Virtual Private Network

FAQs

Issue 01

Date 2025-06-25





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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process.* For details about this process, visit the following web page:

https://www.huawei.com/en/psirt/vul-response-process

For vulnerability information, enterprise customers can visit the following web page:

https://securitybulletin.huawei.com/enterprise/en/security-advisory

FAQs - S2C Enterprise Edition VPN

1.1 Popular Questions

1.1.1 What Devices Can Be Connected to Huawei Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- 1. Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

□ NOTE

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- The following products can connect to the cloud through VPNs:
 - Devices: Huawei firewalls and access routers (ARs), Hillstone firewalls, and Check Point firewalls
 - Cloud services: Alibaba Cloud, Tencent Cloud, Amazon Web Services (AWS), and Microsoft Azure
 - Software: strongSwan
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.

Most enterprise-class routers and firewalls support the IPsec protocol.

Some devices support IPsec VPN only after you purchase required software licenses.
 Your on-premises data center administrator can check with the device vendor whether a license is required based on the device model.

1.1.2 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 1-1 VPN negotiation parameters

Protocol	Parameter	Value
IKE	Authentication Algorithm	 MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-128 (default value) AES-192 AES-256 AES-256-GCM-16
	DH Algorithm	 Group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 14 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 15 (default value) Group 15 Group 19 Group 20 Group 21
	Version	 v1 (not recommended due to security risks) v2 (default value)

Protocol	Parameter	Value
	Lifetime (s)	86400 (default value) Unit: second Value range: 60 to 604800
	Local ID	 IP Address The local IP address is automatically displayed as the EIP of the VPN gateway, removing the need to manually configure it. FQDN By default, the local ID type is IP address and the local ID value is the EIP of the VPN gateway.
	Customer ID	 IP Address FQDN By default, the customer ID type is IP address and the customer ID value is the public IP address of the customer gateway.
IPsec	Authentication Algorithm	 SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 AES-128 (default value) AES-192 AES-256 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-256-GCM-16

Protocol	Parameter	Value
	PFS	Disable (not recommended due to security risks)
		DH group 1 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 2 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 5 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 14 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 15 (default value)
		DH group 16
		DH group 19
		DH group 20
		DH group 21
	Transfer Protocol	ESP (default value)
	Lifetime (s)	3600 (default value) Unit: second
		Value range: 30 to 604800

□ NOTE

- Perfect Forward Secrecy (PFS) is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. When PFS is enabled, an additional DH exchange will be performed during IPsec SA negotiation to generate a new IPsec SA key, improving IPsec SA security.
- For security purposes, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the gateway device in your on-premises data center and the PFS settings on both ends are the same. Otherwise, the negotiation will fail.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB on the cloud side and cannot be changed for the VPN gateway. This parameter is not involved in negotiation and has no impact on the establishment of an IPsec SA.

1.1.3 Can I Deploy an Application on the Cloud and a Database in an On-premises Data Center and Connect Them Through a VPN Gateway?

Yes.

A VPN connects a VPC and an on-premises data center.

After a VPN is set up, service traffic can be transmitted between the VPC and on-premises data center. For an application server on the cloud, access to an on-premises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check the network latency and packet loss rate to ensure smooth service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

1.1.4 Can I Visit Websites Across International Borders Using a VPN?

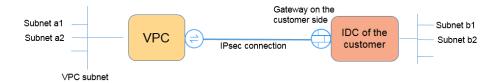
No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

1.1.5 What Is a VPN Connection? How Do I Set the Number of VPN Connections When Buying a VPN Gateway?

A VPN connection is an IPsec connection established between a VPN gateway and an independent public IP address of an on-premises data center. You can configure multiple local subnets (VPC subnets) and customer subnets (on-premises subnets) for one VPN connection.

The number of VPN connections to be created is determined by the number of onpremises data centers. Each VPN connection can connect a VPC to only one onpremises data center.



In the preceding figure, if subnets a1 and a2 on the cloud need to communicate with subnets b1 and b2 on the on-premises network, you only need to create one VPN connection, with source CIDR blocks set to a1 and a2 and destination CIDR blocks set to b1 and b2.

1.1.6 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

To view the status of a VPN connection, click in the **Monitoring** column of the VPN connection.

1.1.7 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

◯ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

1.1.8 What Are the Differences Between IPsec VPN and SSL VPN in Application Scenarios and Connection Modes?

Application Scenarios

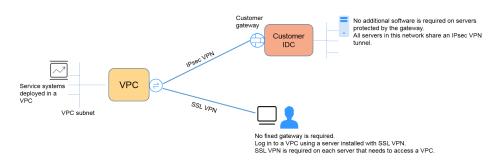
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to enable them to complete IPsec VPN negotiation.

SSL VPN requires a specific client program installed on hosts. Users need to enter usernames and password to connect the hosts to SSL servers.



Ⅲ NOTE

IPsec VPN and SSL VPN are supported.

1.1.9 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

1.1.10 How Will I Be Charged for My Use of a VPN? Will I Be Charged for VPN Gateway EIPs?

VPNs are billed by the following items on a yearly/monthly or pay-per-use basis.

- VPN gateway
- VPN connection

By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. You can purchase additional VPN connections if required.

• EIP bandwidth of a VPN gateway

The VPN gateway bandwidth can be billed by traffic or bandwidth.

- a. A yearly/monthly VPN gateway can only be billed by bandwidth. The price of a yearly/monthly VPN gateway includes the price of the VPN connections that can be created for the gateway and the bandwidth price.
- b. The billing cycle of the pay-per-use billing mode is 1 hour. When you create a pay-per-use VPN gateway, the system prompts you to create VPN connections. By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. If more connection groups are required, you need to purchase them.

\mathbf{m}	INOT	E

A VPN gateway cannot share a bandwidth with an EIP bound to an ECS.

1.1.11 Can the EIP of a VPN Gateway Be Retained After the VPN Gateway Is Deleted?

If a pay-per-use EIP is bound to a pay-per-use VPN gateway, deleting the VPN gateway will also delete the bound EIP.

To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway.

1.1.12 What VPN Resources Can Be Monitored?

VPN gateway

The following bandwidth information of a VPN gateway IP address can be monitored: inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view the monitoring information, click in the **Gateway IP Address** column in the VPN gateway list.

VPN connection

The following information about a VPN connection can be monitored: VPN connection status, average link round-trip time (RTT), maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate.

To monitor average link RTT, maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate, click the VPN connection name and click **Add** in the **Health Check** area on the **Summary** tab page to add health check items. Private network metrics can be configured only when the VPN connection uses the static routing mode and the NQA function is enabled.

To view the status of a VPN connection, click $\stackrel{\triangle}{=}$ in the **Monitoring** column of the VPN connection.

1.1.13 In Which Direction Is the VPN Bandwidth Limited? What Is the Unit of Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited as follows:

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is greater than 10 Mbit/s, the bandwidth in the inbound direction is the same as the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

1.1.14 How Is the Network Speed of a VPN Connection Tested?

Test environment: A VPN connection has been created. ECSs have been created on the local subnets of VPCs at the two ends of the VPN connection. The ECSs can ping each other.

When the bandwidth of a purchased VPN gateway is 200 Mbit/s:

1. When the ECSs at the two ends of the VPN connection run Windows, iPerf3 and FileZilla (a free FTP application for file upload and download) are used to test the network speed. The test result is 180 Mbit/s, meeting requirements.

