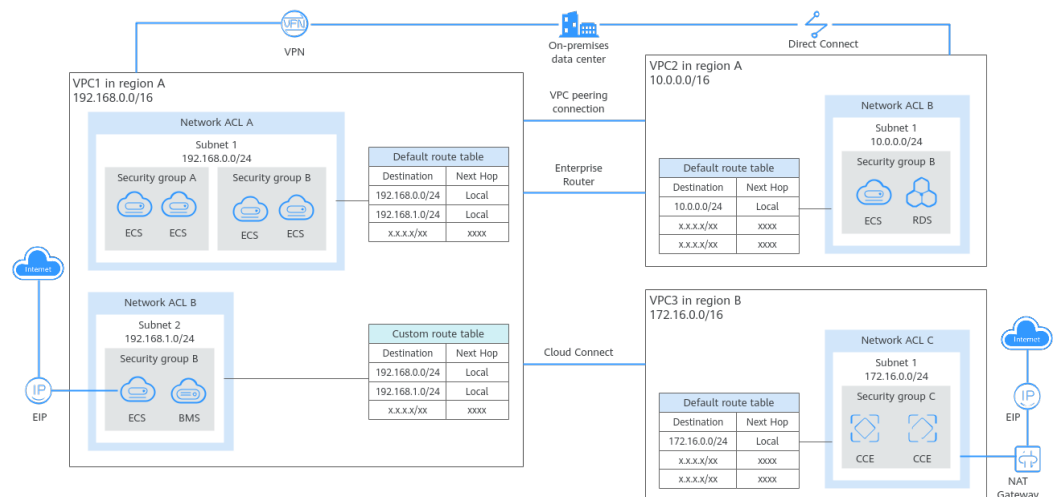


Table 2-1 Architecture description

Item	Brief	Details
VPC basics	<p>A VPC is a logically isolated virtual private network. You can define a CIDR block for each VPC and add one or more subnets. You can also configure route tables to control where the traffic from your subnet is directed.</p> <p>VPCs are logically isolated from each other, but subnets in a VPC can communicate with each other by default.</p>	<ul style="list-style-type: none"> • IPv4 CIDR block: When creating a VPC, you need to specify an IPv4 CIDR block for it. Supported IPv4 CIDR blocks are 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24. • Subnet: You can divide a VPC into one or more subnets as required to deploy your instances (such as cloud servers, containers, and databases). Private IP addresses are then assigned to your instances from the subnets where they are running. For more information, see Subnet. • Route table: Each VPC comes with a default route table that allows communications between subnets in a VPC. You can add routes to the default route table or create a route table to control traffic. For details, see Route Tables and Routes.
VPC security	<p>Security groups and network ACLs protect the cloud resources deployed in a VPC.</p>	<ul style="list-style-type: none"> • Security groups protect instances. You can add inbound and outbound rule to protect all the resources in a security group. For details about security groups, see Security Groups and Security Group Rules. • Network ACLs protect associated subnets. You can add inbound and outbound rule to protect all the resources in a subnet. For details, see Network ACL Overview. <p>Network ACLs protect subnets, while security groups protect instances in a subnet. If both security group and network ACL rules are configured, traffic matches network ACL rules first and then security group rules. For details, see What Is Access Control?</p>

Item	Brief	Details
VPC connectivity	<p>You can combine VPC and other networking services to build networks to meet different requirements.</p> <ul style="list-style-type: none">• Use VPC peering connections or an enterprise router to connect different VPCs in the same region.• Use Cloud Connect to connect VPCs in different regions.• Use an EIP or NAT gateway to allow the instances in a VPC to the Internet.• Use Direct Connect or VPN to connect an on-premises data center to VPCs.	<ul style="list-style-type: none">• Connecting VPCs in the same region<ul style="list-style-type: none">– VPC peering connections: connect VPCs in a region in the same account or different accounts. For details, see VPC Peering Connection Overview.– Enterprise routers: connect multiple VPCs in the same region, as a high-performance centralized router. For details, see What Is an Enterprise Router?<p>VPC peering connections are free of charge, while enterprise routers are not free. Compared with VPC peering connections, enterprise routers simplify the network structure and make it easy for scale-out and O&M.</p>• Connecting VPCs across regions Cloud Connect: connects VPCs in different regions to quickly build cross-region networks. For details, see What Is Cloud Connect?• Connecting a VPC to the Internet<ul style="list-style-type: none">– EIPs: enable your cloud resources to communicate with the Internet. For details, see What Is Elastic IP?– Public NAT gateways: enables instances (such as ECSs or BMSs) in a VPC to share an EIP to communicate with the Internet. A public NAT gateway supports up to 20 Gbit/s of bandwidth. For details, see What Is NAT Gateway?• Connecting an on-premises data center to a VPC<ul style="list-style-type: none">– Direct Connect: allows you to establish a stable, high-speed, low-latency, secure, and dedicated network connection

Item	Brief	Details
		<p>that connects your on-premises data center to the cloud. Direct Connect helps you build a flexible, scalable hybrid cloud computing environment. For details, see What Is Direct Connect?</p> <ul style="list-style-type: none"> - VPN: establishes a secure, encrypted communication tunnel between your on-premises data center and your VPC. For details, see What Is Virtual Private Network? <p>Compared with Direct Connect, VPN is cost-effective and can be quickly deployed.</p>
VPC O&M	VPC flow logs and traffic mirroring track traffic in a VPC for network O&M.	<ul style="list-style-type: none"> • VPC flow logs: records traffic to and from a VPC in real time. VPC flow logs help you monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification. For details, see VPC Flow Log Overview. • Traffic mirroring: mirrors traffic that meets a mirror filter from an elastic network interface to a destination, where you can use the mirrored traffic for inspection, audit analysis, and troubleshooting. For details, see Traffic Mirroring Overview.

2.2 Which CIDR Blocks Are Available for the VPC Service?

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

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

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

Table 2-2 VPC CIDR blocks (RFC 1918)

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

In addition to the preceding addresses, you can create a VPC with a publicly routable CIDR block that falls outside of the private IPv4 address ranges specified in RFC 1918. However, the reserved system and public CIDR blocks listed in [Table 2-3](#) must be excluded:

Table 2-3 Reserved system and public CIDR blocks

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

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

2.3 How Many VPCs Can I Create?

By default, you can create a maximum of five VPCs in a region.

Quotas vary by your account type and service level. You can go to the [console](#) to check your quotas.

If the current quota cannot meet service requirements, [submit a service ticket](#) to increase the quota.

2.4 Can Subnets Communicate with Each Other?

- Different VPCs cannot communicate with each other, so subnets in different VPCs are isolated from each other.

To connect subnets in different VPCs, you can:

- Use a VPC peering connection or an enterprise router to connect different VPCs in the same region.

For details about peering connections, see [VPC Peering Connection Overview](#).

For details about enterprise routers, see [What Is an Enterprise Router?](#)

- Use Cloud Connect to connect VPCs in different regions.

For details about cloud connections, see [What Is Cloud Connect?](#)

- Subnets in the same VPC can communicate with each other by default. If network ACLs and security groups are used to protect network security, communications between subnets may be denied by these rules.
 - Network ACL: If you associate subnets with different network ACLs and do not add inbound and outbound allow rules, communications between these subnets would fail.
 - Security group: If you associate instances (such as ECSs) in a subnet with different security groups and do not add inbound and outbound allow rules, communications between these instances would fail.

If both network ACLs and security groups are configured, traffic preferentially matches the network ACL rules. For details, see [Table 2-4](#).

Figure 2-2 Communications between subnets in a VPC




Table 2-4 Communication scenarios

Scenario	Access Control Configuration	Description
Between subnets	No network ACLs associated Instances associated with the same security group	<ul style="list-style-type: none"> Subnets 3 and 4 are not associated with a network ACL, so they can communicate with each other. ECS07, ECS08, ECS09, and ECS10 are associated with the same security group (security group A), so they can communicate with each other.
	Subnet associated with the same network ACL Instances associated with different security groups	<ul style="list-style-type: none"> Subnets 1 and 2 are associated with the same network ACL (network ACL A), so they can communicate with each other. ECS01 and ECS02 in subnet 1 are associated with security group A, and ECS05 and ECS06 in subnet 2 are associated with security group B. If security groups A and B have no allow rules, ECSs in the two security groups cannot communicate with each other. For example, ECS01 and ECS05 cannot communicate with each other.
	Subnet associated with different network ACLs	<p>Subnet 1 is associated with network ACL A, and subnet 5 is associated with network ACL B. If network ACLs A and B have no allow rules, subnet 1 and subnet 5 cannot communicate with each other.</p> <p>As a result, ECSs in subnets 1 and 5 are blocked from each other even they are in the same security group. For example, ECS01 and ECS 11 cannot communicate with each other.</p>
Within a subnet	Instances associated with different security groups	<p>ECS01 and ECS02 in subnet 1 are associated with security group A, and ECS03 and ECS04 are associated with security group B. If security groups A and B have no allow rules, ECSs in the two security groups cannot communicate with each other even they are in the same subnet (subnet 1). For example, ECS01 and ECS03 cannot communicate with each other.</p>

2.5 What Subnet CIDR Blocks Are Available?

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.

The subnet CIDR block must be within the VPC CIDR block. Subnet CIDR blocks in a VPC must be unique. A subnet mask can be between the netmask of its VPC CIDR block and /29 netmask. If a VPC CIDR block is 10.0.0.0/16, its subnet mask can be between 16 to 29.

For example, if the CIDR block of a VPC is 10.0.0.0/16, you can specify 10.0.0.0/24 for a subnet in this VPC, 10.0.1.0/24 for the second subnet, and 10.0.2.0/24 for the third subnet.

2.6 Can I Change the CIDR Block of a Subnet?

After a subnet is created, the CIDR block cannot be changed.

Properly plan subnet CIDR blocks in advance.

- The subnet CIDR block size cannot be too small. Ensure that the number of available IP addresses in the subnet meets service requirements. Remember that the first and last three addresses in a subnet CIDR block are reserved for system use. For example, in subnet 10.0.0.0/24, 10.0.0.1 is the gateway address, 10.0.0.253 is the system interface address, 10.0.0.254 is used by DHCP, and 10.0.0.255 is the broadcast address.
- The subnet CIDR block cannot be too large, either. If you use a CIDR block that is too large, you may not have enough CIDR blocks available later for new subnets, which can be a problem when you want to scale out services.
- Avoid CIDR block conflicts if you need to connect two VPCs or connect a VPC to an on-premises data center.

2.7 How Many Subnets Can I Create?

By default, you can create a maximum of 100 subnets in a region.

Quotas vary by your account type and service level. You can go to the [console](#) to check your quotas.

If the current quota cannot meet service requirements, [submit a service ticket](#) to increase the quota.

2.8 How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?

Scenario

After you change the DHCP lease time on the console, the change will be applied when the DHCP lease of an instance (such as ECS) is renewed. You can manually

renew the lease or wait for the system to do it. Renewing lease will not change the IP address of the instance.

- If you want the change to be applied immediately, manually renew the lease. Manual renewal will interrupt services.
- If you do not want the change to be applied immediately, wait for the system to renew the DHCP lease automatically. The system renews the lease for the first time when half of the current lease time is left. If the renewal fails, the system attempts to renew the lease again when one eighth of the current lease time is left. If the renewal still fails, the IP address of the instance will be released after the current lease expires. To prevent the IP address from being released, manually renew the lease.

For details about the DHCP lease renewal, see [Table 2-5](#).

Table 2-5 DHCP lease renewal

Applied Immediately	How to Renew	Description
Yes	Manually	<p>You can manually renew the DHCP lease of the instance by referring to Viewing and Renewing the DHCP Lease (Windows ECS) or Viewing and Updating the DHCP Lease (Linux ECS).</p> <p>You can also restart the instance to make the new DHCP release to be applied immediately.</p> <p>NOTICE If you renew the DHCP lease manually, the IP address of the instance will be released, and a new IP address will be assigned when the new lease is applied. This may cause service interruption.</p>

Applied Immediately	How to Renew	Description
No	Automatically	<p>You can wait for the new lease to be applied automatically.</p> <ul style="list-style-type: none">• First attempt: The system renews the lease for the first time when half of the current lease time is left. If the renewal succeeds, the new lease is applied.• Second attempt: If the first attempt fails, the system attempts to renew the lease again when one eighth of the current lease time is left. If the renewal succeeds, the new lease is applied. If the second attempt still fails, the IP address of the instance is released after the current DHCP lease expires. <p>Suppose the DHCP lease time of an ECS is 30 days, and the lease will expire on January 30, 2024. If you change the DHCP lease time to 10 days on January 2, 2024:</p> <ul style="list-style-type: none">• First attempt: The system will renew the lease on January 15, 2024 when half of the current lease time is left. If the attempt succeeds, the lease of the ECS will expire on January 25, 2024. When half of the new lease time is left, the lease will be renewed on January 20, 2024.• Second attempt: If the first attempt fails, the system attempts to renew the lease again on January 26, 2024 when one eighth of the current lease time is left. If the attempt succeeds, the lease will expire on February 5, 2024. If the second attempt still fails, the IP address of the ECS will be released on January 30, 2024.

Viewing and Renewing the DHCP Lease (Windows ECS)

1. After you change the DHCP lease time on the console, log in to the ECS whose lease you want to renew.
2. Enter **cmd** in the search box to open the CLI.
3. View the expiration time of the current DHCP lease:
ipconfig /all
4. Renew the DHCP lease:
ipconfig /release && ipconfig /renew
5. Check the new DHCP lease expiration time:
ipconfig /all

Viewing and Updating the DHCP Lease (Linux ECS)

1. After you change the DHCP lease time on the console, log in to the ECS whose lease you want to renew.
2. Check whether the client that provides the DHCP service is **dhclient**:

ps -ef | grep dhclient

- If information similar to the following is displayed, the **dhclient** process exists and the client is **dhclient**. The **lease** file following the **-lf** parameter contains lease information.

```
[root@ecs-A ~]# ps -ef | grep dhclient
root      580   526   0 18:49 ?        00:00:00 /sbin/dhclient -d -q -sf /usr/libexec/nm-dhcp-helper -
pf /var/run/dhclient-eth0.pid -lf /var/lib/NetworkManager/
dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0.lease -cf /var/lib/NetworkManager/
dhclient-eth0.conf eth0
root      1512  1470   0 18:50 pts/0    00:00:00 grep --color=auto dhclient
```

- If the **dhclient** process does not exist, this procedure may not be applicable. In this case, you need to search for the operation commands of the corresponding DHCP client.

3. View the latest DHCP lease information in the **lease** file obtained in [2](#):

cat lease File name

Example command:

cat /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0.lease

Information similar to the following is displayed. The **lease** file contains historical DHCP lease information, and the information following the last **lease** is about the latest DHCP lease.

```
[root@ecs-A ~]# cat /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0.lease
```

```
lease {
  interface "eth0";
  fixed-address 172.16.0.54;
  option subnet-mask 255.255.255.0;
  option dhcp-lease-time 108000000;
  option routers 172.16.0.1;
  option dhcp-message-type 5;
  option dhcp-server-identifier 172.16.0.254;
  option domain-name-servers 100.125.1.250,100.125.64.250;
  option interface-mtu 1500;
  option dhcp-renewal-time 54000000;
  option dhcp-rebinding-time 94500000;
  option rfc3442-classless-static-routes 0,172,16,0,1,32,169,254,169,254,172,16,0,1;
  option broadcast-address 172.16.0.255;
  option host-name "host-172-16-0-54";
  option domain-name "openstacklocal";
  renew 3 2025/06/18 21:46:42;
  rebind 3 2027/01/20 04:46:44;
  expire 5 2027/06/25 10:46:44;
}
```

```
lease {
  interface "eth0";
  fixed-address 172.16.0.54;
  option subnet-mask 255.255.255.0;
  option routers 172.16.0.1;
  option dhcp-lease-time 108000000;
  option dhcp-message-type 5;
  option domain-name-servers 100.125.1.250,100.125.64.250;
  option dhcp-server-identifier 172.16.0.254;
  option interface-mtu 1500;
  option dhcp-renewal-time 54000000;
  option broadcast-address 172.16.0.255;
```



```
option rfc3442-classless-static-routes 0,172,16,0,1,32,169,254,169,254,172,16,0,1;
option dhcp-rebinding-time 94500000;
option host-name "host-172-16-0-54";
option domain-name "openstacklocal";
renew 3 2025/08/20 23:57:15;
rebind 3 2027/01/20 04:50:00;
expire 5 2027/06/25 10:50:00;
}
```

4. Release the IP address of the ECS:

```
dhclient -r
```

5. Obtain the new DHCP lease:

```
killall dhclient && systemctl restart NetworkManager
```

6. View the latest DHCP lease information in the **lease** file obtained in [2](#):

```
cat lease File name
```

Example command:

```
cat /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0.lease
```

Information similar to the following is displayed. The **lease** file contains historical DHCP lease information, and the information following the last **lease** is about the latest DHCP lease.

```
[root@ecs-A ~]# cat /var/lib/NetworkManager/dhclient-5fb06bd0-0bb0-7ffb-45f1-d6edd65f3e03-eth0.lease
```

```
lease {
  interface "eth0";
  fixed-address 172.16.0.54;
  option subnet-mask 255.255.255.0;
  option dhcp-lease-time 108000000;
  option routers 172.16.0.1;
  option dhcp-message-type 5;
  option dhcp-server-identifier 172.16.0.254;
  option domain-name-servers 100.125.1.250,100.125.64.250;
  option interface-mtu 1500;
  option dhcp-renewal-time 54000000;
  option dhcp-rebinding-time 94500000;
  option rfc3442-classless-static-routes 0,172,16,0,1,32,169,254,169,254,172,16,0,1;
  option broadcast-address 172.16.0.255;
  option host-name "host-172-16-0-54";
  option domain-name "openstacklocal";
  renew 3 2025/08/20 23:57:15;
  rebind 3 2027/01/20 04:50:00;
  expire 5 2027/06/25 10:50:00;
}
lease {
  interface "eth0";
  fixed-address 172.16.0.54;
  option subnet-mask 255.255.255.0;
  option routers 172.16.0.1;
  option dhcp-lease-time 108000000;
  option dhcp-message-type 5;
  option domain-name-servers 100.125.1.250,100.125.64.250;
  option dhcp-server-identifier 172.16.0.254;
  option interface-mtu 1500;
  option dhcp-renewal-time 54000000;
  option broadcast-address 172.16.0.255;
  option rfc3442-classless-static-routes 0,172,16,0,1,32,169,254,169,254,172,16,0,1;
  option dhcp-rebinding-time 94500000;
  option host-name "host-172-16-0-54";
  option domain-name "openstacklocal";
  renew 4 2025/07/03 00:34:04;
  rebind 3 2027/01/20 04:52:43;
  expire 5 2027/06/25 10:52:43;
}Sub-eni
```

2.9 How Can I Make a Domain Name in a Subnet Take Effect Immediately After Being Changed?

When you create a subnet, you can also configure domain names. To access a domain name, you only need to enter the domain name prefix and ECSs in the subnet will automatically match the domain name suffix. After a subnet is created, you can change the configured domain names. [Table 2-6](#) describes how to make the change to take effect.

