

Virtual Private Cloud

FAQs

Issue 43
Date 2020-07-23



Copyright © Huawei Technologies Co., Ltd. 2020. 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	1
1.1 What Is a Quota?	1
2 Billing and Payments	3
2.1 Will I Be Charged for Using the VPC Service?	3
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 ?	6
3 VPC and Subnet	7
3.1 What Is Virtual Private Cloud?	7
3.2 Which CIDR Blocks Are Available to 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 What Do I Do If a Subnet Cannot Be Deleted Because It Is Being Used by Other Resources?	10
3.9 What Do I Do If the ECS IP Address Is Lost After the ECS System Time Has Been Changed?	10
3.10 How Do I Make the Changed DHCP Lease Time of a Subnet Take Effect Immediately?	11
3.11 Can I Change the VPC of an ECS?	12
3.12 How Do I Switch to a Private DNS Server?	12
4 EIP	14
4.1 How Do I Assign or Retrieve a Specific EIP?	14
4.2 What Are the Differences Between EIP, Private IP Address, Floating IP Address, and Virtual IP Address?	14
4.3 How Do I Access the Internet Using an EIP Bound to an Extension NIC?	15
4.4 What Are the Differences Between the Primary and Extension NICs of ECSs?	16
4.5 What Are EIPs?	16
4.6 Can an EIP That Uses Dedicated Bandwidth Be Changed to Use Shared Bandwidth?	17
4.7 How Many ECSs Can One EIP Be Assigned to?	17
4.8 How Do I Access an ECS from the Internet After an EIP Is Bound to the ECS?	17
4.9 What Is the EIP Assignment Policy?	18
4.10 Can I Bind an EIP to an ECS, to Another ECS?	18

4.11 Will the EIP Bound to an ECS Be Changed After the ECS Is Stopped and Then Started?.....	18
4.12 Can I Assign a Specific EIP?.....	18
4.13 Will an EIP Be Changed After I Assign It?.....	18
4.14 How Do I Query the Region of My EIPs?.....	19
4.15 Can I Transfer an EIP to Another Account?.....	19
4.16 Can an EIP Be Bound to a Cloud Resource in Another Region?.....	19
5 Bandwidth.....	20
5.1 What Are Inbound Bandwidth and Outbound Bandwidth?.....	20
5.2 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?.....	21
5.3 How Do I Check Whether the Bandwidth Exceeds the Limit?.....	21
5.4 What Are the Differences Between EIP Bandwidth and Private Network Bandwidth?.....	23
5.5 What Is the Bandwidth Size Range?.....	24
5.6 What Bandwidth Types Does the VPC Service Support?.....	24
5.7 How Can I Use Shared Bandwidth?.....	24
5.8 How Many EIPs Can Use the Same Shared Bandwidth?.....	24
5.9 Can I Increase My Yearly/Monthly Bandwidth Size and Then Decrease It?.....	24
5.10 Can I Use Multiple Bandwidth Add-On Packages with Overlapping Validity Periods?.....	24
5.11 What Is the Relationship Between Bandwidth and Upload/Download Rate?.....	25
5.12 What Are the Differences Between Static BGP and Dynamic BGP?.....	25
5.13 What Is Enhanced 95th Percentile Bandwidth Billing?.....	26
6 Connectivity.....	29
6.1 Does a VPN Allow for Communication Between Two VPCs?.....	29
6.2 Why Cannot I Access Public Websites Through Domain Names or Access Internal Domain Names in the Cloud When My ECS Has Multiple NICs?.....	29
6.3 What Are the Limitations of VPC Peering?.....	30
6.4 What Do I Do If VPCs in a VPC Peering Connection Cannot Communicate with Each Other?.....	30
6.5 How Many VPC Peering Connections Can I Have?.....	31
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?.....	31
6.7 What Do I Do If Intermittent Interruption Occurs When a Local Host Accesses a Website Built on an ECS?.....	31
6.8 What Do I Do If ECSs Using Private IP Addresses in the Same Subnet Only Support One-Way Communication?.....	32
6.9 What Do I Do If Two ECSs in the Same VPC Cannot Communicate with Each Other or Packet Loss Occurs When They Communicate?.....	33
6.10 What Do I Do If a Virtual IP Address Cannot Be Pinged After It Is Bound to an ECS NIC?.....	35
6.11 How Do I Handle the Cloud-init Connection Failure?.....	39
6.12 How Do I Handle EIP Connection Failure?.....	43
6.13 How Do I Handle the IB Network Failure?.....	47
6.14 How Do I Handle the VPC Peering Connection Failure?.....	48
6.15 How Do I Handle the Layer 2 or Layer 3 Network Communication Failure?.....	52
6.16 How Do I Handle the BMS Network Failure?.....	52

6.17 How Do I Handle the ECS IP Address Obtaining Failure?.....	54
6.18 How Do I Handle the VPN or Direct Connect Connection Network Failure?.....	55
6.19 What Do I Do If My Server Can Be Accessed from the Internet But Cannot Access the Internet?.....	57
6.20 Can a VPC Peering Connection Be Deployed Across Regions?.....	58
6.21 Is a VPC Peering Connection Charged?.....	58
6.22 Why the IPv6 Address Manually Configured for an ECS Cannot Be Used for Communication?.....	58
6.23 What Devices Can Connect to a L2CG on HUAWEI CLOUD?.....	59
6.24 Why Is the Layer 2 Connection in the Not Connected State Even After Its Configuration Is Complete?	59
6.25 Why Is the Communication Between the Cloud and On-premises Servers Unavailable Even When the Layer 2 Connection Status Is Connected?.....	59
7 Routing.....	60
7.1 How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?.....	60
7.2 Can a Route Table Span Multiple VPCs?.....	62
7.3 How Many Routes Can Exist in a Route Table?.....	62
7.4 What Are the Limitations of a Route Table?.....	62
7.5 Will a Route Table be Billed?.....	62
7.6 Are There Different Routing Priorities for Direct Connect Connections and Custom Routes in the Same VPC?.....	63
7.7 Are There Different Routing Priorities of the VPN and Custom Routes in the Same VPC?.....	63
8 Security.....	64
8.1 Are the Security Group Rules Considered as the Same If All Their Parameters Except the Description Are the Same?.....	64
8.2 What Should I Do Before Deleting a Security Group?.....	64
8.3 What Do I Do If Outbound Access Through TCP Port 25 Is Restricted?.....	64
8.4 Can I Change the Security Group of an ECS?.....	65
8.5 How Many Security Groups Can Each User Have?.....	65
8.6 Is the Security Group Service Charged?.....	66
8.7 How Do I Configure a Security Group for Multi-Channel Protocols?.....	66
8.8 How Many Network ACLs Can a User Have?.....	66
8.9 Does a Security Group Rule or a Network ACL Rule Immediately Take Effect for Its Original Traffic After It Is Modified?.....	66
8.10 What Do I Do If Some Ports in the Public Cloud System Are Inaccessible?.....	67
8.11 Why the Access from a Specified IP Address Is Still Allowed After a Network ACL Rule that Denies the Access from this Specified IP Address Has Been Added?.....	67

1 General

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?


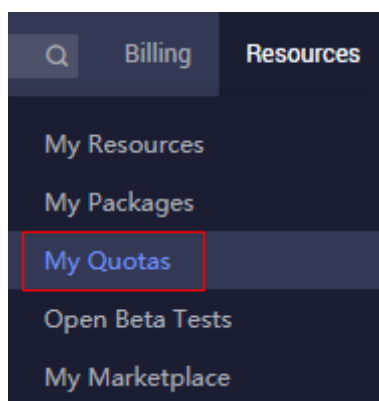
1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper right corner of the page, choose **Resources > My Quotas**.
The **Service Quota** page is displayed.

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.
2. 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.
5. 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 the standard pricing.

2.2 How Is an EIP Billed?

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

The EIP price varies according to the billing mode.

Table 2-1 EIP billing details

Billing Mode	Billed By	EIP Retention Fee	Bandwidth Price	Public Network Traffic Price
Yearly/ Monthly	Bandwidth	-	√	-
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.	√	-
	Traffic		-	√

NOTE

"-" indicates that the fee will not be included in the bill. "√" indicates that the fee will be included in the bill.

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.

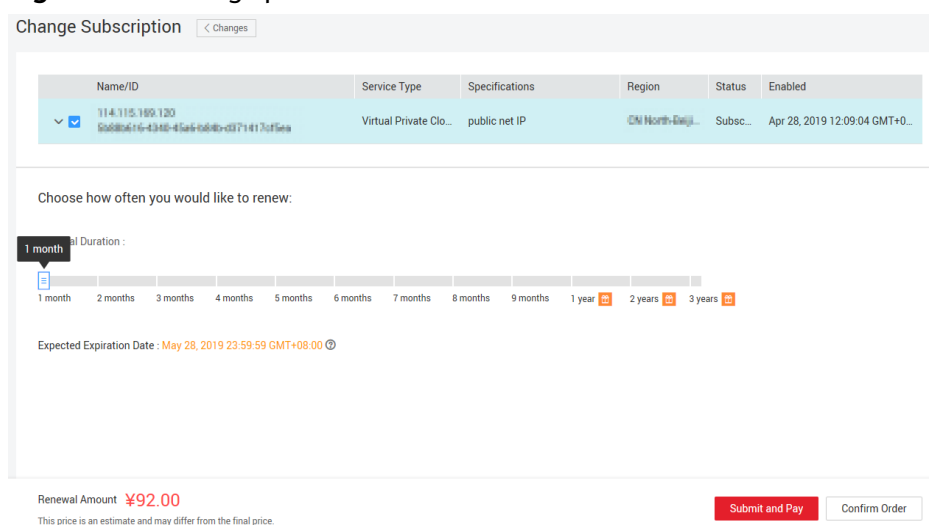
1. Log in to the management console.
2. Under **Network**, click **Elastic IP**.
3. On the displayed page, search for the pay-per-use EIP whose billing mode is to be changed.
4. Locate the row that contains the target EIP and click **Change Billing Mode** in the **Operation** column.