The TCP-based FTP protocol has a congestion control mechanism, and the IPsec protocol adds new headers to original packets. As such, it is normal in the industry, to have a network speed deviation of about 10%.

Figure 1-1 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 client.

ID] Interval

perf Done.

0.00-10.01 sec 0.00-10.01 sec

sender

receiver

C:\Windows>iperf3 -c 10.1.0.180 -i 1 -w 1M Connecting to host 10.1.0.180, port 5201 41 local 192.168.0.75 port 49212 connected to 10.1.0.180 port 5201 Transfer Bandwidth ID1 Interval 0.00-1.01 sec 17.1 MBytes 142 Mbits/sec 1.01-2.00 sec 30.0 MBytes 253 Mbits/sec 2.00-3.01 19.8 MBytes 165 Mbits/sec 194 Mbits/sec sec 23.2 MBytes 3.01-4.01 4.01-5.00 18.9 MBytes 5.00-6.01 sec 26.2 MBytes 219 Mbits/sec 18.4 MBytes 153 Mhits/sec 6.01-7.01 sec sec 23.2 MBytes sec 21.1 MBytes 7.01-8.01 195 Mbits/sec 180 Mbits/sec 8.01-9.00 sec 9.00-10.01 sec 21.0 MButes 174 Mbits/sec

Figure 1-1 Test result for 200 Mbit/s bandwidth (iPerf3 client)

Figure 1-2 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 server.

Bandwidth

183 Mbits/sec

183 Mbits/sec

Figure 1-2 Test result for 200 Mbit/s bandwidth (iPerf3 server)

Iransfer

219 MBytes

219 MBytes

- 2. When the ECSs at the two ends of the VPN connection run CentOS 7, iPerf3 is used to test the network speed. The test result is 180 Mbit/s, meeting requirements.
- 3. When the ECS functioning as a server runs CentOS 7 and the ECS functioning as a client runs Windows, iPerf3 and FileZilla are used to test the network speed. The test result is 20 Mbit/s, failing to meet requirements.

This is because TCP implementations on Windows and Linux are different.

Figure 1-3 shows the result of using iPerf3 to test the network speed between two ECSs running different operating systems.

Figure 1-3 Test result on iPerf3

```
:\Windows>iperf3 -c 10.1.0.182 -i 1 -w 1M
onnecting to host 10.1.0.182, port 5201
 41 local 192.168.0.75 port 49426 connected to 10.1.0.182 port 5201
ID] Interval
                         Transfer
                                       Bandwidth
                                      36.7 Mbits/sec
37.7 Mbits/sec
      0.00-1.00
 4]
                         4.38 MBytes
                   sec
       1.00-2.00
                         4.50 MBytes
                                       43.0 Mbits/sec
                         5.12 MBytes
                         1.75
                              MBytes
                                      14.7
                                            Mbits/sec
                   sec
                         2.12
                              MBytes
                                            Mbits/sec
                   sec
                         3.25
                   sec
                              MBytes
                                            Mbits/sec
                         2.12 MBytes
                                            Mbits/sec
                   sec
                         1.25
                              MBytes
                                       10.5
                         2.25 MBytes
                                       18.9 Mbits/sec
       9.00-10.00
                         2.38 MBytes
                                       19.9 Mbits/sec
IDI
    Interval
                         Iransfer
                                       Bandwidth
       0.00-10.00
                   sec
                         29.1 MBytes
                                      24.4 Mbits/sec
                                                                         sender
       0.00-10.00
                         28.2 MBytes
                                      23.6 Mbits/sec
                                                                         receiver
```

When the bandwidth of a purchased VPN gateway is 1000 Mbit/s:

◯ NOTE

Some regions support only 300 Mbit/s bandwidth by default. If higher bandwidth is required, select 300 Mbit/s bandwidth and then **submit a service ticket** for capacity expansion.

The VPN gateway bandwidth is shared by all of its VPN connections. To fully use the large bandwidth of 1000 Mbit/s, deploy multiple ECSs with high specifications as the forwarding performance of a single ECS is limited. ECSs with their NICs supporting the bandwidth of 2 Gbit/s or higher are recommended.

Conclusions: Based on the preceding test results, bandwidths of VPN gateways meet requirements. To fully use your purchased bandwidth, you are advised to use servers running the same operating system and using NICs meeting certain requirements at the two ends of a VPN connection.

1.1.15 Can a VPN Billed by Traffic Use a Shared Data Package?

Yes.

The VPN service fee includes the EIP fee. An EIP can use a shared data package.

1.1.16 How Do I Change the Billing Mode of a VPN Gateway from Pay-Per-Use to Yearly/Monthly?

Prerequisites

- A pay-per-use VPN gateway is billed by bandwidth.
- To change the billing mode of a VPN gateway billed by traffic from pay-peruse to yearly/monthly, first change the VPN gateway from being billed by traffic to being billed by bandwidth and then from pay-per-use to yearly/ monthly.

Procedure

Perform the following operations:

- 1. Log in to the management console.
- 2. Click in the upper left corner and select the desired region and project.
- 3. Click Service List and choose Networking > Virtual Private Network.
- 4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise VPN Gateways**.
- 5. On the **VPN Gateways** page, locate the row that contains the target VPN gateway, choose **More** > **Change Billing Mode** in the **Operation** column.
- 6. In the Change Billing Mode dialog box, click OK.

The billing mode of a VPN gateway cannot be changed from yearly/monthly to payper-use. The resource quotas of a yearly/monthly VPN gateway can be decreased upon a renewal.

- 7. Confirm the VPN gateway information, set a renewal duration, and click **Pay**.
- 8. On the payment page, confirm the order information, select a coupon or discount, select a payment method, and click **Pay**.

Changing the billing mode of a VPN gateway from pay-per-use to yearly/monthly will not affect your services.

1.1.17 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

- VPC
 - A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
 - You can use the VPN service to securely access ECSs in a VPC.
- VPN gateway
 - A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have multiple VPN gateways, and one VPN gateway can have multiple VPN connections.
 - With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.
- VPN connection

A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

□ NOTE

The number of VPN connections is irrelevant to the number of local subnets or the number of customer subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

1.1.18 What Are a Customer Gateway and a Customer Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a customer subnet and a customer gateway, respectively.

A customer gateway's IP address is a public IP address.

1.1.19 How Many VPN Connections Do I Need to Connect Multiple On-premises Servers to the Cloud?

VPN uses the IPsec technology to connect your on-premises data center to a VPC on the cloud. As such, the number of VPN connections is related to the number of data centers where the servers to be connected to the cloud are located, but not to the number of servers.

Two EIPs can be bound to a VPN gateway for communication with a customer gateway.

- If an on-premises data center has only one egress gateway, all servers or hosts in the data center connect to the Internet through this gateway. In this case, you need to configure a VPN connection group consisting of two VPN connections. That is, configure a VPN connection for each of the two EIPs of the VPN gateway to communicate with the egress gateway in the onpremises data center.
- If an on-premises data center has two egress gateways, the servers or user hosts in the data center connect to the Internet through the two egress gateways. In this case, you need to configure two VPN connection groups, each of which consisting of two VPN connections. That is, configure a VPN connection for each of the two EIPs of each VPN gateway to communicate with both egress gateways in the on-premises data center.