Table 2-6 Domain name effective policies

ECS	Effective Policy
New ECSs in the subnet	ECSs newly added to a subnet will use the new domain names automatically. No additional configuration is required.
Existing ECSs in a subnet	To use the new domain names, you can use either of the following method: <ul style="list-style-type: none">• Restart the ECS.• Restart the DHCP Client service: service dhcpd restart• Restart the network service: service network restart <p>NOTE The command for updating the DHCP configuration depends on the ECS OS. The commands here are only for your reference.</p>

2.10 Why Can't I Delete My VPCs and Subnets?

If VPCs and subnets are being used by other resources, you need to delete these resources first based on the prompts on the console before deleting the VPCs and subnets. This following provides detailed deletion prompts and corresponding deletion guide.

- [Deleting Subnets](#)
- [Deleting VPCs](#)

NOTICE

VPCs and subnets are free of charge.

If your VPCs and subnets cannot be deleted, [submit a service ticket](#).

Deleting Subnets

You can refer to [Table 2-7](#) to delete subnets.

Table 2-7 Deleting subnets

Prompts	Cause	Solution
You do not have permission to perform this operation.	Your account does not have permissions to delete subnets.	Contact the account administrator to grant permissions to your account and then delete the subnet. Permissions Management
Delete custom routes from the associated route table of the subnet and then delete the subnet.	The route table has custom routes with the following as the next hop type: <ul style="list-style-type: none"> • Server • Extension NIC • Virtual IP address • NAT gateway 	Delete the custom routes from the route table and then delete the subnet. <ol style="list-style-type: none"> 1. Viewing the Route Table Associated with a Subnet 2. Deleting a Route
Release any virtual IP addresses configured in the subnet and then delete the subnet.	The subnet has virtual IP addresses configured.	Release the virtual IP addresses from the subnet and then delete the subnet. Releasing a Virtual IP Address
Release any private IP addresses configured in the subnet and then delete the subnet.	The subnet has virtual IP addresses that are not used by any instance.	On the IP Addresses tab, release these private IP addresses that are not required and then delete the subnet. <ol style="list-style-type: none"> 1. Viewing IP Addresses in a Subnet 2. In the private IP address list, locate the IP address that is not being used and click Release in the Operation column. <p>NOTICE If you want to release an in-use private IP address, you need to delete the resource that uses the IP address first.</p>

Prompts	Cause	Solution
Delete the resource (ECS or load balancer) that is using the subnet and then delete the subnet.	The subnet is being used by an ECS or a load balancer.	Delete the ECS or load balancer and then delete the subnet. Viewing and Deleting Resources in a Subnet
Delete the load balancer that is using the subnet and then delete the subnet.	The subnet is being used by a load balancer.	Delete the load balancer and then delete the subnet. Viewing and Deleting Resources in a Subnet
Delete the NAT gateway that is using the subnet and then delete the subnet.	The subnet is being used by a NAT gateway.	Delete the NAT gateway and then delete the subnet. Viewing and Deleting Resources in a Subnet
Delete the resource that is using the subnet and then delete the subnet.	The subnet is being used by cloud resources.	On the IP Addresses tab, view the usage of the IP address, find the resource that is using the IP address, delete the resource, and delete the subnet. <ol style="list-style-type: none"> Viewing IP Addresses in a Subnet Locate resource based on the usage of the IP address by referring to Searching for Cloud Resources. Delete the resource and then delete the subnet.

Deleting VPCs

Before deleting a VPC, ensure that all subnets in the VPC have been deleted. You can refer to [Table 2-8](#) to delete VPCs.

Table 2-8 Deleting VPCs

Prompts	Cause	Solution
You do not have permission to perform this operation.	Your account does not have permissions to delete VPCs.	Contact the account administrator to grant permissions to your account and then delete the VPC. Permissions Management
Delete the VPC endpoint service or the route configured for the service from the VPC route table and then delete the VPC.	The VPC route table has custom routes.	Delete the custom routes and then delete the VPC. 1. In the VPC list, locate the row that contains the VPC and click the number in the Route Tables column. The route table list is displayed. 2. Deleting a Route
	The VPC is being used by a VPC endpoint service.	Search for the VPC endpoint service on the VPC endpoint service console and delete it. Deleting a VPC Endpoint Service
This VPC has VPC endpoint services configured. Delete the services first.	The VPC is being used by VPC endpoint services.	Search for the VPC endpoint service on the VPC endpoint service console and delete it. Deleting a VPC Endpoint Service
This VPC cannot be deleted because it has associated resources.	The VPC is being used by the following resources: <ul style="list-style-type: none"> Subnet VPC peering connection Custom route table 	Click the resource name hyperlink as prompted to delete the resource. <ul style="list-style-type: none"> Table 2-7 Deleting a VPC Peering Connection Deleting a Route Table
Delete the virtual gateway that is using the VPC and then delete the VPC.	The VPC is being used by a Direct Connect virtual gateway.	On the Direct Connect console, locate the virtual gateway and delete it. Deleting a Virtual Gateway
Delete the VPN gateway that is using the VPC and then delete the VPC.	The VPC is being used by a VPN gateway.	On the VPN console, locate the VPN gateway and delete it. Deleting a VPN Gateway

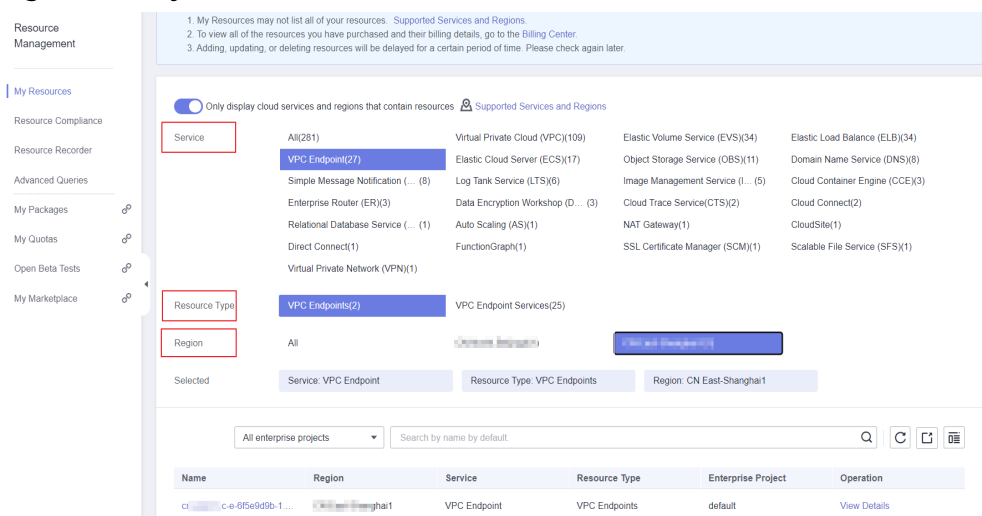
Prompts	Cause	Solution
Remove the VPC from the cloud connection and then delete the VPC.	The VPC is being used by a Cloud Connect connection.	On the Cloud Connect console, locate the connection and remove the VPC from it. Removing a Network Instance
Delete all custom security groups in this region and then delete this last VPC.	In the current region, this is the last VPC and there are custom security groups. NOTICE You only need to delete the custom security groups. The default security group does not affect the deletion of VPCs.	Delete all custom security groups and then delete the VPC. Deleting a Security Group
Release all EIPs in this region and then delete this last VPC.	In the current region, this is the last VPC and there are EIPs.	Release all EIPs in this region and then delete this last VPC. Releasing an EIP

Searching for Cloud Resources

1. Log in to the management console.
2. In the upper right corner of the console, choose **More > Resources > My Resources**.

The **My Resources** page is displayed.

Figure 2-3 My Resources



3. On the **My Resources** page, set search criteria to quickly find the resources in the subnet.
 - **Service:** Select a service that has resources in subnets.
Table 2-9 lists some common resources. If you have other resources, check them.
 - **Resource Type:** Check the types of resources.
 - **Region:** Select the region where the VPC and subnet are located to filter resources in the same region. VPCs and subnets can be used only by resources from their same region.

Table 2-9 Common resources in subnets

Product Category	Service
Compute	Elastic Cloud Server (ECS)
	Bare Metal Server (BMS)
	Cloud Container Engine (CCE)
	Cloud Container Instance (CCI)
Containers	Application Service Mesh (ASM)
Networking	Elastic Load Balance (ELB)
	NAT Gateway
	VPC Endpoint (VPCEP)
Databases	GaussDB
	Relational Database Service (RDS)
	Document Database Service (DDS)
	GaussDB NoSQL
	Distributed Database Middleware (DDM)
Applications	Distributed Cache Service (DCS) <ul style="list-style-type: none">● Redis instance● Memcached instance
	Distributed Message Service (DMS) <ul style="list-style-type: none">● Kafka instance● RabbitMQ instance
EI	MapReduce Service (MRS)
	Data Warehouse Service (DWS)
	Cloud Search Service (CSS)

If you cannot delete a subnet even after deleting all the resources it contains, [submit a service ticket](#).

2.11 Can I Change the VPC of an ECS?

Yes.

You can click **Change VPC** in the **Operation** column on the **Elastic Cloud Server** page.

For details, see [Changing a VPC](#).

2.12 Why Is the ECS IP Address Released After the System Time Is Changed?

Cause

If the new ECS system time is later than the DHCP lease expiration time, the system will attempt to renew the DHCP lease. If the renewal fails, the IP address will be released.

Solution

If you need to set the ECS system time to a later time than the DHCP lease expiration time, assign a static IP address to the ECS.

2.13 How Do I Change the DNS Server Address of an ECS?

Scenarios

This section describes how to change the DNS server address of an ECS and make the new DNS server address take effect immediately on the ECS.

The required operations are as follows:

1. [Viewing the DNS Server Addresses for an ECS](#)
2. [Changing the DNS Server Addresses for a VPC Subnet](#)
3. [Changing DNS Servers for the ECS](#)

Background

ECSs use private DNS servers for domain name resolution in VPCs. The private DNS servers do not affect the domain name resolution for the ECSs to access the Internet. Additionally, you can use the private DNS servers to directly access the private IP addresses of cloud services, such as OBS and SMN. Compared with the access through the Internet, this access features high performance and low latency.

Before private domain names are available, VPC subnets use the public DNS server (114.114.114.114). To allow ECSs in these VPCs to access private domain names,


you can change the public DNS server to the private DNS servers configured for the VPC subnets. For instructions about how to obtain a private DNS server address, see [What Are the Private DNS Servers Provided by the Huawei Cloud DNS Service?](#)

Viewing the DNS Server Addresses for an ECS

1. Log in to the management console.
2. Choose **Computing** > **Elastic Cloud Server**.
The **Elastic Cloud Server** page is displayed.
3. In the ECS list, click the name of the target ECS.
4. On the ECS details page, click the VPC name.
The **Virtual Private Cloud** page is displayed.
5. Locate the VPC and click the number in the **Subnets** column.
The **Subnets** page is displayed.
6. Click the name of the subnet.
In the **Gateway and DNS Information** area, view the DNS server addresses used by the ECS.

Changing the DNS Server Addresses for a VPC Subnet

If the ECS uses default public DNS server addresses, change them to Huawei Cloud private DNS server addresses.

1. In the **Gateway and DNS Information** area, click  next to **DNS Server Address**.
2. Change the DNS server addresses to private DNS server addresses.

Changing DNS Servers for the ECS

After you change the DNS server addresses of a VPC subnet, the DNS server addresses of ECSs in the subnet will not take effect immediately.

To make the DNS server addresses take effect immediately, do as follows:

- Restart the OS. The ECS will then obtain the new DNS server addresses from the DHCP server.

NOTICE

Restarting the OS will interrupt services on the ECS. Perform this operation during off-peak hours.

Alternatively, wait for the DHCP lease to expire. After the lease time expires, the DHCP server allocates another IP address and updates the DNS server addresses to the ECS.

- Obtain the new DNS server addresses.
 - a. Log in to the ECS.
For details, see [How Do I Log In to My ECS?](#)

- b. Run the following command to view the DNS server address of the ECS:

```
cat /etc/resolv.conf
```

If information similar to the following is displayed, 114.114.114.114 is the DNS server address of the ECS.

```
[root@ecs-01 ~]# cat /etc/resolv.conf
# Generated by NetworkManager
search openstacklocal
nameserver 114.114.114.114
options timeout:1 single-request-reopen
```

- c. Run the following command to check whether the **dhclient** process exists:

```
ps -ef | grep dhclient | grep -v grep
```

If information similar to the following is displayed, no process exists (CentOS 8.1 is used as an example).

In this case, run the **dhclient** command to start the process and check whether the **dhclient** process exists.

```
[root@ecs-01 ~]# ps -ef | grep dhclient | grep -v grep
[root@ecs-01 ~]# dhclient
[root@ecs-01 ~]# ps -ef | grep dhclient | grep -v grep
root      5712      1  0 09:52 ?        00:00:00 dhclient
```

If information similar to the following is displayed, process exists (CentOS 7.2 is used as an example).

```
[root@ecs-01 ~]# ps -ef | grep dhclient | grep -v grep
root      651  477  0 18:36 ?        00:00:00 /sbin/dhclient -d -q -sf /usr/libexec/om-dhcp-helper -pf /var/run/dhclient-eth0.
pid -lf /var/lib/NetworkManager/dhclient-5f86bd0-0bb9-7ffb-45f1-d6edd65f3e03-eth0.lease -cf /var/lib/NetworkManager/dhclient-eth0.conf eth0
```

- d. Run the following command to release the current DNS server address:

```
dhclient -r
```

- e. Run the following command to restart the **dhclient** process and obtain new DNS server addresses:

```
dhclient
```

- f. Run the following command to view the new DNS server addresses of the ECS:

```
cat /etc/resolv.conf
```

If information similar to the following is displayed, 100.125.1.250 and 100.125.64.250 are the new DNS server addresses of the ECS.

```
[root@ecs-01 ~]# dhclient -r
[root@ecs-01 ~]# dhclient
[root@ecs-01 ~]# cat /etc/resolv.conf
options timeout:1 single-request-reopen
; generated by /usr/sbin/dhclient-script
search openstacklocal
nameserver 100.125.1.250
nameserver 100.125.64.250
```

3 EIPs

3.1 How Do I Assign or Retrieve a Specific EIP?

If you want to retrieve an EIP that you have released or assign a specific EIP, you can use APIs by setting the value of `ip_address` to the one that you want to assign. For details, see [Elastic IP API Reference](#).

NOTE

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

3.2 What Are the Differences Between EIP, Private IP Address, and Virtual IP Address?

Different types of IP addresses have different functions.

Figure 3-1 IP address architecture

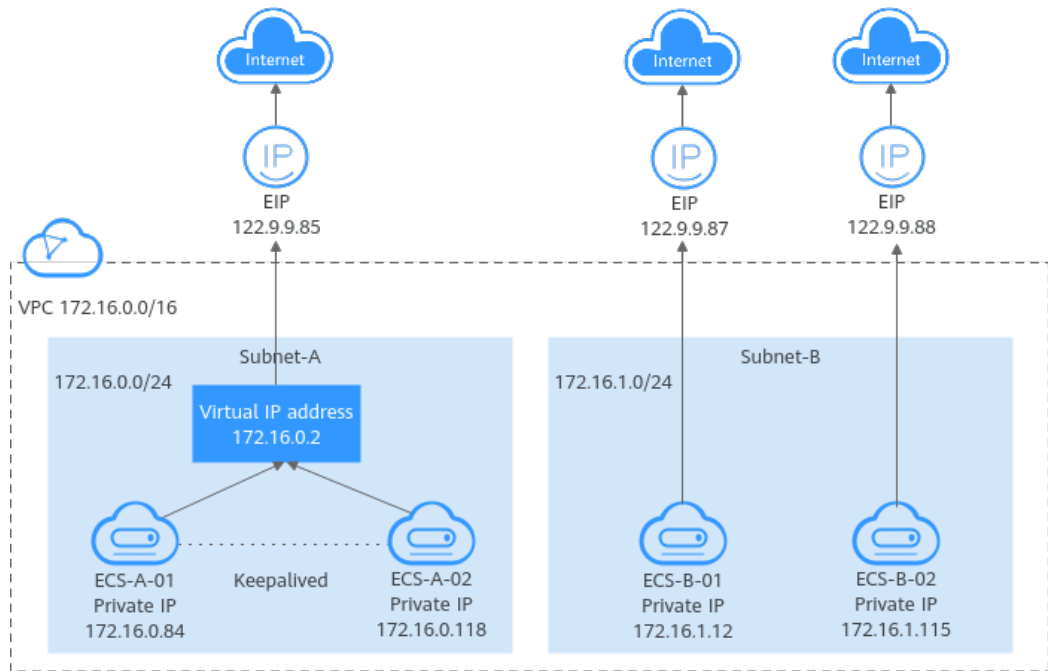


Table 3-1 Functions of different IP address types

IP Address Type	Description	Example Value
Private IP address	Private IP addresses come with your ECSs and belong to the VPC subnets of the ECSs. They are used for private communication on the cloud.	<ul style="list-style-type: none"> Private IP address of ECS-A-01: 172.16.0.84 Private IP address of ECS-B-01: 172.16.1.12

IP Address Type	Description	Example Value
Virtual IP address	<p>A virtual IP address is a private IP address that can be independently assigned from and released to a VPC subnet. You can:</p> <ul style="list-style-type: none"> 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. <p>For more information about virtual IP addresses, see Virtual IP Address Overview. For details about how to set up a high availability cluster, see Using a Virtual IP Address and Keepalived to Set Up a High-Availability Web Cluster.</p>	<p>Bind virtual IP address (172.16.0.2) both ECS-A-01 and ECS-A-02. The active/standby switchover of ECS-A-01 and ECS-A-02 can be implemented by using Keepalived.</p>
EIP	<p>EIPs allow cloud resources to access the Internet. They can be flexibly bound to or unbound from instances.</p> <ul style="list-style-type: none"> You can bind an EIP to a virtual IP address to enable the ECSs with the virtual IP address bound to access the Internet. You can also bind an EIP to the ECSs to enable them to access the Internet. <p>For more information, see EIP Overview.</p>	<ul style="list-style-type: none"> Bind EIP (122.9.9.85) to virtual IP address (172.16.0.2) to enable ECS-A-01 and ECS-A-02 to access the Internet. Bind EIP (122.9.9.87) to ECS-B-01 to enable ECS-B-01 to access the Internet.