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



5. Click **Yes**.
6. Set specifications.

Figure 2-2 Setting specifications



7. Click **Submit and Pay**.

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


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

The billing mode of yearly/monthly EIPs and shared bandwidth can be changed to pay-per-use. The new billing mode takes effect only after the required duration of the yearly/monthly billing mode 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



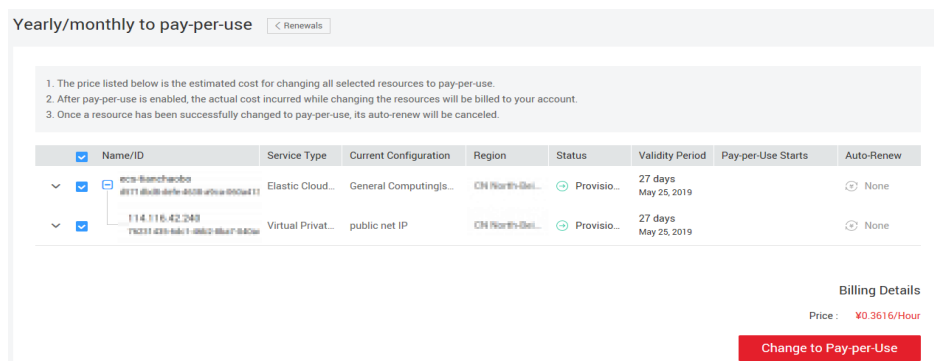
3. In the search box on the right, search for the EIP whose billing mode you want to change.
4. Locate the row that contains the target EIP and click **Change to Pay-per-Use After Expiration** in the **Operation** column.

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



The EIP billed by bandwidth on a yearly/monthly basis can only be changed to an EIP billed by bandwidth on a pay-per-use basis. After you change the billing mode of the EIP to pay-per-use, you can change the EIP billed by bandwidth to be billed by traffic. For details, see [Changing Bandwidth Billing](#).

 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 **Yearly/Monthly**. For details, see [Changing Bandwidth Billing](#).
- 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.

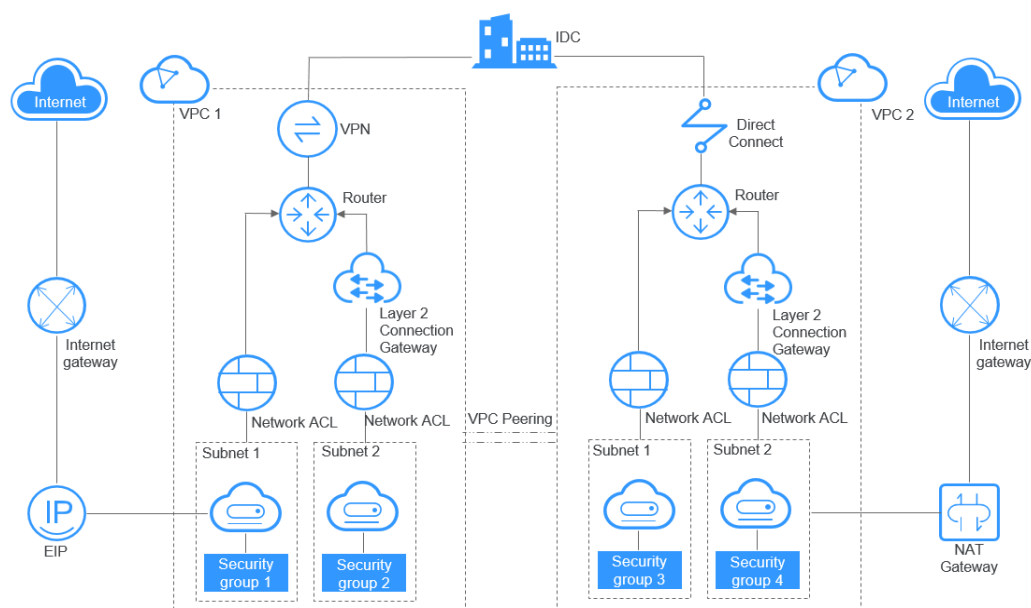
You can create security groups and VPNs, configure IP address ranges, and specify bandwidth sizes in your VPC. With a VPC, you can configure and manage the networks in the VPC, making changes to these networks as needed, quickly and securely. You can also define rules for communication between ECSs in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

The VPC service uses network virtualization technologies, such as link redundancy, distributed gateway clusters, and multi-AZ deployment, to ensure network security, stability, and availability.

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 a VPC is created, you need to 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 are used to ensure the security of cloud resources deployed in a VPC. A security group acts as a virtual firewall to provide access rules for cloud resources that have the same security protection requirements and are mutually trusted in a VPC. For more information, see [Security Group Overview](#). You can associate subnets that have the same traffic control requirements with the same network ACL. 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, Direct Connect, or Layer 2 Connection Gateway can connect your data center to VPCs.

3.2 Which CIDR Blocks Are Available to the VPC Service?

The VPC service supports the following CIDR blocks:

- 10.0.0.0 – 10.255.255.255
- 172.16.0.0 – 172.31.255.255
- 192.168.0.0 – 192.168.255.255

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 quota increase.

3.4 Can Subnets Communicate with Each Other?

Subnets in the same VPC can communicate with each other. 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.

3.5 What Subnet CIDR Blocks Are Available?

The subnet CIDR blocks must be included in the VPC CIDR blocks. The VPC CIDR blocks are **10.0.0.0/8–24**, **172.16.0.0/12–24**, and **192.168.0.0/16–24**. The subnet CIDR blocks must be within these CIDR blocks, and the allowed block size of a subnet is between the netmask of its VPC CIDR block and /29 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 account. If the number of subnets cannot meet your service requirements, [submit a service ticket](#) to request quota increase.

3.8 What Do I Do If a Subnet Cannot Be Deleted Because It 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 network 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 can view all resources of your account on the console homepage and check the resources that are in the subnet to be deleted. You must delete all resources in the subnet before deleting the subnet.

The resources may include:

- ECS
- BMS
- CCE cluster
- CCI instance
- RDS instance
- Workspace
- MRS cluster
- DCS instance
- Load balancer
- VPN
- Private IP address
- Custom route
- NAT gateway
- VPC endpoint and VPC endpoint service

3.9 What Do I Do If the ECS IP Address Is Lost After the ECS System Time Has Been Changed?

Cause: The failure occurs because the time difference between the old and new time is longer than your DHCP lease time. The default DHCP lease time set when you create a subnet on the public cloud system is 365 days (24 hours for subnet created before). If you manually change the ECS system time and the time difference between the old and new time is longer than 365 hours (24 hours for subnet created before), the DHCP lease fails to renew and the ECS IP address will be lost.

Solution: If you do need to change the ECS system time, and the time difference is longer than your DHCP release time, change the ECS IP address assignment mode to be static before changing the ECS system time.

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

Scenarios

After you change the DHCP lease time on the console, the existing ECSs will not use the new DHCP lease until the current lease needs to be renewed. A release is renewed when half of the lease time has passed. For example, if the lease is 30 days set on January 1, the lease will be renewed on January 15.

If you need to make the new DHCP lease time take effect immediately for ECSs in the subnet, refer to the following operations.

Precautions

If you renew the DHCP lease manually, the current IP addresses of ECSs will be released and they have no IP addresses until the new release takes effect and they are assigned with new IP addresses, which may cause service interruption.

You can also directly restart the ECSs to make the new DHCP release take effect immediately.

For a Windows ECS:

1. After you changed the DHCP lease time on the console, log in to the ECS whose lease is to be renewed.
2. Choose **Start > Run**. Type `cmd` to open the DOS operation GUI.
3. Run the `ipconfig /all` command to view the expiration time of the current DHCP lease.
4. Run the `ipconfig /release && ipconfig /renew` command to renew the lease. Run the `ipconfig /all` command again to view the result.

For a Linux ECS:

1. After you changed the DHCP lease time on the console, log in to the ECS whose lease is to be renewed.
2. Run the `ps -ef | grep dhclient` command to check whether the client that provides the DHCP service is **dhclient**. If the process in the following figure exists, the client is **dhclient**. The lease file whose path is next to the `-lf` parameter contains the lease information. If the **dhclient** process does not exist, this document may not be applicable. In this case, you need to search for the operation commands of the corresponding DHCP client.

```
(root@vpc8-subnet-as8-pod8-xen ~)# ps -ef | grep dhclient
root      537   537   0 Jun14  ?        00:00:00 sbin/dhclient -d -q -f /var/lib/isc/dhcp-helper -pf /var/run/dhclient-eth8
pid -lf /var/lib/NetworkManager/dhclient-5f866d8-8bb8-7ffb-45f1-d6cd65f3e83-eth8.lease -cf /var/lib/NetworkManager/dhclient-eth8.conf eth8
```

3. Run the `dhclient -r` command to release the current IP address.

4. Run the **dhclient** command to obtain the new DHCP lease. View the lease file mentioned in the preceding step. You can view the latest DHCP lease. The lease file also stores the historical lease. The last lease in the file is the latest lease.

```
root@vpc0-subnet0-az8-pod0-zen ~# cat /var/lib/NetworkManager/dhclient-5fb86bd8-8bb8-7ffb-45f1-d6edd65f3e83-eth0.lease
lease {
  interface "eth0";
  fixed-address 10.1.1.124;
  option subnet-mask 255.255.255.0;
  option dhcp-lease-time 864880;
  option routers 10.1.1.1;
  option dhcp-message-type 5;
  option dhcp-server-identifier 10.1.1.254;
  option domain-name-servers 10.1.1.254;
  option interface-mtu 1500;
  option dhcp-renewal-time 432880;
  option dhcp-rebinding-time 756880;
  option broadcast-address 10.1.1.255;
  option rfc3442-classless-static-routes 0,10.1.1.1,32,169.254.169.254,10.1.1.254;
  option host-name "host-10-1-1-124";
  renew 2 2019/06/10 14:09:18;
  rebind 0 2019/06/23 01:53:27;
  expire 1 2019/06/24 07:53:27;
}

lease {
  interface "eth0";
  fixed-address 10.1.1.124;
  option subnet-mask 255.255.255.0;
  option routers 10.1.1.1;
  option dhcp-lease-time 4294967295;
  option dhcp-message-type 5;
  option domain-name-servers 10.1.1.254;
  option dhcp-server-identifier 10.1.1.254;
  option interface-mtu 1500;
  option rfc3442-classless-static-routes 0,10.1.1.1,32,169.254.169.254,10.1.1.254;
  option broadcast-address 10.1.1.255;
  option host-name "host-10-1-1-124";
  renew 3 2007/07/02 11:32:40;
  rebind 3 2130/07/16 19:59:09;
  expire 1 2155/07/21 14:46:47;
}
```

3.11 Can I Change the VPC of an ECS?

Yes.

