Virtual Private Cloud

FAQs

Issue 30

Date 2021-03-24





Copyright © Huawei Technologies Co., Ltd. 2021. All rights reserved.

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

Trademarks and Permissions

HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

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

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

Contents

1 General Questions	1
1.1 What Is a Quota?	1
2 Billing and Payments	3
2.1 Will I Be Charged for Using the VPC Service?	
2.2 How Is an EIP Billed?	3
2.3 How Do I Change the Billing Mode?	4
2.4 How Do I Change the Bandwidth Billing Option from Bandwidth to Traffic or from Traffic to Bandwidth ?	
3 VPC and Subnet	7
3.1 What Is Virtual Private Cloud?	7
3.2 Which CIDR Blocks Are Available for the VPC Service?	9
3.3 How Many VPCs Can I Create?	9
3.4 Can Subnets Communicate with Each Other?	9
3.5 What Subnet CIDR Blocks Are Available?	9
3.6 Can I Modify the CIDR Block of a Subnet?	9
3.7 How Many Subnets Can I Create?	10
3.8 How Can I Delete a Subnet That Is Being Used by Other Resources?	10
3.9 How Do I Switch to a Private DNS Server?	10
4 EIP	12
4.1 How Do I Assign or Retrieve a Specific EIP?	12
4.2 What Are the Differences Between EIP, Private IP Address, Floating IP Address, and Virtual IP	
4.3 How Do I Access the Internet Using an EIP Bound to an Extension NIC?	13
4.4 What Are the Differences Between the Primary and Extension NICs of ECSs?	14
4.5 Can an EIP That Uses Dedicated Bandwidth Be Changed to Use Shared Bandwidth?	14
4.6 Can I Bind an EIP to Multiple ECSs?	14
4.7 How Do I Access an ECS from the Internet After an EIP Is Bound to the ECS?	14
4.8 What Is the EIP Assignment Policy?	
4.9 Can I Bind an EIP to an ECS, to Another ECS?	15
4.10 Does an EIP Change Over Time?	
4.11 Can I Assign a Specific EIP?	16
4.12 How Do I Query the Region of My EIPs?	
4.13 Can a Bandwidth Be Used by Multiple Accounts?	16

4.14 How Do I Change an EIP for an Instance?	16
4.15 Can I Bind an EIP to a Cloud Resource in Another Region?	19
5 Bandwidth	20
5.1 What Are Inbound Bandwidth and Outbound Bandwidth?	
5.2 How Do I Know If My Used Bandwidth Exceeds the Limit?	
5.3 What Is the Bandwidth Size Range?	
5.4 What Bandwidth Types Are Available?	
5.5 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth? Can a De Bandwidth Be Changed to a Shared Bandwidth or the Other Way Around?	
5.6 How Do I Buy a Shared Bandwidth?	23
5.7 Is There a Limit to the Number of EIPs That Can Be Added to Each Shared Bandwidth?	24
5.8 Can I Increase My Bandwidth Billed on Yearly/Monthly Basis and Then Decrease It?	24
5.9 What Is the Relationship Between Bandwidth and Upload/Download Rate?	24
5.10 What Are the Differences Between Static BGP and Dynamic BGP?	24
6 Connectivity	26
6.1 Does a VPN Allow Communication Between Two VPCs?	
6.2 Why Is Internet or Internal Domain Names in the Cloud Inaccessible Through Domain Names \ My ECS Has Multiple NICs?	Vhen
6.3 What Are the Constraints Related to VPC Peering?	27
6.4 Why Does Communication Fail Between VPCs That Are Connected by a VPC Peering Connectio	n?28
6.5 How Many VPC Peering Connections Can I Create?	32
6.6 What Are the Priorities of the Custom Route and EIP If Both Are Configured for an ECS to Enab ECS to Access the Internet?	
6.7 Why Does Intermittent Interruption Occur When a Local Host Accesses a Website Built on an E	
6.8 Why Do ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communi	cation?
6.9 Why Does Communication Fail Between Two ECSs in the Same VPC or Packet Loss Occur Wher Communicate?	n They
6.10 Why Cannot the Virtual IP Address Be Pinged After It Is Bound to an ECS NIC?	
6.11 Why Does My ECS Fail to Use Cloud-init?	
6.12 Why Does Internet Access Fail Even If My ECS Is Bound with an EIP?	46
6.13 How Do I Handle the IB Network Failure?	50
6.14 Why Does My ECS Fail to Communicate at a Layer 2 or Layer 3 Network?	52
6.15 How Do I Handle the BMS Network Failure?	54
6.16 Why Does My ECS Fail to Obtain an IP Address?	55
6.17 How Do I Handle the VPN or Direct Connect Connection Network Failure?	57
6.18 Why Does My Server Can Be Accessed from the Internet But Cannot Access the Internet?	59
6.19 Can I Use a VPC Peering Connection to Connect VPCs in Different Regions?	61
6.20 Will I Be Billed for Using a VPC Peering Connection?	61
6.21 What Switches Can Connect to a L2CG on HUAWEI CLOUD?	62
6.22 Why Is the Layer 2 Connection in the Not Connected State Even After Its Configuration Is Con	•

6.23 Why Is Communication Between the Cloud and On-premises Servers Unavailable Even When th Layer 2 Connection Status Is Connected?	
6.24 Why Can't I Access Websites Using IPv6 Addresses After IPv4/IPv6 Dual Stack Is Configured?	62
7 Routing	64
7.1 How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?	64
7.2 Why Can't I Ping an ECS with Two NICs Configured?	68
7.3 Can a Route Table Span Multiple VPCs?	69
7.4 How Many Routes Can a Route Table Contain?	69
7.5 Are There Any Restrictions on Using a Route Table?	69
7.6 Will a Route Table Be Billed?	70
7.7 Do the Same Routing Priorities Apply to Direct Connect Connections and Custom Routes in the S VPC?	
7.8 Are There Different Routing Priorities of the VPN and Custom Routes in the Same VPC?	70
8 Security	71
8.1 Are the Security Group Rules Considered the Same If All Parameters Except Their Description Are Same?	
8.2 What Are the Requirements for Deleting a Security Group?	71
8.3 Why Is Outbound Access Through TCP Port 25 Restricted?	72
8.4 Can I Change the Security Group of an ECS?	73
8.5 How Many Security Groups Can I Have?	73
8.6 Will a Security Group Be Billed?	73
8.7 How Do I Configure a Security Group for Multi-Channel Protocols?	73
8.8 How Many Network ACLs Can I Create?	73
8.9 Does a Security Group Rule or a Network ACL Rule Immediately Take Effect for Its Original Traffi After It Is Modified?	
8.10 Why Are Some Ports in the Public Cloud System Inaccessible?	74
8.11 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?	e
8.12 What Do My Security Group Rules Not Take Effect?	75

1 General Questions

1.1 What Is a Quota?

What Is a Quota?

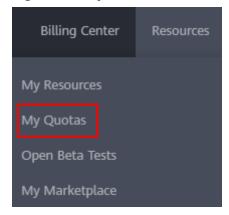
A quota limits the quantity of a resource available to users, thereby preventing spikes in the usage of the resource. For example, a VPC quota limits the number of VPCs that can be created.

You can also request for an increase in quota if an existing quota cannot meet your service requirements.

How Do I View My Quotas?

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- In the upper right corner of the page, choose Resources > My Quotas.
 The Service Quota page is displayed.

Figure 1-1 My Quotas



4. View the used and total quota of each type of resources on the displayed page.

If a quota cannot meet service requirements, apply for a higher quota.

How Do I Apply for a Higher Quota?

- 1. Log in to the management console.
- In the upper right corner of the page, choose Resources > My Quotas.
 The Service Quota page is displayed.

Figure 1-2 My Quotas



- 3. Click Increase Quota.
- 4. On the **Create Service Ticket** page, configure parameters as required. In **Problem Description** area, fill in the content and reason for adjustment.
- After all necessary parameters are configured, select I have read and agree to the Tenant Authorization Letter and Privacy Statement and click Submit.

2 Billing and Payments

2.1 Will I Be Charged for Using the VPC Service?

The VPC service is free of charge. However, EIP and bandwidth used together with a VPC will be billed based on standard pricing.

2.2 How Is an EIP Billed?

EIPs can be billed on a yearly/monthly or pay-per-use basis.

Table 2-1 EIP billing details

Billing Mode	Billed By	EIP Retention Fee	Bandwidth Price	Public Network Traffic Price
Yearly/ Monthly	Bandwidth	-	Included	Not included
Pay-per-use	Bandwidth	EIP retention fee is not included if the EIP is bound to an ECS, BMS, or load balancer. EIP retention fee is included if the EIP is unbound but not released.	Included	Not included
	Traffic		Not included	Included

- "Not included" indicates that the fee will not be included in the bill. "Included" indicates that the fee will be included in the bill.
- For details about the EIP pricing, see Product Pricing Details.

2.3 How Do I Change the Billing Mode?

Changing the Billing Mode from Pay-per-Use to Yearly/Monthly

You can change the billing mode of pay-per-use EIPs and shared bandwidth billed by bandwidth to yearly/monthly. After the change is successful, the new billing mode will take effect immediately.

You can change the billing mode on the EIP console. Do as follows to change the billing mode of an EIP from pay-per-use to yearly/monthly.

■ NOTE

The billing mode of an EIP that is billed by traffic on a pay-per-use basis cannot be directly changed to yearly/monthly. Change the EIP to be billed by bandwidth and then change its billing mode to yearly/monthly.

- 1. Log in to the management console.
- 2. Under Network, click Elastic IP.
- 3. On the displayed page, search for the pay-per-use EIP whose billing mode is to be changed.
- 4. Locate the row that contains the target EIP and click **Change Billing Mode** in the **Operation** column.

Figure 2-1 Changing the billing mode on the EIP console



- 5. Click Yes.
- 6. Set specifications.

Change Subscription (Investigation Service Supplemental Service Supplementary Suppleme

Figure 2-2 Setting specifications

Click Submit and Pay.

You can also select multiple EIPs and click **Change Billing Mode** above the EIP list to change the billing mode of all selected EIPs at the same time.

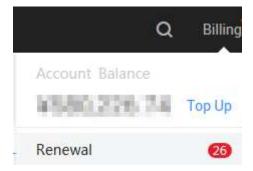
Changing the Billing Mode from Yearly/Monthly to Pay-per-Use

The billing mode of yearly/monthly EIPs and shared bandwidths can be changed to pay-per-use. The new billing mode takes effect only after the validity period of the EIPs or bandwidths expires.

The billing mode of an EIP can be changed from yearly/monthly to pay-per-use in the billing center. Do as follows to change the billing mode of an EIP from yearly/monthly to pay-per-use:

- 1. Log in to the management console.
- 2. Choose **Billing** > **Renewal**.

Figure 2-3 Renewal



- 3. In the search box on the right, search for the EIP whose billing mode you want to change.
- 4. Locate the row that contains the target EIP and click **Change to Pay-per-Use After Expiration** in the **Operation** column.