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

Yes.

You can change the dedicated bandwidth used by a pay-per-use EIP to a shared bandwidth.

You cannot change the dedicated bandwidth used by a yearly/monthly EIP to a shared bandwidth.

3.4 How Many ECSs Can I Bind an EIP To?

An EIP can be bound to only one ECS.

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

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

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

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

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

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

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

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

3.6 What Is the EIP Assignment Policy?

By default, EIPs are assigned randomly.

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

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

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

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

Yes.

You can unbind the EIP from the original ECS. For details, see [Unbinding an EIP from an Instance](#).

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

If you want to change an EIP for your ECS, refer to [Changing an EIP](#).

3.8 Can I Buy a Specific EIP?

By default, EIPs are assigned randomly.

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

You can assign a specific EIP only by calling an API. For details, see [Assigning an EIP](#).

3.9 How Do I Query the Region of My EIPs?

You can visit <https://en.ipip.net/ip.html> to query the region of your EIPs.

- The region of an EIP identified using a third-party website may be different from the region that the EIP belongs to.
- If the region identified using another third-party website is different from the one identified using <https://en.ipip.net/ip.html>, use the region identified using <https://en.ipip.net/ip.html>.
- If the region identified using <https://en.ipip.net/ip.html> is different from the one you selected when purchasing the EIP, use the region you had selected during EIP purchase.

NOTE

The geographical location of an EIP purchased in CN North-Ulanqab1 is Beijing.

- If your service is adversely affected because the region of your EIP cannot be determined, [submit a service ticket](#).

To know more about the region of EIPs, [submit a service ticket](#).

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

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

1. Unbind an EIP.
 - a. Go to the [EIP list](#) page.
 - b. On the displayed page, locate the row that contains the target EIP, and click **Unbind**.
 - c. Click **Yes**.
2. Assign an EIP.

 **NOTE**

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

- a. Go to the [EIP list](#) page.
 - b. On the displayed page, click **Buy EIP**.
 - c. Set the parameters as prompted.
 - d. Click **Next**.
3. Bind the new EIP to the ECS.
 - a. Go to the [EIP list](#) page.
 - b. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
 - c. Select the desired ECS.
 - d. Click **OK**.
 4. Release the EIP that is unbound.

 **NOTE**

If an unbound EIP is no longer required, you can release it. If you do not release an unbound EIP, it will continue to be billed.

- a. Go to the [EIP list](#) page.
- b. In the EIP list, locate the row that contains the EIP, and choose **More > Release** in the **Operation** column.
- c. Click **Yes**.

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

1. Unbind an EIP.
 - a. Log in to the management console.
 - b. Click **Service List**. Under **Networking**, click **Elastic Load Balance**.
 - c. In the load balancer list, locate the target load balancer and choose **More > Unbind EIP** in the **Operation** column.

- d. Click **Yes**.
2. Assign an EIP by referring to [2](#).

 **NOTE**

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

3. Bind the new EIP to the load balancer.
 - a. Log in to the management console.
 - b. Click **Service List**. Under **Networking**, click **Elastic Load Balance**.
 - c. In the load balancer list, locate the target load balancer and choose **More > Bind EIP** in the **Operation** column.
 - d. In the **Bind EIP** dialog box, select the EIP to be bound and click **OK**.
4. Release the EIP that was replaced. For details, see [4](#).

 **NOTE**

If an unbound EIP is no longer required, you can release it. If you do not release an unbound EIP, it will continue to be billed.

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

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

 **NOTE**

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

2. Modify an SNAT rule.

For details, see [Modifying an SNAT Rule](#). In the EIP list, select the new EIP and deselect the existing EIP.
3. Modify a DNAT rule.

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

 **NOTE**

If an unbound EIP is no longer required, you can release it. If you do not release an unbound EIP, it will continue to be billed.

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

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

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

For example, EIPs in CN-Hong Kong cannot be bound to the cloud resources in AP-Singapore.

3.12 Can I Change the Region of an EIP?

The region of an EIP cannot be changed.

The following is an example:

If you assign an EIP in a region but need an EIP in another region, you cannot directly change the region of the assigned EIP. Instead, you have to assign an EIP in the desired region.

4 VPC Peering Connections

4.1 How Many VPC Peering Connections Can I Create in an Account?

To check how many VPC peering connections you can create in a region, you can [go to the console](#).

- Number of VPC peering connections that you can create in each region between VPCs in the same account: subject to the actual quota
- Number of VPC peering connections that you can create in each region between VPCs in different accounts: Accepted VPC peering connections use the quotas of both accounts. To-be-accepted VPC peering connections only use the quotas of accounts that request the connections.

An account can create VPC peering connections with different accounts if the account has enough quota.

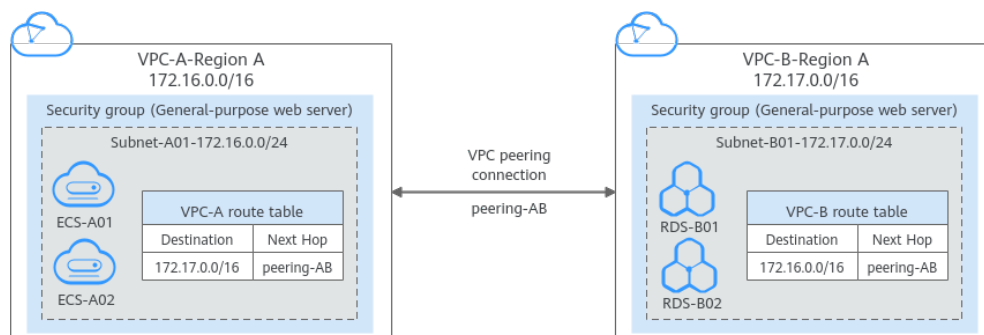
4.2 Can a VPC Peering Connection Connect VPCs in Different Regions?

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

A VPC peering connection can only connect VPCs in the same region. If your VPCs are in different regions, use [Cloud Connect](#).

Figure 4-1 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 4-1 Two VPCs connected by a VPC peering connection

4.3 Why Did Communication Fail Between VPCs That Were Connected by a VPC Peering Connection?

Symptom

After a VPC peering connection is created, the local and peer VPCs cannot communicate with each other.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Figure 4-2 Troubleshooting process

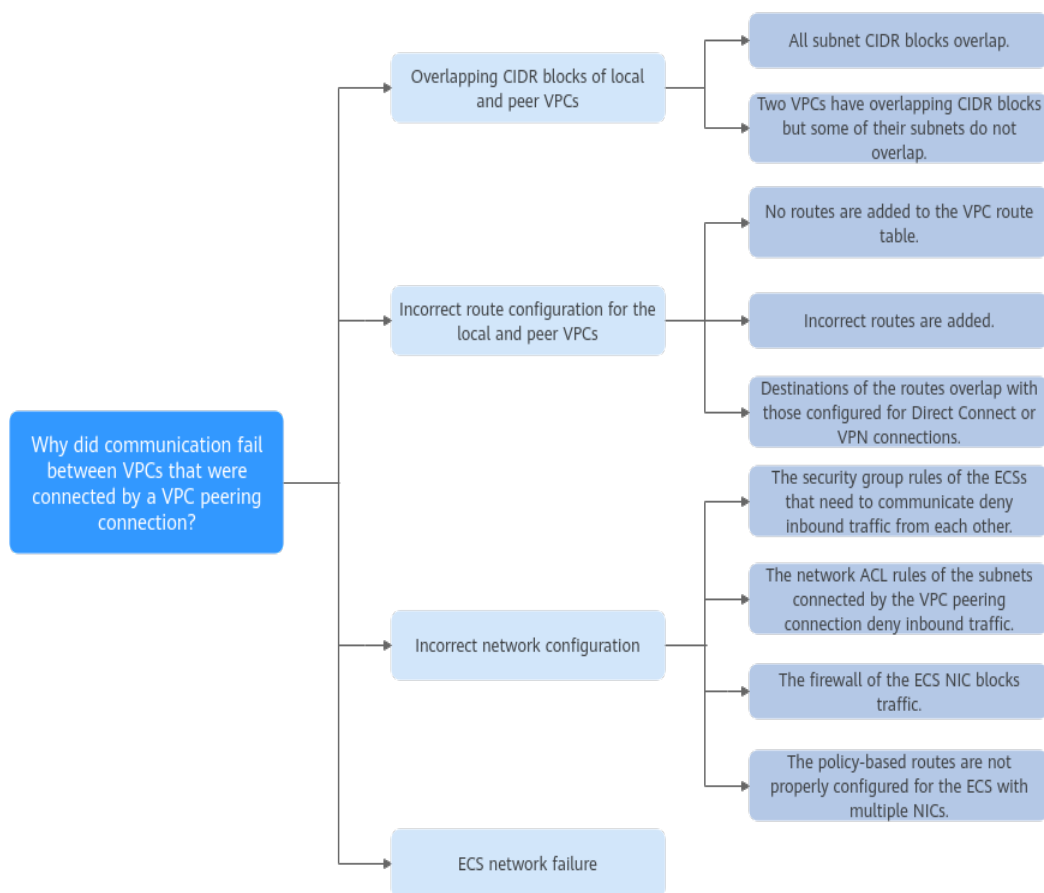


Table 4-1 Possible causes and solutions

No.	Possible Cause	Solution
1	Overlapping CIDR blocks of local and peer VPCs <ul style="list-style-type: none"> All their subnet CIDR blocks overlap. Some of their subnet CIDR blocks overlap. 	Refer to Overlapping CIDR Blocks of Local and Peer VPCs .
2	Incorrect route configuration for the local and peer VPCs <ul style="list-style-type: none"> No routes are added. Incorrect routes are added. Destinations of the routes overlap with that configured for Direct Connect or VPN connections. 	Refer to Incorrect Route Configuration for Local and Peer VPCs .

No.	Possible Cause	Solution
3	Incorrect network configuration <ul style="list-style-type: none"> The security group rules of the ECSs that need to communicate deny inbound traffic from each other. The firewall of the ECS NIC blocks traffic. The network ACL rules of the subnets connected by the VPC peering connection deny inbound traffic. Check the policy-based routing configuration of an ECS with multiple NICs. 	Refer to Incorrect Network Configuration .
4	ECS network failure	Refer to ECS Network Failure .

NOTICE

If the problem persists, [submit a service ticket](#).

Overlapping CIDR Blocks of Local and Peer VPCs

If the CIDR blocks of VPCs connected by a VPC peering connection overlap, the connection may not take effect due to route conflicts.

Table 4-2 Overlapping CIDR blocks of local and peer VPCs

Scenario	Description	Solution
VPCs with overlapping CIDR blocks also include subnets that overlap.	As shown in Figure 4-3 , the CIDR blocks of VPC-A and VPC-B overlap, and all their subnets overlap. <ul style="list-style-type: none"> Overlapping CIDR blocks of VPC-A and VPC-B: 10.0.0.0/16 Overlapping CIDR blocks of Subnet-A01 in VPC-A and Subnet-B01 in VPC-B: 10.0.0.0/24 Overlapping CIDR blocks of Subnet-A02 in VPC-A and Subnet-B02 in VPC-B: 10.0.1.0/24 	VPC-A and VPC-B cannot be connected using a VPC peering connection. Replan the network.

Scenario	Description	Solution
Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.	<p>As shown in Figure 4-4, the CIDR blocks of VPC-A and VPC-B overlap, and some of their subnets overlap.</p> <ul style="list-style-type: none"> Overlapping CIDR blocks of VPC-A and VPC-B: 10.0.0.0/16 Overlapping CIDR blocks of Subnet-A01 in VPC-A and Subnet-B01 in VPC-B: 10.0.0.0/24 CIDR blocks of Subnet-A02 in VPC-A and Subnet-B02 in VPC-B do not overlap. 	<ul style="list-style-type: none"> A VPC peering connection cannot connect the entire VPCs, VPC-A and VPC-B. A connection can connect their subnets (Subnet-A02 and Subnet-B02) that do not overlap. For details, see Figure 4-5.

Figure 4-3 Networking diagram (IPv4)

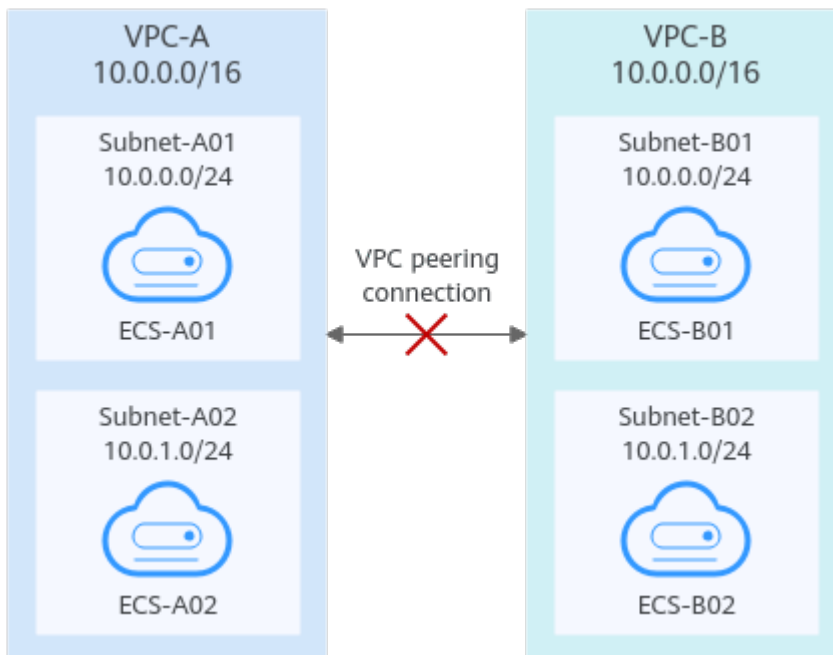
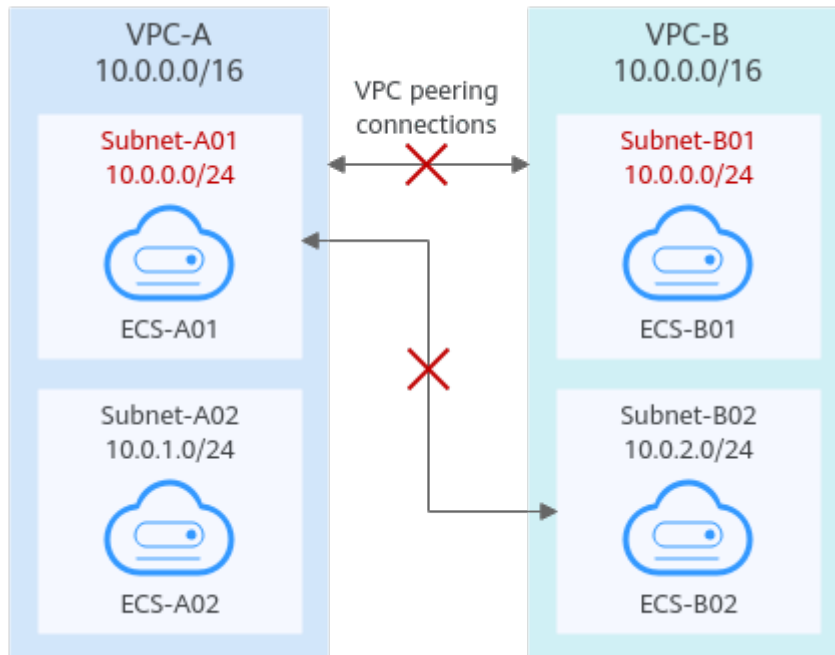


Figure 4-4 Networking diagram (IPv4)



If CIDR blocks of VPCs overlap and some of their subnets overlap, you can create a VPC peering connection between their subnets with non-overlapping CIDR blocks. [Figure 4-5](#) shows the networking diagram of connecting Subnet-A02 and Subnet-B02. [Table 4-3](#) describes the routes required.

Figure 4-5 Networking diagram (IPv4)

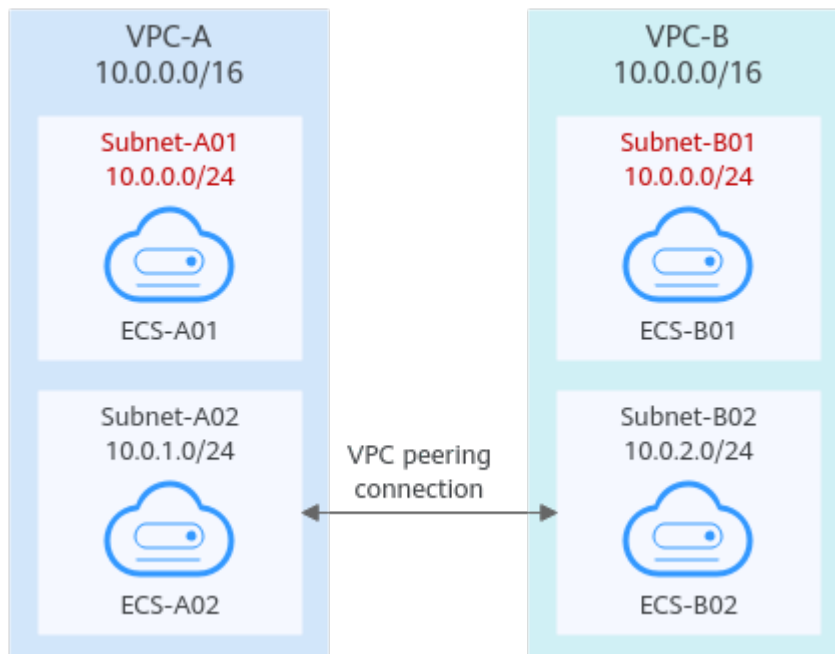


Table 4-3 Routes required for the VPC peering connection between Subnet-A02 and Subnet-B02

Route Table	Destination	Next Hop	Description
VPC-A route table	10.0.2.0/24	Peering-AB	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
VPC-B route table	10.0.1.0/24	Peering-AB	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AB as the next hop.