You can click **Change VPC** in the **Operation** column on the **Elastic Cloud Server** page. For details, see [Changing a VPC](#).

3.12 How Do I Switch to a Private DNS Server?

Private DNS servers are configured for VPC subnets by default. ECSs in the subnets can use private DNS servers to access internal addresses of other cloud services, such as OBS and SMN, without going through the Internet, and to request domain names on the Internet.

For VPCs created earlier, 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 ones 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?](#)

To switch the DNS server of an ECS to a private DNS server, you need to check the DNS server addresses of the ECS, change the DNS servers for the VPC subnet where the ECS resides, and update the DNS server addresses of ECS.

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.

6. 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 Server Addresses for a VPC Subnet

If the subnet of the ECS is not using a private DNS server address, you need to do as follows:

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

For example, in the CN North-Beijing1 region, you need to 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

After you change the DNS server addresses of the subnet, the DNS server addresses of the ECS are not updated immediately.

You can use either of the following methods to update the DNS server address:

- 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. Therefore, perform this operation during off-peak hours.

Alternatively, wait for the DHCP lease time to end, which lasts for 24 hours by default. The ECS will then update the IP address and DNS server addresses with the DHCP server.

- Manually change DNS configurations of the ECS.

If the DHCP function is disabled on the ECS, you need to manually update DNS configurations.

For example, in a Linux OS, change DNS configurations in 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. For details, see [EIP Overview](#).

A private IP address is used by the internal network of the public cloud for internal communications. It cannot be accessed over 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 is not allocated to an ECS NIC. 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 ~]# route -n
Kernel IP routing table
Destination      Gateway         Genmask        Flags Metric Ref    Use Iface
0.0.0.0          192.168.11.1   0.0.0.0        UG    0      0      0 eth0
169.254.0.0      0.0.0.0        255.255.0.0    U     1002   0      0 eth0
169.254.0.0      0.0.0.0        255.255.0.0    U     1003   0      0 eth1
169.254.169.254 192.168.11.1   255.255.255.255 UGH   0      0      0 eth0
192.168.11.0     0.0.0.0        255.255.255.0  U     0      0      0 eth0
192.168.17.0     0.0.0.0        255.255.255.0  U     0      0      0 eth1
[root@ecs-b926 ~]#
```

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

Figure 4-2 Viewing NIC information

```
[root@ecs-b926 ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.11.42 netmask 255.255.255.0 broadcast 192.168.11.255
    inet6 fe80::f816:3eff:fe7:1c44 prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:f7:1c:44 txqueuelen 1000 (Ethernet)
    RX packets 127 bytes 21633 (21.1 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 258 bytes 22412 (21.8 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
    inet 192.168.17.191 netmask 255.255.255.0 broadcast 192.168.17.255
    inet6 fe80::f816:3eff:fe1c:b57f prefixlen 64 scopeid 0x20<link>
    ether fa:16:3e:1c:b5:7f txqueuelen 1000 (Ethernet)
    RX packets 11 bytes 1283 (1.2 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 12 bytes 1388 (1.3 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