Figure 2-4 Changing the billing mode to pay-per-use



5. In the page that is displayed, click the **Change to Pay-per-Use** button.

Figure 2-5 Confirming the change



■ NOTE

The EIP remains the same after the billing mode is changed.

2.4 How Do I Change the Bandwidth Billing Option from Bandwidth to Traffic or from Traffic to Bandwidth?

- The billing option can be changed only when the billing mode is **Pay-per-use**. For details, see **Modifying EIP Bandwidth**.
- A yearly/monthly resource can only be billed by bandwidth.

3 VPC and Subnet

3.1 What Is Virtual Private Cloud?

The Virtual Private Cloud (VPC) service enables you to provision logically isolated, configurable, and manageable virtual networks for cloud servers, cloud containers, and cloud databases, improving cloud service security and simplifying network deployment.

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

Product Architecture

The product architecture consists of the VPC components, security features, and VPC connectivity options.

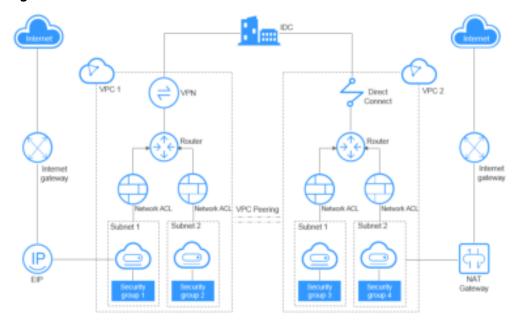


Figure 3-1 Architecture

VPC Components

Each VPC consists of a private CIDR block, route tables, and at least one subnet.

- Private CIDR block: When creating a VPC, you need to specify the private CIDR block used by the VPC. The VPC service supports the following CIDR blocks: 10.0.0.0 10.255.255.255, 172.16.0.0 172.31.255.255, and 192.168.0.0 192.168.255.255
- Subnet: Cloud resources, such as ECSs and databases, must be deployed in subnets. After you create a VPC, divide the VPC into one or more subnets.
 Each subnet must be within the VPC. For more information, see Subnet.
- Route table: When you create a VPC, the system automatically generates a
 default route table. The route table ensures that all subnets in the VPC can
 communicate with each other. If the routes in the default route table cannot
 meet application requirements (for example, an ECS without an elastic IP
 address (EIP) bound needs to access the Internet), you can create a custom
 route table. For more information, see Example Custom Route in a VPC and
 Example Custom Route Outside a VPC.

Security Features

Security groups and network ACLs ensure the security of cloud resources deployed in a VPC. A security group acts as a virtual firewall to provide access rules for instances that have the same security requirements and are mutually trusted in a VPC. For more information, see **Security Group Overview**. A network ACL can be associated with subnets that have the same access control requirements. You can add inbound and outbound rules to precisely control inbound and outbound traffic at the subnet level. For more information, see **Network ACL Overview**.

VPC Connectivity

HUAWEI CLOUD provides multiple VPC connectivity options to meet diverse requirements. For details, see **Application Scenarios**.

- VPC Peering allows two VPCs in the same region to communicate with each other using private IP addresses.
- Elastic IP or NAT Gateway allows ECSs in a VPC to communicate with the Internet.
- Virtual Private Network (VPN), Cloud Connect, or Direct Connect can connect a VPC to your data center.

3.2 Which CIDR Blocks Are Available for the VPC Service?

The VPC service supports the following CIDR blocks:

- 10.0.0.0/8-24
- 172.16.0.0/12-24
- 192.168.0.0/16-24

3.3 How Many VPCs Can I Create?

By default, you can create a maximum of five VPCs in your account. If the number of VPCs cannot meet your service requirements, **submit a service ticket** to request a quota increase.

3.4 Can Subnets Communicate with Each Other?

Subnets in the same VPC can communicate with each other while subnets in different VPCs cannot communicate with each other by default. However, you can create VPC peering connections to enable subnets in different VPCs to communicate with each other.

If a subnet is associated with a network ACL, configure network ACL rules to allow communication between subnets.

3.5 What Subnet CIDR Blocks Are Available?

A subnet CIDR block must be included in its VPC CIDR block. Supported VPC CIDR blocks are 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24. The allowed block size of a subnet is between the netmask of its VPC CIDR block and the /28 netmask.

3.6 Can I Modify the CIDR Block of a Subnet?

You can modify the CIDR block of a subnet only when you are creating the subnet. After the subnet is created, you cannot modify its CIDR block.

3.7 How Many Subnets Can I Create?

By default, you can create a maximum of 100 subnets in your cloud account. If the number of subnets cannot meet your service requirements, **submit a service ticket** to request a quota increase.

3.8 How Can I Delete a Subnet That Is Being Used by Other Resources?

The VPC service allows you to create private, isolated virtual networks. In a VPC, you can manage private IP address ranges, subnets, and gateways. ECSs, BMSs, databases, and some other applications can use subnets created in VPCs.

A subnet cannot be deleted if it is being used by other resources. You must delete all resources in the subnet before you can delete the subnet.

You can view all resources of your account on the console homepage and check the resources that are in the subnet you want to delete.

The resources may include:

- ECS
- CCI instance
- Load balancer
- VPN
- Private IP address
- Custom route
- NAT gateway
- VPC endpoint and VPC endpoint service

If you cannot delete a subnet even after deleting all the resources it contains, submit a service ticket.

3.9 How Do I Switch to a Private DNS Server?

ECSs use private DNS servers for domain name resolution in VPCs. ECSs in a VPC can access the Internet using public domain names and other cloud services like OBS and SMN through private DNS servers, with no need to connect to the Internet.

For VPCs created earlier before private domain names are available, a public DNS server (114.114.114.114) is configured. 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 Server Addresses Provided by the DNS Service?

Perform the operations provided in this section to change the public DNS servers to private DNS servers.

Checking the DNS Server Addresses of an ECS

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

Changing the DNS Servers for a VPC Subnet

If the ECS uses default public DNS servers, change them to private DNS servers provided by the DNS service.

- In the Gateway and DNS Information area, click next to DNS Server Address.
- 2. Change the DNS server addresses to private DNS server addresses. For example, in the CN North-Beijing1 region, change the DNS server addresses of a VPC subnet to **100.125.1.250** and **100.125.21.250**.

Updating the DNS Server Addresses for the ECS

New DNS server addresses will not take effect immediately on the ECS.

The DNS server addresses needs to be updated first. There are two ways to do this:

• 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, which takes 24 hours by default. After the lease time expires, the DHCP server allocates another IP address and updates the DNS server addresses to the ECS.

Manually change the DNS configurations on the ECS.
 If DHCP is disabled on the ECS, manually update DNS configurations.
 For example, if the ECS is running Linux, change the DNS configurations by editing the /etc/resolv.conf file.

4 EIP

4.1 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. When assigning an EIP, set the value of **ip_address** to the IP address 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.

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

An EIP is an IP address that can be accessed over the Internet. Each EIP can be used by only one ECS at a time.

A private IP address is used on the private network of the public cloud for private communications. It cannot be reached from the Internet.

A floating IP address is similar to an EIP. They are both public IP addresses that are used to connect to the Internet, but a floating IP address API cannot be used to configure bandwidth parameters. For details, see **Floating IP Address**.

A virtual IP address can be shared among multiple ECSs. A virtual IP address is used for active/standby switchover of ECSs for higher availability. If the active ECS becomes faulty and cannot provide services, the virtual IP address is dynamically re-assigned to the standby ECS so services can continue uninterrupted. For details, see Virtual IP Address Overview.

4.3 How Do I Access the Internet Using an EIP Bound to an Extension NIC?

1. After an EIP is bound to an extension NIC, log in to the ECS and run the **route** command to query the route.

You can run **route --help** to learn more about the **route** command.

Figure 4-1 Viewing route information

```
[root@ecs-b926 ~1# route -n
Kernel IP routing table
                Gateway
Destination
                                                  Flags Metric Ref
                                                                       Use Iface
                                 Genmask
                192.168.11.1
                                 0.0.0.0
0.0.0.0
                                                  UG
                                                                         0 eth0
                                                        1002
169.254.0.0
                0.0.0.0
                                 255.255.0.0
                                                  U
                                                                0
                                                                         0 eth0
169.254.0.0
                                 255.255.0.0
                                                        1003
                0.0.0.0
                                                  ш
                                                                Я
                                                                         0 eth1
169.254.169.254 192.168.11.1
                                 255.255.255.255 UGH
                                                        0
                                                                0
                                                                         0 eth0
192.168.11.0
                0.0.0.0
                                 255.255.255.0
                                                  Ш
                                                        ø
                                                                ø
                                                                         0 eth0
192.168.17.0
                0.0.0.0
                                 255.255.255.0
                                                                0
                                                                         0 eth1
Croot@ecs-b926
                1]#
```