NOTICE

- If a VPC peering connection between VPCs with overlapping CIDR blocks does not take effect, see [Unsupported VPC Peering Configurations](#).
- If two VPCs want to use their IPv6 CIDR blocks for communication through a VPC peering connection but their IPv4 CIDR blocks or subnets overlap, the connection is not usable.

Incorrect Route Configuration for Local and Peer VPCs

Check the routes in the route tables of the local and peer VPCs by referring to [Viewing Routes Configured for a VPC Peering Connection](#). [Table 4-4](#) lists the items that you need to check.

Table 4-4 Route check items

Item	Solution
Check whether routes are added to the route tables of the local and peer VPCs.	If routes are not added, add routes by referring to: <ul style="list-style-type: none">• Creating a VPC Peering Connection with Another VPC in Your Account• Creating a VPC Peering Connection with a VPC in Another Account

Item	Solution
<p>Check the destinations of routes added to the route tables of the local and peer VPCs.</p> <ul style="list-style-type: none">• In the route table of the local VPC, check whether the route destination is the CIDR block, subnet CIDR block, or related private IP address of the peer VPC.• In the route table of the peer VPC, check whether the route destination is the CIDR block, subnet CIDR block, or related private IP address of the local VPC.	<p>If the route destination is incorrect, change it by referring to Modifying Routes Configured for a VPC Peering Connection.</p>
<p>Destinations of the routes overlap with that configured for Direct Connect or VPN connections.</p>	<p>Check whether any of the VPCs connected by the VPC peering connection also has a VPN or Direct Connect connection connected. If they do, check the destinations of their routes.</p> <p>If the destinations of the routes overlap, the VPC peering connection does not take effect. In this case, replan the network connection.</p>

Incorrect Network Configuration

1. Check whether the security group rules of the ECSs that need to communicate with each other allow inbound traffic from each other. For details, see [Viewing the Details of a Security Group](#).
 - Check the security group rules by referring to Step 3: Configure Security Group Rules for Instances in Local and Peer VPCs in [Using a VPC Peering Connection to Connect Two VPCs in the Same Account](#).
 - Check the security group rules by referring to Step 4: Configure Security Group Rules for Instances in Local and Peer VPCs in [Using a VPC Peering Connection Connect Two VPCs in Different Accounts](#).
2. Check whether the firewall of the ECS NIC blocks traffic.
If the firewall blocks traffic, configure the firewall to allow inbound traffic.
3. Check whether network ACL rules of the subnets connected by the VPC peering connection deny inbound traffic.
If the network ACL rules deny inbound traffic, configure rules to allow the traffic.
4. If an ECS has more than one NIC, check whether correct policy-based routes have been configured for the ECS and packets with different source IP addresses match their own routes from each NIC.
If an ECS has two NICs (eth0 and eth1):

- IP address of eth0: 192.168.1.10; subnet gateway: 192.168.1.1
- IP address of eth1: 192.168.2.10; subnet gateway: 192.168.2.1

Command format:

- **ping -l** *IP address of eth0 Subnet gateway address of eth0*
- **ping -l** *IP address of eth1 Subnet gateway address of eth1*

Run the following commands:

- **ping -I 192.168.1.10 192.168.1.1**
- **ping -I 192.168.2.10 192.168.2.1**

If the network communication is normal, the routes of the NICs are correctly configured.

Otherwise, you need to configure policy-based routing for the ECS with multiple NICs by referring to [How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?](#)

ECS Network Failure

1. Log in to the ECS.

Multiple methods are available for logging in to an ECS. For details, see [Logging In to an ECS](#).

2. Check whether the ECS NIC has an IP address assigned.


- Linux ECS: Use the **ifconfig** or **ip address** command to view the IP address of the NIC.
- Windows ECS: In the search box, enter **cmd** and press **Enter**. In the displayed command prompt, run the **ipconfig** command.

If the ECS NIC has no IP address assigned, see [Why Does My ECS Fail to Obtain an IP Address?](#)

3. Check whether the subnet gateway of the ECS can be pinged.

- a. In the ECS list, click the ECS name.
The ECS details page is displayed.
- b. On the ECS details page, click the hyperlink of VPC.
The **Virtual Private Cloud** page is displayed.
- c. In the VPC list, locate the target VPC and click the number in the **Subnets** column.
The **Subnets** page is displayed.
- d. In the subnet list, click the subnet name.
The subnet details page is displayed.
- e. Click the **IP Addresses** tab and view the gateway address of the subnet.

Figure 4-6 Gateway address



- f. Check whether the gateway communication is normal:

ping *Subnet gateway address*

Example command: **ping 172.17.0.1**

If the gateway address cannot be pinged, see [Why Does My ECS Fail to Communicate at a Layer 2 or Layer 3 Network?](#)

5 Virtual IP Addresses

5.1 Why Can't the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?

Symptom

After you bind a virtual IP address to an ECS NIC, you cannot ping the virtual IP address.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 5-1 Troubleshooting




Table 5-1 Troubleshooting

Possible Cause	Solution
Virtual IP address of the ECS NIC	See Virtual IP Address of the ECS NIC
Virtual IP address of the internal NIC of the ECS	See Virtual IP Address of the Internal NIC of the ECS

Possible Cause	Solution
Security group or network ACL configuration	See Security Group or Network ACL Configuration

Virtual IP Address of the ECS NIC

Check whether the source/destination check of the NIC is disabled and whether a virtual IP address is bound to the NIC.

1. Log in to the management console.
2. Click **Service List** and click **Elastic Cloud Server** under **Computing**.
3. In the ECS list, click the name of the ECS.
4. On the displayed ECS details page, click the **Network Interfaces** tab.
5. Ensure that **Source/Destination Check** is disabled.
6. Ensure that an IP address is displayed for **Virtual IP Address** on the NIC details page.

If there is no virtual IP address, click **Manage Virtual IP Address**. On the displayed **IP Addresses** tab, click **Assign Virtual IP Address**.

NOTE

To check whether a virtual IP address has been configured, **ifconfig** will not work. Use **ip address** instead. For more information, see [Binding a Virtual IP Address to an EIP or ECS](#).

Virtual IP Address of the Internal NIC of the ECS

The following uses Linux and Windows ECSs as examples to describe how to check whether an ECS NIC has a virtual IP address.

For a Linux ECS:

1. Check if there is a NIC **ethX:X**
ifconfig

Figure 5-2 Checking for NIC **ethX:X**

```
[root@scy ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.1.2 netmask 255.255.255.0 broadcast 192.168.1.255
    inet6 fe80::f816:3eff:fe4d:5b98 prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
    RX packets 77399 bytes 5101164 (4.8 MiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 68798 bytes 8090922 (7.7 MiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth0:1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.1.137 netmask 255.255.255.0 broadcast 192.168.1.255
    ether fa:16:3e:4d:5b:98 txqueuelen 1000 (Ethernet)
```

The command output in the preceding figure contains a NIC **ethX:X**. **192.168.1.137** is its virtual IP address.

- If the NIC **ethX:X** is there, the ECS NIC is correctly configured.
 - If the NIC **ethX:X** cannot be found, perform the following operations:
2. If the command output does not contain a NIC **ethX:X**, switch to the **/etc/sysconfig/network-scripts** directory:
cd /etc/sysconfig/network-scripts
 3. Run the following command to create and then modify the **ifcfg-eth0:1** file:
vi ifcfg-eth0:1
Add the following NIC information to the file:

```
BOOTPROTO=static
DEVICE=eth0:1
HWADDR=fa:16:3e:4d:5b:98
IPADDR=192.168.1.137
GATEWAY=192.168.1.1
NETMASK=255.255.255.0
ONBOOT=yes
ONPARENT=yes
```
 4. Press **Esc**, enter **:wq!**, and save the file and exit.
 5. Restart the ECS and run the **ifconfig** command to check whether the virtual IP address has been configured for the ECS.

For a Windows ECS:

1. In the **Start** menu, open the Windows command line window and run the following command to check whether the virtual IP address has been configured:

ipconfig /all

Figure 5-3 Checking whether the virtual IP address has been configured

```
C:\Users\Administrator>ipconfig /all

Windows IP Configuration

Host Name . . . . . : dst-win
Primary Dns Suffix . . . . . :
Node Type . . . . . : Hybrid
IP Routing Enabled. . . . . : No
WINS Proxy Enabled. . . . . : No


Ethernet adapter Ethernet 5:

Connection-specific DNS Suffix . . :
Description . . . . . : Red Hat VirtIO Ethernet Adapter #2
Physical Address. . . . . : FA-16-3E-83-B2-73
DHCP Enabled. . . . . : No
Autoconfiguration Enabled . . . . : Yes
Link-local IPv6 Address . . . . . : fe80::6182:a265:10bc:134e%3(Preferred)
IPv4 Address. . . . . : 192.168.10.41(Preferred)
Subnet Mask . . . . . : 255.255.255.0
IPv4 Address. . . . . : 192.168.10.137(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : 192.168.10.1
DHCPv6 IAID . . . . . : 184161854
DHCPv6 Client DUID. . . . . : 00-01-00-01-21-9F-1A-85-52-54-00-A6-AD-AC
DNS Servers . . . . . : 100.125.1.250
                        114.114.114.114
NetBIOS over Tcpip. . . . . : Enabled
```


In the preceding command output, check whether the value of **IPv4 Address** (192.168.10.137) is the IP address of the ECS NIC.

- If yes, the virtual IP address has been configured for the ECS NIC.

- If no, perform the following operations:
- 2. Choose **Control Panel > Network and Internet > Network Connections**. Right-click the corresponding local connection and then click **Properties**.
- 3. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.
- 4. Click **Properties**.
- 5. Select **Use the following IP address**, and set **IP address** to the private IP address displayed in [Figure 5-3](#). For example, 192.168.10.41.

Figure 5-4 Configuring a private IP address

- 6. Click **Advanced**.
- 7. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address configured in [Figure 5-3](#). For example, 192.168.10.137.

Figure 5-5 Configuring virtual IP address

Security Group or Network ACL Configuration

Check whether the ECS security groups and the network ACLs associated with the subnet used by the ECS NIC are blocking traffic.

1. On the ECS details page, click the **Security Groups** tab and confirm that required security group rules have been configured for the virtual IP address. If the required security group rules have not been configured, click **Change Security Group** or **Modify Security Group Rule** to change the security group or modify the security group rules.
2. Click **Service List**. Under **Networking**, click **Virtual Private Cloud**. In the navigation pane on the left of the network console, click **Network ACLs** and check whether the network ACL rules associated with the subnet used by the ECS NIC are blocking access to the virtual IP address.

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

5.2 How Do I Bind a Virtual IP Address on Huawei Cloud to a Server in an On-premises Data Center?

Prerequisites

- You have assigned virtual IP addresses. For details, see [Assigning a Virtual IP Address](#).
- You have created a Layer 2 connection for the subnet where the virtual IP address resides. For details, see [Buying an Enterprise Switch](#).

Procedure

1. Log in to the management console.
2. On the console homepage, under **Networking**, click **Virtual Private Cloud**.
3. In the navigation tree on the left, choose **Enterprise Switch**.
4. Click **Manage Virtual IP Address** on the right of **Layer 2 Connection Topology**.
5. In the IP address list, locate row that contains the target virtual IP address and click **Bind to Instance** in the **Operation** column.
6. On the **Bind to Instance** page, set **Instance Type** to **Layer 2 Connection**, select the target Layer 2 connection, and click **OK**.

5.3 Why Is the Network Disconnected Between Servers Using a Virtual IP Address After an Active/Standby Switchover?

For an HA cluster using virtual IP addresses and Keepalived, if you find that the network between the client and the server is disconnected after an active/standby switchover, the possible cause is that the switchover is performed manually. As a result, the ARP table on the client is not updated, you can perform the following operations to update the ARP table:

1. Log in to the client.
2. Update the ARP table on the client.
 - Method 1: Trigger the client to learn the new MAC address corresponding to the virtual IP address:
ping *Virtual IP address*
Example command: **ping 192.168.3.22**
 - Method 2: Clear the residual entries in the ARP table of the virtual IP address to trigger the client to learn the new ARP table:
arp -d *Virtual IP address*
Command example: **arp -d 192.168.3.22**

6 Bandwidth

6.1 What Are Inbound Bandwidth and Outbound Bandwidth?

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

Public bandwidth is the bandwidth consumed when data is transferred between Huawei Cloud instances and the Internet. Public bandwidth is classified into inbound bandwidth and outbound bandwidth. For details the outbound bandwidth and inbound bandwidth, see [Table 6-1](#).

Figure 6-1 Inbound bandwidth and outbound bandwidth




Table 6-1 Inbound bandwidth and outbound bandwidth

Type	Description
Outbound bandwidth	<p>Bandwidth consumed when data is transferred from Huawei Cloud to the Internet. For example, the outbound bandwidth is used when ECSs provide services accessible from the Internet and FTP clients download resources from the ECSs. Outbound bandwidth means the same thing as upstream bandwidth on the Cloud Eye console.</p> <p>Huawei Cloud only charges for the outbound bandwidth.</p> <p>NOTE</p> <ul style="list-style-type: none">• To view the bandwidth usage, see Viewing Metrics.• To view bandwidth billing details, see Bills.
Inbound bandwidth	<p>Bandwidth consumed when data is transferred from the Internet to Huawei Cloud. For example, the inbound bandwidth is used when resources are downloaded from the Internet to ECSs and FTP clients upload resources to the ECSs. Inbound bandwidth means the same thing as downstream bandwidth on the Cloud Eye console.</p> <p>The maximum inbound bandwidth depends on the size of the outbound bandwidth.</p> <ul style="list-style-type: none">• If your purchased bandwidth is less than or equal to 10 Mbit/s, the inbound bandwidth will be 10 Mbit/s, and the outbound bandwidth will be the same as the purchased bandwidth.• If your purchased bandwidth is greater than 10 Mbit/s, the outbound and inbound bandwidth will be the same as the purchased bandwidth. <p>The preceding bandwidth limit is not applicable in CN North-Beijing1 and CN East-Shanghai2.</p>

6.2 What Are the Differences Among Static BGP, Dynamic BGP, and Premium BGP?

Border Gateway Protocol (BGP) is a routing protocol used between autonomous systems (ASs). BGP is the only protocol that can process many connections between unrelated routing domains. The EIP service connects to networks provided by China Unicom, China Telecom, China Mobile, and other carriers.

When assigning an EIP, you can select from the following EIP types:

- Static BGP routes are manually configured by network carriers.
- Dynamic BGP provides automatic failover and chooses the optimal path when a network connection fails.
- Premium BGP chooses the optimal path and ensures low-latency and high-quality networks. BGP is used to interconnect with lines of multiple mainstream carriers. Public network connections that feature low latency and

high quality are directly established between the Chinese mainland and Hong Kong (China). (Premium BGP is available only in **CN-Hong Kong**.)

The EIP type cannot be changed. For example, dynamic BGP EIPs cannot be changed to static BGP EIPs because they are in different IP address pools.

For details about static BGP, dynamic BGP, and premium BGP and their differences, see [Table 6-2](#).

Table 6-2 Differences among static BGP, dynamic BGP, and premium BGP

Item	Static BGP	Dynamic BGP	Premium BGP
Definition	Static routes are manually configured and must be manually reconfigured anytime when the network topology or link status changes.	Dynamic BGP provides automatic failover and chooses the optimal path based on the real-time network conditions as well as preset policies.	Premium BGP chooses the optimal path and ensures low-latency and high-quality networks. BGP is used to interconnect with lines of multiple mainstream carriers. Public network connections that feature low latency and high quality are directly established between Chinese mainland and Hong Kong (China).
Assurance	When changes occur on a network that uses static BGP, the manual configuration takes some time and high availability cannot be guaranteed. If you select static BGP, your application system must have disaster recovery setups in place.	When a fault occurs on a carrier's link, dynamic BGP will quickly select another optimal path to take over services, ensuring service availability. Currently, carriers in China that support dynamic BGP routing include China Telecom, China Mobile, China Unicom, China Education and Research Network (CERNET), National Radio and Television Administration, and Dr. Peng Group.	Premium BGP has the same assurance capability as that of dynamic BGP. In addition, premium BGP ensures higher network quality and lower latency. Currently, mainstream carriers in Hong Kong (China) are supported.

Item	Static BGP	Dynamic BGP	Premium BGP
Advantages	This is a more cost-effective option that allows resources to access the Internet over a single carrier network. The routes can be manually configured.	The BGP public network egress supports switchover across domains within seconds, providing your users with high-speed and secure networks.	<ul style="list-style-type: none">• Premium BGP chooses the optimal path for access from the abroad.• It allows users in the Chinese mainland to access cross-border applications faster.
Service availability	99%	99.95%	99.95%
Billing	Their price from least to most expensive: static BGP, dynamic BGP, and premium BGP. For details, see EIP Pricing Details .		

 NOTE

For more information about service availability, see [Huawei Cloud Service Level Agreement](#).

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

Symptom

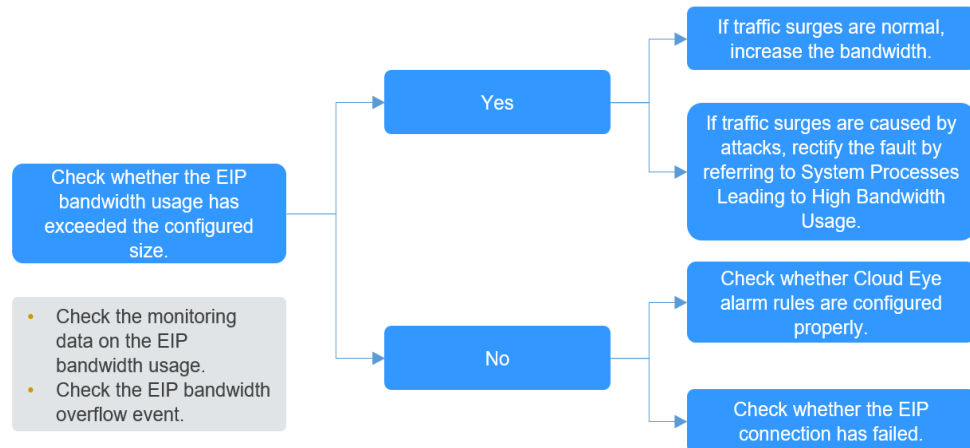
The bandwidth size configured when you buy a dedicated or shared bandwidth defines the maximum amount of outbound bandwidth supported. If an ECS running your web applications cannot be accessed smoothly from the Internet, check whether the bandwidth of the EIP bound to the ECS is greater than the configured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

 NOTE

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

Troubleshooting

Troubleshoot the issue by following the procedure described below. If the problem persists, [submit a service ticket](#).