lo: flags=73<UP,LOOPBACK,RUNNING> mtu 65536
    inet 127.0.0.1 netmask 255.0.0.0
    inet6 ::1 prefixlen 128 scopeid 0x10<host>
    loop txqueuelen 1 (Local Loopback)
    RX packets 51 bytes 12818 (11.7 KiB)
    RX errors 0 dropped 0 overruns 0 frame 0
    TX packets 51 bytes 12818 (11.7 KiB)
    TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
```

3. Configure the extension NIC to 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

 NOTE

This operation will interrupt ECS communication. Exercise caution when performing this operation. It is recommended that you perform the configuration by following step 4.

- b. 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 network segment (for example, *xx.xx.0.0/16*) through the extension NIC:

You can configure the network segment 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 What Are EIPs?

An EIP is a static public IP address that can be accessed over the Internet. Bandwidth will be assigned together with an EIP. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers. Various billing modes are provided to meet diverse service requirements.

Each EIP can be used by only one cloud resource at a time.

Figure 4-3 Accessing the Internet using an EIP

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

Yes. A pay-per-use EIP that uses the dedicated bandwidth can be changed to use the shared bandwidth. However, a yearly/monthly EIP that uses the dedicated bandwidth cannot be changed to use the shared bandwidth.

4.7 How Many ECSs Can One EIP Be Assigned to?

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

Multiple ECSs can share the same EIP, but the EIP and the ECSs must be in the same region. To enable ECSs across AZs in a VPC to share an EIP, you can use a NAT gateway. For more information, see [NAT Gateway User Guide](#).

4.8 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.

4.9 What Is the EIP Assignment Policy?

By default, an EIP is assigned randomly. If you have released EIPs, the system preferentially assigns EIPs from the ones you released in the last 24 hours.

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

Yes.

Unbind the EIP from the ECS. For details, see [Unbinding or Releasing an EIP](#).

Then, bind the EIP to the target ECS. For details, see [Binding an EIP to Cloud Resources](#).

4.11 Will the EIP Bound to an ECS Be Changed After the ECS Is Stopped and Then Started?

The EIP bound to an ECS will not be changed after the ECS is stopped and then started.

Stopping and starting an ECS does not affect its EIP.

4.12 Can I Assign a Specific EIP?

By default, an EIP is assigned randomly. If you used to release EIPs, the system preferentially assigns an EIP from what you released.

To assign a specific EIP, you can call APIs. For details, see [Assigning an EIP](#).

4.13 Will an EIP Be Changed After I Assign It?

EIPs will not change after they are assigned.

4.14 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 IP address you searched through a third-party website may be different from the actual region to which the IP address belongs.
- If the region searched through a third-party website is different from that searched on <https://en.ipip.net/?origin=CN>, use the region searched on <https://en.ipip.net/?origin=CN>.
- If the region searched on <https://en.ipip.net/?origin=CN> is different from the region you selected when purchasing the EIP, use the region you selected when purchasing the EIP.
- If your service is adversely affected due to the inconsistency between the region of the EIP you searched on the third-party database and the region of the EIP you selected when buying the EIP, [submit a service ticket](#).

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

4.15 Can I Transfer an EIP to Another Account?

EIPs **cannot** be transferred across accounts.

4.16 Can an EIP Be Bound to a Cloud Resource in Another Region?

No. EIPs can only be bound to cloud resources 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. If the bandwidth in the outbound direction is less than 100 Mbit/s, the bandwidth in the inbound direction will be 100 Mbit/s. If the bandwidth in the outbound direction is greater than 100 Mbit/s, the bandwidth in the inbound direction will be the same as that in the outbound direction.

 NOTE

- If you have selected the enhanced 95th percentile bandwidth option, the bandwidth will be billed based on the average bandwidth in the inbound and outbound directions.
- Inbound bandwidth limits of EIPs purchased after July 31, 2020 00:00:00 GMT+08:00 in Chinese mainland regions are adjusted as follows:
 - If your purchased bandwidth is less than or equal to 10 Mbit/s, the bandwidth in the inbound direction will be 10 Mbit/s, and the bandwidth in the outbound direction will be the same as the purchased bandwidth.
 - If your purchased bandwidth is greater than 10 Mbit/s, the bandwidth both in the outbound and inbound directions will be the same as the purchased bandwidth.

5.2 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 and a NAT gateway.

Shared bandwidth: The bandwidth can be shared by multiple EIPs. You can add multiple pay-per-use EIPs to the bandwidth. Adding an EIP to or removing an EIP from a shared bandwidth does not affect running services.

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

- After an EIP is added to a shared bandwidth, the EIP will use the shared bandwidth.
- After an EIP is removed from the shared bandwidth, the EIP will use the dedicated bandwidth.

5.3 How Do I Check Whether the Bandwidth Exceeds the Limit?

Symptom

The bandwidth size configured when buying a dedicated or shared bandwidth is the upper limit of the outbound bandwidth. If the web applications that depend on the Internet freeze, check whether the dedicated bandwidth of the EIP bound to the ECS is greater than the configured bandwidth size.

 NOTE


If the bandwidth exceeds the configured bandwidth size, packet loss may occur. To ensure normal service running, it is recommended that you monitor the bandwidth.

Method 1

Query the historical dedicated bandwidth usage of an EIP.

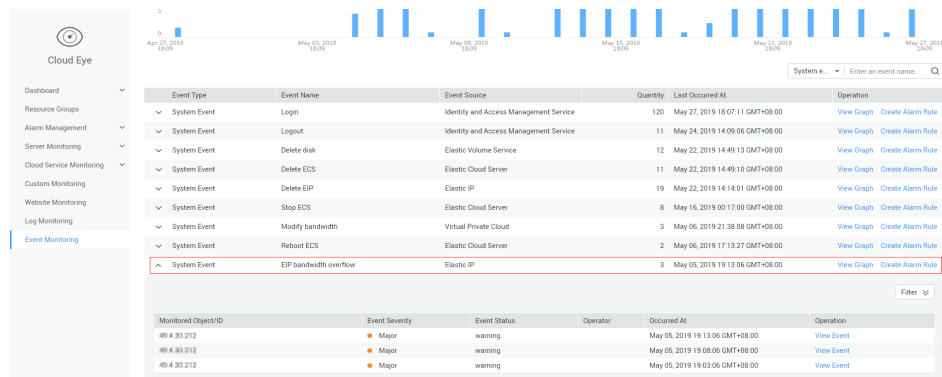
1. Log in to the management console and view the EIP bound to the ECS.

Figure 5-2 Viewing the EIP



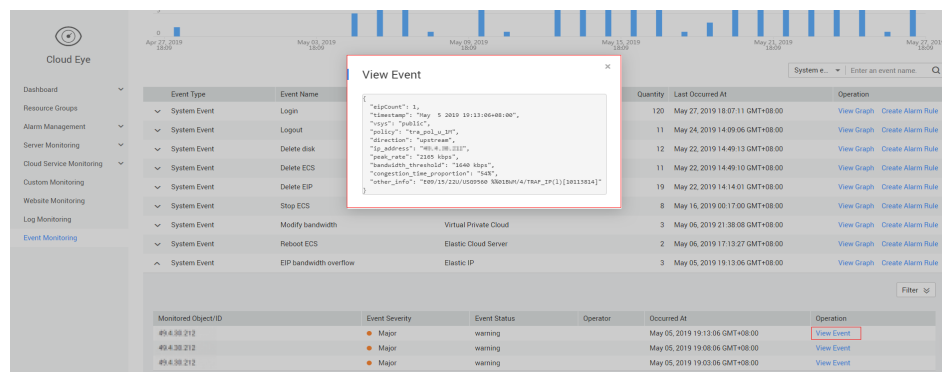
2. On the **Cloud Eye** console, click **Event Monitoring**.

Figure 5-3 Event Monitoring



3. Click **View Event**.

Figure 5-4 View Event



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

If the event **EIP bandwidth overflow** is displayed, the dedicated bandwidth of the EIP has exceeded the limit. To ensure stable service running, increase the bandwidth. For details, see [Modifying EIP Bandwidth](#).

NOTE


- EIP bandwidth overflow is only launched in regions **CN North-Beijing1**, **CN East-Shanghai2**, and **CN South-Guangzhou** regions.
- Cloud Eye does not display the shared bandwidth details on the **Event Monitoring** page. For details, see *Cloud Eye User Guide*.

Method 2

Create an alarm rule to generate alarms when the bandwidth exceeds the limit.

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

Figure 5-5 Alarm Rules



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

Figure 5-6 Create Alarm Rule

Specify Rule Name

* Name: alarm-6etk

Description: [Empty text area]

Alarm Content

* Resource Type: Elastic IP and Bandwid...

* Dimension: Bandwidths (selected), Elastic IPs

* Monitoring Scope: Resource groups, Specific resources (selected)

When the actual bandwidth exceeds the limit, an alarm is generated and the system automatically sends a notification.

The basic alarm function is free of charge. Simple Message Notification (SMN) sends you alarm notifications, and this service will be billed. For details, see *Cloud Eye User Guide*.

5.4 What Are the Differences Between EIP Bandwidth and Private Network Bandwidth?

The EIP bandwidth is used by ECSs to access the Internet through EIPs. The EIP bandwidth displays network resource usage and can be used for service metering.

The private network bandwidth is used for ECS communication in the public cloud system. The private network bandwidth and PPS of an ECS are determined based on ECS specifications. For details, see [ECS Types](#).


5.5 What Is the Bandwidth Size Range?

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

5.6 What Bandwidth Types Does the VPC Service Support?

The VPC service supports dedicated bandwidth and shared bandwidth. The dedicated bandwidth can only be used by one EIP, whereas the shared bandwidth can be used by multiple EIPs.

5.7 How Can I Use Shared Bandwidth?

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

5.8 How Many EIPs Can Use the Same Shared Bandwidth?

A shared bandwidth can be used by up to 20 EIPs. If you want to add more EIPs to the same shared bandwidth, [submit a service ticket to](#) request quota increase.

5.9 Can I Increase My Yearly/Monthly Bandwidth Size and Then Decrease It?

The increased bandwidth size takes effect immediately. The decreased bandwidth size takes effect in the subsequent billing cycle after a successful renewal. For details, see [Modifying EIP Bandwidth](#).

5.10 Can I Use Multiple Bandwidth Add-On Packages with Overlapping Validity Periods?

No.

A bandwidth add-on package is used to temporarily increase the maximum shared or dedicated bandwidth of a yearly/monthly EIP. You can buy multiple bandwidth

add-on packages for the same bandwidth, but the validity periods of the packages cannot overlap. For details, see [Bandwidth Add-On Package Overview](#) and [Buying a Bandwidth Add-On Package](#).

5.11 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.

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

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

Table 5-1 Differences between static BGP and dynamic BGP

Aspect	Static BGP	Dynamic BGP
Definition	Routes are manually configured by carriers. If the network topology or link status changes, carriers must manually modify static routes in the route table.	When changes occur on a network using dynamic BGP, routing protocols provide automatic, real-time optimization of network configurations, ensuring network stability and optimal user experience.
Assurance	When changes occur on a network using static BGP, carriers cannot adjust network configurations in real time to ensure optimal user experience. NOTE If you select static BGP, your application system must have the disaster recovery capability.	The multi-path BGP can detect the status of the access line and carrier's internal networks. When an internal fault occurs on the carrier network, another network will take over services quickly to ensure service availability.
Service availability	99%	99.95%

 NOTE

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

5.13 What Is Enhanced 95th Percentile Bandwidth Billing?

The enhanced 95th percentile bandwidth billing mode allows you to use more bandwidth after you pay for the baseline bandwidth. You are billed based on the required duration and the bandwidth volume after eliminating some top sampled usage from a billing period.

Prerequisite

To use the enhanced 95th percentile billing mode, the following requirements must be met:

- Your level is greater than or equal to V4.
- You can select this billing mode when purchasing a shared bandwidth.
- The minimum bandwidth you can purchase is 300 Mbit/s.

Pricing Details

Billing formula: Monthly peak bandwidth x Monthly peak bandwidth price x Shared bandwidth in-use days/Calendar days of a month.

Billing cycle: Bills are generated for each calendar month.

 NOTE

Billing mode: The enhanced 95th percentile bandwidth is billed on a pay-per-use basis, and does not require prepayment. The monthly fee is settled at the end of each calendar month.

Monthly peak bandwidth price: The price is preset and does not vary according to the number of days in a calendar month.

Shared bandwidth in-use days: Calculated based on the actual duration when the shared bandwidth is used. For example, if you apply for a shared bandwidth at 12:00, the shared bandwidth use time is half a day.

 NOTE

Shared bandwidth in-use days = Number of collected bandwidth values in a month/288.

Monthly peak bandwidth: The monthly peak bandwidth is calculated based on the enhanced 95 percentile billing mode in which some peak bandwidth values discarded. It must be higher than the monthly baseline bandwidth.

Daily baseline bandwidth: Daily baseline bandwidth = Baseline percentage x Shared bandwidth. The baseline percentage is 20%.

In enhanced 95 percentile billing mode, the shared bandwidth can be adjusted in real time, and the adjustment takes effect immediately. After the shared bandwidth is adjusted, the baseline bandwidth changes accordingly. The daily

baseline bandwidth is calculated based on the maximum baseline bandwidth set for a day. For example, if the bandwidth is adjusted from 100 Mbit/s to 300 Mbit/s and then to 200 Mbit/s during a day, the daily baseline bandwidth is 60 Mbit/s (300 x 20%).


Monthly baseline bandwidth: The bandwidth can be adjusted frequently to suit the traffic requirements during a month.

The monthly baseline bandwidth is calculated based on the following formula (only the integer is retained in the calculated result): Monthly baseline bandwidth = (Baseline bandwidth 1 x Number of days using baseline bandwidth 1 + Baseline bandwidth 2 x Number of days using baseline bandwidth 2 + ... + Baseline bandwidth n x Number of days using baseline bandwidth n)/Number of days using all baseline bandwidths in a month.

Monthly peak bandwidth calculation method

- Daily peak bandwidth
 - Inbound and outbound traffic data is collected every five minutes.
 - The averages of both inbound bandwidth and outbound bandwidth within five minutes are calculated, and the larger one is used as the bandwidth for that collection.
 - After all meter readings within a day are obtained, they are sequenced in descending order. The top four peak bandwidth values are discarded, and the fifth peak bandwidth value is used as the daily peak bandwidth.

Figure 5-7 Daily peak bandwidth




NOTE

If less than five peak bandwidth values are obtained in a day, the last value in the peak bandwidth sequence is used as the daily peak bandwidth. The daily peak bandwidth is an integer (any fractional parts are discarded).

- Monthly peak bandwidth

At the end of each month, the daily peak bandwidth values are sequenced in descending order. The average of the top five daily peak bandwidth values is the monthly peak bandwidth (only the integer is retained and fractional parts are discarded).

Figure 5-8 Monthly peak bandwidth



- The daily peak bandwidth values of a month are sequenced in descending order.
- The average of the top five daily peak bandwidth values is the monthly peak bandwidth.

NOTE

If less than five daily peak bandwidth values are obtained, the average of all the daily peak bandwidth values in the month is the monthly peak bandwidth. The monthly peak bandwidth is an integer (any fractional parts are discarded).

The larger value between the baseline bandwidth of the month and the average of daily peak bandwidth values is used as the monthly peak bandwidth. If the baseline bandwidth of a month is greater than the average of the daily peak bandwidth values, the monthly peak bandwidth is equal to the baseline bandwidth of the month. Otherwise, the monthly peak bandwidth is equal to the average of daily peak bandwidth values.

6 Connectivity

6.1 Does a VPN Allow for 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 Cannot I Access Public Websites Through Domain Names or Access Internal Domain Names in the Cloud When My ECS Has Multiple NICs?

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

You can rectify this fault 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 Limitations of 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 the same two VPCs at the same time.
- You cannot create a VPC peering connection between VPCs in different regions.
- You cannot use the EIPs in a VPC of a VPC peering connection to access resources in the other VPC. For example, VPC A is peered with VPC B, VPC B has EIPs that can be used to access the Internet, you cannot use EIPs in VPC B to access the Internet from VPC A.
- To request a VPC peering connection with a VPC of another tenant, the peer tenant 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 to activate 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, routes with the same destinations cannot be added in VPC A.
- To ensure security, do not accept VPC peering connections from unknown tenants.
- 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 be automatically deleted immediately, including routes added for the VPC peering connection.
- Currently, the route table of a VPC takes effect for all subnets in the VPC. You cannot add a route table dedicated for a specific subnet. The route preference is as follows: direct route > VPC peering connection route > custom route.
- If two VPCs in a VPC peering connection have overlapping CIDR blocks, the peering connection can only enable communication between two subnets in the two VPCs. If subnets in the two VPCs of a VPC peering connection have overlapping CIDR blocks, the peering connection will not take effect. When creating a VPC peering connection, ensure that the VPCs involved do not contain overlapping subnets.
- You cannot delete a VPC for which VPC peering connection routes have been configured.

6.4 What Do I Do If VPCs in a VPC Peering Connection Cannot Communicate with Each Other?

1. Check whether a VPC peering connection has been successfully created for the two VPCs. Confirm the IDs of the VPCs in the VPC peering connection.

2. Check whether routes that point to the CIDR block (or portion of the CIDR block) of the other VPC have been configured.
3. Check whether routes configured for the VPC peering connection are correct. If VPCs in a VPC peering connections have overlapping CIDR blocks, you can only add routes to enable communication between two subnets in the two VPCs.
4. Check whether the VPCs in the VPC peering connection contain overlapping subnets.
5. Check whether required security group rules have been configured for the ECSs that need to communicate with each other and whether restriction rules have been added to the iptables or firewall used by the ECSs.
6. If a message indicating that this route already exists is displayed when you add routes for a VPC peering connection, check whether the route's destination IP addresses of the VPN, Direct Connect, and VPC peering connections already exist.
7. If the route's destination IP addresses of a VPC peering connection overlap with those of a Direct Connect or VPN connection, the route may be invalid.
8. Check whether the ECS has multiple NICs. If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure a policy-based route in the ECS. For details, see [How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?](#)
9. If VPCs in a VPC peering connection cannot communicate with each other after all these possible faults have been rectified, contact customer service.

6.5 How Many VPC Peering Connections Can I Have?

A tenant can have a maximum of 50 VPC peering connections in one region. Accepted VPC peering connections consume the quota of both owners of a VPC peering connection. A VPC peering connection consumes the quota of only the requester (tenant of the local VPC).

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.

6.7 What Do I Do If Intermittent Interruption Occurs When a Local Host Accesses a Website Built on an ECS?

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

Fault Locating

1. Check the local network of the user.

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

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