1.1.20 Does a VPN Allow for Communications Between Two VPCs?

- If the two VPCs are in the same region, use a VPC peering connection to connect them.
- If the two VPCs are in different regions, use a VPN to connect them. The operations are as follows:
 - a. Create a VPN gateway for each VPC, and create a VPN connection between the two VPN gateways.
 - b. For the VPN connection, set the customer gateway to the EIP of the peer VPN gateway.

- c. For the VPN connection, set the customer subnet to the subnet of the peer VPC.
- d. Set the same pre-shared keys (PSKs) and algorithms for the two VPCs.

1.1.21 What Are the Impacts of a VPN on an On-premises Network? What Are the Changes to the Route for Accessing an ECS?

When configuring a VPN, you need to perform the following operations on the gateway in your on-premises data center:

- Configure IKE and IPsec policies.
- Set the connection mode to route-based or policy-based.
- Check the route configuration on the gateway to ensure that traffic destined for a VPC can be routed to the correct outbound interface (interface having an IPsec policy bound).

1.1.22 Can I Connect a Network with Two Egresses to a VPC Through Two VPN Connections?

Yes

1.1.23 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes of disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- Dead Peer Detection (DPD) is not configured on the device in your onpremises data center.
- Configuration is modified when the VPN connection is in use.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- At the two ends of the VPN connection, the local and customer subnet configurations are reversed.
- SA lifetime settings at both ends of the VPN connection are the same.
- DPD is enabled on the on-premises gateway device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.

- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway device is large enough for the VPN connection.
- VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway device.

1.1.24 What Do I Do If a VPN Connection Fails to Be Established?

- Log in to the management console, and choose Virtual Private Network > Enterprise - VPN Connections.
- In the VPN connection list, locate the target VPN connection, and choose More > Modify Policy Settings on the right to view IKE and IPsec policies of the VPN connection.
- 3. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.

 If the IKE SA has been set up in phase 1 but no IPsec SA has been established in phase 2, the IPsec policies at both ends of the VPN connection may be
- 4. Check whether the ACL configurations are correct.

inconsistent.

If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

5. Ping the two ends of the VPN connection from each other to check whether the VPN connection is normal.

1.1.25 Can EIPs Be Used as VPN Gateway IP Addresses?

In Enterprise Edition VPN, EIPs can be used as VPN gateway IP addresses.

When creating a VPN gateway, you can bind EIPs as the gateway IP addresses.

1.1.26 Why Is a VPN Connection Always in Not Connected State After Its Configuration Is Complete?

The configuration may be incorrect.

- At the two ends (cloud and on-premises data center) of the VPN connection, ensure that the pre-shared keys (PSKs) and negotiation information are consistent, the local and remote subnets are reversed, and the local and remote gateways are also reversed.
- 2. Ensure that routes, NAT, and security policies are correctly configured on the device in your on-premises data center.

1.1.27 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to configure policy rules (ACL rules) for a VPN connection on the management console only when **VPN Type** is set to **Policy-based**.

1.1.28 How Do I Determine Which EIP Is Used for Transmitting Service Traffic That Leaves the Cloud?

- If the HA mode of a VPN gateway is active/standby:
 The outgoing traffic from the VPN gateway to the customer subnet is preferentially transmitted through the VPN connection set up between the customer subnet and the active EIP.
- If the HA mode of a VPN gateway is active-active:
 - When Associate With is set to Enterprise Router, the outgoing traffic from the VPN gateway to the customer subnet is load balanced among all VPN connections set up with the customer subnet.
 - When Associate With is set to VPC, the outgoing traffic from the VPN gateway to the customer subnet is preferentially forwarded through the VPN connection that is first set up between the customer subnet and an EIP.
- You can perform the following operations to check the VPN connection through which traffic leaves the cloud:
 - a. Log in to the management console.
 - b. Click in the upper left corner of the page, select a region, and choose Management & Governance > Cloud Eye.
 - Choose Cloud Service Monitoring from the navigation tree. The Cloud Service Dashboards page is displayed.
 - d. Click Virtual Private Network VPN in the Dashboard column. The Details page is displayed.
 - e. Choose **S2C VPN Connection**, click the **Resources** tab, and click **View Metric** in the **Operation** column of the target VPN connection.
 - Check the metrics of the VPN connection. If the value of the **Traffic Send Rate** metric is not 0, the traffic is transmitted through this connection.

1.2 General Consulting

1.2.1 What Are the Typical Scenarios of IPsec VPN?

A VPN is a point-to-point connection that implements private network access between two points.

- Applicable scenarios:
 - A VPN is created between different regions to enable cross-region VPC communications.

- A VPN is created between VPCs and another public cloud, for example, Alibaba Cloud.
- A VPN is created between a VPC and your on-premises data center.
- A VPN hub is used together with VPC peering connections and Cloud Connect connections to enable communications between an on-premises data center and multiple VPCs on the cloud.
- A VPN is used together with source NAT to enable access to specific IP addresses across clouds.
- A VPN can be used between the cloud and your home network that uses PPPoE dial-up.
- A VPN can be used between the cloud and 4G/5G routers.
- A VPN can be used between the cloud and your personal terminals.
- Not applicable scenarios:
 - A VPN cannot be used to connect VPCs in the same region. It is recommended that you use VPC peering connections to enable communications between VPCs in the same region.

1.2.2 What Are a VPC, a VPN Gateway, and a VPN Connection?

VPC enables you to create private, isolated virtual networks. You can use VPN to securely access ECSs in VPCs.

A VPN gateway is an egress gateway for a VPC. With a VPN gateway, you can create a secure, reliable, and encrypted connection between a VPC and an onpremises data center or between two VPCs in different regions.

A VPN connection is a secure and reliable IPsec encrypted communications tunnel established between a VPN gateway and the customer gateway in an on-premises data center.

To create a VPN on the cloud, perform the following operations:

- 1. Create a VPN gateway. You need to specify the VPC to be connected, as well as the bandwidth and EIPs of the VPN gateway.
- 2. Create a VPN connection. You need to specify the gateway EIP used to connect to the customer gateway, subnets, and negotiation policies.

1.2.3 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

- VPC
 - A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
 - You can use the VPN service to securely access ECSs in a VPC.

VPN gateway

- A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have multiple VPN gateways, and one VPN gateway can have multiple VPN connections.
- With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.

VPN connection

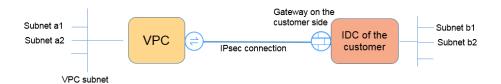
A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

The number of VPN connections is irrelevant to the number of local subnets or the number of customer subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

1.2.4 What Is a VPN Connection? How Do I Set the Number of VPN Connections When Buying a VPN Gateway?

A VPN connection is an IPsec connection established between a VPN gateway and an independent public IP address of an on-premises data center. You can configure multiple local subnets (VPC subnets) and customer subnets (on-premises subnets) for one VPN connection.

The number of VPN connections to be created is determined by the number of onpremises data centers. Each VPN connection can connect a VPC to only one onpremises data center.



Ⅲ NOTE

In the preceding figure, if subnets a1 and a2 on the cloud need to communicate with subnets b1 and b2 on the on-premises network, you only need to create one VPN connection, with source CIDR blocks set to a1 and a2 and destination CIDR blocks set to b1 and b2.

1.2.5 What Are a Customer Gateway and a Customer Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a customer subnet and a customer gateway, respectively.

A customer gateway's IP address is a public IP address.