Figure 6-2 Troubleshooting procedure**Step 1 Check whether the EIP bandwidth usage has exceeded the configured size.**

1. Check the monitoring data on the EIP bandwidth usage.
Check whether the inbound bandwidth and outbound bandwidth usage have exceeded the amount purchased. For details, see [Exporting Monitoring Data](#).
2. Check **EIP bandwidth overflow** event.
For details about how to check the event, see [a](#).
If you have not configured EIP bandwidth overflow events, configure one by referring to [solution 2](#). If there is packet loss or access delay, you can view **EIP bandwidth overflow** event on the **Event Monitoring** page.

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

Step 2 Fix the excessive bandwidth usage issue.

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

Check whether the sudden increase in bandwidth is normal.

1. If it is normal, [increase the bandwidth](#).
2. If it is not normal, for example, the sudden increase in bandwidth was caused by online attacks, rectify the fault by referring to [System Processes Leading to High Bandwidth Usage](#).

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

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

- Check whether Cloud Eye alarm rules are configured properly.
If there are alarm rules for bandwidth usage on the Cloud Eye console, where the outbound bandwidth limit or the alarm period is set too small, the system may generate excessive alarms. You can refer to [Improper Cloud Eye Alarm Rules](#) to fix the problem.

- Check whether the EIP connection has failed.
If an ECS with an EIP bound cannot access the Internet, you can refer to [Why Can't My ECS Access the Internet Even After an EIP Is Bound?](#)

----End

System Processes Leading to High Bandwidth Usage

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

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

- [Why Is My Windows ECS Running Slowly?](#)
- [Why Is My Linux ECS Running Slowly?](#)

Improper Cloud Eye Alarm Rules

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

- Solution 1: Create a more appropriate bandwidth alarm rule.
If there are alarm rules for bandwidth usage on the Cloud Eye console, where the outbound bandwidth limit or the alarm period is set too small, the system may generate excessive alarms. You need to set an appropriate alarm rule based on your purchased bandwidth. For example, if your purchased bandwidth is 5 Mbit/s, you can create an alarm rule to report an alarm if the outbound bandwidth reaches 4.8 Mbit/s for three consecutive periods. To create an alarm rule:
 - a. Log in to the management console, under **Management & Governance**, click **Cloud Eye**. On the **Cloud Eye** console, choose **Alarm Management > Alarm Rules**.
 - b. Click **Create Alarm Rule** and configure an alarm rule to generate alarms when the bandwidth usage exceeds the configured limit.
- Solution 2: Configure **EIP bandwidth overflow** events.

NOTE

The **Event Monitoring** page only displays EIP status. It does not display the shared bandwidth limit.

To configure **EIP bandwidth overflow** events:

- a. Log in to the management console, under **Management & Governance**, click **Cloud Eye**. On the **Cloud Eye** console, choose **Event Monitoring**.
- b. Click **Create Alarm Rule** and configure an alarm rule to generate alarms when the EIP bandwidth usage exceeds the limit.

After the configuration, you can view the usage details of the EIP dedicated bandwidth on the **Event Monitoring** page when there are packet loss or data transfer delays.

To check the EIP bandwidth overflow history, perform the following steps:

- a. On the **Cloud Eye** console, click **Event Monitoring**.
- b. On the **Event Monitoring** page, locate the target monitoring event and click **View Graph** in the **Operation** column.
- c. On the system event list page, locate the target monitored object and click **View Event** in the **Operation** column to view the bandwidth overflow details.

If the event **EIP bandwidth overflow** is not displayed, the usage of the dedicated EIP bandwidth did not exceed the preset limit.

If the event **EIP bandwidth overflow** is displayed, the usage of the dedicated EIP bandwidth exceeded the limit. To ensure stability and high availability of your workload, [you can increase your bandwidth](#).

You will not be billed for Cloud Eye alarms, but if you enable SMN to send out alarm notifications, this will incur charges. For details, see the [Cloud Eye User Guide](#).

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

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

Public Bandwidth

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

Public bandwidth is classified into inbound bandwidth and outbound bandwidth.

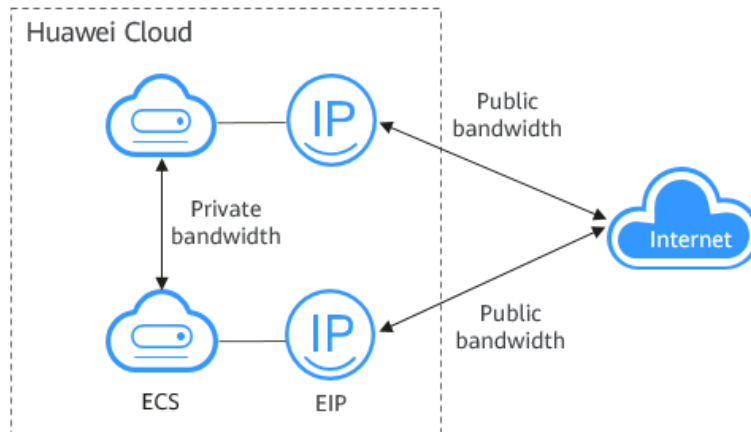
Inbound bandwidth is the bandwidth consumed when data is transferred from the Internet to Huawei Cloud. For example, when resources are downloaded from the Internet to ECSs, that consumes inbound bandwidth.

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

Private Bandwidth

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

For details, see [ECS Types](#).

Figure 6-3 Public bandwidth and private bandwidth

6.5 What Bandwidth Types Are Available?

There are dedicated or shared bandwidths.

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

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

6.6 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth?

A dedicated bandwidth can only be used by one EIP that is bound to one cloud resource, such as an ECS, a NAT gateway, or a load balancer.

A shared bandwidth can be shared by multiple pay-per-use EIPs. Adding an EIP to or removing an EIP from a shared bandwidth does not affect your services.

A dedicated bandwidth cannot be changed to a shared bandwidth or the other way around. You can purchase a shared bandwidth for your pay-per-use EIPs.

- After you add an EIP to a shared bandwidth, the EIP will use the shared bandwidth.
- After you remove an EIP from a shared bandwidth, the EIP will use the dedicated bandwidth.

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

A shared bandwidth can be used by multiple EIPs.

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

If the current quota cannot meet service requirements, [submit a service ticket](#) to increase the quota.

6.8 Can I Increase a Yearly/Monthly Bandwidth and Decrease It Later?

The bandwidth can be decreased after it is increased.

For details about how to adjust the bandwidth, see [Modifying an EIP Bandwidth](#).

- Increasing the bandwidth: The change is applied immediately.
You need to pay the price difference.
- Decrease bandwidth upon renewal: The change is not applied immediately.
You need to select a new bandwidth size and a renewal duration. The change will be applied in the first billing cycle after a successful renewal.
- Decreasing the bandwidth immediately: The change is applied immediately.
The price difference between the original and the new bandwidth will be refunded.

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

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

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

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

7 Connectivity

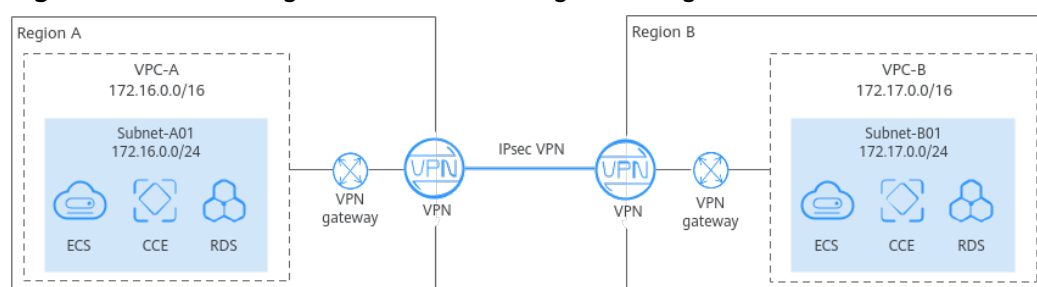
7.1 Does a VPN Allow Communication Between Two VPCs?

If the two VPCs are in the same region, you can use a VPC peering connection to enable communication between them.

If the two VPCs are in different regions, you can use a VPN to enable communication between the VPCs. The CIDR blocks of the two VPCs are the local and remote subnets, respectively.

In [Figure 7-1](#), there is a VPC in each region: VPC-A in region A and VPC-B in region B. Each VPC is connected to a VPN connection. The two VPCs can communicate with each other through an encrypted channel on the Internet. VPN can be enabled fast and is cost-effective.

Figure 7-1 Connecting VPCs in different regions using VPN



7.2 Why Cannot I Access Public Websites Through Domain Names or Access Internal Domain Names on the Cloud When My ECS Has Multiple NICs?

When an ECS has more than one NIC, if different DNS server addresses are configured for the subnets used by the NICs, the ECS cannot access the Internet or domain names on the cloud.

You can resolve this issue by configuring the same DNS server address for the subnets used by the same ECS. You can perform the following steps to modify DNS server addresses of subnets in a VPC:

1. Log in to the management console.
2. On the console homepage, under **Networking**, click **Virtual Private Cloud**.
3. In the navigation pane on the left, click **Virtual Private Cloud**.
4. On the **Virtual Private Cloud** page, locate the VPC for which a subnet is to be modified and click the VPC name.
5. In the subnet list, locate the row that contains the subnet to be modified, click **Modify**. On the displayed page, change the DNS server address as prompted.
6. Click **OK**.

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

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

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

If an ECS in the VPC has an EIP bound, the VPC route table will have a policy-based route with 0.0.0.0/0 as the destination, which has a higher priority than its custom route. In this case, traffic is forwarded to the EIP and cannot reach the NAT gateway.

7.4 Why Are There Intermittent Interruptions When a Local Host Accesses a Website Built on an ECS?

Symptom

After you build a website on an ECS, some users occasionally are unable to access the website through the local network.

Troubleshooting

1. Check the local network of the user.
If the local host communicates with the ECS using NAT, this problem may occur.
2. Run the following command to check whether **tcp_tw_recycle** is enabled on the ECS:
sysctl -a|grep tcp_tw_recycle
If the value of **tcp_tw_recycle** is **1**, the function is enabled.
3. Run the following command to check the number of lost packets of the ECS:
cat /proc/net/netstat | awk '/TcpExt/ { print \$21,\$22 }'

If the value of **ListenDrops** is not **0**, there is packet loss, that is, the network is faulty.

Procedure

This problem can be solved by modifying the kernel parameters of the ECS.

- Run the following command to temporarily modify the parameters (the parameters will change back after a restart):
sysctl -w net.ipv4.tcp_tw_recycle=0
- Perform the following operations to permanently modify the parameters:
 - a. Run the following command and modify the **/etc/sysctl.conf** file:
vi /etc/sysctl.conf
Add the following content to the file:
`net.ipv4.tcp_tw_recycle=0`
 - b. Press **Esc**, enter **:wq!**, and save the file and exit.
 - c. Run the following command to make the modification take effect:
sysctl -p

7.5 Why Do ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communication?

Symptom

Two ECSs (**ecs01** and **ecs02**) are in the same subnet in a VPC. Their IP addresses are 192.168.1.141 and 192.168.1.40.

The **ecs01** can ping **ecs02** through a private IP address successfully, but **ecs02** cannot ping **ecs01** through a private IP address.

Troubleshooting

1. Ping **ecs01** from **ecs02** through the EIP. If **ecs01** can be pinged, the NIC of **ecs01** is working properly.
2. Run the **arp -n** command on **ecs02** to check whether the command output contains the MAC address of **ecs01**. If the command output does not contain the MAC address of **ecs01**, **ecs02** fails to learn the MAC address of **ecs01** when using the private IP address to ping **ecs01**.
3. Run the **ip a** command on **ecs01** to check the NIC configuration of **ecs01**. The following figure shows an example.

Figure 7-2 Viewing **ecs01** NIC configuration

```
[root@bd-slave1 ~]# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP qlen 1000
    link/ether fa:16:3e:62:1d:d5 brd ff:ff:ff:ff:ff:ff
    inet 192.168.1.141/24 brd 192.168.1.255 scope global eth0
    inet 192.168.1.40/32 scope global eth0
    inet6 fe80::f816:3eff:fe62:1dd5/64 scope link
        valid_lft forever preferred_lft forever
```

The IP address 192.168.1.40/32 should not be configured based on the command output. As a result, **ecs01** fails to send packets to **ecs02**.

Procedure

Modify the NIC configuration of **ecs01**. Run the following command to delete the redundant IP address, for example, 192.168.1.40/32, configured on the NIC **eth0**:

```
ip a del 192.168.1.40/32 dev eth0
```

7.6 Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur When They Communicate?

Symptom

Two ECSs in the same VPC cannot communicate with each other or there is packet loss when they communicate.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 7-3 Troubleshooting

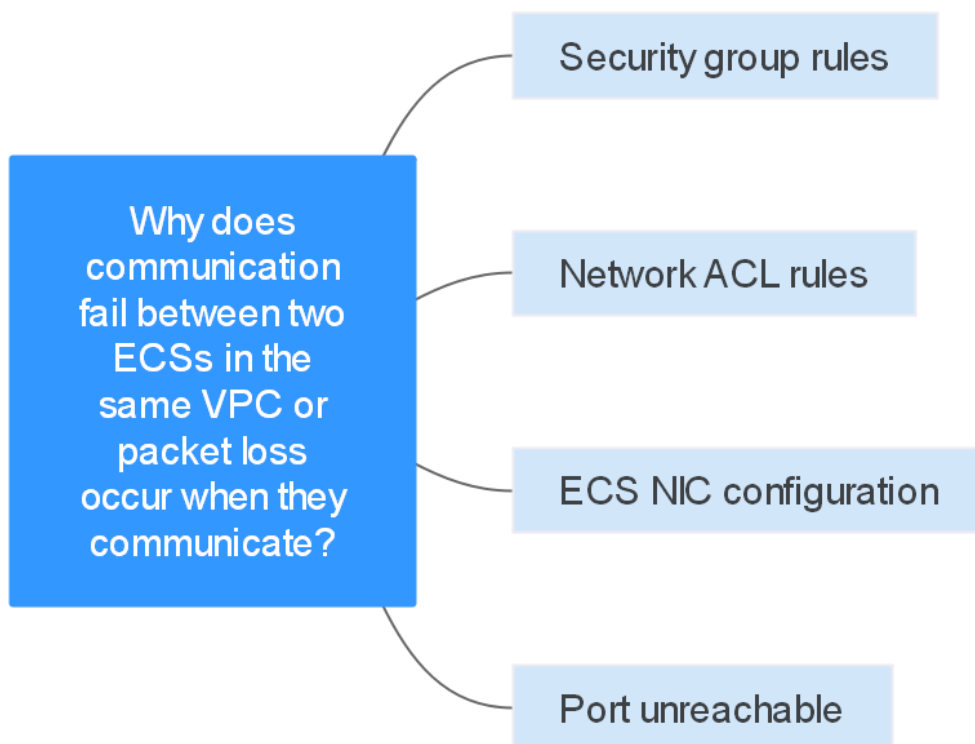


Table 7-1 Troubleshooting

Possible Cause	Solution
Security group rules	See Security Group Rules
Network ACL rules	See Network ACL Rules
ECS NIC configuration	See ECS NIC Configuration
Port unreachable	See Port Unreachable

Security Group Rules

Check whether the ECS NIC security group allows the outbound and inbound ICMP traffic.

Take the inbound direction as an example. The security group rules must contain at least one of the following rules.

Figure 7-4 Inbound security group rule

Protocol & Port	Type	Source	Description	Operation
All	IPv4	0.0.0.0/0	--	Modify Replicate Delete
ICMP: All	IPv4	0.0.0.0/0	--	Modify Replicate Delete

If packets of other protocols are tested, configure the security group rules to allow the corresponding protocol traffic. For example, if UDP packets are tested, check whether the security group allows the inbound UDP traffic.

Network ACL Rules

1. Check whether the subnet of the ECS NIC has an associated network ACL.
2. Check the network ACL status in the network ACL list.
 - If **Disabled** is displayed in the **Status** column, the network ACL has been disabled. Go to [3](#).
 - If **Enabled** is displayed in the **Status** column, the network ACL has been enabled. Go to [4](#).
3. Click the network ACL name and configure rules on the **Inbound Rules** and **Outbound Rules** tabs to allow the ICMP traffic.
4. If the network ACL is disabled, all packets in the inbound and outbound directions are discarded by default. In this case, delete the network ACL or enable the network ACL and allow the ICMP traffic.

ECS NIC Configuration

The following procedure uses a Linux ECS as an example. For a Windows ECS, check the firewall restrictions.

1. Check whether multiple NICs are configured for the ECS. If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policy-based routing for the ECS.

For details, see [How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?](#)

2. Log in to the ECS and run the following command to check whether the NIC has been created and obtained a private IP address. If there is no NIC information or the private IP address cannot be obtained, contact technical support.

ifconfig

Figure 7-5 NIC IP address

```
root@ecs-acl ~]# ifconfig
eth0  Link encap:Ethernet  HWaddr FA:16:3E:BC:B7:81
       inet addr:192.168.72.289  Bcast:192.168.72.255  Mask:255.255.255.0
       inet6 addr: fe80::f816:3eff:febc:b781/64  Scope:Link
       UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
       RX packets:881  errors:0  dropped:0  overruns:0  frame:0
       TX packets:547  errors:0  dropped:0  overruns:0  carrier:0
       collisions:0  txqueuelen:1000
       RX bytes:49684 (48.4 KiB)  TX bytes:44454 (43.4 KiB)
       Interrupt:46
```

3. If the CPU usage exceeds 80%, the ECS communication may be adversely affected. Run the following command to check whether the CPU usage of the ECS is too high:

top

4. Run the following command to check whether the ECS has any restrictions on security group rules:

iptables-save

5. Run the following command to check whether the `/etc/hosts.deny` file contains the IP addresses that limit communication:

vi /etc/hosts.deny

If the `hosts.deny` file contains the IP address of another ECS, delete the IP address from the `hosts.deny` file and save the file.