```
sysctl -a|grep tcp_tw_recycle
```

The value of **tcp_tw_recycle** is **1**, indicating 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 faulty.

Troubleshooting 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 What Do I Do If 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.

The ECS **ecs02** private IP address can be pinged from ECS **ecs01**, but ECS **ecs01** private IP address cannot be pinged from ECS **ecs02**.

Fault Locating

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**, the ECS **ecs02** fails to learn the ECS **ecs01** MAC address 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 6-1 Viewing ECS **ecs01** NIC configuration

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

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

Troubleshooting 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 What Do I Do If Two ECSs in the Same VPC Cannot Communicate with Each Other or Packet Loss Occurs When They Communicate?

Fault Locating

1. Check security group rules.
2. Check network ACLs.
3. Check the NIC information of ECSs.
4. Check the disconnected ports.

Troubleshooting Procedure

1. Check 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 any of the following rules.

Figure 6-2 Inbound security group rule

Transfer Direction	Type	Protocol	Port Range/ICMP Type	Remote End	Operation
Inbound	IPv4	Any	Any	0.0.0.0/0 ?	Delete
Inbound	IPv4	ICMP	Any	0.0.0.0/0 ?	Delete

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

2. Check network ACLs.
 - a. Check whether the ECS NIC is in the associated subnet of the network ACL.
 - b. Check the network ACL status in the network ACL list.
 - If **Disable** is displayed in the **Operation** column, the network ACL has been enabled. Go to [2.c](#).
 - If **Enable** is displayed in the **Operation** column, the network ACL has been disabled. Go to [2.d](#).
 - c. Click the network ACL name and configure rules on the **Inbound** and **Outbound** tabs to allow the ICMP traffic.
 - d. When 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.
3. Check the NIC information of the ECS. (The following procedure uses a Linux ECS as an example. For a Windows ECS, check the firewall restrictions.)
 - a. Check whether multiple NICs are configured for the ECS. If the ECS has multiple NICs and the EIP is bound to an extension NIC, configure policy-based routing for the ECS. For details, see [How Do I Configure Policy-Based Routing for ECSs with Multiple NICs?](#)
 - b. 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 6-3 NIC IP address

```
(root@ecs-acl ~)# ifconfig
eth8      Link encap:Ethernet  HWaddr FA:16:3E:BC:B7:81
          inet addr:192.168.72.289  Bcast:192.168.72.255  Mask:255.255.0
          inet6 addr: fe80::f016:3eff:febc:b701/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:881 errors:0 dropped:0 overruns:0 frame:0
          TX packets:547 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:49684 (48.4 KiB)  TX bytes:44454 (43.4 KiB)
          Interrupt:46
```

- c. 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

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

iptables-save

- e. 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.

4. Check the disconnected ports.
 - a. If the special port of the ECS cannot be accessed, check whether the security group rules and network ACL rules enable the port.
 - b. On the Linux ECS, run the following command to check whether the ECS listens to the port: If the port is not listened, the ECS communication may be adversely affected.
netstat -na | grep <Port number>

6.10 What Do I Do If a Virtual IP Address Cannot Be Pinged After It Is Bound to an ECS NIC?

Fault Locating

1. Check whether the source/destination check function of the NIC is disabled and whether the virtual IP address has been bound to the NIC.
2. Check whether the ECS NIC subinterfaces are successfully created.
3. Check whether the ECS security groups and the network ACL rules associated with the subnets used by the ECS NICs block traffic.

Troubleshooting Procedure

1. Check whether the source/destination check function of the NIC is disabled and whether the virtual IP address has been bound to the NIC.
 - a. Log in to the management console.
 - b. Click **Service List** and click **Elastic Cloud Server** under **Computing**.
 - c. In the ECS list, click the name of the target ECS.
 - d. On the displayed ECS details page, click the **NICs** tab.
 - e. Check that **Source/Destination Check** is disabled.
 - f. 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.
2. Check whether an ECS has been configured with a virtual IP address.

This following uses Linux and Windows ECSs as examples to describe how to check whether an ECS has been configured with a virtual IP address.

For a Linux ECS:

- a. Run the following command on the ECS and check whether NIC **ethX:X** exists:
ifconfig

Figure 6-4 Checking for NIC **ethX:X**

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

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

The command output in the preceding figure contains the NIC of the **ethX:X** type. **192.168.1.137** is the virtual IP address in **1**.

- If yes, the sub-interface of the ECS NIC has been created properly.
 - If no, perform the following operations:
- b. If the command output does not contain the NIC of the **ethX:X** type, run the following command to switch to the **/etc/sysconfig/network-scripts** directory:

```
cd /etc/sysconfig/network-scripts
```

- c. 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:

Figure 6-5 Adding NIC information