1.2.6 How Do I Plan CIDR Blocks for Access to a VPC Through a VPN Connection?

- The CIDR blocks of a VPC cannot conflict with on-premises CIDR blocks.
- To avoid conflicts with cloud service addresses, do not use 127.0.0.0/8, 169.254.0.0/16, 224.0.0.0/3, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 for your on-premises network.

If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.

1.2.7 Is an IPsec VPN Connection Automatically Established?

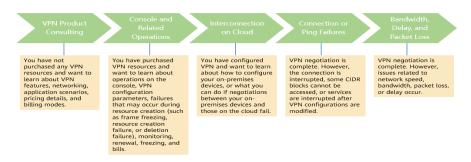
Yes. An IPsec VPN connection is automatically established.

1.2.8 What Types of VPN Service Tickets Are There? How Do I Create a VPN Service Ticket?

- 1. Log in to the management console.
- 2. Choose **Service Tickets** > **Create Service Ticket** in the upper right corner.
- 3. Search for "VPN" and choose Virtual Private Network (VPN).
- 4. Select an issue category.

When you **submit a service ticket**, select an issue category to facilitate problem handling.

Figure 1-4 Issue category and classification basis



1.2.9 What Devices Can Be Connected to Huawei Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- 1. Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- The following products can connect to the cloud through VPNs:
 - Devices: Huawei firewalls and access routers (ARs), Hillstone firewalls, and Check Point firewalls
 - Cloud services: Alibaba Cloud, Tencent Cloud, Amazon Web Services (AWS), and Microsoft Azure
 - Software: strongSwan
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.
 - Most enterprise-class routers and firewalls support the IPsec protocol.
- Some devices support IPsec VPN only after you purchase required software licenses. Your on-premises data center administrator can check with the device vendor whether a license is required based on the device model.

1.2.10 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 1-2 VPN negotiation parameters

Protocol	Parameter	Value
IKE	Version	 v1 (v1 has low security. If the device supports v2, v2 is recommended.) v2 (default value)
	Negotiation Mode	Main (default value) Aggressive
		Aggressive
	Authentication Algorithm	 MD5 (This algorithm is insecure. Exercise caution when using this algorithm.)
		 SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.)
		SHA2-256 (default value)
		• SHA2-384
		• SHA2-512

Protocol	Parameter	Value
	Encryption Algorithm	 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-128 (default value) AES-192 (This algorithm is insecure. Exercise caution when using this algorithm.) AES-256 (This algorithm is insecure. Exercise caution when using this algorithm.)
	DH Algorithm	 Group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 14 (default value) Group 15 Group 16 Group 20 Group 21
	Lifetime (s)	86400 (default value) Unit: second Value range: 60 to 604800
	Local ID	 IP Address The local IP address is automatically displayed as the EIP of the VPN gateway, removing the need to manually configure it. FQDN By default, the local ID type is IP address and the local ID value is the EIP of the VPN gateway.

Protocol	Parameter	Value
	Customer ID	 IP Address FQDN By default, the customer ID type is IP address and the customer ID value is the public IP address of the customer gateway.
IPsec	Authentication Algorithm	 SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 AES-128 (default value) AES-192 (This algorithm is insecure. Exercise caution when using this algorithm.) AES-256 (This algorithm is insecure. Exercise caution when using this algorithm.) 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-128-GCM-16 AES-256-GCM-16

Protocol	Parameter	Value
	PFS	DH group 1 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 2 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 5 (This algorithm is insecure. Exercise caution when using this algorithm.)
		DH group 14 (default value)
		DH group 15
		DH group 16
		DH group 19
		DH group 20
		DH group 21
		Disable (not recommended due to security risks)
	Transfer Protocol	ESP (default value)
	Lifetime (s)	3600 (default value)
		Unit: second
		Value range: 30 to 604800

□ NOTE

- Perfect Forward Secrecy (PFS) is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. When PFS is enabled, an additional DH exchange will be performed during IPsec SA negotiation to generate a new IPsec SA key, improving IPsec SA security.
- For security purposes, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the gateway device in your on-premises data center and the PFS settings on both ends are the same. Otherwise, the negotiation will fail.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB on the cloud side and cannot be changed for the VPN. This parameter is not involved in negotiation and has no impact on the establishment of an IPsec SA.

1.2.11 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for

creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

1.2.12 How Do I Allow Specific Hosts to Access a VPC Subnet Through a Created VPN Connection?

Restrictions in the on-premises data center:

- Access control policies on the VPN device
- ACL rules on the router or switch

Restrictions at the cloud side:

- Security group rules that permit access only from specified IP addresses
- ACL rules



You are advised not to change the local or customer subnet to control access.

1.2.13 What VPN Resources Can Be Monitored?

VPN gateway

The following bandwidth information of a VPN gateway IP address can be monitored: inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view the monitoring information, click $\stackrel{\square}{=}$ in the **Gateway IP Address** column in the VPN gateway list.

VPN connection

The following information about a VPN connection can be monitored: VPN connection status, average link round-trip time (RTT), maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate.

To monitor average link RTT, maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate, click the VPN connection name and click **Add** in the **Health Check** area on the **Summary** tab page to add health check items. Private network metrics can be configured only when the VPN connection uses the static routing mode and the NQA function is enabled.

To view the status of a VPN connection, click $\stackrel{\square}{=}$ in the **Monitoring** column of the VPN connection.

1.2.14 Can EIPs Be Used as VPN Gateway IP Addresses?

Yes.

When creating a VPN gateway, you can bind EIPs as the gateway IP addresses.

1.2.15 Do I Need to Purchase EIPs for Hosts to Communicate with Each Other Through a VPN?

If your on-premises hosts need to access an ECS on the cloud through a VPN, you do not need to purchase any EIPs for the ECS.

If an ECS needs to provide services accessible from the Internet, you need to purchase an EIP for the ECS.

1.2.16 Are SSL VPNs Supported?

Currently, SSL VPNs are supported.

1.2.17 How Long Does It Take for Delivered VPN Configurations to Take Effect?

It takes 1–5 minutes for the VPN configurations to take effect.

□ NOTE

After VPN configurations take effect, configure your gateway device on your on-premises network to complete tunnel negotiation with the VPN gateway.

1.2.18 Does VPN Support IPv6?

Yes.

Currently, VPN supports both IPv4 and IPv6.

1.2.19 How Do I Determine My VPN Bandwidth?

Consider the following when you determine the bandwidth:

- Amount of data transmitted over a VPN tunnel in a period of time (Reserve enough bandwidth to prevent link congestion.)
- Egress bandwidths at the two ends of a VPN connection: The egress bandwidth at the cloud side must be less than that at the on-premises side.

1.2.20 Does a VPN Connection Support SM Series Cryptographic Algorithms?

Yes.

Use the algorithms provided on the management console for VPN negotiation. Additionally, ensure that the two ends of a VPN connection use the same algorithms.

1.2.21 Which IKE Version Should I Select When I Create a VPN Connection?

IKEv2 is recommended because IKEv1 is not secure. In addition, IKEv2 outperforms IKEv1 in connection negotiation and establishment, authentication methods, dead peer detection (DPD) timeout processing, and security association (SA) timeout processing.

IKEv2 will be widely used, and IKEv1 will gradually phase out.