Port Unreachable

1. If a port of the ECS cannot be reached, check whether the security group rules and network ACL rules enable the port.
2. On the Linux ECS, run the following command to check whether the ECS listens on the port: If the ECS does not listen on the port, the ECS communication may be adversely affected.

netstat -na | grep <Port number>

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

7.7 Why Can't My ECS Use Cloud-init?


Symptom

An ECS is not able to use cloud-init.

Troubleshooting

Figure 7-6 shows the process for an ECS to obtain metadata using the cloud-init.

Figure 7-6 Process for obtaining metadata



Check the following possible causes.

Figure 7-7 Possible causes

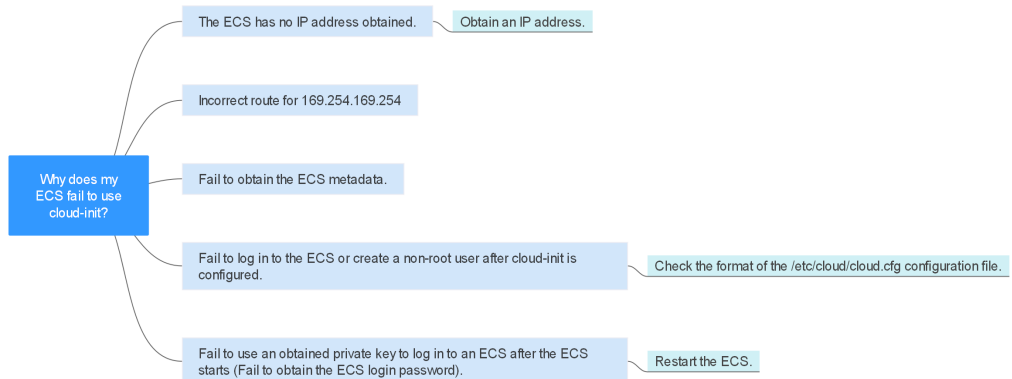


Table 7-2 Possible causes

Possible Cause	Solution
The ECS has no IP address obtained.	See The ECS Has Not Obtained IP Address
Incorrect route for 169.254.169.254	See Incorrect Route for 169.254.169.254
Fail to obtain the ECS metadata.	See Did Not Obtain the ECS Metadata

Possible Cause	Solution
Fail to log in to the ECS or create a non-root user after cloud-init is configured.	Check the format of the <code>/etc/cloud/cloud.cfg</code> configuration file. For details, see Cannot Log in to the ECS or Create a Non-root User After Cloud-init Is Configured .
Fail to use an obtained private key to log in to an ECS after the ECS starts (Fail to obtain the ECS login password).	Restart the ECS and try again.

The ECS Has Not Obtained IP Address

Check whether the ECS has obtained an IP address.

If no IP address is obtained, run the `dhclient` command to obtain the IP address (this command varies depending on the ECS OSs). Alternatively, you can run the `ifdown ethx` command to disable the network port and then run the `ifup ethx` command to enable it to allow the ECS NIC to automatically obtain an IP address again.

Figure 7-8 ECS IP address

```
-bash-4.1# ifconfig
eth0      Link encap:Ethernet  HWaddr FA:16:3E:BD:36:DD
          inet addr:192.168.1.200  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:febd:36dd/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:73008 errors:0 dropped:0 overruns:0 frame:0
          TX packets:26295 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:4162713 (3.9 MiB)  TX bytes:2336476 (2.2 MiB)
          Interrupt:35

eth1      Link encap:Ethernet  HWaddr FA:16:3E:A9:C7:1D
          inet addr:192.168.1.179  Bcast:192.168.1.255  Mask:255.255.255.0
          inet6 addr: fe80::f816:3eff:fea9:c71d/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:45026 errors:0 dropped:0 overruns:0 frame:0
          TX packets:12244 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1270534 (1.2 MiB)  TX bytes:4178924 (3.9 MiB)
          Interrupt:34

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:65536  Metric:1
          RX packets:1 errors:0 dropped:0 overruns:0 frame:0
          TX packets:1 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:28 (28.0 b)  TX bytes:28 (28.0 b)
```

Incorrect Route for 169.254.169.254

Ping **169.254.169.254/32** from the ECS. If the IP address cannot be pinged, perform the following steps:

1. Check the exact route configured on the ECS for IP address **169.254.169.254/32**.

In most cases, the next hop of the exact route for IP address **169.254.169.254/32** is the same as that of the default route for the IP address.

Figure 7-9 Route for IP address 169.254.169.254/32

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

2. If there is no exact route for IP address **169.254.169.254/32**, the cause is as follows:

Images with CentOS 5 OSs are not compatible with cloud-init. To use cloud-init, select a different OS.

3. If the next hop of the exact route for IP address **169.254.169.254/32** is different from that of the default route for the IP address:
 - If the ECS was created before cloud-init was enabled, run **service network restart** to obtain the correct route.
 - If the ECS is newly created, [submit a service ticket](#) or contact technical support.

Did Not Obtain the ECS Metadata

Run the following command on the ECS to obtain the metadata:

```
curl http://169.254.169.254/openstack/latest/meta_data.json
```

If information similar to that shown in [Figure 7-10](#) is displayed, the ECS successfully obtains the metadata.

Figure 7-10 Command output

```
-bash-4.1# curl http://169.254.169.254/openstack/latest/meta_data.json
{"random_seed": "rTUrSdIEh6A jUKLnvg51U8S0pH6xC78MFRTEwI0munBNyqos6q/EsAEJondF8iJkMDG0TzbcTb815HntS9X
XHu61u+y8fAeykba j60Av8KHMPgDv6Xdf hKu6gy jCr jXn5hUFvqfZ/yaJ3LrAE jB8N j59hI +umbP i8oYc2WzYmTqW jXYRNwpmqJM
s lKYm0CLuFbwYoZak1y27/WEUzDU0Q16pRkkwNwFaCn/rQQ/hHd+3UwSjBArsgUeowCTp5oxixL iCJzSSHAKz41UviziRxaYwm0go
iTfTopvZTwmYEk lFmkZsy7h6PP0kgm jgPn+1kZf0qqht lVpyRr2pw4aPaeZa4z7QX1RtnwT7MUGUbeaB5/1PDUE1J/GJpoH1/+z
rDye lA09Cs0G1UFuELadyDcrWA4k42f0o7dDmEjDm lNnEBeeqa5r7Eohb04KTImzi+3nb10Q jPq/S7J+mfM/UoZEJH0bZE4uw1Aj
2nhvy/pc6ho7fQKbx0C78f biPh59CKyF0WB35nNJ/CZNNBTd3UdG25SQ701FnA+NtbDeo8+g05 iFLwWeww0G5BLc jm1f jh9+mqot
+5ae6Zcexds11fscqm8 jwCnCimthJ lY6mbxu +6Fm9XpLDopDFrRtBUcRSnt IK67JprBSRppc+4shMyiuKY1JOTUJYQYDBUzB7F3o
=", "uuid": "53ebb737-ddc5-4303-9fac-aa72b00b101a", "availability_zone": "eu-de-02", "hostname": "ec
s-gjm-55eb.novalocal", "launch_index": 0, "meta": {"metering.image_id": "98721f93-722f-4386-a975-3cb
df1abf56d", "metering.imagetype": "gold", "metering.resourcespeccode": "c2.large.oracle", "metering.
cloudServiceType": "sys.service.type.ec2", "image_name": "AutoC_DTC_OEL_6.8", "metering.resourcetype":
"1", "os_bit": "64", "opc_id": "120b71c7-94ac-45b8-8ed6-30aafc8fbd8a", "os_type": "Linux", "charg
ing_mode": "0"}, "project_id": "efdf974f549b4eaab05c3903ddd2ab0e", "name": "ecs-gjm-55eb"}-bash-4.1#
```

Cannot Log in to the ECS or Create a Non-root User After Cloud-init Is Configured

Check whether the `/etc/cloud/cloud.cfg` configuration file format is correct. For details, see the file format requirements for different Linux distributions. The following figure shows an example `/etc/cloud/cloud.cfg` configuration file for Ubuntu.

Figure 7-11 Configuration file

```
system_info:
# This will affect which distro class gets used
distro: rhel
# Default user name + that default users groups (if added/used)
default_user:
  name: linux // Specifies the username for login.
  lock_passwd: False // The value False indicates that the password login mode is enabled. For some OSs, the value 0 indicates that the password login mode is enabled.
  gecos: Cloud User
  groups: users // Specifies whether the user will be added to a group. This parameter is optional. The groups parameter value must be an existing group under /etc/group in the system.
  passwd: $6$I63DBVKK$Zh41chiJR7NuZvtJHsYBQJIg5RoQCRL5IX2Hsgj2s5JwXITKU01we8WYcwbze aS2VNrRnNo28vmaxCyU6LwoD0
  sudo: ["ALL=(ALL) NOPASSWD:ALL"] // Specifies that all permissions of user root will be granted to the user.
  shell: /bin/bash // Specifies that the bash shell is used.
# Other config here will be given to the distro class and/or path classes
paths:
  cloud_dir: /var/lib/cloud/
  templates_dir: /etc/cloud/templates/
  ssh_svcname: sshd
```

Obtained Private Key Cannot Be Used to Log in to an ECS After the ECS Starts (Failed to Obtain the ECS Login Password)

Restart the ECS to rectify the fault.

Submitting a Service Ticket

If the EIP still fails to use cloud-init after performing the preceding steps, [submit a service ticket](#).

Provide the following information to the technical support engineer.

Item	Description	Example	Value
VPC CIDR block	Required for customer gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	Example: 120b71c7-94ac-45b8-8ed6-30aafc8fbd8a	N/A
CIDR block of subnet 1 (can be the same as the VPC CIDR block)	N/A	Example: 10.0.1.0/24	N/A
ECS ID	N/A	N/A	N/A
ECS IP address	N/A	Example: 192.168.1.192/24	N/A

Item	Description	Example	Value
ECS route information	N/A	N/A	N/A

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

Symptom

An ECS with an EIP bound cannot access the Internet.

Troubleshooting


Checking Whether EIPs Are Blocked or Frozen

- Check whether the EIP is blocked. For details, see [How Do I Unblock an EIP?](#)
- Check whether the EIP is frozen. For details, see [Why My EIPs Are Frozen? How Do I Unfreeze My EIPs?](#)

Checking EIP Connectivity


[Figure 7-12](#) shows the networking diagram for an ECS to access the Internet using an EIP.

Figure 7-12 EIP network diagram



Locate the fault based on the following procedure.

Figure 7-13 Troubleshooting procedure



1. **Step 1: Check Whether the ECS Is Running Properly**
2. **Step 2: Check Whether the Network Configuration of the ECS Is Correct**
3. **Step 3: Check Whether an EIP Has Been Assigned and Bound to the ECS**
4. **Step 4: Check Whether an EIP Is Bound to the Primary NIC of the ECS**
5. **Step 5: Check Whether Required Security Group Rules Have Been Configured.**
6. **Step 6: Check Whether Traffic from the ECS Subnet Is Blocked**

Step 1: Check Whether the ECS Is Running Properly

Check the ECS status.

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

Figure 7-14 ECS status

Name/ID	Monit...	Se...	Status	AZ	Specifications/Image	OS Type	IP Address	Billing Mode	Enterprise Pr...
697 4f9b-...			Stopped Locked...	AZ3	2 vCPUs 4 GiB t5.large.2 CCE_images_HCE20-Node-2...	Linux	99 (Privat...	Pay-per-use Created on Apr 11, 2024...	default
822 4ece-b...			Stopped Locked...	AZ1	2 vCPUs 4 GiB s7.large.2 CCE_images_HCE20-Node-2...	Linux	30 (Privat...	Pay-per-use Created on Apr 11, 2024...	default
388 499d-...			Stopped Locked...	AZ3	4 vCPUs 8 GiB t5.xlarge.2 CCE_images_HCE20-Node-2...	Linux	45 (EIP)... 8 (Private...	Pay-per-use Created on Apr 10, 2024...	default
c... 42ct1-...			Running	AZ4	2 vCPUs 4 GiB c7.large.2 CentOS 7.8 64bit	Linux	4 (Private...	Pay-per-use Created on Jan 22, 2024...	default

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

1. Check whether the ECS NIC has an IP address assigned.

Log in to the ECS, and run **ifconfig** or **ip address** to check the ECS NIC IP address.

If both the primary and extension NICs of an ECS have an EIP bound, check whether the ECS has policy-based routes configured. If policy-based routes are not configured, refer to [Configuring Policy-based Routes for a Linux ECS with Multiple NICs \(IPv4/IPv6\)](#).

If the ECS runs Windows, run **ipconfig**.

2. Check whether the ECS NIC has a virtual IP address.

Log in to the ECS, and run **ifconfig** or **ip address** to check whether the ECS NIC has a virtual IP address. If the ECS NIC has no virtual IP address, run the **ip addr add virtual IP address eth0** command to configure an IP address for the ECS NIC.

Figure 7-15 Virtual IP address of a NIC

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

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

Figure 7-16 Default route


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

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

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

The ECS shown in [Figure 7-17](#) has no EIP bound. It only has a private IP address bound.

Figure 7-17 EIP status



Name/ID	Monito...	Sec...	Status	AZ	Specifications/Image	OS Type	IP Address
ecs- b3b9*			Running	AZ3	1 vCPU 2 GiB s6.medium.2 CentOS 7.5 64bit	Linux	192.168.x.x (Private IP)

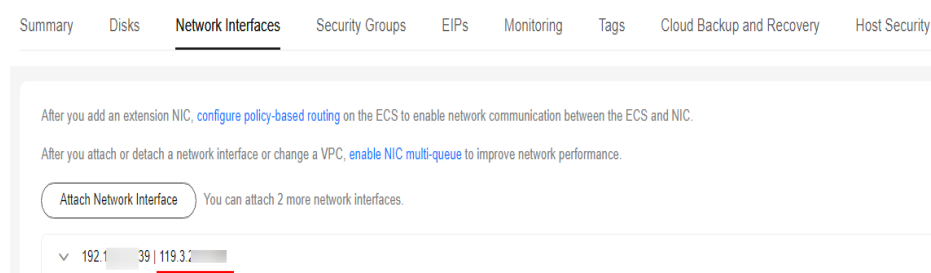
Step 4: Check Whether an EIP Is Bound to the Primary NIC of the ECS

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

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

As shown in the following figure, the EIP is bound to the primary NIC.

Figure 7-18 Checking whether the EIP is bound to the primary NIC of the ECS



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

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

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

Step 6: Check Whether Traffic from the ECS Subnet Is Blocked

Check whether the network ACL of the NIC subnet blocks certain traffic from the subnet.

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

Submitting a Service Ticket

If the EIP still cannot communicate with the Internet after you perform all the steps above, [submit a service ticket](#).

Provide the following information to technical support.

Item	Description	Example	Value
VPC CIDR block	Required for gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	Example: 120b71c7-94ac-45b8-8ed6-30aafc8fbdba	N/A
CIDR block of subnet 1 (can be the same as the VPC CIDR block)	N/A	Example: 10.0.1.0/24	N/A
ECS ID	N/A	N/A	N/A
ECS IP address	N/A	Example: 192.168.1.192/24	N/A
ECS route information	N/A	N/A	N/A
EIP	Required for the ECS to access the Internet	Example: 10.154.55.175	N/A
EIP bandwidth	Maximum bandwidth size used by the ECS to access the Internet	Example: 1 Mbit/s	N/A
EIP ID	N/A	Example: b556c80e-6345-4003-b512-4e6086abbd48	N/A

7.9 Why Is My ECS Unable to Communicate at a Layer 2 or Layer 3 Network?

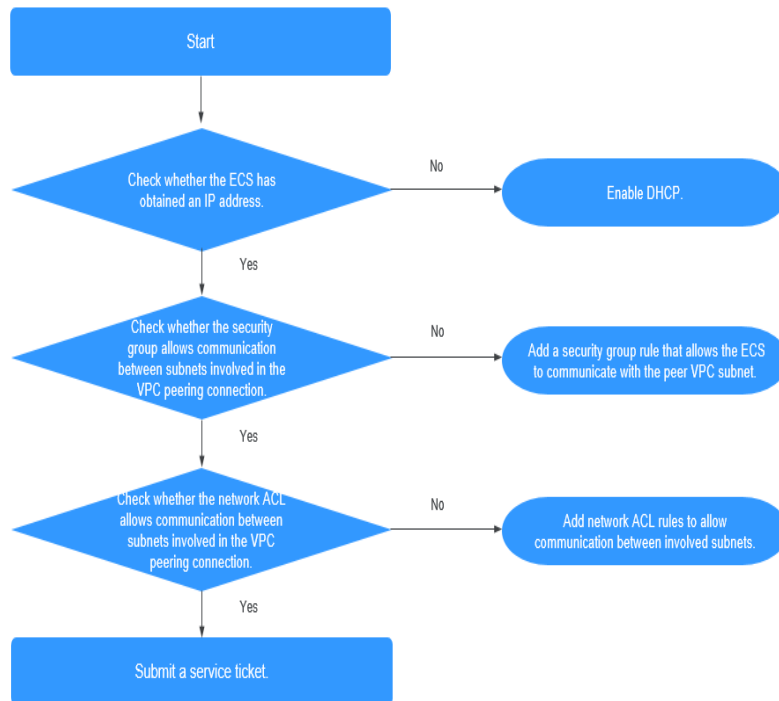
Symptom

An ECS cannot ping the gateway of the subnet where the ECS resides.

Troubleshooting

Locate the fault based on the following procedure.

Figure 7-19 Troubleshooting procedure



1. **Checking Whether the ECS Has Obtained an IP Address**
2. **Checking Whether the Security Group Allows Communication Between Subnets Involved in the VPC Peering Connection**
3. **Checking Whether the Network ACL Allows Communication Between Subnets Involved in the VPC Peering Connection**

Checking Whether the ECS Has Obtained an IP Address

Log in to the ECS, and run **ifconfig** or **ip address** to check the ECS NIC IP address. If an ECS runs Windows, use **ipconfig**.