```
BOOTPROTO=static
DEVICE=eth0:1
HWADDR=fa:16:3e:4d:5b:98
IPADDR=192.168.1.137
GATEWAY=192.168.1.1
NETMASK=255.255.255.0
ONBOOT=yes
ONPARENT=yes
```

- d. Press **Esc**, enter **:wq!**, and save the file and exit.
- e. 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:

- a. 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 6-6 Checking whether the virtual IP address has been configured

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

Windows IP Configuration

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

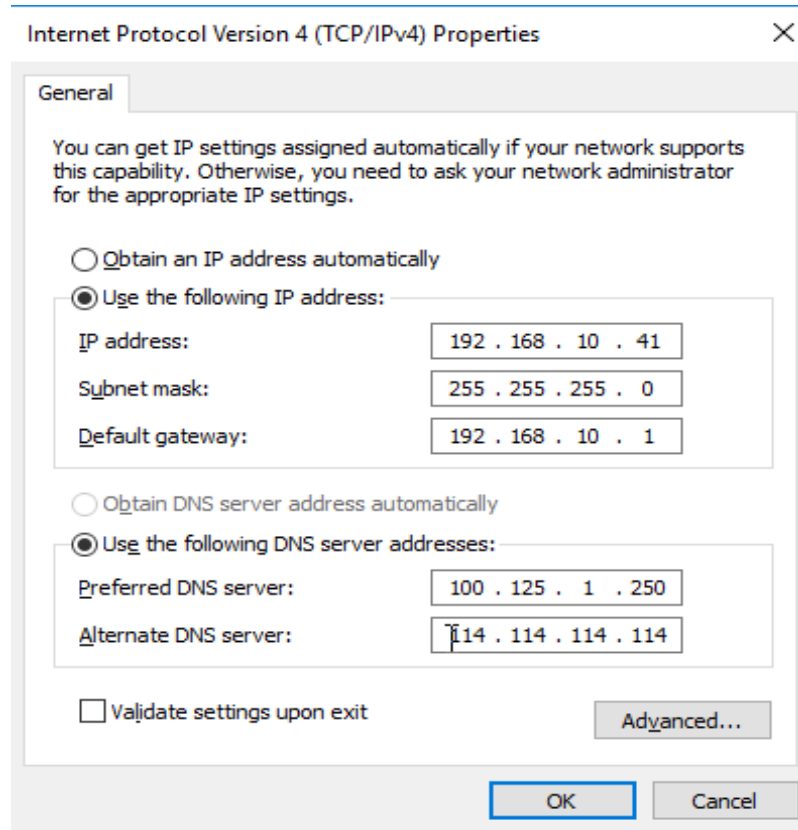
Ethernet adapter Ethernet 5:

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

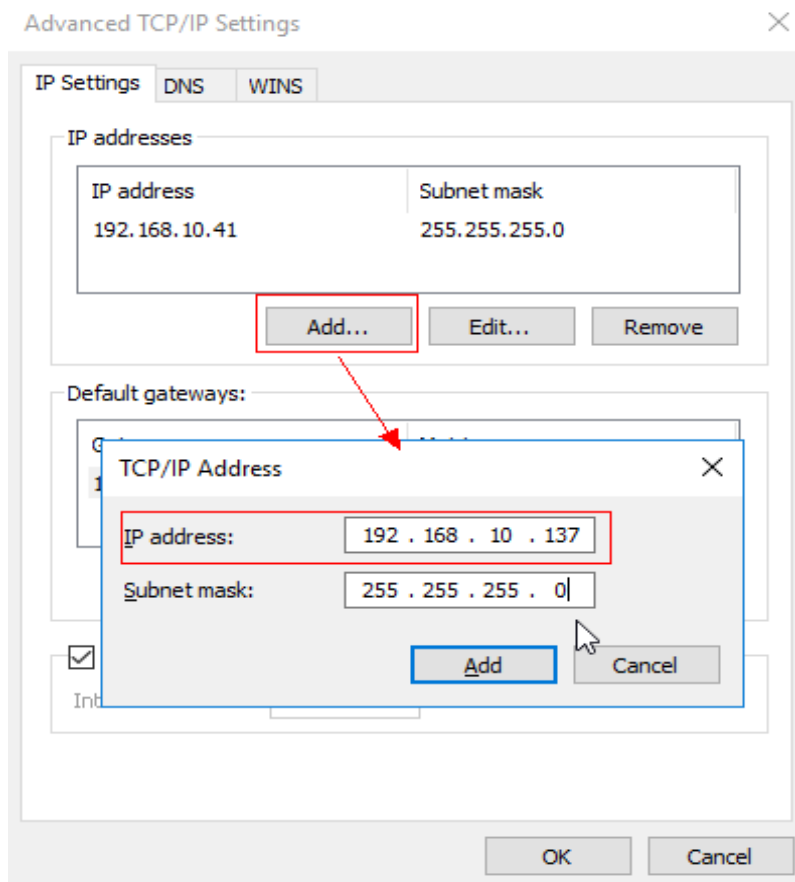
In the preceding command output, check whether the value of **IPv4 Address** is the virtual IP address 192.168.10.137 in **1**.

- If yes, the virtual IP address has been configured for the ECS NIC.
- If no, perform the following operations:
 - b. Choose **Control Panel > Network and Internet > Network Connections**. Right-click the corresponding local connection and then click **Properties**.
 - c. On the **Network** tab page, select **Internet Protocol Version 4 (TCP/IPv4)**.
 - d. Click **Properties**.
 - e. Select **Use the following IP address**, and set **IP address** to the private IP address displayed in **Figure 6-6**. For example, 192.168.10.41.

Figure 6-7 Configuring private IP address



- f. Click **Advanced**.
- g. On the **IP Settings** tab, click **Add** in the **IP addresses** area. Add the virtual IP address configured in **1**. For example, 192.168.10.137.

Figure 6-8 Configuring virtual IP address

3. Check whether the ECS security groups and the network ACL rules associated with the subnets used by the ECS NICs block traffic.
 - a. 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.
 - b. Click **Service List**. Under **Network**, click **Virtual Private Cloud**. In the navigation pane on the left of the network console, click **Network ACL** and check whether the network ACL rules associated with the subnets used by the ECS NICs block access to the virtual IP address.

6.11 How Do I Handle the Cloud-init Connection Failure?

Cloud-init Network

Figure 6-9 shows the process for an ECS to obtain metadata using the cloud-init.

Figure 6-9 Process for obtaining metadata



Troubleshooting Procedure

1. 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 6-10 ECS IP address

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

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

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

2. Ping IP address **169.254.169.254/32** from the ECS. If the IP address cannot be pinged, perform the following steps:
 - a. 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 6-11 Route for IP address **169.254.169.254/32**

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

- b. 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.
 - c. 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 the cloud-init function is enabled, run the **service network restart** command to obtain the correct route.

- If the ECS was created after the cloud-init function is enabled, go to step 6.
3. The cloud-init connection failure may also be caused by the metadata obtaining fault on the ECS.

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

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

If information similar to that shown in [Figure 6-12](#) is displayed, the ECS successfully obtains the metadata.

Figure 6-12 Command output

```
-bash-4.1# curl http://169.254.169.254/openstack/latest/meta_data.json
{"random_seed": "rTVrsD1Eh6AjUKLnvg51U8S0pH6xC70MFRTelW10munBNyqos6q/EsAEJondF8iJkMDG0TzbCTb815HntS9X
XHu61u+y8fAeylxAj60Ae8KHMPgDu6Xdfhku6gyjCrjXn5hUFvqfZ/yaJ3LrAEjB8Nj59hI+umbPi8oYc2WzYmTqWjXYRNwpmqJM
s1KYm0CLuFbwYo2aK1y27WEUZDU0Q1GpRkkaWwFaCN/rQQ/hHd+3UwSJbArsgUeolWCTp5oxixLiCJzSSHAKz41UiziRuxYam0go
iTFtopvZTwmYEk1FmkZsy7h6PPDkgmJgPn+1kZf0qqhtlvpjRrZpw4aPAeZa4z7QX1RtmwT7MjyGUbea85/1PDUE1J/GJpoH1/+z
rDye1A09CsDG1UFuELadYDcrW44k42f0o7dDmEgDmInnEBeega5r7Eohb04KTimzi+3nb10QjPq/S7J+mFMUoZEJHObZE4u1Ajj
Znhwv/pc6ho7fQKbx0C78fbipH59CKyF0W835mNJ/CZNNBTA3UdG25SQ70IFnA+NtbDeo8+g05iFLvweu0G5Blcjm1fjh9+mqot
+5ae6ZcexUs1ifscqm0jwCnCimthJLYGmbxu+6Fm9XpLDopDPrRtBUcRSntIK67JprBSRppc+4smJyiuKY1JOTUJYQYDBU2B7F3o
=", "uid": "53ebb737-ddc5-4303-9fac-aa72b00b101a", "availability_zone": "eu-de-02", "hostname": "ecs-
g-jm-55eb.nova.local", "launch_index": 0, "meta": {"metering.image_id": "98721f93-722f-4386-a975-3cb
df1abf56d", "metering.imagetype": "gold", "metering.resourcecode": "c2.large.oracle", "metering.
cloudServiceType": "sys.service.type.ec2", "image_name": "autoC_OTC_OEL_6.0", "metering.resourceType
": "1", "os_bit": "64", "vpc_id": "120b71c7-94ac-45b8-8ed6-30aafc8fbdba", "os_type": "Linux", "charg
ing_mode": "0"}, "project_id": "efdf974f549b4eaab05c3903ddd2ab0e", "name": "ecs-g-jm-55eb"}-bash-4.1#
```

4. If you cannot log in to the ECS or cannot create a non-root user after cloud-init is configured and when the service is running properly, the cause may be as follows:

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 6-13 Configuration file

```
system_info:
# This will affect which distro class gets used
distro: rhel
# Default user name + that default users groups (if added/used)
default_user:
name: linux // Specifies the username for login.
lock_passwd: False // The value False indicates that the password login mode is enabled. For some OSs, the value 0 indicates that the password login mode is enabled.
gecos: Cloud User
groups: users // Specifies whether the user will be added to a group. This parameter is optional. The groups parameter value must be an existing group under /etc/group in the system.
passwd: $6$I63DBVXK$Zh41chiJR7NuZvtJHsYDQJIg5RoQCRL51X2Hsgj2s5JwXI7XU01we8WYcwbza52VHpRa#028vmxcCyU6LwoD0
sudo: ["ALL=(ALL) NOPASSWD:ALL"] // Specifies that all permissions of user root will be granted to the user.
shell: /bin/bash // Specifies that the bash shell is used.
# Other config here will be given to the distro class and/or path classes
paths:
cloud_dir: /var/lib/cloud/
templates_dir: /etc/cloud/templates/
ssh_svcname: sshd
```

5. If you cannot use an obtained private key to log in to an ECS after the ECS starts or you cannot obtain the ECS login password, restart the ECS to rectify the fault.
6. If you still cannot use the cloud-init function after the preceding steps, contact technical support for assistance.

You need to provide the technical support engineer with the following information.


Item	Description	Example	Your Value
VPC CIDR block	Required for customer gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	Example: 120b71c7-94ac-45b8-8ed6-30aafc8fbdba	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

6.12 How Do I Handle EIP Connection Failure?

EIP Network

[Figure 6-14](#) shows the process for an ECS to access the Internet using an EIP.

Figure 6-14 EIP network



Fault Locating

The possible causes to the EIP connection failure are as follows:

- The ECS is not running properly.
- The internal network configuration of the ECS is incorrect.
- No EIP is bound to the ECS.
- The EIP is not bound to the primary NIC of the ECS.
- Required security group rules are not configured for the ECS.
- Required packets are discarded by the firewall.

Troubleshooting Procedure

1. Check whether the ECS is running properly.
If the ECS state is not **Running**, start or restart the ECS.

Figure 6-15 ECS status

Name ID	AZ	Status	Specifications/Image	Private IP Address	EIP	Operation
ecs-gim-55eb 53eb0737-d4d5-4303-9fac-aa72b00	eu-de-02	Running	2 vCPUs 4 GB AutoC_OTC_OEL_6.8	192.168.1.200	-	Remote Login More

2. Check the ECS internal network configuration.
 - a. Confirm that 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.
 The **ipconfig** command applies only to Windows ECSs.
 - b. Check whether the 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 192.168.1.192/24 eth0** command to configure IP address **192.168.1.192/24** for the NIC.

Figure 6-16 NIC virtual IP address

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

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

Figure 6-17 Default route

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

3. Check whether the EIP has been assigned and bound to the ECS. (If the EIP has not been assigned or bound to the ECS, assign an EIP and bind it to the ECS.)

Figure 6-18 EIP status

Name ID	AZ	Status	Specifications/Image	Private IP Address	EIP	Operation
ecs-gm-f55eb 53ebb737-ddc5-4303-9fac-aa72b001	eu-nl-02	Running	2 vCPUs 4 GB AutoC_OTC_OEL_6.8	192.168.1.200	-	Remote Login More

4. Check whether the EIP is bound to the primary NIC of the ECS.

Figure 6-19 EIP binding status

Name ID	AZ	Status	Specifications/Image	Private IP Address	EIP	Operation
ecs-gm-f55eb 53ebb737-ddc5-4303-9fac-aa72b001	eu-nl-02	Running	2 vCPUs 4 GB AutoC_OTC_OEL_6.8	192.168.1.200	192.168.1.1	Remote Login More

5. Check whether required security group rules have been configured.

Configure security group rules 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.)

6. Check whether traffic filtering has been configured on the firewall for the subnet used by the ECS NIC.

If you can configure the firewall on the VPC console, confirm that the firewall rules allow traffic from the subnet used by the ECS to pass through.

7. Contact technical support.

If you still cannot properly use the EIP after the preceding steps, contact technical support for assistance.

You need to provide the technical support engineer with the following information.

Item	Description	Example	Your Value
VPC CIDR block	Required for customer gateway configuration	Example: 10.0.0.0/16	N/A
VPC ID	N/A	Example: 120b71c7-94ac-45b8-8ed6-30aafc8fbdba	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 customer ECS to access the Internet	Example: 10.154.55.175	N/A
EIP bandwidth	Maximum bandwidth size used by the customer 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

1. 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 6-20 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 6-21 Checking the firewall