Introduction to IKEv1 and IKEv2

- As a hybrid protocol, IKEv1 brings some security and performance defects due to its complexity. As such, it has become a bottleneck in the IPsec system.
- IKEv2 addresses the issues of IKEv1 while retaining basic functions of IKEv1.
 IKEv2 is more simplified, efficient, secure, and robust than IKEv1. Additionally, IKEv2 is defined by RFC 4306 in a single document, whereas IKEv1 are defined in multiple documents. By minimizing core functions and default password algorithms, IKEv2 greatly improves interoperability between different IPsec VPNs.

Security Risks of IKEv1

- The cryptographic algorithms supported by IKEv1 have not been updated for more than 10 years. In addition, IKEv1 does not support strong cryptographic algorithms such as AES-GCM and ChaCha20-Poly1305. For IKEv1, the E (Encryption) bit in the ISALMP header specifies that the payloads following the ISALMP header are encrypted, but any data integrity verification of those payloads is handled by a separate hash payload. This separation of encryption from data integrity protection prevents the use of authenticated encryption (AES-GCM) with IKEv1.
- IKEv1 is vulnerable to DoS amplification attacks and half-open connection attacks. After responding to spoofed packets, the responder maintains initiator-responder relationships, consuming a large number of system resources.
 - This defect is inherent to IKEv1 and is addressed in IKEv2.
- The aggressive mode of IKEv1 is not secure. In this mode, information packets are not encrypted, posing risks of information leakage. There are also bruteforce attacks targeting at the aggressive mode, such as man-in-the-middle attacks.

Differences Between IKEv1 and IKEv2

Negotiation process

IKEv1 is complex and consumes a large amount of bandwidth. IKEv1 SA negotiation consists of two phases. In IKEv1 phase 1, an IKE SA is established in either main mode or aggressive mode. Main mode requires three exchanges between peers totaling six ISAKMP messages, whereas aggressive mode requires two exchanges totaling three ISAKMP messages. Aggressive mode is faster, but does not provide identity protection for peers as key exchange and identity authentication are performed simultaneously. In IKEv1 phase 2, IPsec SAs are established through three ISAKMP messages in quick mode.

Compared with IKEv1, IKEv2 simplifies the SA negotiation process. IKEv2 requires only two exchanges, totaling four messages, to establish an IKE SA and a pair of IPsec SAs. To create multiple pairs of IPsec SAs, only one additional exchange is needed for each additional pair of SAs.

For IKEv1 negotiation, its main mode involves nine (6+3) messages, and its aggressive mode involves six (3+3) messages. In contrast, IKEv2 negotiation requires only four (2+2) messages.

Authentication methods

- Only IKEv1 (requiring an encryption card) supports digital envelope authentication (HSS-DE).
- IKEv2 supports Extensible Authentication Protocol (EAP) authentication.
 IKEv2 can use an AAA server to remotely authenticate mobile and PC users and assign private IP addresses to these users. IKEv1 does not provide this function and must use L2TP to assign private IP addresses.
- Only IKEv2 supports IKE SA integrity algorithms.

DPD timeout processing

- Only IKEv1 supports the retry-interval parameter. If a device sends a
 DPD packet but receives no reply within the specified retry-interval, the
 device records a DPD failure event. When the number of DPD failure
 events reaches 5, both the IKE SA and IPsec SAs are deleted. IKE SA
 negotiation will start again only when there is traffic to be transmitted
 over the IPsec tunnel.
- In IKEv2, the retransmission interval increases from 1, 2, 4, 8, 16, 32 to 64, in seconds. If no reply is received within eight consecutive transmissions, the peer end is considered dead, and the IKE SA and IPsec SAs are deleted.

• IKE SA timeout processing and IPsec SA timeout processing

In IKEv2, the IKE SA soft lifetime is 9/10 of the IKE SA hard lifetime plus or minus a random number. This reduces the likelihood that two ends initiate renegotiation simultaneously. Therefore, you do not manually set the soft lifetime in IKEv2.

Advantages of IKEv2 over IKEv1

- Simplifies the SA negotiation process, improving efficiency.
- Fixes many cryptographic security vulnerabilities, improving security.
- Supports EAP authentication, improving authentication flexibility and scalability.
 - EAP is an authentication protocol that supports multiple authentication methods. The biggest advantage of EAP is its scalability. That is, new authentication methods can be added without changing the original authentication system. EAP authentication has been widely used in dial-up access networks.
- Employs an Encrypted Payload on basis of ESP. This payload contains both an encryption algorithm and a data integrity algorithm. AES-GCM ensures confidentiality, integrity, and authentication, and works well with IKEv2.

1.2.22 How Many Bits Do the DH Groups Used by VPN Have?

Diffie-Hellman (DH) groups determine the strength of the key used in the key exchange process. Higher DH group numbers are usually more secure, but more time is required to calculate the key.

Table 1-3 lists the number of bits corresponding to the DH groups used by VPN.

Table 1-3 Number of bits corresponding to each DH group

DH Group	Modulus
1	768 bits
2	1024 bits
5	1536 bits
14	2048 bits
15	3072 bits
16	4096 bits
19	256 bits (ECP)
20	384 bits (ECP)
21	521 bits (ECP)

□ NOTE

The following DH algorithms have security risks and are not recommended: DH group 1, DH group 2, and DH group 5.

1.2.23 Can I Visit Websites Across International Borders Using a VPN?

No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

1.2.24 Can I Deploy an Application on the Cloud and a Database in an On-premises Data Center and Connect Them Through a VPN?

Yes.

A VPN connects a VPC and an on-premises data center.

After a VPN is set up, service traffic can be transmitted between the VPC and onpremises data center. For an application server on the cloud, access to an onpremises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check the network latency and packet loss rate to ensure smooth service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

1.2.25 What Are the Differences Between IPsec VPN and SSL VPN in Application Scenarios and Connection Modes?

Application Scenarios

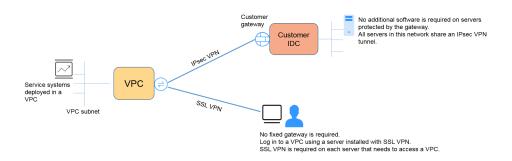
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to enable them to complete IPsec VPN negotiation.

SSL VPN requires a specific client program installed on hosts. Users need to enter usernames and password to connect the hosts to SSL servers.



□ NOTE

IPsec VPN and SSL VPN are supported.

1.2.26 How Will I Be Charged for My Use of a VPN? Will I Be Charged for VPN Gateway EIPs?

VPNs are billed by the following items on a yearly/monthly or pay-per-use basis.

- VPN gateway
- VPN connection

By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. You can purchase additional VPN connections if required.

• EIP bandwidth of a VPN gateway

The VPN gateway bandwidth can be billed by traffic or bandwidth.

- a. A yearly/monthly VPN gateway can only be billed by bandwidth. The price of a yearly/monthly VPN gateway includes the price of the VPN connections that can be created for the gateway and the bandwidth price.
- b. The billing cycle of the pay-per-use billing mode is 1 hour. When you create a pay-per-use VPN gateway, the system prompts you to create VPN connections. By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. If more connection groups are required, you need to purchase them.

∩ NOTE

A VPN gateway cannot share a bandwidth with an EIP bound to an ECS.

1.2.27 What Are the Differences Between Billing the VPN Gateway EIP Bandwidth by Bandwidth and by Traffic?

The VPN gateway EIP bandwidth can be billed by bandwidth or by traffic.

The differences are as follows:

- Billed by bandwidth: The billing cycle is 1 hour. The generated fee depends on the bandwidth.
- Billed by traffic: The fee is calculated based on the outgoing traffic of a VPC generated every hour, which is not affected by the bandwidth.