2. Run the **ifconfig** command to view NIC information.

Figure 4-2 Viewing NIC information

```
root@ecs-b926 ~1# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 192.168.11.42 netmask 255.255.255.0 broadcast 192.168.11.255
         inet6 fe80::f816:3eff:fef7:1c44 prefixien 64 scopeid 0x20<link> ether fa:16:3e:f7:1c:44 txqueuelen 1000 (Ethernet)
         RX packets 127 bytes 21633 (21.1 KiB)
         RX errors 0 dropped 0 overruns 0 frame 0
         TX packets 258 bytes 22412 (21.8 KiB)
         TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
         inet 192.168.17.191 netmask 255.255.255.0 broadcast 192.168.17.255 inet6 fe88::f816:3eff:fe1c:b57f prefixlen 64 scopeid 9x20link> ether fa:16:3e:1c:b5:7f txqueuelen 1000 (Ethernet)
         RX packets 11 bytes 1283 (1.2 KiB)
         RX errors 0 dropped 0 overruns 0 frame 0 TX packets 12 bytes 1388 (1.3 KiB)
         TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
         inet 127.0.0.1 netmask 255.0.0.0
         inet6 ::1 prefixlen 128 scopeid 0x10<host>
         loop txqueuelen 1 (Local Loopback)
         RX packets 51 bytes 12018 (11.7 KiB)
         RX errors 0 dropped 0 overruns 0 frame 0
         TX packets 51 bytes 12018 (11.7 KiB)
         TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

- 3. Enable access to the Internet through the extension NIC by default.
 - a. Run the following command to delete the default route of the primary NIC:

route del 0.0.0.0 192.168.11.1 dev eth0

This operation will interrupt ECS communication. It is recommended that you perform the configuration by following step 4.

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

route add default gw 192.168.17.1

4. Configure Internet access from the extension NIC based on your destination address.

Run the following command to configure access to a specified CIDR block (for example, xx.xx.0.0/16) through the extension NIC:

You can configure the CIDR block as required.

route add -net xx.xx.0.0 netmask 255.255.0.0 gw 192.168.17.1

4.4 What Are the Differences Between the Primary and Extension NICs of ECSs?

The differences are as follows:

- Generally, the OS default routes preferentially use the primary NICs. If the OS
 default routes use the extension NICs, network communication will be
 interrupted. Then, you can check the route configuration to rectify the
 network communication error.
- Primary NICs can communicate with the public service zone (zone where PaaS and DNS services are deployed). Extension NICs cannot communicate this zone.

4.5 Can an EIP That Uses Dedicated Bandwidth Be Changed to Use Shared Bandwidth?

No. An EIP that uses a dedicated bandwidth cannot be changed to use a shared bandwidth.

In addition, an EIP that uses a shared bandwidth cannot be changed to use a dedicated bandwidth.

4.6 Can I Bind an EIP to Multiple ECSs?

Each EIP can be bound to only one ECS at a time.

Multiple ECSs cannot share the same EIP. An ECS and its bound EIP must be in the same region. If you want multiple ECSs in the same VPC to share an EIP, you have to use a NAT gateway. For more information, see **NAT Gateway User Guide**.

4.7 How Do I Access an ECS from the Internet After an EIP Is Bound to the ECS?

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 the Linux OS and TCP traffic from

port 3389 through RDP to the Windows OS). 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 the ECS needs to be accessible over the Internet and the IP address used to
 access the ECS over the Internet has been configured on the ECS, or the ECS
 does not need to be accessible over the Internet, set Source to the IP address
 range containing the IP address that is allowed to access the ECS over the
 Internet.
- If the ECS needs to be accessible over the Internet and the IP address used to
 access the ECS over the Internet has not been configured on the ECS, it is
 recommended that you retain the default setting 0.0.0.0/0 for Source, and
 then set Port Range to improve network security.
- Allocate ECSs that have different Internet access policies to different security groups.

◯ NOTE

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

4.8 What Is the EIP Assignment Policy?

By default, EIPs are assigned randomly.

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

If you want an EIP that you released 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 want.

4.9 Can I Bind an EIP to an ECS, to Another ECS?

Yes.

Unbind the EIP from the current ECS. For details, see **Unbinding or Releasing an EIP**.

Then, bind the EIP to another ECS. For details, see **Binding an EIP to Cloud Resources**.

Another related operation is to change the EIP associated with an ECS.

For details, see Changing an EIP.

4.10 Does an EIP Change Over Time?

EIPs will not be changed after they are assigned.

Stopping and starting an ECS does not affect its EIP.

An EIP will be released if it expires or if the EIP owner's account is in arrears.

4.11 Can I Assign a Specific EIP?

By default, EIPs are assigned randomly. If you have released EIPs before, the system preferentially assigns an EIP from what you released.

Certain APIs need to be called to assign specific EIPs. For details, see **Assigning an EIP**.

4.12 How Do I Query the Region of My EIPs?

You can visit https://en.ipip.net/?origin=CN 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/?origin=CN, use the region identified using https://en.ipip.net/?origin=CN.
- If the region identified using https://en.ipip.net/?origin=CN is different from the one you selected when purchasing the EIP, use the region you had selected during EIP purchase.
- 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.

4.13 Can a Bandwidth Be Used by Multiple Accounts?

A bandwidth cannot be shared between different accounts. Each account can use and manage only its own EIP bandwidths.

4.14 How Do I Change an EIP for an Instance?

Scenario 1: Changing an EIP for an ECS

- 1. Unbind an EIP.
 - a. Log in to the management console.
 - b. On the console homepage, under **Network**, click **Elastic IP**.
 - c. On the displayed page, locate the row that contains the target EIP, and click **Unbind**.
 - d. Click Yes.
- 2. Assign an EIP.
 - a. Log in to the management console.
 - b. On the console homepage, under **Network**, click **Elastic IP**.
 - c. On the displayed page, click **Buy EIP**.

- d. Set the parameters as prompted.
- e. Click Next.
- 3. Bind the new EIP to the ECS.
 - a. On the **EIPs** page, locate the row that contains the target EIP, and click **Bind**.
 - b. Select the desired ECS.
 - c. Click **OK**.
- 4. Release the EIP that has been replaced.
 - a. Release a single EIP.
 - i. Log in to the management console.
 - ii. On the console homepage, under Network, click Elastic IP.
 - iii. In the EIP list, locate the row that contains the target EIP, and click **Release**.
 - iv. Click Yes.
 - b. Unbind multiple EIPs at a time.
 - i. Log in to the management console.
 - ii. On the console homepage, under Network, click Elastic IP.
 - iii. In the EIP list, select the EIPs to be unbound.
 - iv. Click **Unbind** above the EIP list.
 - v. Click Yes.

Scenario 2: Changing an EIP for a Load Balancer

- 1. Unbind an EIP.
 - a. Log in to the management console.
 - b. Click **Service List**. Under **Network**, 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. For details, see 2.
- Bind the new EIP to the load balancer.
 - a. Log in to the management console.
 - b. Click Service List. Under Network, 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 has been replaced. For details, see 4.

Scenario 3: Changing an EIP for a NAT Gateway

- 1. Assign an EIP. For details, see 2.
- 2. Modify an SNAT rule.

For details about how to modify an SNAT rule, see **Modifying an SNAT Rule**. In the EIP area, select the newly assigned EIP and deselect the original EIP

(ensure that the deselected EIP belongs to the IP address range on Telefonica Open Cloud).

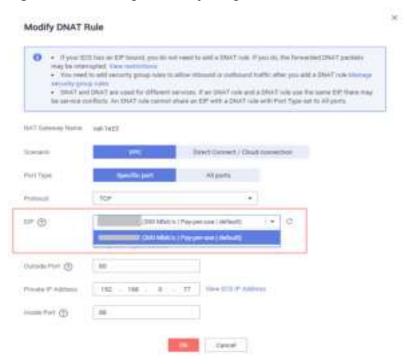
Figure 4-3 Selecting the newly assigned EIP



3. Modify a DNAT rule.

For details about how to modify a DNAT rule, see **Modifying a DNAT Rule**. In the EIP area, select the newly assigned EIP (ensure that the original EIP belongs to the IP address range on Telefonica Open Cloud).

Figure 4-4 Selecting the newly assigned EIP



4. Release the EIP that has been replaced. For details, see 4.

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

No. EIPs and their associated cloud resources must be in the same region. For example, an EIP in the **CN North-Beijing1** region cannot be bound to a resource in the **CN North-Beijing4** region.

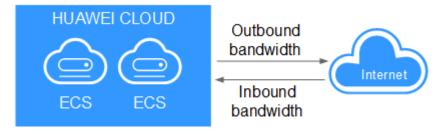
5 Bandwidth

5.1 What Are Inbound Bandwidth and Outbound Bandwidth?

Inbound bandwidth: refers to the bandwidth consumed when data is transferred from the Internet to HUAWEI CLOUD. For example, resources are downloaded from the Internet to ECSs in the cloud.

Outbound bandwidth: refers to the bandwidth consumed when data is transferred from HUAWEI CLOUD to the Internet. For example, the ECSs in the cloud provide services accessible from the Internet and external users download resources from the ECSs.

Figure 5-1 Inbound bandwidth and outbound bandwidth



HUAWEI CLOUD only bills for the outbound bandwidth.

Ⅲ NOTE

Inbound and outbound bandwidths have been adjusted as follows since July 31, 2020 00:00:00 GMT+08:00:

- If your purchased or modified 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 or modified bandwidth.
- If your purchased or modified bandwidth is greater than 10 Mbit/s, both the inbound and the outbound bandwidth will be the same as the purchased or modified bandwidth.

5.2 How Do I Know If My Used Bandwidth Exceeds the Limit?

Symptom

The bandwidth size configured when you buy a dedicated or shared bandwidth is the upper limit of the outbound bandwidth. If the traffic of your web application bound for the Internet is not transferred smoothly, check whether the outbound bandwidth of the EIP bound to the ECS is greater than the configured bandwidth size.

□ NOTE

If the outbound bandwidth exceeds the configured bandwidth size, packet loss may occur. To prevent data loss, it is recommended that you monitor the bandwidth.

Troubleshooting

The following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 5-2 Troubleshooting

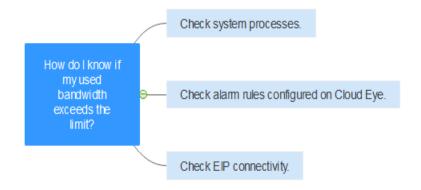


Table 5-1 Troubleshooting

Possible Cause	Solution
System processes leading to high bandwidth	See System Processes Leading to High Bandwidth Usage
Improper Cloud Eye alarm rules	See Improper Cloud Eye Alarm Rules
EIP connection failure	See Why Does Internet Access Fail Even If My ECS Is Bound with an EIP?

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 become inaccessible unexpectedly.

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

- Troubleshooting High Bandwidth or CPU Usage of a Windows ECS
- Troubleshooting High Bandwidth or CPU Usage of a Linux ECS

Improper Cloud Eye Alarm Rules

In the case that you have created alarm rules for bandwidth usage on the Cloud Eye console, if the outbound bandwidth limit or the alarm period is set too small, the system may generate alarms frequently.

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 when the maximum outbound bandwidth is greater than or equal to 4.8 Mbit/s in three consecutive periods. You can also **increase your bandwidth**.

 Log in to the management console, under Management & Deployment, click Cloud Eye. On the Cloud Eye console, choose Alarm Management > Alarm Rules.

Figure 5-3 Alarm Rules



2. Click **Create Alarm Rule** and configure an alarm rule to generate alarms when the bandwidth exceeds the limit.

Figure 5-4 Creating an alarm rule



Submitting a Service Ticket

If the problem persists, **submit a service ticket**.

5.3 What Is the Bandwidth Size Range?

The bandwidth ranges from 1 Mbit/s to 2000 Mbit/s.

The bandwidth in regions **LA-Mexico City1** and **LA-Sao Paulo1** ranges from 1 Mbit/s to 1000 Mbit/s.

5.4 What Bandwidth Types Are Available?

There are dedicated bandwidth and shared bandwidth. A dedicated bandwidth can only be used by one EIP, whereas a shared bandwidth can be used by multiple EIPs.

5.5 What Are the Differences Between a Dedicated Bandwidth and a Shared Bandwidth? Can a Dedicated Bandwidth Be Changed to a Shared Bandwidth or the Other Way Around?

Dedicated bandwidth: The bandwidth can only be used by one EIP and the EIP can only be used by one cloud resource, such as an ECS, a NAT gateway, or a load balancer.

Shared bandwidth: The 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 workloads.

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.

5.6 How Do I Buy a Shared Bandwidth?

- 1. Log in to the management console.
- 2. On the console homepage, under **Network**, click **Virtual Private Cloud**.