```
[root@test2 ~]# service firewalld status
Redirecting to /bin/systemctl status firewalld.service
• firewalld.service - firewalld - dynamic firewall daemon
  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 ExecStart=/usr/sbin/firewalld --nofork --nopid $FIREWALLD_ARGS (code=exited, status=0/SUCCESS)
  Main PID: 861 (code=exited, status=0/SUCCESS)

Jan 02 06:04:39 ecs-g00200264-h2-0002.novalocal systemd[1]: Starting firewalld - dynamic firewall daemon...
Jan 02 06:04:39 ecs-g00200264-h2-0002.novalocal systemd[1]: Started firewalld - dynamic firewall daemon.
Jan 02 20:27:35 test2 systemd[1]: Stopping firewalld - dynamic firewall daemon...
Jan 02 20:27:36 test2 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, it is found that no IP address has been assigned to the ECS InfiniBand (IB) port.

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

Figure 6-22 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.

2. 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.

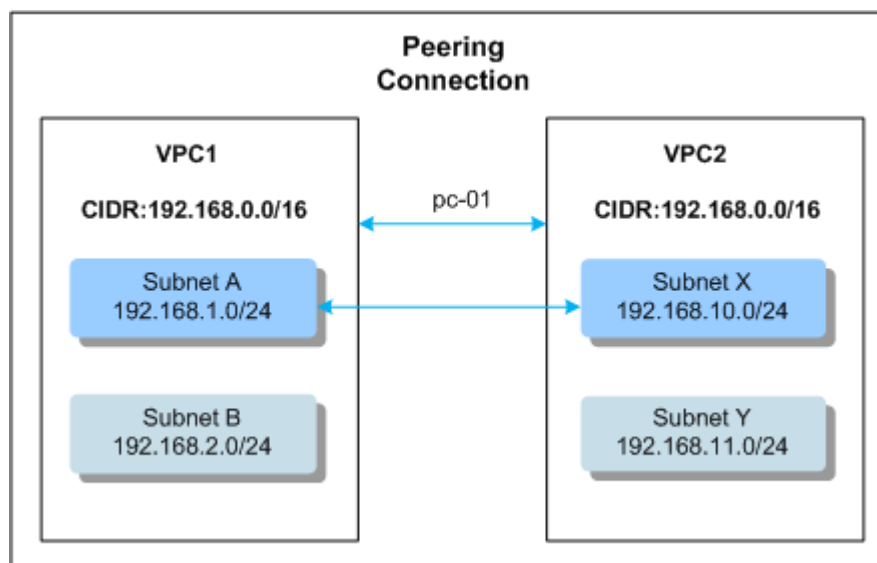
Item	Description	Example	Your Value
VPC1 ID	VPC 1 ID	Example: fef65559-c154-4229-afc4-9ad0314437ea	N/A
VM1 ID	ID of ECS 1 in VPC 1	Example: f7619b12-3683-4203-9271-f34f283cd740	N/A
VM2 ID	ID of ECS 2 in VPC 1	Example: f75df766-68aa-4ef3-a493-06cdc26ac37a	N/A

6.14 How Do I Handle the VPC Peering Connection Failure?

VPC Peering Connection Network

[Figure 6-23](#) shows the VPC peering connection network.

Figure 6-23 VPC peering connection network



Routes are required to enable communication between Subnet A in VPC1 and Subnet X in VPC2 in the figure. [Figure 6-24](#) shows the route table configuration.

Figure 6-24 VPC peering connection route table

VPC1		VPC2	
VPC Peering Route Table		VPC Peering Route Table	
Destination	Next Hop	Destination	Next Hop
192.168.10.0/24	pc-01	192.168.1.0/24	pc-01

Checking ECS Basic Network Functions

1. Confirm that 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.
The **ipconfig** command applies only to Windows ECSs.
2. Ping the gateway address of the subnet from the ECS to check the ECS communication with external networks.
Obtain the gateway address from the VPC details page on the console. In most cases, the gateway address is *xxx.xxx.xxx.1*. Ping the gateway address to check the communication. If the gateway address cannot be pinged, troubleshoot the layer 2 and layer 3 networks.

Checking VPC Network Configuration

1. Confirm that the security group configuration of the ECS NIC is correct.

Obtain the security group used by the ECS NIC from the ECS details page. The security group rule that allows the ECS to access the peer subnet has been configured for the security group. For example, you must configure security group rules described in [Figure 6-25](#) for the NICs of all ECSs in VPC 1 in [Figure 6-23](#).

Figure 6-25 Security group configuration

Direction	Type	Protocol	Port Range/ICMP Type	Remote End	Operation
Inbound	IPv4	Any	Any	192.168.10.0/24	Delete

2. Confirm that the firewall for the subnet used by the ECS NIC does not block required traffic.

If you can configure the firewall on the VPC console, confirm that the firewall rules allow traffic from the subnets used by the VPC peering connection to pass through.

3. 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 IP addresses can be pinged, the policy-based routing configured for the two NICs is correct.

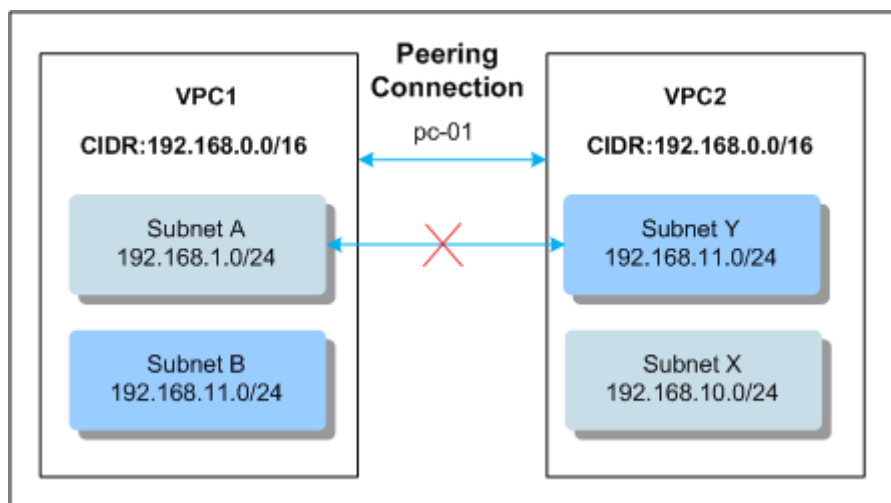
Checking VPC Peering Connection Configuration

1. The VPC peering connection described in [Figure 6-23](#) is used as an example to show how to check the configuration. Check whether correct routes have been added to the VPC peering connection. For example, the destination of the route for VPC 1 must be the subnet CIDR block in VPC 2.

Add local and peer routes on the VPC peering connection page. The VPC peering connection works properly after the routes are correctly configured.

2. Check VPC 1 and VPC 2 for subnets that conflict with the subnets involved in the VPC peering connection. 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 6-26](#) shows the invalid VPC peering connection.

Figure 6-26 Invalid VPC peering connection example



O&M Operations That Require Assistance

If the VPC peering connection failure cannot be rectified after you perform the preceding operations, contact technical support.

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	Your Value
VPC1 ID	VPC 1 ID	N/A
VPC2 ID	VPC 2 ID	N/A
VM1 ID	ID of the ECS in VPC 1	N/A
VM2 ID	ID of the ECS in VPC 2	N/A
Subnet1 ID	ID of the subnet used by ECS 1	N/A
Subnet2 ID	ID of the subnet used by ECS 2	N/A
IP1	ECS 1 IP address	N/A
IP2	ECS 2 IP address	N/A

NOTE

You can add `-t` to the end of the ping command to enable the Windows ECS to continuously send ICMP packets.

6.15 How Do I Handle the Layer 2 or Layer 3 Network Communication Failure?

1. Check 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 information. The **ipconfig** command applies only to Windows ECSs.
2. 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.
3. If the ECS has an IP address, check whether the security group rules and firewall rules allow the required traffic to pass through.
 - a. Check the security group rules.
Obtain the security group used by the ECS NIC from the ECS details page. The security group rule that allows the ECS to access the required subnet has been configured for the security group.
 - b. Check the firewall rules.
If you can configure the firewall on the VPC console, confirm that the firewall rules allow traffic from the required subnets to pass through.

6.16 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 6-27 Checking for bond

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

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

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

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

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

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

```
route -n
```

Figure 6-28 Checking BMS route information

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

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

Figure 6-29 Checking the default route

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

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

Figure 6-30 Checking the route for IP address range **169.254.169.254**

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

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

3. If BMSs in a VPC cannot communicate with each other or the BMS with an EIP cannot access the Internet, rectify the failure based on the related ECS 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.

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

6.17 How Do I Handle the ECS IP Address Obtaining Failure?

1. 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 4.
2. Run the following command to check whether the **dhclient** process exists:
ps -ef | grep dhclient
3. If the **dhclient** process does not exist, log in to the ECS and try restarting the ECS NIC or sending 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.
4. 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.

5. Check the ECS **messages** log in the **/var/log/messages** directory to troubleshoot the ECS.

Search for the NIC MAC address and check whether processes that cause failures in obtaining IP addresses over DHCP exists.

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


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.

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

VPN Network

Figure 6-31 shows your network, the customer gateway, the VPN, and the VPC.

Figure 6-31 VPN network



Customer Self-Check Guidance

1. Provide your network information.

Obtain information listed in [Table 6-1](#). This table lists example values. You can determine the actual values based on the example values. You must obtain all actual values of your project.

NOTE

You can print this table and fill in your values.