If you select the more cost-effective yearly/monthly billing mode, VPN gateways can only be billed by bandwidth.

1.2.28 Can a VPN Billed by Traffic Use a Shared Data Package?

Yes.

The VPN service fee includes the EIP fee. An EIP can use a shared data package.

1.2.29 Can the EIP of a VPN Gateway Be Retained After the VPN Gateway Is Deleted?

If a pay-per-use EIP is bound to a pay-per-use VPN gateway, deleting the VPN gateway will also delete the bound EIP.

To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway.

1.2.30 Where Can I Add Routes to Customer Subnets on the VPN Console?

When a VPN connection is created, routes to customer subnets are automatically delivered.

1.2.31 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

To view the status of a VPN connection, click in the **Monitoring** column of the VPN connection.

1.2.32 What Do I Do If a VPN Connection Fails to Be Established?

- Log in to the management console and choose Virtual Private Network > Enterprise - VPN Connections.
- In the VPN connection list, locate the target VPN connection, and choose
 More > Modify Policy Settings on the right to view IKE and IPsec policies of
 the VPN connection.
- 3. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.
 If the IKE SA has been set up in phase 1 but no IPsec SA has been established in phase 2, the IPsec policies at both ends of the VPN connection may be inconsistent.
- 4. Check whether the ACL configurations are correct.
 - If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

5. Ping the two ends of the VPN connection from each other to check whether the VPN connection is normal.

1.2.33 In Which Direction Is the VPN Bandwidth Limited? What Is the Unit of Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited as follows:

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is greater than 10 Mbit/s, the bandwidth in the inbound direction is the same as the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

1.2.34 Can I Restore a VPN Gateway or VPN Connection That Is Incorrectly Deleted?

- A yearly/monthly VPN gateway or VPN connection cannot be restored.
- A pay-per-use VPN gateway can be restored only when the following conditions are met:
 - The VPN gateway was deleted within 24 hours.
 - Both EIPs bound to the VPN gateway have not been unbounded.
 - The VPC or enterprise router interconnected with the VPN gateway is available. If the VPC or enterprise router does not exist, restore the VPC or enterprise router first.
 - Your account is normal and is not in arrears or frozen.
- A VPN connection of a pay-per-use VPN gateway can be restored only when the following conditions are met:
 - The VPN gateway and customer gateway are available. If one of them does not exist, restore it first.
 - Your account is normal and is not in arrears or frozen.

The health check configuration of a pay-per-use VPN connection cannot be restored even after the VPN connection is restored. Therefore, you need to reconfigure the health check function.

1.2.35 Can the Specification of a VPN Gateway Be Changed (for Example, from Professional 1 to Professional 2)?

You can change the specification of a VPN gateway on the VPN gateway page. The following specification changes are subject to the console.

- The specification of a VPN gateway can be changed between Basic and Professional 1.
- The specification of a VPN gateway can be changed between Professional 1 and Professional 2.
- The specification of a VPN gateway cannot be changed from Professional 1 supporting access via non-fixed IP addresses to Professional 1, from Professional 2 supporting access via non-fixed IP addresses to Professional 2, or from Professional 3 supporting access via non-fixed IP addresses to Professional 3.
- When Network Type is set to Public network and Billing Mode is set to Yearly/Monthly, the specification of a VPN gateway can be changed from Professional 1 to Professional 1 supporting access via non-fixed IP addresses, from Professional 2 to Professional 2 supporting access via non-fixed IP addresses, or from Professional 3 to Professional 3 supporting access via non-fixed IP addresses.

1.2.36 Can I Upgrade Classic VPN to Enterprise Edition VPN?

Yes.

If you have such an upgrade requirement, submit a service ticket.

1.2.37 Can Pay-per-Use EIPs Be Bound to a Yearly/Monthly VPN Gateway? What About EIPs Using a Shared Data Package?

Both pay-per-use EIPs and EIPs using a shared data package can be bound to yearly/monthly VPN gateways. When you subscribe to pay-per-use EIPs billed by traffic and buy a shared data package, the EIPs will use the shared data package. After the package quota is used up or if the package expires, the EIPs will continue to be billed on a pay-per-use basis. Note the EIP type (static BGP or dynamic BGP) during the binding.

1.3 Networking and Application Scenarios

1.3.1 Can I Visit Websites Across International Borders Using a VPN?

No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

1.3.2 Can I Deploy an Application on the Cloud and a Database in an On-premises Data Center and Connect Them Through a VPN?

Yes

A VPN connects a VPC and an on-premises data center.

After a VPN is set up, service traffic can be transmitted between the VPC and onpremises data center. For an application server on the cloud, access to an onpremises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check the network latency and packet loss rate to ensure smooth service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

1.3.3 How Many VPN Connections Do I Need to Connect Multiple On-premises Servers to the Cloud?

VPN uses the IPsec technology to connect your on-premises data center to a VPC on the cloud. As such, the number of VPN connections is related to the number of data centers where the servers to be connected to the cloud are located, but not to the number of servers.

Two EIPs can be bound to a VPN gateway for communication with a customer gateway.

- If an on-premises data center has only one egress gateway, all servers or hosts in the data center connect to the Internet through this gateway. In this case, you need to configure a VPN connection group consisting of two VPN connections. That is, configure a VPN connection for each of the two EIPs of the VPN gateway to communicate with the egress gateway in the onpremises data center.
- If an on-premises data center has two egress gateways, the servers or user hosts in the data center connect to the Internet through the two egress gateways. In this case, you need to configure two VPN connection groups, each of which consisting of two VPN connections. That is, configure a VPN connection for each of the two EIPs of each VPN gateway to communicate with both egress gateways in the on-premises data center.

1.3.4 What Are the Differences Between IPsec VPN and SSL VPN in Application Scenarios and Connection Modes?

Application Scenarios

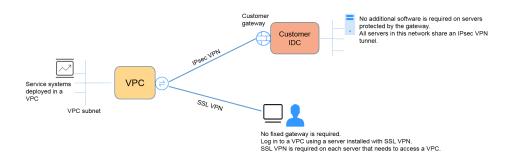
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to enable them to complete IPsec VPN negotiation.

SSL VPN requires a specific client program installed on hosts. Users need to enter usernames and password to connect the hosts to SSL servers.



□ NOTE

IPsec VPN and SSL VPN are supported.

1.3.5 Does a VPN Allow for Communications Between Two VPCs?

- If the two VPCs are in the same region, use a VPC peering connection to connect them.
- If the two VPCs are in different regions, use a VPN to connect them. The operations are as follows:
 - a. Create a VPN gateway for each VPC, and create a VPN connection between the two VPN gateways.
 - b. For the VPN connection, set the customer gateway to the EIP of the peer VPN gateway.
 - c. For the VPN connection, set the customer subnet to the subnet of the peer VPC.
 - d. Set the same pre-shared keys (PSKs) and algorithms for the two VPCs.

1.3.6 What Are the Impacts of a VPN on an On-premises Network? What Are the Changes to the Route for Accessing an ECS?

When configuring a VPN, you need to perform the following operations on the gateway in your on-premises data center:

- Configure IKE and IPsec policies.
- Configure a VPN connection in static routing, BGP routing, or policy-based mode.
- Check the route configuration on the gateway to ensure that traffic destined for a VPC can be routed to the correct outbound interface (interface having an IPsec policy bound).