- 3. In the navigation pane on the left, choose **Elastic IP and Bandwidth** > **Shared Bandwidths**.
- 4. In the upper right corner, click **Buy Shared Bandwidth**. On the displayed page, configure parameters as prompted to buy a shared bandwidth.

5.7 Is There a Limit to the Number of EIPs That Can Be Added to Each Shared Bandwidth?

A maximum of 20 EIPs can be added to each shared bandwidth. If you want to add more EIPs to each shared bandwidth, **submit a service ticket** to request a quota increase.

5.8 Can I Increase My Bandwidth Billed on Yearly/ Monthly Basis and Then Decrease It?

You can increase bandwidth for a yearly/monthly EIP any time you want to, and the change takes effect immediately. But you can only decrease it when you renew the EIP subscription, and the decreased bandwidth will take effect in the new billing cycle. For details, see **Modifying EIP Bandwidth**.

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

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

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

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) in consideration of losses, such as computer performance, network device quality, resource usage, and network peak hours.

5.10 What Are the Differences Between Static BGP and Dynamic BGP?

The differences between static BGP and dynamic BGP are as follows:

Table 5-2 Differences between static BGP and dynamic BGP

Aspect	Static BGP	Dynamic BGP
Definition	Static routes are manually configured and must be manually reconfigured any time 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.

Aspect	Static BGP	Dynamic BGP
Assuranc e	When changes occur on a network that uses static BGP, the manual configuration takes some time and high availability cannot be guaranteed.	When a fault occurs on a carrier's link, dynamic BGP will quickly select another optimal path to take over services, ensuring service availability.
	NOTE If you select static BGP, your application system must have disaster recovery setups in place.	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.
Service availabilit y	99%	99.95%
Billing	Their price from least to most expensive: static BGP, dynamic BGP. For details, see EIP Pricing Details .	

□ NOTE

For more information about service availability, see **Huawei Cloud Service Level Agreement**.

6 Connectivity

6.1 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.

6.2 Why Is Internet or Internal Domain Names in the Cloud Inaccessible Through Domain Names When My ECS Has Multiple NICs?

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

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

- 1. Log in to the management console.
- 2. On the console homepage, under **Network**, click **Virtual Private Cloud**.
- 3. In the 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**.

6.3 What Are the Constraints Related to VPC Peering?

- VPC peering connections created between VPCs that have overlapping subnet CIDR blocks may not take effect.
- You cannot have more than one VPC peering connection between any two VPCs at the same time.
- You cannot create a VPC peering connection between VPCs in different regions.
- Even if VPC 1 and VPC 2 are connected using a VPC peering connection, ECSs in VPC 2 cannot access the Internet through the EIP of VPC 1. If you want to allow the ECSs in VPC 2 to access the Internet through the EIP of VPC 1, you can use a NAT gateway service or configure an SNAT server. For details, see Having an ECS Without a Public IP Address Access the Internet.
- If you request a VPC peering connection with a VPC of another account, the peer account must accept the request to activate the connection. If you request a VPC peering connection with a VPC of your own, the system automatically accepts the request and activates the connection.
- After a VPC peering connection is established, the local and peer tenants must add routes in the local and peer VPCs to enable communication between the two VPCs.
- VPC A is peered with both VPC B and VPC C. If VPC B and VPC C have overlapping CIDR blocks, you cannot configure routes with the same destinations for VPC A.
- To ensure security, do not accept VPC peering connections from unknown accounts.
- Either owner of a VPC in a peering connection can delete the VPC peering connection at any time. If a VPC peering connection is deleted by one of its owners, all information about this connection will also be deleted immediately, including routes added for the VPC peering connection.
- If VPCs connected by a VPC peering connection have overlapping CIDR blocks, the connection can only enable communication between specific (nonoverlapping) subnets in the VPCs. If subnets in the two VPCs of a VPC peering connection have overlapping CIDR blocks, the peering connection will not take effect. When you create a VPC peering connection, ensure that the VPCs involved do not contain overlapping subnets.
- You cannot delete a VPC that has VPC peering connection routes configured.
- A VPC peering connection can be created between VPCs in same region even if one is created on the HUAWEI CLOUD Chinese Mainland console and another on the HUAWEI CLOUD international console.

6.4 Why Does Communication Fail Between VPCs That Are Connected by a VPC Peering Connection?

Symptom

Two VPCs cannot communicate with each other after you create a VPC peering connection between them.

Troubleshooting

The following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 6-1 Troubleshooting



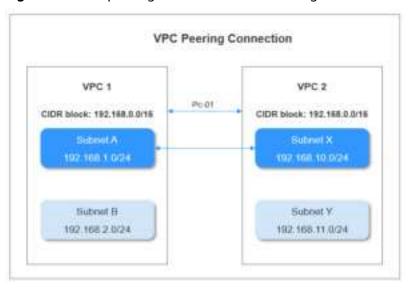
Table 6-1 Troubleshooting

Possible Cause	Solution
Incorrect VPC Peering Connection Configuration	See Incorrect VPC Peering Connection Configuration
Incorrect Network Configuration	See Incorrect Network Configuration
ECS Communication Failure	See ECS Communication Failure

Possible Cause	Solution
Route Conflicts Between VPC Peering and Direct Connect or Between VPC Peering and VPN	Replan the network connection scheme.
Route Already Exists	Replan the network connection scheme.

Incorrect VPC Peering Connection Configuration

Figure 6-2 VPC peering connection network diagram



Add routes to enable communication between Subnet A in VPC 1 and Subnet X in VPC 2. Figure 6-3 shows the route table configurations.

Figure 6-3 VPC peering connection route table



Figure 6-2 is used as an example. Perform the following operations:

1. Check whether a VPC peering connection has been successfully created for the two VPCs, especially, whether the VPC IDs are correctly configured.

- If the VPC peering connection is not correctly configured, create it again.
- 2. Check whether routes have been configured for the VPC peering connection. For example, the destination of the route for VPC 1 must be the subnet CIDR block in VPC 2.
 - If the routes of the VPC peering connection are incorrect, add local and peer routes on the VPC peering connection details page. The VPC peering connection works properly only after the routes are correctly configured.
- 3. Check whether VPC 1 and VPC 2 have overlapping subnets. For example, if VPC 1 and VPC 2 each has a subnet with the same CIDR block, such as 192.168.11.0/24, the VPC peering connection will become invalid.

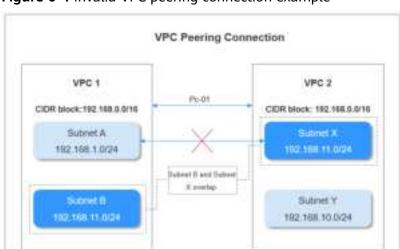


Figure 6-4 Invalid VPC peering connection example

Incorrect Network Configuration

1. Check whether the security group of the ECS NIC is correctly configured. You can view the security group on the ECS details page. Check whether a security group rule that allows the ECS to communicate with the peer VPC subnet has been configured. For example, a security group rule described in Figure 6-5 has to be configured for the NICs of all ECSs in VPC 1.

Figure 6-5 Security group configuration



- 2. Check whether traffic filtering has been configured on the firewall associated with the subnet to which the ECS NIC belongs. If the required traffic is blocked, configure firewall rules to allow the traffic.
- 3. Check whether the traffic between the subnets involved in the VPC peering connection is blocked by the network ACLs. If the required traffic is blocked, configure network ACL rules to allow the traffic.
- 4. If the ECS has more than one NIC, ensure that correct policy-based routing has been configured for the ECS and that packets with different source IP addresses match their own rules.

For example, if the IP address of eth0 is 192.168.1.10/24, and that of eth1 is 192.168.2.10/24, 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 two IP addresses can be pinged, the policy-based routing configured for the two NICs is correct.

Otherwise, you need to configure policy-based routing for the ECS with multiple NICs. For details, see **How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?**

ECS Communication Failure

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

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

If an ECS runs the Window OS, run the **ipconfig** command.

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

2. Ping the gateway address of the subnet to which the ECS belongs to check the ECS communication.

Obtain the gateway address from the VPC details page on the console. In most cases, the gateway address is in the format xxx. xxx. xxx. 1. Ping the gateway address to check the communication. If the ping operation for the gateway address fails, see Why Does My ECS Fail to Communicate at a Layer 2 or Layer 3 Network?

Route Conflicts Between VPC Peering and Direct Connect or Between VPC Peering and VPN

Check whether any of the VPC connected by the VPC peering connection have a VPN or Direct Connect connection connected. If yes, check the next hop destination of their routes.

If the route destination of the VPC peering connection overlaps with that of a Direct Connect or VPN connection, the route may be invalid.

Route Already Exists

If a message indicating that this route already exists is displayed when you add a route for a VPC peering connection, check whether the destination of a VPN, Direct Connect, or VPC peering connection route already exists. If the destination already exists, the VPC peering connection cannot take effect.

Submitting a Service Ticket

If the problem persists, **submit a service ticket**.

You need to ping the ECS at one side of the VPC peering connection from another ECS at the other side of the VPC peering connection to send ICMP packets and provide the technical support engineer with the following information:

Item	Description	Value
VPC1 ID	ID of VPC 1	-
VPC2 ID	ID of VPC 2	-
VM1 ID	ID of the ECS in VPC 1	-
VM2 ID	ID of the ECS in VPC 2	-
Subnet1 ID	ID of the subnet used by ECS 1	-
Subnet2 ID	ID of the subnet used by ECS 2	-
IP1	ECS 1 IP address	-
IP2	ECS 2 IP address	-

Ⅲ NOTE

You can add - ${\bf t}$ to the end of the ping command to enable the Windows ECS to continuously send ICMP packets.

6.5 How Many VPC Peering Connections Can I Create?

You can create a maximum of 50 VPC peering connections in one region. Accepted VPC peering connections consume the quota of both the owners of a VPC peering connection. A VPC peering connection in the pending approval state consumes the quota of only the requester.

6.6 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. That is, if both are configured for an ECS to enable Internet access, the EIP will be used preferentially.

6.7 Why Does Intermittent Interruption Occur When a Local Host Accesses a Website Built on an ECS?

Symptom

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

Troubleshooting

Check the local network of the user.

If the local host communicates with the ECS using NAT, this problem may occur.

Run the following command to check whether tcp_tw_recycle is enabled on the ECS:

sysctl -a|grep tcp_tw_recycle

faulty.

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, packet loss occurs, that is, the network is

Procedure

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

• Run the following command to temporarily modifying the parameters (the modification becomes invalid after the ECS is restarted):

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

6.8 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, respectively.

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

Troubleshooting

- 1. Ping ECS **ecs01** from ECS **ecs02** through the EIP. If ECS **ecs01** can be pinged, the NIC of ECS **ecs01** is working properly.
- 2. Run the **arp -n** command on ECS **ecs02** to check whether the command output contains the MAC address of ECS **ecs01**. If the command output does

- not contain the MAC address of ECS **ecs01**, ECS **ecs02** fails to learn the MAC address of ECS **ecs01** when using the private IP address to ping ECS **ecs01**.
- 3. Run the **ip a** command on ecs01 to check the NIC configuration of ECS **ecs01**. The following figure shows an example.

Figure 6-6 Viewing ECS ecs01 NIC configuration