If the ECS does not have an IP address, check whether DHCP has been enabled for the required subnet.


Switch to the subnet details page and check whether the DHCP function has been enabled.

For details, see [Why Does My ECS Fail to Obtain an IP Address?](#)

Checking Whether the Security Group Allows Communication Between Subnets Involved in the VPC Peering Connection

You can view the security group on the ECS details page. Check whether a security group rule has been configured to allow the ECS to communicate with the peer VPC subnet.

Figure 7-20 Security group rule



Checking Whether the Network ACL Allows Communication Between Subnets Involved in the VPC Peering Connection

In the navigation pane on the left of the VPC console, choose **Network ACLs**. On the displayed page, select the network ACL associated with the subnets of the VPC peering connection. On the network ACL details page, check whether network ACL rules allow communication between the subnets involved in the VPC peering connection.

Figure 7-21 Network ACL rule

The screenshot shows the AWS Management Console interface for Network ACL Rules. A table lists the rules:

Priority	Status	Type	Action	Protocol	Source	Source Port Range	Destination	Destination Port Range	Description	Operation
1	Enabled	IPv4	Allow	All	192.168.10.0/24	All	0.0.0.0	All	--	Modify Delete More
*	Enabled	--	Deny	All	0.0.0.0	All	0.0.0.0	All	--	Modify Delete More

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

7.10 How Do I Handle a BMS Network Failure?

1. Run the following command to check whether the BMS network ports have been bonded:

```
ifconfig
```

Figure 7-22 Checking for bond

```
[root@bms2 rhel]# ifconfig
bond0      Link encap:Ethernet  HWaddr FA:16:3E:E9:B0:8A
           inet addr:192.168.2.46  Bcast:192.168.2.255  Mask:255.255.255.0
           inet6 addr: fe80::f816:3eff:fee9:b08a/64 Scope:Link
           UP BROADCAST RUNNING PROMISC MASTER MULTICAST  MTU:8888  Metric:1
           RX packets:188108 errors:0 dropped:0 overruns:0 frame:0
           TX packets:112119 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:0
           RX bytes:42689694 (40.7 MiB)  TX bytes:82939564 (79.0 MiB)

bond0.2966 Link encap:Ethernet  HWaddr FA:16:3E:60:9C:CF
           inet addr:192.168.4.113  Bcast:192.168.4.255  Mask:255.255.255.0
           inet6 addr: fe80::f816:3eff:fe60:9ccf/64 Scope:Link
           UP BROADCAST RUNNING MULTICAST  MTU:8888  Metric:1
           RX packets:12 errors:0 dropped:0 overruns:0 frame:0
           TX packets:12 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:0
           RX bytes:660 (660.0 b)  TX bytes:720 (720.0 b)

eth0       Link encap:Ethernet  HWaddr FA:16:3E:E9:B0:8A
           UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
           RX packets:174667 errors:0 dropped:0 overruns:0 frame:0
           TX packets:112119 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:1000
           RX bytes:41874228 (39.9 MiB)  TX bytes:82939564 (79.0 MiB)

eth1       Link encap:Ethernet  HWaddr FA:16:3E:E9:B0:8A
           UP BROADCAST RUNNING SLAVE MULTICAST  MTU:8888  Metric:1
           RX packets:13441 errors:0 dropped:0 overruns:0 frame:0
           TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
           collisions:0 txqueuelen:1000
           RX bytes:815466 (796.3 KiB)  TX bytes:0 (0.0 b)
```

If no bonding information is obtained, the BMS network ports are not bonded. Contact technical support.

- Run the following command to check whether the BMS route information is correct:

```
route -n
```

Figure 7-23 Checking BMS route information

```
[root@bms2 rhel]# route -n
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
169.254.169.254 192.168.2.1 255.255.255.255 UGH 0 0 0 bond0
192.168.4.0 0.0.0.0 255.255.255.0 U 0 0 0 bond0.2966
192.168.2.0 0.0.0.0 255.255.255.0 U 0 0 0 bond0
169.254.0.0 0.0.0.0 255.255.0.0 U 1006 0 0 bond0
169.254.0.0 0.0.0.0 255.255.0.0 U 1007 0 0 bond0.2966
0.0.0.0 192.168.2.1 0.0.0.0 UG 0 0 0 bond0
[root@bms2 rhel]# █
```

Check whether the default route (with a destination of 0.0.0.0/0) exists.

Figure 7-24 Checking the default route

```
0.0.0.0 192.168.2.1 0.0.0.0 UG 0 0 0 bond0
[root@bms2 rhel]# █
```

Check whether a route to **169.254.169.254** exists.

Figure 7-25 Checking the route for IP address range **169.254.169.254**

```
Destination Gateway Genmask Flags Metric Ref Use Iface
169.254.169.254 192.168.2.1 255.255.255.255 UGH 0 0 0 bond0
```

If the required routes are not there, contact technical support.

- If BMSs in a VPC cannot communicate with each other or a BMS with an EIP bound cannot access the Internet, rectify the failure based on the related FAQ.
- If the failure cannot be rectified after you perform these operations, contact technical support.

Obtain the VPC and BMS information on the management console and provide the technical support engineer with the following information.

Item	Description	Example	Value
VPC 1 ID	VPC 1 ID	Example: fef65559-c154-4229-afc4-9ad0314437ea	N/A
BMS 1 ID	ID of BMS 1 in VPC 1	Example: f7619b12-3683-4203-9271-f34f283cd740	N/A
BMS 2 ID	ID of BMS 2 in VPC 1	Example: f75df766-68aa-4ef3-a493-06cdc26ac37a	N/A

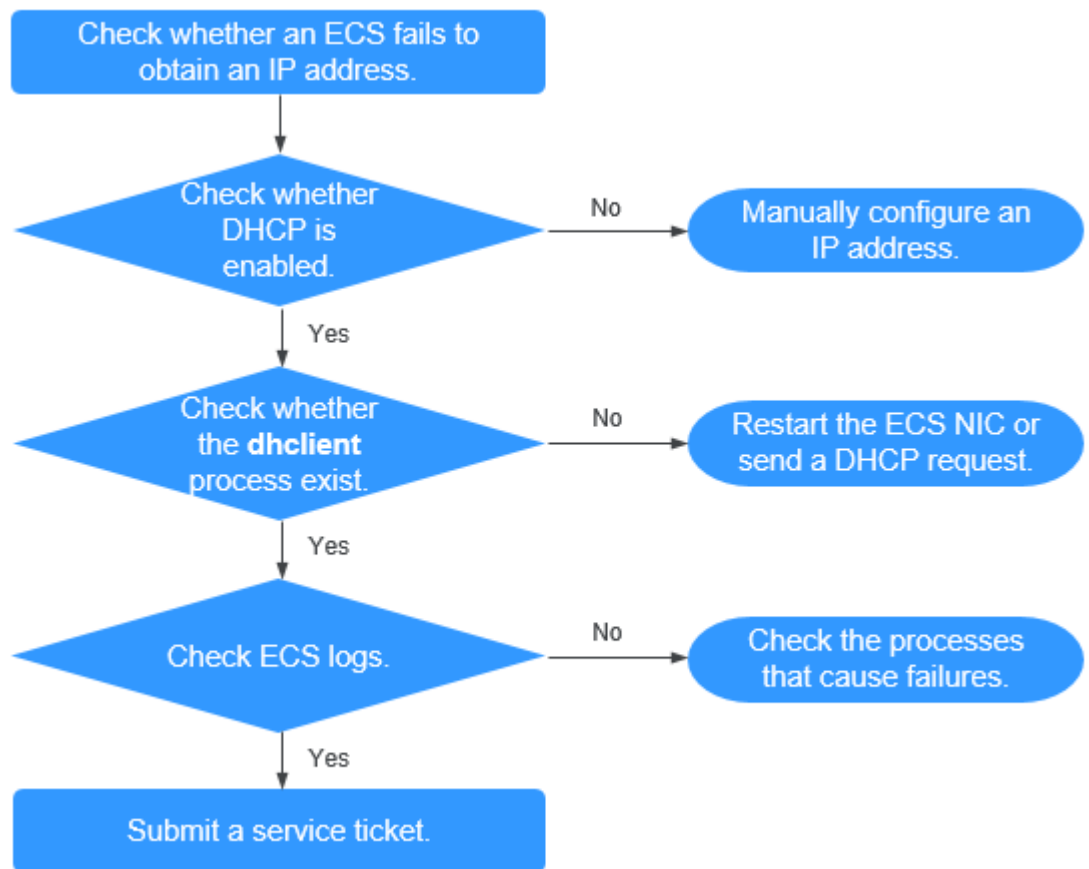
7.11 Why Does My ECS Fail to Obtain an IP Address?

Symptom

The private IP address of the ECS fails to be obtained.

Troubleshooting

Locate the fault based on the following procedure.

Figure 7-26 Troubleshooting process

1. [Checking Whether DHCP Is Enabled](#)
2. [Checking Whether the dhclient Process Exist](#)
3. [Checking ECS Logs](#)

Checking Whether DHCP Is Enabled

Check whether the DHCP function of the subnet is enabled (enabled by default).

Switch to the subnet details page. If DHCP is disabled, you must manually configure a static IP address for the ECS by referring to step [3](#).

Checking Whether the dhclient Process Exist

1. Check whether the **dhclient** process exists:
ps -ef | grep dhclient
2. If the **dhclient** process does not exist, log in to the ECS and restart the ECS NIC or send a DHCP request.
 - Linux:
Run the **dhclient ethx** command. If **dhclient** commands are supported, run the **ifdown ethx;ifup ethx** command. In the command, *ethx* indicates the ECS NIC, for example, **eth0** and **eth1**.

- Windows:
Disconnect the network connection and connect it.
- 3. If the DHCP client does not send requests for a long time, for example, the fault occurs again after the NIC restarts, you can use the following method to configure the static IP address.
 - Linux:
 - i. Run the following command to open the `/etc/sysconfig/network-scripts/ifcfg-eth0` file:
vi /etc/sysconfig/network-scripts/ifcfg-eth0
 - ii. Modify the following configuration items in the `/etc/sysconfig/network-scripts/ifcfg-eth0` file.
BOOTPROTO=static
IPADDR=192.168.1.100 #IP address
NETMASK=255.255.255.0 #Subnet mask
GATEWAY=192.168.1.1 #Gateway address
 - iii. Run the following command to restart the network service:
service network restart
 - Windows:
On the **Local Area Connection Status** tab, click **Properties**. In the displayed area, select **Internet Protocol Version 4 (TCP/IPv4)** and click **Properties**. In the displayed area, enter the IP address, subnet mask, and the default gateway address.

Checking ECS Logs

Check the ECS **messages** log in the `/var/log/messages` directory.

Search for the NIC MAC address and check whether there are any processes causing failures in obtaining IP addresses over DHCP.

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).


Provide customer service with the ECS ID, the ID of the subnet used by the ECS, and the ID of the VPC used by the ECS.

7.12 How Do I Handle a VPN or Direct Connect Connection Network Failure?

VPN Network

[Figure 7-27](#) shows your network, the customer gateway, the VPN, and the VPC.

Figure 7-27 VPN network



Customer Self-Check Guidance

1. Provide your network information.

Obtain information listed in [Table 7-3](#). This table lists example values. You can determine the actual values based on the example values. You must obtain all actual values of your project.

NOTE

You can print this table and fill in your values.

Table 7-3 Network information

Item	Description	Example	Value
VPC CIDR block	Required for customer gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	N/A	N/A

Item	Description	Example	Value
CIDR block of subnet 1 (can be the same as the VPC CIDR block)	N/A	Example: 10.0.1.0/24	N/A
ECS ID	N/A	N/A	N/A
Customer gateway type (for example, Cisco)	N/A	N/A	N/A
Public IP address used by the customer gateway	N/A	The value must be static.	N/A

2. Provide your gateway configuration information.

You can check the gateway connectivity issues based on the following steps:

You must take the IKE, IPsec, ACL rules, and route selection into consideration. You can rectify the failure in any desired sequence. However, it is recommended that you check for the failure in the following sequence: IKE, IPsec, ACL rules, and route selection.

- a. Obtain the IKE policy used by your gateway device.
- b. Obtain the IPsec policy used by your gateway device.
- c. Obtain the ACL rule used by your gateway device.
- d. Check whether your gateway device can communicate with the gateway devices on the cloud.

 **NOTE**

The commands used on different gateway devices are different. You can run the commands based on your gateway device (such as Cisco, H3C, AR, or Fortinet device) to obtain the preceding required information.

O&M Operations That Require Assistance

You must send communication requests from the ECSs to the remote device.

Method:

Log in to an ECS and ping an IP address in your on-premises data center.

7.13 Why Can My Server Be Accessed from the Internet But Cannot Access the Internet?

Symptom

The server can be accessed from, but cannot access the Internet.

Troubleshooting

Check the following possible causes.

Figure 7-28 Possible causes

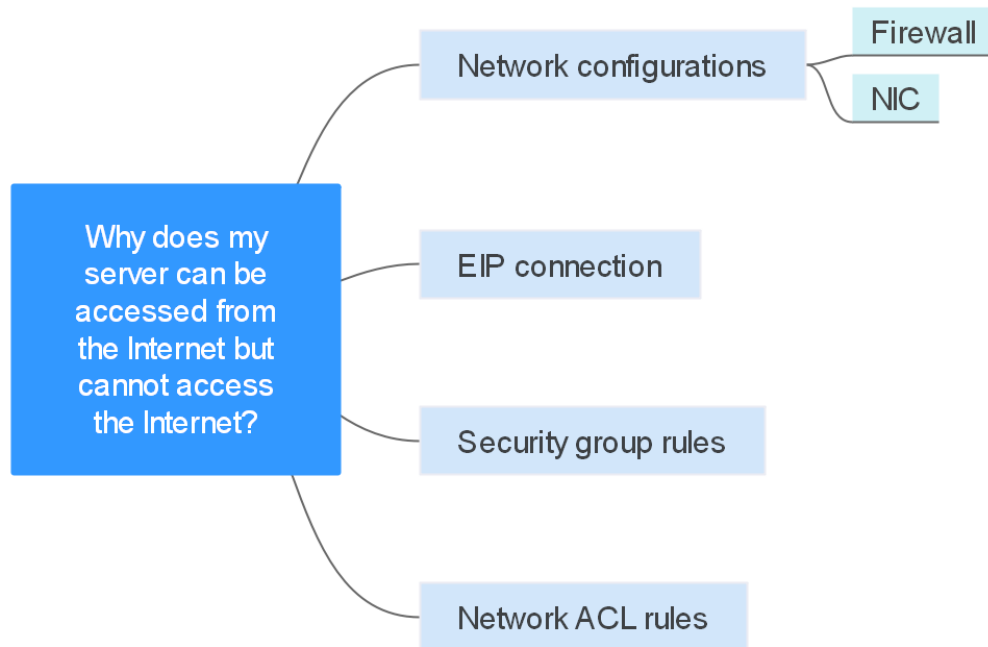


Table 7-4 Possible causes

Possible Cause	Solution
Network configurations	Check firewall and NIC configurations. See Network Configurations .
EIP connection	See Why Does Internet Access Fail Even If My ECS Is Bound with an EIP?
Security group rules	See Security Group Rules
Network ACL rules	See Network ACL Rules

Network Configurations

- Firewall
 - Disable firewall rules for the ECS and check if Internet connectivity is restored:
 - Linux ECS: [Checking the Firewall Configuration](#).
 - Windows ECS: [Checking the Firewall Configuration](#).
- NIC
 - Check the NIC and DNS configurations.
 - Linux ECS: [Checking the NIC Configuration](#).

- Windows ECS: [Checking the NIC Configuration](#).

Security Group Rules

Check if there is a security group rule for the server denying the outbound traffic.

By default, a security group allows all outbound traffic. If the outbound traffic is denied, you can [configure security group rules](#) or click **Allow Common Ports**.

Network ACL Rules

Check whether the network ACL of the subnet that the server belongs to denies the outbound traffic.

By default, a network ACL denies all outbound traffic. You need to add an outbound rule with **Action** set to **Allow** to the network ACL associated with the server.

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

7.14 Why Can't I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?


Symptom

You have enabled IPv4/IPv6 dual stack for an ECS, but the ECS cannot access websites using IPv6 addresses.

Troubleshooting

- Check whether the IPv4/IPv6 dual stack is correctly configured and whether the dual-stack NIC of the ECS has obtained an IPv6 address.
- Check whether the obtained IPv6 address of the dual-stack NIC has been added to a shared bandwidth.
- If the ECS has multiple NICs, check whether policy-based routes have been configured for these NICs.

Figure 7-29 NIC details



Solution

- When you buy an ECS, select **Automatically-assigned IPv6 address** for **Network**.

If an IPv6 address fails to be automatically assigned or the selected image does not support automatic IPv6 address allocation, manually obtain the IPv6 address by referring to [Dynamically Assigning IPv6 Addresses](#).

- By default, dynamic IPv6 address assignment is enabled for Windows public images.
- Before enabling dynamic IPv6 address assignment for a Linux public image, check whether IPv6 is supported first.

Currently, all Linux public images support IPv6. By default, dynamic IPv6 address assignment is enabled for Ubuntu 16. For other Linux public images, you need to enable this function.

- By default, IPv6 addresses can only be used for private network communication. If you want to use an IPv6 address to access the Internet or want it to be accessed by IPv6 clients on the Internet, you need to add the IPv6 address to a shared bandwidth. For details, see [Buy a Shared Bandwidth and Add the IPv6 Address to It](#).

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

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

If your ECS runs Linux, refer to [Configuring Policy-based Routes for a Linux ECS with Multiple NICs \(IPv4/IPv6\)](#).

If your ECS runs Windows, refer to [Configuring Policy-based Routes for a Windows ECS with Multiple NICs \(IPv4/IPv6\)](#).

7.15 Why Does My ECS Fail to Communicate with Other After It Has Firewall Installed?

Symptom

An ECS has a single NIC and failed to communicate with others after the ECS has a firewall installed. An example scenario is as follows:

In a VPC, there are three ECSs. Services are deployed on ECS 1 and ECS 2, and a third-party firewall is installed on ECS X. Traffic from ECS 1 and ECS 2 needs to be filtered by the firewall of ECS X.