1.3.7 What Configurations Are Required at Both Ends of a VPN That Connects an On-premises Data Center to a VPC?

To implement the VPN interconnection, create a VPN on the cloud and configure the VPN device in the on-premises data center.

- Create a VPN on the cloud.
 - Buy a VPN gateway, and configure the billing mode, bandwidth, and interconnected VPC.
 - Create a customer gateway and configure the routing mode.
 - Buy a VPN connection, and configure the gateway IP addresses and subnets at both ends, as well as negotiation policies.
- Configure the VPN device in the on-premises data center.
 - a. Configure the public IP address used by the on-premises data center to connect to the cloud, and complete the configurations of IPsec negotiation phase 1 and phase 2 on the VPN device.
 - b. Configure routes, NAT, and security policies on the VPN device.

1.3.8 Can I Connect a Network with Two Egresses to a VPC Through Two VPN Connections?

Yes.

1.3.9 Can I Connect Two VPCs in the Same Region Through a VPN?

No.

You can use a VPC peering connection or Cloud Connect connection to connect two VPCs in the same region.

1.3.10 How Can I Connect Two VPCs in the Same Region?

You can use a VPC peering connection or Cloud Connect connection to connect two VPCs in the same region. VPC peering can only connect VPCs in the same region; Cloud Connect can also connect VPCs in different regions.

1.3.11 How Do I Enable Communications Between Two VPCs and an On-premises Network?

Network Topology

IDC-VPC 1-VPC 2



IDC indicates an on-premises data center. A VPN connection is established between VPC 1 and the IDC.

Procedure

1. Check whether the two VPCs are in the same region.

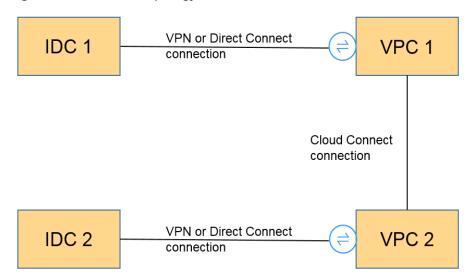
- **FAQs**
- If so, use a VPC peering connection or Cloud Connect connection to connect the two VPCs. Such a connection is free of charge.
- If not, use a Cloud Connect connection to connect the two VPCs. You need to pay for the Cloud Connect bandwidth.
- 2. Establish a VPN connection between the IDC and one VPC (VPC 1 in this example).

In the on-premises data center, set subnets of VPC 1 and VPC 2 as customer subnets. The local subnet of VPC 1 must contain the subnet connected through a VPC peering connection or Cloud Connect connection. The subnet route of the VPC peering connection or Cloud Connect connection needs to destine for the on-premises subnet.

1.3.12 How Do I Connect Four Subnets?

Figure 1-5 shows the network topology.

Figure 1-5 Network topology



- Use a VPN connection or Direct Connect connection to connect IDC 1 to VPC
 1.
- 2. Use a Cloud Connect connection to connect VPC 1 to VPC 2. (You can also use a VPC peering connection to connect VPC 1 to VPC 2 if they are in the same region.)
- 3. Use a VPN connection or Direct Connect connection to connect VPC 2 to IDC 2
- 4. Update VPN subnets, Cloud Connect subnet routes, and Direct Connect subnet routes. Then, the four subnets are reachable to reach other.

1.3.13 Do I Need Two VPN Connections to Connect Four Subnets of Two Regions If Each Region Has Two Subnets?

No.

Only one VPN connection is required between two regions. The subnets can all be added to the VPN connection.

In this scenario, if you attempt to create a second VPN connection, the management console displays a message indicating that a conflict occurs because the two connections have the same customer gateway address.

1.3.14 Can I Access OBS Through a VPN?

Yes.

- 1. With the help of the VPC endpoint service, you can access OBS through a VPN. You need to create two VPC endpoints for the private DNS server and OBS, respectively.
- 2. Configure the private DNS server and routes in your on-premises data center.

1.3.15 How Do I Connect My Personal Computer to the Cloud Through a VPN?

Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.

To use VPN, on-premises devices must support the standard IPsec protocol.

1.3.16 How Do I Access ECSs at Home When My Enterprise Network Has Been Connected to the Cloud Through a VPN?

A VPN is an IPsec VPN that connects an on-premises LAN to a VPC on the cloud. Your home network is not a part of your enterprise LAN, so you cannot directly connect to the VPC on the cloud at home.

If your host at home needs to access VPC resources on the cloud, your host can directly access the EIP of the corresponding service. Alternatively, your host can connect to the LAN of your enterprise through SSL VPN (if supported), and then access VPC resources on the cloud through the LAN.

1.3.17 How Do I Establish a VPN Connection Temporarily If No IPsec-Capable On-Premises Device Is Available After I Purchase a VPN Gateway and VPN Connection?

To establish a VPN connection with the cloud, you must have an on-premises device that supports the standard IPsec protocol and have a fixed public IP address.

If the preceding requirements are not met, you can install third-party IPsec software on a host to temporarily connect to the cloud.

Recommended third-party IPsec software includes strongSwan, Openswan, and TheGreenBow. For details about the interconnection, see **Administrator Guide**.

1.3.18 How Do I Select a Proper Region on the Cloud When I Buy a VPN Gateway?

You can select a VPC in any region when you buy a VPN gateway.

It is recommended that you select the region nearest to your on-premises data center to minimize the impact of the Internet on the VPN.

- To connect to multiple VPCs in the same region, you can use VPN and Direct Connect.
- To connect to multiple VPCs in different regions, you can use VPN and Cloud Connect.

1.4 Billing and Payments

1.4.1 How Will I Be Charged for My Use of a VPN? Will I Be Charged for VPN Gateway EIPs?

VPNs are billed by the following items on a yearly/monthly or pay-per-use basis.

- VPN gateway
- VPN connection

By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. You can purchase additional VPN connections if required.

• EIP bandwidth of a VPN gateway

The VPN gateway bandwidth can be billed by traffic or bandwidth.

- a. A yearly/monthly VPN gateway can only be billed by bandwidth. The price of a yearly/monthly VPN gateway includes the price of the VPN connections that can be created for the gateway and the bandwidth price.
- b. The billing cycle of the pay-per-use billing mode is 1 hour. When you create a pay-per-use VPN gateway, the system prompts you to create VPN connections. By default, 10 VPN connection groups are included free of charge with the purchase of a VPN gateway. If more connection groups are required, you need to purchase them.

A VPN gateway cannot share a bandwidth with an EIP bound to an ECS.

1.4.2 What Are the Differences Between Billing the VPN Gateway EIP Bandwidth by Bandwidth and by Traffic?

The VPN gateway EIP bandwidth can be billed by bandwidth or by traffic.

The differences are as follows:

- Billed by bandwidth: The billing cycle is 1 hour. The generated fee depends on the bandwidth.
- Billed by traffic: The fee is calculated based on the outgoing traffic of a VPC generated every hour, which is not affected by the bandwidth.

If you select the more cost-effective yearly/monthly billing mode, VPN gateways can only be billed by bandwidth.

1.4.3 Can a VPN Billed by Traffic Use a Shared Data Package?

In Enterprise Edition VPN, EIPs can be used as VPN gateway IP addresses.

The VPN service fee includes the EIP fee. An EIP can use a shared data package.