```
[root@bd-slavel ~l# ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
  link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00:00
  inet 127.0.0.18 scope host lo
  inet6 ::1/128 scope host
    valid_lft forever preferred lft forever
2: eth0: <BPDADCAST,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
  inet 192.168.1.141/24 brd 192.168.1.255 scope global eth0
  inet6 fe80::f816:3eff:fe02:ldd5/64 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, ECS **ecs01** fails to send packets to ECS **ecs02**.

Procedure

Modify the NIC configuration of ECS **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

6.9 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 packet loss occurs when they communicate.

Troubleshooting

The following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 6-7 Troubleshooting

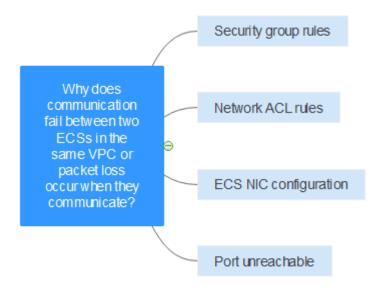


Table 6-2 Troubleshooting

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 Internet Control Message Protocol (ICMP) traffic.

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

Figure 6-8 Inbound security group rule



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

Network ACL Rules

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

ECS NIC Configuration

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

- Check whether multiple NICs are configured for the ECS. If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policybased routing for the ECS. For details, see How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?
- Log in to the ECS and run the following command to check whether the NIC
 has been created and obtained a private IP address. If there is no NIC
 information or the private IP address cannot be obtained, contact technical
 support.

ifconfig

Figure 6-9 NIC IP address

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

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.

6.10 Why Cannot 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 following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 6-10 Troubleshooting

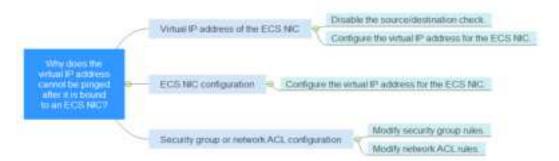


Table 6-3 Troubleshooting

Possible Cause	Solution
Virtual IP address of the ECS NIC	See Virtual IP Address of the ECS NIC
ECS NIC configuration	See ECS NIC Configuration
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 the 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 **NICs** 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 no IP address is displayed for **Virtual IP Address**, click **Manage Virtual IP Address** and configure an IP address.

∩ NOTE

To check whether the virtual IP address has been configured, you can only run the **ip** address command. For details, see **Binding a Virtual IP Address to an EIP or ECS**.

ECS NIC Configuration

The following uses Linux and Windows ECSs as examples to describe how to check whether an ECS NIC has been correctly configured.

For a Linux ECS:

 Run the following command on the ECS to check whether NIC ethX:X exists: ifconfig

Figure 6-11 Checking for NIC ethX:X

```
[rvotMacy ] # ifcoming
with: flags=4163/UP, SNCADCAST, NUMNING, NULTICAST) win 1500
Inst 102.168.1.2 netwark 255.265.285.0 benedicast 192.168.2.255
(net6 fa80;:SSCADCAST, NUMNING, NULTICAST) win 1500
wither fallo:Se:4d:5h:98 txpususien 1000 (Bithernet)
NI packets 77390 bytes 5101364 (4.8 MiS)
NI errors 0 dropped 0 overrans 0 frume 0
TX packets 68798 bytes 6000022 (7.7 NIS)
TX errors 0 dropped 0 overrans 0 carrier 0 calligions 0
with0:1: flags=4163/UP, SNCADCAST, NUMNING, NULTICAST; with 1500
Inst 192.168.1.137 setwark 255.256.258.0 kroedcast 192.168.1.255
wither fail0:3e:4d:Sh:98 txpususien 1000 (Ethernet)
```

The command output in the preceding figure contains the NIC **eth***X:X.* **192.168.1.137** is the virtual IP address of the ECS NIC.

- If NIC ethX:X exists, the ECS NIC is correctly configured.
- If NIC ethX:X does not exist, perform the following operations:
- 2. If the command output does not contain the NIC **eth** *X:X*, run the following command to 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
NETMADK=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 6-12 Checking whether the virtual IP address has been configured

```
Users\Administrator>ipconfig /all
indows IP Configuration
 Host Name .
                                   : dst-win
 Primary Dns Suffix
 Node Type . . . . .
IP Routing Enabled.
WINS Proxy Enabled.
                       . . . . . : Hybrid
thernet adapter Ethernet 5:
 Connection-specific DNS Suffix
 Description . . . . . . . . . . . .
                                   : Red Hat VirtIO Ethernet Adapter #2
 Physical Address. . . . . . . . .
                                     FA-16-3E-83-B2-73
 Link-local IPv6 Address . . . . . : fe80::6182:a265:10bc:134e%3(Preferred)
  192.168.10.41(Preferred)
  IPv4 Address. . . . . . . . . . : 192.168.10.137(Preferred)
                                   : 255.255.255.0
: 192.168.10.1
  Subnet Mask . . . . . . . . . . . .
  Default Gateway . . . . . .
  DHCPv6 IAID .
                                     184161854
                                     80-81-80-81-21-9F-1A-85-52-54-88-A6-AD-AC
  DHCPv6 Client DUID. . .
  DNS Servers . . . .
                                     100.125.1.250
                                     114.114.114.114
```

In the preceding command output, check whether the value of **IPv4 Address** is the virtual IP address 192.168.10.137 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 6-12**. For example, 192.168.10.41.

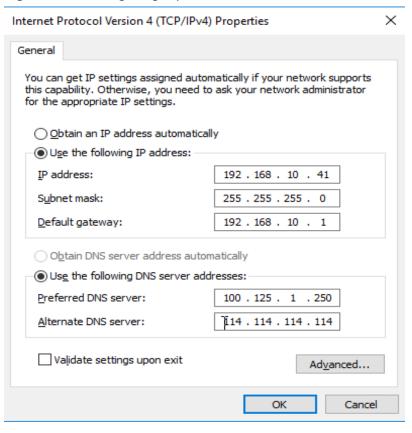


Figure 6-13 Configuring a private IP address

- 6. Click Advanced.
- 7. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address configured in **Figure 6-12**. For example, 192.168.10.137.

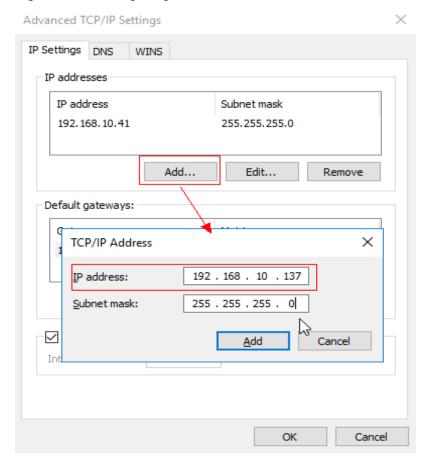


Figure 6-14 Configuring virtual IP address

Security Group or Network ACL Configuration

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

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

Submitting a Service Ticket

If the problem persists, submit a service ticket.

6.11 Why Does My ECS Fail to Use Cloud-init?

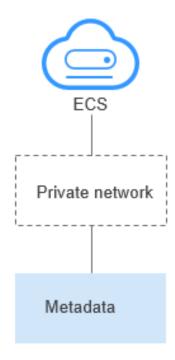
Symptom

An ECS fails to use cloud-init.

Troubleshooting

Figure 6-15 shows the process for an ECS to obtain metadata using the cloud-init.

Figure 6-15 Process for obtaining metadata



Check the following possible causes.

Figure 6-16 Possible causes



Table 6-4 Possible causes

Possible Cause	Solution
The ECS has no IP address obtained.	See The ECS Has No IP Address Obtained
Incorrect route for 169.254.169.254	See Incorrect Route for 169.254.169.254
Fail to obtain the ECS metadata.	See Failing to Obtain the ECS Metadata
Fail to log in to the ECS or create a non-root user after cloud-init is configured.	Check the format of the /etc/cloud/cloud.cfg configuration file. For details, see Failing to 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 No IP Address Obtained

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 6-17 ECS IP address

```
bash-4.1# if config
ethØ
            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
            inetb addr: febB::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
           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
eth1
            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 IP address **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 6-18 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
192.168.1.8/24 dev eth0 proto kernel scope link src 192.168.1.200
192.168.1.8/24 dev eth1 proto kernel scope link src 192.168.1.179
169.254.0.8/16 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 do not support the cloud-init function. To use this function, change the ECS 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, handle the issue based on the following information:
 - If the ECS was created before cloud-init is enabled, run the service network restart command to obtain the correct route.

 If the ECS is newly created, submit a service ticket or contact technical support.

Failing to 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 6-19** is displayed, the ECS successfully obtains the metadata.

Figure 6-19 Command output

```
-bash-1.18 curl http://169.254.169.254/openstack/latest/weta_data.json
("random_seed": "rTV-sblEh64 jUC.mq511038pH5cC78FFRTeWilloundPupostq/Es8EJondF8iJkMGGTzbCTh815HH4S1X
Ob6in-yBfheyib6 j6Oh4608FFgDv6C7fFRTeWilloundPupostq/Es8EJondF8iJkMGGTzbCTh815HH4S1X
Ob6in-yBfheyib6 j6Oh4608FFgDv6C7fFRTeWilloundPupostq/Es8EJondF8iJkMGGTzbCTh815HH4S1X
Ob6in-yBfheyib6 j6Oh4608FFgDv6C7fDkgmjgPr-1k2F8qptlvq/LAMSibhrsqVeokTy5sxis.iCl2SSHWG+1ViZiH8uh4m8qp
of ITTtopv2TwhYEkIFwkZsy7h6FFDkgmjgPr-1k2F8qptlVpyBr2pH4h7beZa4z7gCh8thWT7HUyGUbeaU5/IPDUE1JJGDp0H1/-vz
eDgc1n09Cs0G1VPuELadVDc-Wrlk-4Z2F0DYBMJShHJCZSMTDSUB0SSQ701FnA+NtbDeoB-gB5iFLvWexw8G5BLcjm1fjh0-wqot
vSae6Zeedis1ffscqu8jbCh6inthJ1YGMSon+6F#SQLEppDFrRUBUERSH11867JprES8ppc-4srhygiuMYIJUTUJYQMBUERFF3o
"," unid": "SJebb737-ddc5-4380-9fac-aa72b48b1B1a", "availabilineme": "eu-de-d2", "hostamme": "eu-d
```

Failing to 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 posed by Linux OS providers. The following figure shows the example /etc/cloud/cloud.cfg configuration file for the Ubuntu OSs.

Figure 6-20 Configuration file

```
System_info:

8 This will affect which distro class gets used distro: the!

8 Default user name + that default users groups (if added/used) default user:
name: Jinux // Spedie the unmover briogs.

1ock_passwd: False // The vole hole inducte that the powered had model to encided for some 00, the vole 0 inducts that the powered had model to encided for some 00, the vole 0 inducts that the powered had model to encided genes: Cloud User
groups: arexx //Spedie whether the new will be added to a more. The powered are entire that are entired to a south of the some powered and the on entire many under failures in the other.

passwd: f6515050VXXIIIA61shiJITFM2VIJHSTF0JIZ5NOVXXIIIMSgjDSSJWXIIIX00Jwe8FTcwbweeSSVNpRaMoSPwaxxCy961woB0
passwd: f7AL1=(AL1) ND0FASSWD:ALL*] //Spedie that all permission of one not will be quoted to the som.

thell: /bin/bash //Spedie that the bost desire all as and/or path classes
pather
cloud_dir: /war/lib/cloud/
templates_dir: /etc/cloud/templates/
ssh_rvcame: schd
```

Failing to Use an Obtained Private Key to Log in to an ECS After the ECS Starts (Failing 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-45b8-8ed6-30 aafc8fbdba	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	-

6.12 Why Does Internet Access Fail Even If My ECS Is Bound with an EIP?

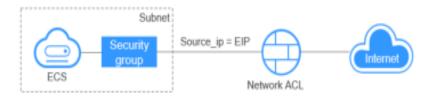
Symptom

You have an ECS that has an EIP bound, but the ECS cannot access the Internet.

Troubleshooting

Figure 6-21 shows the process for an ECS to access the Internet using an EIP.

Figure 6-21 EIP network diagram



Locate the fault based on the following procedure.



Figure 6-22 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 the 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 whether the ECS is running properly.

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

Figure 6-23 ECS status



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

- Check whether the ECS NIC has an IP address assigned.
 Log in to the ECS, and run the ifconfig or ip address command to check the ECS NIC IP address.
 - If an ECS runs the Window OS, run the **ipconfig** command.
- Check whether the virtual IP address is correctly configured on the ECS NIC.
 Log in to the ECS, and run the ifconfig or ip address command to check the ECS NIC IP address. If the ECS NIC does not have an IP address configured, run a command to configure an IP address for the ECS NIC. For example, run the ip addr add virtual IP address eth0 command to configure IP address 192.168.1.192/24 for the NIC.

Figure 6-24 Virtual IP address of a NIC

Check whether the ECS NIC has a default route. If no default route exists, run the **ip route add** command to add the default route.

Figure 6-25 Default route

```
192.168.1.8/24 dev eth8 proto kernel scope link src 192.168.1.288
192.168.1.8/24 dev eth1 proto kernel scope link src 192.168.1.179
169.254.8.8/16 dev eth8 scope link metric 1882
default via 192.168.1.1 dev eth8 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 6-26** has no EIP bound and only has a private IP address bound.

Figure 6-26 EIP status

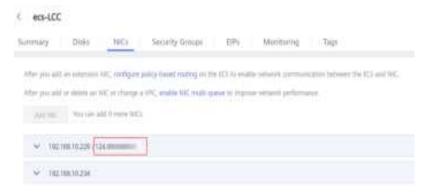


Step 4: Check Whether the EIP Is Bound to the Primary NIC of the ECS

Check whether the EIP is bound to the primary NIC of the ECS. If the EIP is not bound to the primary NIC of the ECS, bind it.

You can view the NIC details by clicking the **NICs** tab on the ECS details page. By default, the first record in the list is the primary NIC and the EIP is bound to the primary NIC as shown in the following figure.

Figure 6-27 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 traffic filtering by network ACL has been configured to block certain traffic from the subnet used by the ECS NIC.

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 fails to 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

Item	Description	Example	Value
VPC ID	N/A	Example: 120b71c7-94ac-45b8-8e d6-30aafc8fbdba	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

6.13 How Do I Handle the IB Network Failure?

RDMA Communication Failure Between Two IB ECSs

Check whether the Pkeys on the two ECSs are consistent.
 Run the following command to check for the Pkeys allocated to the ECSs: cat /sys/class/infiniband/mlx5_0/ports/1/pkeys/* | grep -v "0x0000"

Figure 6-28 Checking Pkey consistency

[root@test2 ~]# cat /sys/class/infiniband/mlx5_0/ports/1/pkeys/* | grep -v "0x0000" 0x8ee5 0x7fff

- If only one Pkey is obtained, contact technical support.
- If two Pkeys are obtained, ensure that the two Pkeys on the two ECSs are the same.
- 2. Run the following command to check whether the firewall is disabled:

service firewalld status

Figure 6-29 Checking the firewall

```
[root@test2 -]# Dervice firewalld status firewalld.service

* firewalld.service - firewalld - dynamic firewalld.service

* firewalld.service - firewalld - dynamic firewalld.service; enabled; vendor preset; enabled)

Loaded: loaded (/usr/lib/systemd/system/firewalld.service; enabled; vendor preset; enabled)

Active: inactive (dead) since Tue 2018-01-02 20:27:36 EST; 10h ago

Docs: man:firewalld(1)

Process: 861 Exec@tatt=/usr/sbin/firewalld --mofork --nopid SFIREWALLD_ARGS (code=exited, status=0/SUCCESS)

Main FID: 861 (code=exited, status=0/SUCCESS)

Jan 02 06:04:35 ecs=g00200266-h2-0002.novalocal systemd[1]: Starting firewalld - dynamic firewall daemon...
Jan 02 00:27:35 est2 systemd[1]: Stopping firewalld - dynamic firewall daemon...
Jan 02 20:27:36 tast2 systemd[1]: Stopped firewalld - dynamic firewall daemon...
```

If the firewall is not disabled, run the following command to disable it: service firewalld stop

3. Check whether the RDMA communication command is correct.

Run the following command on ECS 1 (client):

ib_write_lat -x 0 --pkey_index 0 192.168.0.218

Run the following command on ECS 2 (server):

ib_write_lat -x 0 --pkey_index 0

No IP Address for the ECS IB Port

After you run the **ifconfig** command, the command output shows that no IP address has been assigned to the ECS InfiniBand (IB) port.

Run the following command to check for the Pkey:
 cat /sys/class/infiniband/mlx5_0/ports/1/pkeys/* | grep -v "0x0000"

Figure 6-30 Checking Pkey

```
[root@test2 ~]# cat /sys/class/infiniband/mlx5_0/ports/1/pkeys/* | grep -v "0x0000"
0x8ee5
0x7fff
```

If only one Pkey is obtained, contact technical support.

Run the following command to assign an IP address to the ECS IB port: dhclient ib0

If no command output is displayed, the IP address cannot be obtained using DHCP.

3. Contact technical support.

After you have performed the preceding steps, if the IB network still cannot be used for communication or the IB port still cannot obtain an IP address, contact technical support for assistance and provide the technical support engineer with the following information.

Ite m	Descripti on	Example	Value
VP C1 ID	VPC 1 ID	Example: fef65559- c154-4229- afc4-9ad0314437ea	N/A

Ite m	Descripti on	Example	Value
VM	ID of ECS	Example:	N/A
1	1 in VPC	f7619b12-3683-4203-9	
ID	1	271-f34f283cd740	
VM	ID of ECS	Example:	N/A
2	2 in VPC	f75df766-68aa-4ef3-	
ID	1	a493-06cdc26ac37a	

6.14 Why Does My ECS Fail to Communicate at a Layer 2 or Layer 3 Network?

Symptom

An ECS fails to ping the gateway of the subnet where the ECS resides.

Troubleshooting

Locate the fault based on the following procedure.



Figure 6-31 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 the **ifconfig** or **ip address** command to check the ECS NIC IP address. If an ECS runs the Window OS, run the **ipconfig** command.

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 that allows the ECS to communicate with the peer VPC subnet has been configured.

Figure 6-32 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 the communication between the subnets involved in the VPC peering connection.

Figure 6-33 Network ACL rule



Submitting a Service Ticket

If the problem persists, **submit a service ticket**.

6.15 How Do I Handle the BMS Network Failure?

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

ifconfig

Figure 6-34 Checking for bond

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

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

route -n

Figure 6-35 Checking BMS route information

```
[root@bms2 rhel]# route =n
Rernel 19 routing table
Destination Gateway Germask Flags Metric Ref Use Iface
169.254.169.254 192.168.2.1 255.255.255 UGM 0 0 0 bond0
192.168.4.0 0.0.0.0 255.255.255.0 U 0 0 0 bond0
192.169.2.0 0.0.0.0 255.255.255.0 U 0 0 0 bond0
169.254.0.0 0.0.0.0 255.255.050 U 1006 0 0 bond0
169.254.0.0 0.0.0.0 255.255.050 U 1006 0 0 bond0
169.254.0.0 100.0.0 255.255.050 U 1007 0 0 bond0
169.254.0.0 100.0.0 255.255.00 U 1007 0 0 bond0
169.254.0.0 100.0.0 0 0 0 bond0
```

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

Figure 6-36 Checking the default route



Check whether a route to 169.254.169.254 exists.

Figure 6-37 Checking the route for IP address range 169.254.169.254

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

If required routes do not exist, contact technical support engineers.

- 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.
- 4. If the failure cannot be rectified after you perform the preceding operations, contact technical support.

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

Ite m	Descript ion	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: f7619b12-3683-4203-92 71-f34f283cd740	N/A
BMS 2 ID	ID of BMS 2 in VPC 1	Example: f75df766-68aa-4ef3- a493-06cdc26ac37a	N/A

6.16 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.

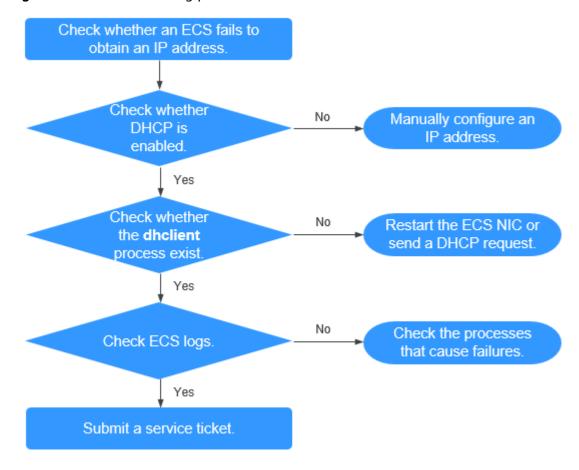


Figure 6-38 Troubleshooting process

- 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

- Run the following command to 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 OS:

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 OS:

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 OS:
 - 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 OS:

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 any processes that cause failures in obtaining IP addresses over DHCP exist.

Submitting a Service Ticket

If the problem persists, submit a service ticket.

Provide the 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.

6.17 How Do I Handle the VPN or Direct Connect Connection Network Failure?

VPN Network

Figure 6-39 shows your network, the customer gateway, the VPN, and the VPC.

Customer Network

Customer
Customer
Customer
Customer
Virtual Private Gatereary

Hurwei VFC
Virtual Private Closef

Figure 6-39 VPN network

Customer Self-Check Guidance

1. Provide your network information.

Obtain information listed in **Table 6-5**. 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.