Fault Locating

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Table 7-5 Fault locating

Possible Cause	Solution
Security group rules	See Whether Security Group Rules Are Configured
Source/destination check	See Whether Source/Destination Check Is Disabled
VPC custom routes	See Whether VPC Custom Routes Are Added

Whether Security Group Rules Are Configured



Subnets in the same VPC can communicate with each other. If your service ECS cannot communicate with the ECS that has firewall installed, check whether they are in the same security group.

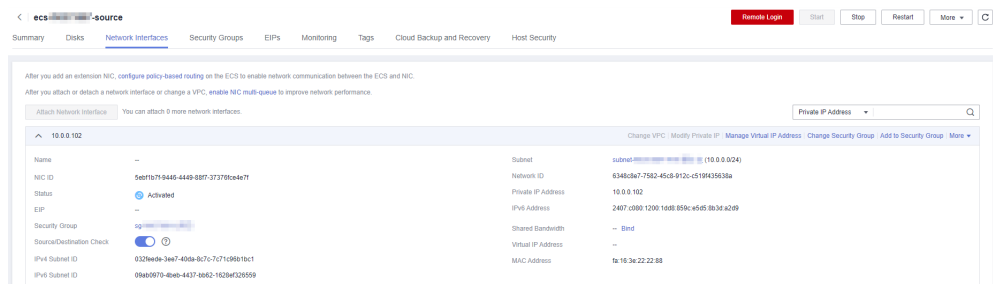
If the ECSs are in different security groups, you need to add rules to the security groups to allow access from each other.

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

Whether Source/Destination Check Is Disabled

Check whether the source/destination check function is disabled on the NIC of the ECS with firewall installed. If the function is not disabled, perform the following operations to disable it:

- Log in to the management console.
- Click  in the upper left corner and select the desired region and project.
- Click **Service List** and choose **Compute > Elastic Cloud Server**.
- In the ECS list, click the target ECS name.
The **Summary** tab page of the ECS is displayed.
- Click the **Network Interfaces** tab, click  to expand information about the primary NIC, and check whether **Source/Destination Check** is disabled.
If it is not disabled, disable it.



Whether VPC Custom Routes Are Added

Check whether the subnet route table of the service VPC has a route pointing to the ECS with firewall installed.

If there is no such a route, add a custom route with next hop set to ECS and destination set to the ECS with the firewall installed.

For details, see [Adding a Custom Route](#).

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).

8 Routing

8.1 How Do I Configure Policy-Based Routes for an ECS with Multiple NICs?

Background

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

Scenarios

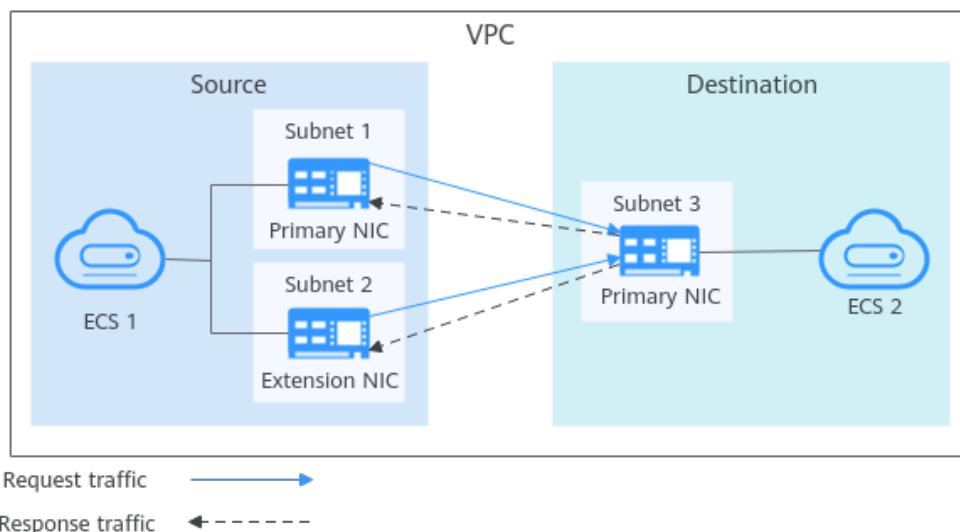
This example describes how to configure policy-based routes for an ECS with two NICs. [Figure 8-1](#) shows the networking. The details are as follows:

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

NOTICE

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

Figure 8-1 Dual-NIC ECS networking



Operation Guide

This document describes how to configure policy-based routes for Linux and Windows ECSs. For details, see [Table 8-1](#).

Table 8-1 Operation instructions

OS Type	IP Address Version	Procedure
Linux	IPv4	Take an ECS running CentOS 8.0 (64-bit) as an example. Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6)
	IPv6	
Windows	IPv4	Take an ECS running Windows Server 2012 (64-bit) as an example. Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/IPv6)
	IPv6	

8.2 Can a Route Table Span Multiple VPCs?

A route table cannot span multiple VPCs.

A route table contains a set of routes that are used to determine where network traffic is directed. A VPC has a default route table and can have multiple custom route tables.

Each subnet in a VPC must be associated with a route table. A subnet can only be associated with one route table at a time, but you can associate multiple subnets in a VPC with the same route table.

8.3 Are There Any Restrictions on Using a Route Table?

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.

9 Security

9.1 Does a Modified Security Group Rule or a Network ACL Rule Take Effect Immediately for Existing Connections?

- Security groups use connection tracking to track traffic to and from instances. If an inbound rule is modified, the modified rule immediately takes effect for the existing traffic. Changes to outbound security group rules do not affect existing persistent connections and take effect only for new connections.
If you add, modify, or delete a security group rule, or add or remove an instance to or from a security group, the inbound connections of all instances in the security group will be automatically cleared.
 - The existing inbound persistent connections will be disconnected. All the new connections will match the new rules.
 - The existing outbound persistent connections will not be disconnected, and the original rule will still be applied. All the new connections will match the new rules.
- Network ACLs use connection tracking to track traffic to and from instances. Changes to inbound and outbound rules do not take effect immediately for the existing traffic.
If you add, modify, or delete a network ACL rule, or associate or disassociate a subnet with or from a network ACL, all the inbound and outbound persistent connections will not be disconnected. New rules will only be applied for the new connections.

NOTICE

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

9.2 Why Is Outbound Access on TCP Port 25 Blocked?

Symptom

You cannot access an external address on TCP port 25. For example, running the **Telnet smtp.***.com 25** command fails.

Cause

For security reasons, TCP port 25 is disabled in the outbound direction by default.

You do not need to enable TCP port 25, unless you want to deploy an email service on the cloud.

 **NOTE**

Outbound traffic over TCP port 25 is only blocked in CN North-Beijing1.

Solution

Use port 465 supported by the third-party email service provider.

9.3 How Do I Know the Instances Associated with a Security Group?

When you create an instance, such as ECS, cloud container, or database, you need to add the instance to a security group. To delete a security group, you must remove all instances from the security group first.

You can perform the following operations to view the instances associated with a security group:

1. In the security group list, locate the row that contains the target security group and click **Manage Instance** in the **Operation** column.

On the **Associated Instances** tab, you can view instances associated with the security group, such as servers and extension NICs.

If there is no instance associated with the security group on the **Associated Instances** tab, but the system still displays a message indicating that the security group has instances associated, perform the following operations:

2. In the upper right corner of the console, choose **Resources > My Resources** and check whether there are resources listed in **Table 9-1** in the same region as the security group.

The table lists some common resources. If you have other resources, check them.

A security group is associated with various instances through ports. You can call the APIs to **query the ports** and then use the ID of the security group to search for the target ports. Then you will get the required information, including **name** (port name), **device_id** (ID of the device to which the port belongs), **device_owner** (device owner), **instance_id** (ID of the instance), and **instance_type** (type of the instance to which the port belongs). The port information helps you check the resources associated with the security group.

If the security group still cannot be deleted after you disassociate all the resources from it, **submit a service ticket**.


Table 9-1 Check list

Product Category	Product/Instance
Databases	GaussDB
	Relational Database Service (RDS)
	Document Database Service (DDS)
	GaussDB NoSQL
	Distributed Database Middleware (DDM)
Applications	Distributed Cache Service (DCS) <ul style="list-style-type: none"> ● Redis instance ● Memcached instance
	Distributed Message Service (DMS) <ul style="list-style-type: none"> ● Kafka instance ● RabbitMQ instance
	API Gateway
EI	MapReduce Service (MRS)
	Data Warehouse Service (DWS)
	Cloud Search Service (CSS)

9.4 Why Can't I Delete a Security Group?

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

Figure 9-1 Default security group



Name	Instances	Rules	Created	Tags	Actions
Sys-WebServer 13e7c11b-959c-464e-875c-e0385aa466da	13	1	Jul 09, 2022 15:29:49 GMT+08:00	default	Manage Rules Manage Instances More
default 3b44184a-8784-4a86-82f7-75a91e0233ae	6	0	Jul 09, 2022 15:29:48 GMT+08:00	default	Manage Rules Manage Instances Clone

- If you want to delete a security group that is associated with instances, such as cloud servers, containers, and databases, you need to remove the instances from the security group first.

If you want to know the instances associated with a security group, refer to [How Do I Know the Instances Associated with a Security Group?](#)

- A security group cannot be deleted if it is used as the source or destination of a rule in another security group.

You need to delete or modify the rule first and delete the security group.

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

9.5 Can I Change the Security Group of an ECS?

Yes.

An ECS must be associated with a security group. You can change the security group of an ECS at any time.

You can log in to the management console to change the security group on the ECS details page.

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

9.6 How Do I Configure a Security Group for Multi-Channel Protocols?

ECS Configuration

The TFTP daemon determines whether a configuration file specifies the port range. If you use a TFTP configuration file that allows the data channel ports to be configurable, it is a good practice to configure a small range of ports that are not listened on.

Security Group Configuration

You can configure port 69 and configure data channel ports used by TFTP for the security group. In RFC1350, the TFTP protocol specifies that ports available to data channels range from 0 to 65535. However, not all these ports are used by the TFTP daemon processes of different applications. You can configure a smaller range of ports for the TFTP daemon.

The following figure provides an example of the security group rule configuration if the ports used by data channels range from 60001 to 60100.

Figure 9-2 Security group rule

Priority	Action	Type	Protocol & Port	Source
100	Allow	IPv4	UDP: 60001-60100	0.0.0.0/0

9.7 Why Are Some Ports of ECSs Inaccessible?

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.

Table 9-2 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 9-2 High-risk ports

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

9.8 Why Is Access from a Specific IP Address Still Allowed After a Network ACL Rule That Denies the Access from the IP Address Has Been Added?

Network ACL rules have priorities. A smaller priority value represents a higher priority. Each network ACL includes a default rule whose priority value is an asterisk (*). Default rules have the lowest priority.

If rules conflict, the rule with the highest priority takes effect.

If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule. For example, if the priority of rule A is 1 but you need rule B to take priority over rule A, insert rule B before rule A. Then, rule B will have a priority of 1 and rule A will be 2. Similarly, if rule B is less important than rule A, insert rule B after rule A.

When a rule that denies access from a specified IP address is added, insert the rules that allow access from all IP addresses at the end. Then, the rule that denies access from the specified IP address will take priority over the other rules and will be effective. For details, see [Adding a Network ACL Rule \(Custom Rule Number\)](#).

9.9 Why Are My Security Group Rules Not Working?

Symptom

After a security group rule is configured for an instance (such as an ECS), the rule does not work. For example, although a security group rule is added to allow a specific IP address to access instances in the security group, the access still fails.

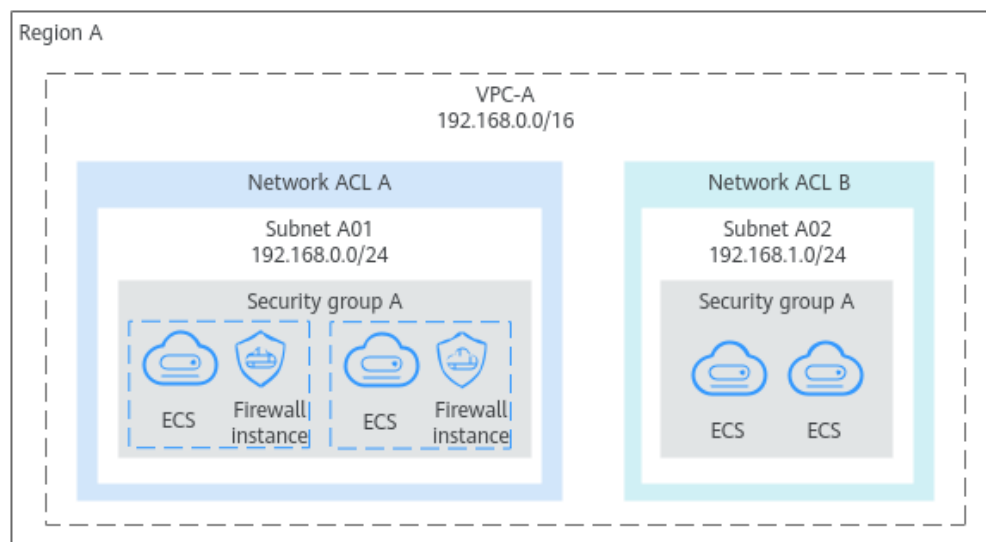
Background

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

- A security group protects the instances in it.
- A network ACL protects the entire subnet. After a subnet is associated with a network ACL, all instances in the subnet are protected by the network ACL.

In addition to the access control provided by the VPC service, you can configure the cloud firewalls to further improve the instance security. [Figure 9-3](#) shows how they work together.

Figure 9-3 Controlling access to a VPC



If your security group rules are not working, they may be configured wrong, or there may be conflicting rules.

Troubleshooting

The issues here are described in order of how likely they are to occur.

Troubleshoot the issue by ruling out the causes described here, one by one.

Figure 9-4 Troubleshooting

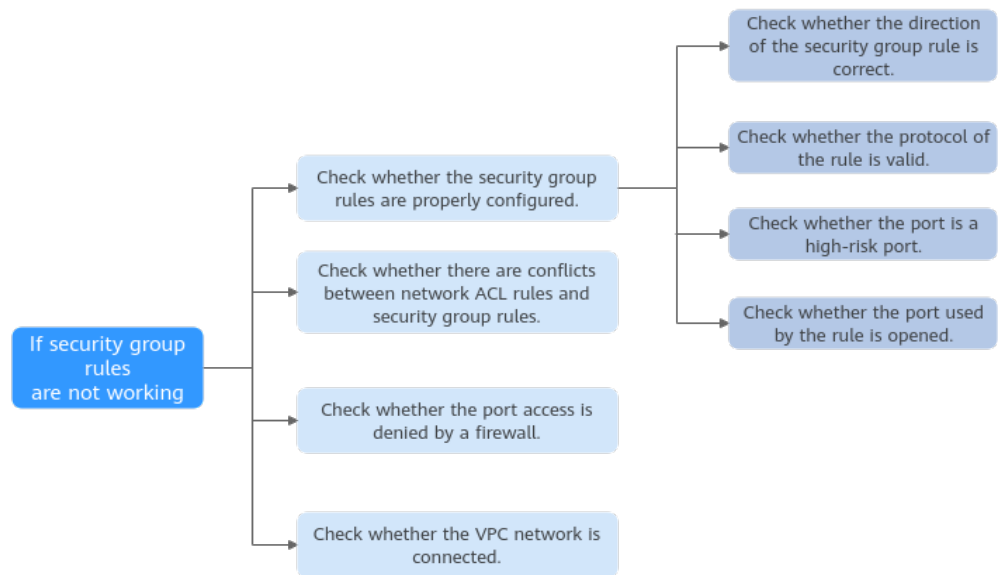


Table 9-3 Troubleshooting

Possible Cause	Solution
Improper security group rule configurations	Incorrect Security Group Rule Configurations
Conflicts between network ACL rules and security group rules	Conflicts Between Network ACL Rules and Security Group Rules
Port access denied by a firewall	Port Access Denied by a Firewall
Network disconnected	Network Disconnectivity Between VPCs

Incorrect Security Group Rule Configurations

If security group rules are improperly configured, instances cannot be protected. Check the security group rules based on the following causes:

- The wrong direction is configured.

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 the incoming traffic to the instances in a security group.
 - Outbound rules control traffic from the instances in a security group for accessing external networks.
- The wrong protocol is configured.

The protocols that can be used to filter traffic can be TCP, UDP, ICMP, or GRE. Select a valid protocol for the security group rule.

3. The configured ports are high-risk ports that are blocked by carriers. These ports cannot be accessed in restricted regions. In this case, change the ports to common ports.

For details about common ports and risky ports, see [Common Ports Used by ECSs](#).

4. The port required by the service is not opened.

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

Check whether the ECS port is opened and whether the configuration takes effect by referring to section "Verifying Security Group Rules" in [Adding a Security Group Rule](#).

After causes are identified, you can [add a security group rule](#) or [modify a security group rule](#) to select the correct direction, protocol, and open the ports.

Conflicts Between Network ACL Rules and Security Group Rules

Security groups protect instances (such as ECSs), while network ACLs protect subnets. If a network ACL rule conflicts with a security group rule, the network ACL rule takes precedence over the security group rule. As a result, the security group rule may not be applied.

For example, if the inbound rule of your security group allows access over port 80 but the network ACL rule denies access over the port, the traffic preferentially matches the network ACL rule.

You can [add a network ACL rule](#) or [modify a network ACL rule](#) to allow traffic from the corresponding protocol port.

Port Access Denied by a Firewall

Both security groups and firewalls protect instances (such as ECSs). Although access to a port is allowed by a security group rule, the ECS firewall may deny the access to the port. In this case, you need to disable the firewall or configure an exception port on the firewall.

For details, see [Disabling a Windows ECS Firewall and Adding a Port Exception on a Windows ECS Firewall](#) or [Disabling a Linux ECS Firewall and Adding a Port Exception on a Linux ECS Firewall](#).

Network Disconnectivity Between VPCs

A security group works only when the network communication is normal. If instances are associated with the same security group but in different VPCs, the instances cannot communicate with each other.

In this case, you can use a VPC peering connection to connect the VPCs so that security groups can control traffic from and to the instances in different VPCs. For details about VPC connectivity, see [Application Scenarios](#).

Submitting a Service Ticket

If the problem persists, [submit a service ticket](#).