Table 6-1 Network information

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

Item	Description	Example	Your Value
Customer gateway type (for example, Cisco)	N/A	N/A	N/A
Public IP address used by the customer gateway	N/A	The value must be static.	N/A

2. Provide your gateway configuration information.

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

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

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

 **NOTE**

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

O&M Operations That Require Assistance

You must send communication requests from the ECSs in the public cloud system to the remote device.

Method:

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

6.19 What Do I Do If 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.

Fault Locating

1. Check on-premises network configurations, such as the firewall and IP address.

2. Check the EIP connection.
3. Check whether the security group of the server or the network ACL of the subnet to which the server belongs denies the outbound traffic.

Solution

1. Check on-premises network configurations, such as the firewall and IP address.
2. Check the EIP connection. For details, see [How Do I Handle EIP Connection Failure?](#)
3. Check whether the security group of the server denies the outbound traffic. By default, the security group allows all outbound traffic. If the outbound traffic is denied, click **Allow Common Ports**.
4. Check whether the network ACL of the subnet to which the server belongs 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 6-32 Allowing outbound traffic

Add Outbound Rule [Learn how to add a network ACL rule.](#)

Network ACL: fw-bc38

Action	Protocol	Source & Destination and Ports	Description	Operation
Allow	All	Source: 0 . 0 . 0 . 0 / 0 Destination: 0 . 0 . 0 . 0 / 0		Replicate Delete

[Add Rule](#) You can add 9 more rules.

OK Cancel

6.20 Can a VPC Peering Connection Be Deployed Across Regions?

No, but it can be deployed across AZs in the same region.

6.21 Is a VPC Peering Connection Charged?

It is free of charge currently.

6.22 Why the IPv6 Address Manually Configured for an ECS Cannot Be Used for Communication?

Currently, manually configured IPv6 addresses cannot be used for communication. You can use a virtual IP address for communication.

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

6.23 What Devices Can Connect to a L2CG on HUAWEI CLOUD?

Switches support VXLAN functions, such as CE6850 and Cisco Nexus 9300.

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

Possible causes and solutions:

1. 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 more information, see [Configuring a Tunnel in Your Data Center](#).
2. The Direct Connect or VPN connection used by the L2CG is not properly configured.
Check the Direct Connect or VPN connection configurations. For more information, see, see [Direct Connect Fault Locating](#) or [VPN Connection or Ping Failure](#).

6.25 Why Is the 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 more information, see [Configuring a Tunnel in Your Data Center](#).

7 Routing

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

If an ECS has multiple NICs, perform the following operations to configure policy-based routing for the ECS to enable the network communication of secondary NICs.

For a Linux ECS:

1. Run the following command and add two route tables (**net1** and **net2**) and their priorities to the **/etc/iproute2/rt_tables** file. The priorities of **net1** and **net2** are **252** and **251**, respectively. A smaller value indicates a higher priority.

vi /etc/iproute2/rt_tables

```
# added for dual net
252 net1
251 net2
```

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

vi /etc/rc.local

The IP addresses of NICs **eth0** and **eth1** are 192.168.1.23 and 192.168.2.4, respectively. The subnet mask is 24 bits. The gateway addresses of NICs **eth0** and **eth1** are 192.168.1.1 and 192.168.2.1, respectively. The information to be added is as follows:

```
# Request IP address for eth1
dhclient eth1
# Add routes
ip route flush table net1
ip route add default via 192.168.1.1 dev eth0 src 192.168.1.23 table net1
ip route add 192.168.1.0/24 dev eth0 src 192.168.1.23 table net1
ip rule add from 192.168.1.23 table net1

ip route flush table net2
ip route add default via 192.168.2.1 dev eth1 src 192.168.2.4 table net2
ip route add 192.168.2.0/24 dev eth1 src 192.168.2.4 table net2
ip rule add from 192.168.2.4 table net2
```

3. Run the following command to add the execute permission for the **rc.local** file:

chmod +x /etc/rc.local

4. Run the **reboot** command to restart the ECS.

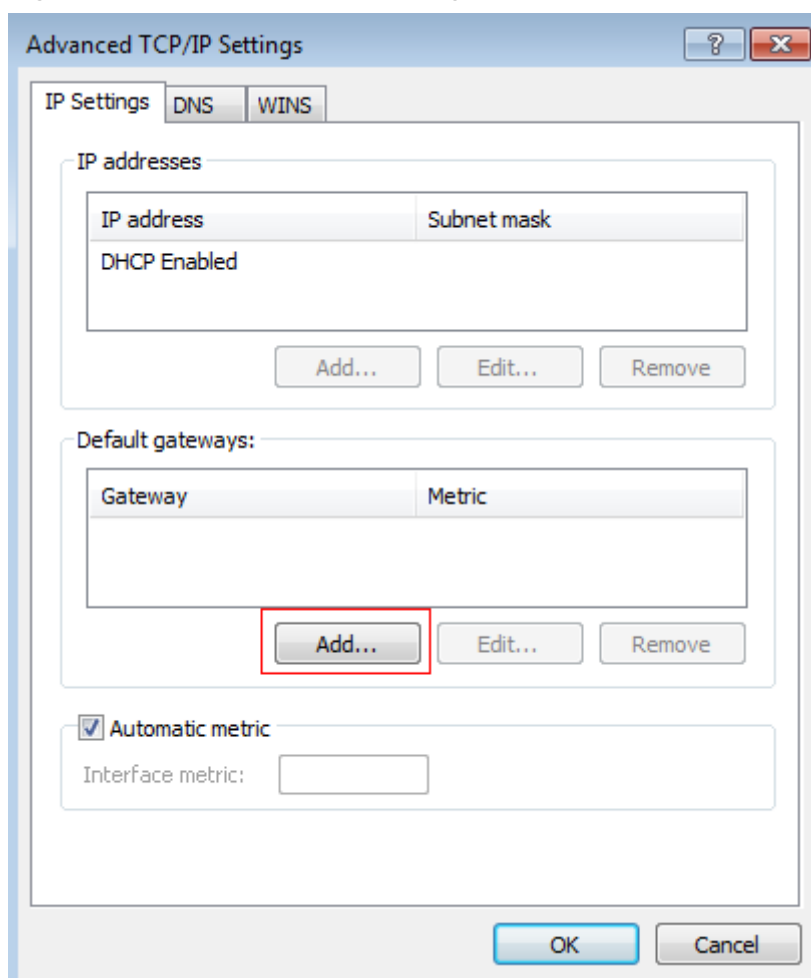
For a Windows ECS:

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

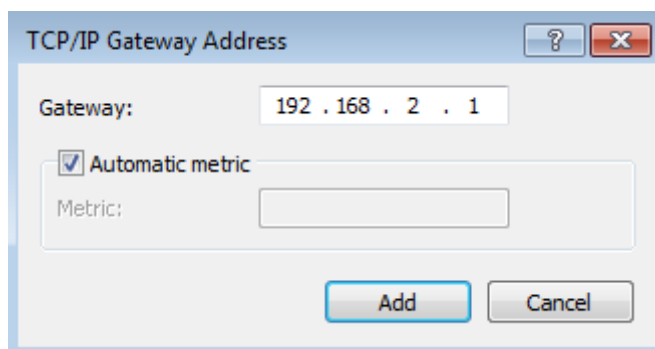
NOTE

Right-click to add NICs based on the site requirements.

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**.

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 Can a Route Table Span Multiple VPCs?

No.

7.3 How Many Routes Can Exist in a Route Table?

Currently, a route table can contain 100 routes.

7.4 What Are the Limitations of a Route Table?

- The ECS providing SNAT can have only one NIC.
- 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 the next hop in a route, EIPs bound with the virtual IP address in the VPC will become invalid.

7.5 Will a Route Table be Billed?

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

7.6 Are There Different Routing Priorities for Direct Connect Connections and Custom Routes in the Same VPC?

No. Direct Connect connections and custom routes are used in different scenarios. Therefore, there is no routing priority competition between them.

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

The routing priority of custom routes and that of VPNs are the same.

8 Security

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

Yes. When you add or import a security group rule with the same parameters except the description as an existing one in the security group, the adding or importing will fail.

8.2 What Should I Do Before 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 any cloud resource, release the corresponding cloud resource or change the security group used by the cloud resource, and then delete the security group.
- If the security group to be deleted has been 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.

8.3 What Do I Do If Outbound Access Through TCP Port 25 Is Restricted?

Symptom

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

Cause

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

If you do not need to deploy the email service on the cloud, disabling the TCP port 25 does not affect your service running.

Solution

- If you want to use an ECS on HUAWEI CLOUD to send emails, you are advised to use Cloud Speedy Mail.
- 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 request, you must agree and guarantee that TCP port 25 is only used to connect 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 for you and will no longer enable it even you request.

1. 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 as prompted.

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 Each User Have?

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

When creating an ECS, you can select multiple security groups (no more than five is recommended).

8.6 Is the Security Group Service Charged?

The security group service is free of charge.

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

ECS Configuration


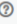
The TFTP daemon determines whether the configuration file specifies the port range. If you use the TFTP configuration file that allows the data channel ports to be configurable, it is a best practice to configure a small range of ports that are not listened on.

Security Group Configuration

You can configure both port 69 and 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

<input type="checkbox"/> Type	Protocol	Port/Range	Source
<input type="checkbox"/> IPv4	All	All	sg-test 
<input type="checkbox"/> IPv4	UDP	60001-60100	0.0.0.0/0 

8.8 How Many Network ACLs Can a User Have?

A user can have 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 more than 20 inbound or outbound rules are configured, 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 is 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. You need to interrupt the original traffic for about 120 seconds for the new rule to take effect for the traffic. To ensure that the traffic is immediately interrupted after the rule is changed, it is recommended that you configure security group rules.

8.10 What Do I Do If Some Ports in the Public Cloud System Are Inaccessible?

Symptom: Users in some 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 not listed in the table for your services.

8.11 Why the Access from a Specified IP Address Is Still Allowed After a Network ACL Rule that Denies the Access from this Specified 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 network ACL 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, put the rules that allow access from all IP addresses to the end. The rule that denies access from the specified IP address takes effect. For details, see [Changing the Sequence of a Network ACL Rule](#).