You can print this table and fill in your values.

Table 6-5 Network information

Item	Description	Example	Valu e
VPC CIDR block	Required for customer gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	N/A	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
Customer gateway type (for example, Cisco)	N/A	N/A	N/A

Item	Description	Example	Valu e
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 in the public cloud system.

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.

6.18 Why Does My Server Can 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 6-40 Possible causes

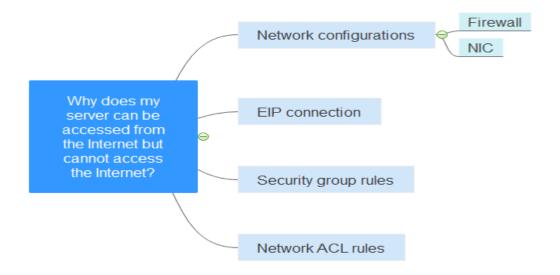


Table 6-6 Possible causes

Possible Cause	Solution	
Network configurations	See Network Configurations	
EIP connection	See Why Does Internet Access Fail Even If My ECS 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 whether the fault is rectified.

- For a Linux ECS, see **Checking the Firewall Configuration**.
- For a Windows ECS, see Checking the Firewall Configuration.
- NIC

Check whether the NIC and DNS configurations.

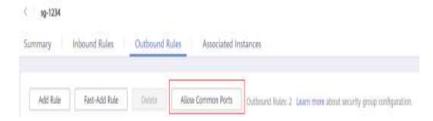
- For a Linux ECS, see Checking the NIC Configuration.
- For a Windows ECS, see Checking the NIC Configuration.

Security Group Rules

Check whether any security group rule of the server denies the outbound traffic.

By default, a security group allows all outbound traffic. If the outbound traffic is denied, **configure security group rules** or click **Allow Common Ports**.

Figure 6-41 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.

Figure 6-42 Allowing outbound traffic



Submitting a Service Ticket

If the problem persists, submit a service ticket.

6.19 Can I Use a VPC Peering Connection to Connect VPCs in Different Regions?

No. You can use a VPC peering connection to connect VPCs in different AZs, but in the same region.

You can use Cloud Connect to enable communication between VPCs in different regions. For details, see **Cloud Connect**.

6.20 Will I Be Billed for Using a VPC Peering Connection?

No. Currently, VPC peering connections are free of charge.

6.21 What Switches Can Connect to a L2CG on HUAWEI CLOUD?

You can use switches, such as CE6850 and Cisco Nexus 9300, which support VXLAN functions.

6.22 Why Is the Layer 2 Connection in the Not Connected State Even After Its Configuration Is Complete?

Possible causes and solutions:

- The VXLAN tunnel of your data center is not properly configured.
 Log in to the switch of your data center and check its tunnel configurations.
 For details, see Configuring a Tunnel Gateway in Your Data Center.
- The Direct Connect connection used by the L2CG is not properly configured.
 Check the Direct Connect connection configurations. For details, see Network and Connectivity.

6.23 Why Is Communication Between the Cloud and On-premises Servers Unavailable Even When the Layer 2 Connection Status Is Connected?

Possible cause: The VXLAN tunnel of your data center is not properly configured.

Solution: Log in to the switch of your data center and check its tunnel configurations. For details, see **Configuring a Tunnel Gateway in Your Data Center**.

6.24 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 the 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.

Figure 6-43 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 the function of automatic IPv6 address allocation, manually obtain the IPv6 address by referring to **Dynamically Assigning IPv6 Addresses**.

■ NOTE

If an ECS is created from a public image:

- 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 and then check whether dynamic IPv6 address assignment has been enabled. Currently, all Linux public images support IPv6, and dynamic IPv6 address assignment is enabled for the Ubuntu 16 OS by default. You do not need to configure dynamic IPv6 address assignment for the Ubuntu 16 OS. For other Linux public images, you need to enable this function.
- By default, IPv6 addresses can only be used for private network communication. If you want to use an IPv6 address to access the Internet or want it to be accessed by IPv6 clients on the Internet, you need to add the IPv6 address to a shared bandwidth. 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.

7 Routing

7.1 How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?

Scenarios

If an ECS has multiple NICs, you can perform the following procedure to configure policy-based routing for the ECS and enable network communication using extension NICs.

Procedure

For a Linux ECS:

 Run the following command to add the priority value and name of the route table for each NIC to the /etc/iproute2/rt_tables file. A smaller priority value represents a higher priority. In this example, 250 and net0 indicate the route table priority value and name of eth0, respectively. 251 and net1 indicate the route table priority value and name of eth1, respectively. If there are multiple NICs, add the route table priority value and name of each NIC one by one.

vi /etc/iproute2/rt_tables

```
# added for dual net
250 net0
```

2. Run the following command to add routing information of each NIC to the /etc/rc.local file:

vi /etc/rc.local

eth0 is used as an example here. If an IPv4 NIC is used, obtain the following information:

IPv4 address (192.168.0.129), subnet (192.168.0.0/24), gateway address (192.168.0.1), and route table added in step **1** (net0)

```
# wait for nics up
sleep 5
# Add v4 routes for eth0
ip route flush table net0
```

```
ip route add default via 192.168.0.1 dev eth0 table net0
ip route add 192.168.0.0/24 dev eth0 table net0
ip rule add from 192.168.0.129 table net0
# Add v4 routes for eth1
ip route flush table net1
ip route add default via 192.168.1.1 dev eth1 table net1
ip route add 192.168.1.0/24 dev eth1 table net1
ip rule add from 192.168.1.138 table net1
```

Before configuring policy-based routing for NICs using IPv6 addresses, ensure that IPv6-related configurations have been performed. For details, see "Linux (Automatic Configuration of IPv6)" in Dynamically Assigning IPv6 Addresses.

eth0 is used as an example here. If an IPv6 NIC is used, obtain the following information:

IPv6 address (2407:c080:802:1be:2233:64bf:b095:54bf), subnet (2407:c080:802:1be::/64), gateway address (fe80::f816:3eff:fef3:20dc), and route table added in step1 is net0

Run the command ip -6 route show| grep default to view the IPv6 gateway address of a NIC.

```
ecs-3c3f ~]# ip -6 route show| grep default
  via fe80::f816:3eff:fef3:20dc dev eth0 proto ra metric 100 pref medium
via fe80::f816:3eff:fe10:5447 dev eth1 proto ra metric 101 pref medium
```

If there are multiple NICs, add their routing information one by one.

```
# Add v6 routes for eth0
ip -6 route flush table net0
ip -6 route add default via fe80::f816:3eff:fef3:20dc dev eth0 table net0
ip -6 route add 2407:c080:802:1be::/64 dev eth0 table net0
ip -6 rule add from 2407:c080:802:1be:2233:64bf:b095:54bf table net0
# Add v6 routes for eth1
ip -6 route flush table net1
ip -6 route add default via fe80::f816:3eff:fe10:5447 dev eth1 table net1
ip -6 route add 2407:c080:802:1bf::/64 dev eth1 table net1
ip -6 rule add from 2407:c080:802:1bf:39ea:bffe:13a2:7a1f table net1
```

Run the following command to add the execute permission for the rc.local

chmod +x /etc/rc.local

- Run the **reboot** command to restart the ECS.
- After the restart, run the following command to check whether the configured routes and route tables take effect.

For IPv4 NICs:

ip rule ip route show table net0 ip route show table net1

```
[root@ecs-3c3f ~]# ip rule
0: from all lookup local
32764: from 192.168.1.138 lookup net1
32765: from 192.168.0.129 lookup net0
32766: from all lookup main
32767: from all lookup default
[root@ecs-3c3f ~]# ip route show table net0
default via 192.168.0.1 dev eth0
192.168.0.0/24 dev eth0 scope link
[root@ecs-3c3f ~]# ip route show table net1
default via 192.168.1.1 dev eth1
192.168.1.0/24 dev eth1 scope link
[root@ecs-3c3f ~]# |
```

For IPv6 NICs:

- ip -6 rule
- ip -6 route show table net0
- ip -6 route show table net1

```
[root@ecs-3c3f -]# ip -6 rule

0; from all lookup local

12764: from 2407:c000:002:15f:39ea:5ffe:13a2:7a1f lookup net1

12765: from 2407:c000:002:15f:39ea:5ffe:13a2:7a1f lookup net0

12765: from 2407:c000:002:15e:2233:64bf:5095:545f lookup net0

12765: from all lookup main

[root@ecs-3c3f +]# ip -6 route show table net0

2407:c000:002:15e::/64 dev eth0 metric 1024 pref medium

default via fe001:fl1f:]eff:fef3:20dc dev eth0 metric 1024 pref medium

[root@ecs-3c3f +]# ip -6 route show table net1

2407:c000:002:15f::/64 dev eth1 metric 1024 pref medium

default via fe00::f016:3eff:fe10:5447 dev eth1 metric 1024 pref medium

[root@ecs-3c3f -]# |
```

6. Specify the source addresses for the test.

For IPv4 addresses:

ping -I 192.168.0.129 xxx

ping -I 192.168.1.138 xxx

For IPv6 addresses:

ping -I 2407:c080:802:1be:2233:64bf:b095:54bf xxx

ping -I 2407:c080:802:1bf:39ea:bffe:13a2:7a1f xxx

For a Windows ECS:

Choose Control Panel > Network and Internet > Network Connections.
 Right-click Local Area Connection 2 and then click Properties.

∩ NOTE

Right-click to add NICs based on the site requirements. If there are multiple NICs, there will be multiple local area connections. Configure them one by one.

- 2. On the Network tab page, select Internet Protocol Version 4 (TCP/IPv4).
- 3. Click **Properties**.
- 4. On the **General** tab page, click **Advanced**.
- 5. On the IP Settings tab, click Add in the Default gateways area.

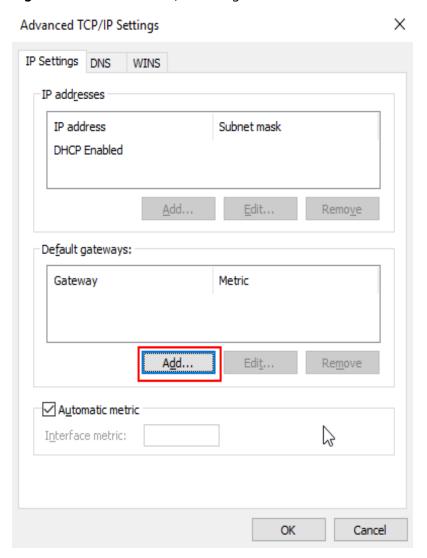
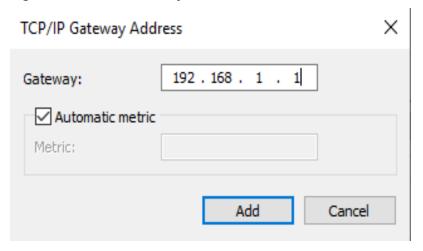


Figure 7-1 Advanced TCP/IP settings

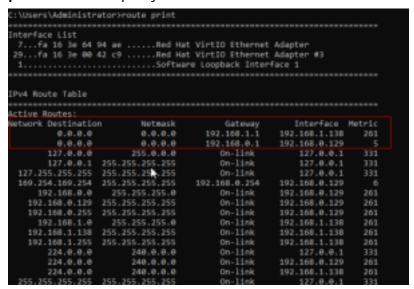
6. Enter the gateway address of the secondary NIC and click **Add**.

Figure 7-2 TCP/IP Gateway Address



7. Click **OK**.

8. Open the command line interface (CLI) of the Windows OS and enter **route print** to view the policy-based routes.



9. Specify the source addresses for the test.

ping -S 192.168.0.129 xxx ping -S 192.168.1.138 xxx

Related Operations

If you want to access the Internet using an extension NIC, see **How Do I Access** the Internet Using an EIP Bound to an Extension NIC?

7.2 Why Can't I Ping an ECS with Two NICs Configured?

Symptom

Your ECS has one primary NIC and one extension NIC in the same subnet. Both the NICs have an EIP bound to access the Internet. The EIP bound to the primary NIC can access the Internet, but that bound to the extension NIC cannot.

Possible Causes

By default, ECSs running CentOS have the reverse path filtering (RP-Filter) enabled. The default route of the ECSs is to forward outgoing traffic through the extension NIC to eth0. However, the system considers that the response data packets should be forwarded from eth1. The system determines that the traffic is received from a wrong NIC and then discards the response packets.

Solution

Configure a policy-based routing rule so that the extension NIC traffic is forwarded from the extension NIC.

1. Run the following command to edit the **rt_tables** file:

vi /etc/iproute2/rt_tables

#
255 local
254 main
253 default
0 unspec
32000 test

Add an alias for the routing table, such as test.

- Save the modification and exit.
- 3. Run the following command to add a route to the **test** table:

ip route add default via *Gateway IP address of the extension NIC* **dev eth1 table** *Name of the routing table*

For example, run the following command:

ip route add default via 192.168.166.1 dev eth1 table test

4. Run the following command to add a policy-based routing rule:

ip rule add from *IP address of the extension NIC* **lookup** *Name of the routing table* **prio** *lower than 32766 but higher than the main table*

For example, run the following command:

ip rule add from 192.168.166.22 lookup test prio 32000

5. Check whether the EIP bound to the extension NIC can access the Internet. If you want to make this rule take effect permanently, add the preceding command to the startup script /etc/rc.local.

7.3 Can a Route Table Span Multiple VPCs?

No.

7.4 How Many Routes Can a Route Table Contain?

Each route table can contain a maximum of 200 routes by default, including routes added for Direct Connect and VPC peering connections.

7.5 Are There Any Restrictions on Using a Route Table?

- The ECS providing SNAT must have the **Unbind IP from MAC** function enabled.
- The destination of each route in a route table must be unique. The next hop must be a private IP address or a virtual IP address in the VPC. Otherwise, the route table will not take effect.
- If a virtual IP address is set to be the next hop in a route, EIPs bound with the virtual IP address in the VPC will become invalid.

7.6 Will a Route Table Be Billed?

The route table function itself is free of charge. However, you are charged for the ECSs and bandwidth that you use together with the route table function.

7.7 Do the Same Routing Priorities Apply to Direct Connect Connections and Custom Routes in the Same VPC?

No. Direct Connect connections and custom routes are used in different scenarios. Therefore, there are different routing priorities for them.

7.8 Are There Different Routing Priorities of the VPN and Custom Routes in the Same VPC?

No. The routing priority of custom routes and that of VPNs are the same.

8 Security

8.1 Are the Security Group Rules Considered the Same If All Parameters Except Their Description Are the Same?

Yes. You cannot add or import a security group rule that has the same parameters but a different description than an existing rule in the security group.

8.2 What Are the Requirements for Deleting a Security Group?

- Before deleting a security group, ensure that the security group is not used by any cloud resource, such as ECS, Relational Database Service (RDS), and Distributed Cache Service (DCS). If the security group is used by a cloud resource, release the cloud resource or change the security group used by the cloud resource, and then delete the security group.
- If the security group you want to delete is associated with rules of another security group (**Source**), delete or modify the associated security group rules, and then delete the security group.

□ NOTE

- The default security group cannot be deleted.
- If a security group is associated with resources other than servers and extension NICs, the security group cannot be deleted.

8.3 Why Is Outbound Access Through TCP Port 25 Restricted?

Symptom

You cannot access an external address using TCP port 25. For example, you cannot run the **Telnet smtp.***.com 25** command.

Cause

By default, TCP port 25 is disabled in the outbound direction for security purposes.

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

This section applies only to the **AP-Hong-Kong** region.

Solution

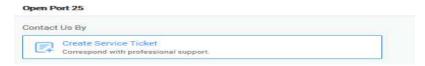
- Use port 465 supported by the third-party email service provider.
- Apply for enabling TCP port 25 in the outbound direction.
 If you must enable TCP port 25 on the ECS for external communications, submit an application.

NOTICE

Before sending the application, you must agree and guarantee that TCP port 25 is only used to connect to third-party Simple Mail Transfer Protocol (SMTP) servers and that emails are sent using the third-party SMTP servers. If you use the EIP specified in the service ticket to directly send emails over SMTP, we will permanently disable TCP port 25 and you will no longer be able to use it or request for it to be enabled.

- On the Create Service Ticket page, choose Products > Elastic Cloud Server.
 For details about how to submit a service ticket, see Submitting a Service Ticket.
- 2. Click Open Port 25 under Select Subtype and click Create Service Ticket.

Figure 8-1 Creating a service ticket



3. On the displayed page, enter the required information.

8.4 Can I Change the Security Group of an ECS?

Yes. Log in to the ECS console, switch to the page showing ECS details, and change the security group of the ECS.

8.5 How Many Security Groups Can I Have?

Each account can have a maximum of 100 security groups and 5000 security group rules.

When you create an ECS, you can select multiple security groups. It is recommended that you select no more than five security groups.

8.6 Will a Security Group Be Billed?

Security groups are free of charge.

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

ECS Configuration

The TFTP daemon determines whether the configuration file specifies the port range. If you use the 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 the 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. Therefore, you can configure a small 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 8-2 Security group rules



8.8 How Many Network ACLs Can I Create?

You can create a maximum of 200 network ACLs. It is recommended that you configure a maximum of 20 inbound or outbound rules for each network ACL. If

you configure more than 20 inbound or outbound rules for a network ACL, the forwarding performance will deteriorate.

8.9 Does a Security Group Rule or a Network ACL Rule Immediately Take Effect for Its Original Traffic After It Is Modified?

- Security groups are stateful. Responses to outbound traffic are allowed to go in to the instance regardless of inbound security group rules, and vice versa. Security groups use connection tracking to track traffic information about traffic to and from instances. If a security group rule is added, deleted, or modified, or an instance in the security group is created or deleted, the connection tracking of all instances in the security group will be automatically cleared. In this case, the inbound or outbound traffic of the instance will be considered as new connections, which need to match the inbound or outbound security group rules to ensure that the rules take effect immediately and the security of incoming traffic.
- A modified network ACL rule will not immediately take effect for its original traffic. It takes about 120 seconds for the new rule to take effect, and traffic will be interrupted during this period. To ensure that the traffic is immediately interrupted after the rule is changed, it is recommended that you configure security group rules.

8.10 Why Are Some Ports in the Public Cloud System Inaccessible?

Symptom: Users in certain areas cannot access some ports in the public cloud system.

Analysis: Ports listed in the following table are high-risk ports and are blocked by default.

Table 8-1 High-risk ports

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

Solution: It is recommended that you use ports that are not listed in the table for your services.

8.11 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 and the priority of rule B is higher than that of rule A, insert rule B before rule A. In this case, the priority of rule B is 1 and that of rule A is 2. Similarly, if the priority of rule B is lower than that of 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 Changing the Sequence of a Network ACL Rule.

8.12 What Do My Security Group Rules Not Take Effect?

Symptom

The security group rules you have configured for an ECS have not taken effect.

Troubleshooting

The following fault causes are sequenced based on their occurrence probability. If the fault persists after you have ruled out a cause, check other causes.

Figure 8-3 Troubleshooting



Table 8-2 Troubleshooting

Possible Cause	Solution
Incorrect Security Group Rule Configurations	See Incorrect Security Group Rule Configuration
Conflicts Between Network ACL Rules and Security Group Rules	See Conflicts Between Network ACL Rules and Security Group Rules
Incorrect ECS Firewall Configurations	See Incorrect ECS Firewall Configurations

Incorrect Security Group Rule Configuration

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

- 1. The direction of a rule is incorrect.
- 2. The protocol of a rule is incorrect.
- 3. The port used in a rule is risky and cannot be accessed. For details about common ports and risky ports, see **Common Ports Used by ECSs**.
- 4. The port used in a rule is not opened. You can perform the following steps to check whether a port is being listened on the server.

For example, you have deployed a website on ECSs. Users need to access your website over TCP (port 80), and you have added the security group rule shown in **Table 8-3**.

Table 8-3 Security group rule

Directio n	Protocol	Port	Source
Inbound	TCP	80	0.0.0.0/0

Linux ECS

To verify the security group rule on a Linux ECS:

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

netstat -an | grep 80

If command output shown in **Figure 8-4** is displayed, TCP port 80 is being listened on.

Figure 8-4 Command output for the Linux ECS



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

Windows ECS

To verify the security group rule on a Windows ECS:

- a. Log in to the ECS.
- b. Choose **Start > Accessories > Command Prompt**.
- c. Run the following command to check whether TCP port 80 is being listened on:

netstat -an | findstr 80

If command output shown in **Figure 8-5** is displayed, TCP port 80 is being listened on

Figure 8-5 Command output for the Windows ECS

TCP 0.0.0.0:80 0.0.0.0:0 LISTENING

- d. Enter http://ECS EIP in the address box of the browser and press Enter.
 If the requested page can be accessed, the security group rule has taken effect.
- ECSs belong to different VPCs. If two ECSs are in the same security group but in different VPCs, the ECSs cannot communicate with each other. To enable communications between the ECSs, use a VPC peering connection to connect the two VPCs. For details about VPC connectivity, see Application Scenarios.

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 operate at the ECS level, whereas network ACLs operate at the subnet level.

For example, if you configure an inbound security group rule to allow access over port 80 and a network ACL rule to deny access over port 80, the security group rule will not take effect.

You can **add a network ACL rule** or **modify a network ACL rule** to allow traffic from the corresponding protocol port.

Incorrect ECS Firewall Configurations

Check whether the firewall of the ECS opens the required ports.

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.

Submitting a Service Ticket

If the problem persists, submit a service ticket.