1.4.4 For How Many VPN Connections Will I Be Charged to Connect VPCs in Different Regions of Huawei Cloud?

VPNs can be used to connect VPCs in different regions. The VPN bandwidth and connections of each region will be billed independently. Therefore, when calculating the estimated fees, you need to check the total number of regions and their connection relationships.

For example, assume that Region A needs to establish a VPN connection with Region B and Region C, respectively. The VPN gateway of Region A has two connections; the VPN gateway of Region B has one connection; and the VPN gateway of Region C has one connection.

In this case, you will be charged for four VPN connections.

1.4.5 How Do I Change the Billing Mode of a VPN Gateway from Pay-Per-Use to Yearly/Monthly?

Procedure

Perform the following operations:

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- 3. Click Service List and choose Networking > Virtual Private Network.
- 4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise VPN Gateways**.
- 5. On the **VPN Gateways** page, locate the row that contains the target VPN gateway, choose **More** > **Change Billing Mode** in the **Operation** column.
- 6. In the Change Billing Mode dialog box, click OK.

\cap	NOT	F

The billing mode of a VPN gateway cannot be changed from yearly/monthly to payper-use. The resource quotas of a yearly/monthly VPN gateway can be decreased upon a renewal.

- 7. Confirm the VPN gateway information, set a renewal duration, and click **Pay**.
- 8. On the payment page, confirm the order information, select a coupon or discount, select a payment method, and click **Pay**.

\cap	NO.	ΓF

Changing the billing mode of a VPN gateway from pay-per-use to yearly/monthly will not affect your services.

1.4.6 Will a Yearly/Monthly VPN Gateway Be Automatically Renewed?

Yes.

Renewal fees will be automatically collected from your balance.

A yearly/monthly VPN gateway needs to be prepaid. To ensure that your connection is normal, top up your account if your balance is not enough.

1.4.7 Can I Unsubscribe from a Yearly/Monthly VPN Gateway?

Yes.

On the **VPN Gateways** page, locate the row that contains the VPN gateway you want to unsubscribe and choose **More** > **Delete** in the **Operation** column. After you unsubscribe from a yearly/monthly VPN gateway, all VPN connections created for the gateway will also be deleted and cannot be recovered.

After the unsubscription, the remaining prepaid fees will be refunded.

1.4.8 When Will My VPN Resources Be Frozen? How Can I Unfreeze the VPN Resources?

- VPN resources billed on a yearly/monthly basis enter the grace period if they expire and are not renewed. During the grace period, you can access and use the resources. If the grace period ends and you have not renewed the resources, the resources enter the retention period and are frozen. Frozen resources are unavailable and cannot be modified or deleted. If the retention period ends and you still have not renewed your resources, they will be released and cannot be restored. To ensure that resources are available, renew them before they expire.
- If pay-per-use VPN resources are in arrears, the resources enter the grace period, during which you can still access and use the resources. If the grace period ends and you have not paid off the arrears, the resources enter the retention period, during which the resources are frozen. Frozen resources are unavailable and cannot be modified or deleted. If the retention period ends and you still have not topped up your account and paid off the arrears, the resources will be released and cannot be restored. To ensure that resources are available, top up your account and pay off the arrears before the resources expire.
- The impact of shared bandwidth freezing on EIPs is subject to the EIP documentation. For details, see Why My EIPs Are Frozen? How Do I Unfreeze My EIPs?.
- Frozen VPN resources will become available after you renew them or top up your account.

1.5 Operations on the Console

1.5.1 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

VPC

- A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
- You can use the VPN service to securely access ECSs in a VPC.

VPN gateway

- A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have multiple VPN gateways, and one VPN gateway can have multiple VPN connections.
- With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.

VPN connection

A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

◯ NOTE

The number of VPN connections is irrelevant to the number of local subnets or the number of customer subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

1.5.2 How Long Does It Take for Delivered VPN Configurations to Take Effect?

It takes 1–5 minutes for the VPN configurations to take effect.

5	$\overline{}$	 	_	_	_
		 NI		П	-
		 w	\sim		_

After VPN configurations take effect, configure your gateway device on your on-premises network to complete tunnel negotiation with the VPN gateway.

1.5.3 Why Is a VPN Connection Always in Not Connected State After Its Configuration Is Complete?

The configuration may be incorrect.

- 1. At the two ends (cloud and on-premises data center) of the VPN connection, ensure that the pre-shared keys (PSKs) and negotiation information are consistent, the local and remote subnets are reversed, and the local and remote gateways are also reversed.
- 2. Ensure that routes, NAT, and security policies are correctly configured on the device in your on-premises data center.

1.5.4 Can the EIP of a VPN Gateway Be Retained After the VPN Gateway Is Deleted?

If a pay-per-use EIP is bound to a pay-per-use VPN gateway, deleting the VPN gateway will also delete the bound EIP.

To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway.

1.5.5 What Information About a Created VPN Can Be Modified and What Information Cannot Be Modified?

- VPN gateway
 - You can modify the following information:
 - Name
 - Local subnet
 - Billing mode (yearly/monthly or pay-per-use)
 - Active and standby EIPs
 - To modify the active or standby EIP, unbind the original EIP and bind a new one.
 - If a VPN connection has been created for an EIP, the EIP cannot be unbound.
 - You can modify EIP attributes, such as the name, type, and bandwidth. For details, see the EIP service documentation.
 - Specification

The supported specifications are subject to those displayed on the management console.

Number of VPN connection groups

The number of VPN connection groups needs to be specified only when **Billing Mode** is set to **Yearly/Monthly**.

- You cannot modify the following information:
 - Region
 - Association mode (VPC or enterprise router)
 - Enterprise router

The associated enterprise router needs to be specified only when **Associate With** is set to **Enterprise Router**.

- VPC
- Interconnection subnet
- BGP ASN
- AZ

- Customer gateway
 - You can modify the following information:
 - Name
 - You cannot modify the following information:
 - BGP ASN

The BGP ASN needs to be specified only when **Routing Mode** is set to **Dynamic (BGP)**.

- Gateway IP address
- VPN connection
 - You can modify the following information:
 - Name
 - Billing mode. Only pay-per-use billing can be changed to yearly/ monthly billing.
 - Local interface address
 - Customer gateway
 - Customer subnet
 - Policy configuration, including IKE and IPsec policies
 - Policy
 - PSK
 - Branch interconnection
 - You cannot modify the following information:
 - VPN gateway
 - EIP
 - VPN type (route-based or policy-based)
 - Routing mode (static or BGP)

The routing mode needs to be specified only when **VPN Type** is set to **Route-based**.

Link detection configuration

The link detection configuration is available only when **VPN Type** is set to **Route-based**.

Policy configuration, including the source and destination CIDR blocks

The policy configuration is available only when **VPN Type** is set to **Policy-based**.

1.5.6 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to configure policy rules (ACL rules) for a VPN connection on the management console only when **VPN Type** is set to **Policy-based**.

1.5.7 What Do I Do If an Exception Occurs When I Add a Customer Subnet During VPN Connection Creation?

Check whether this customer subnet is involved in a route of a VPC peering, Cloud Connect, or Direct Connect connection. If so, a route conflict occurs and you need to delete the route and create a new one to prevent the conflict.

1.5.8 Where Can I Configure Routes to Customer Subnets on the VPN Console?

When a VPN connection is created, routes to customer subnets are automatically delivered.

1.5.9 Can I Call APIs to Manage Huawei Cloud VPN Resources?

Yes.

1.5.10 What Are a Customer Gateway and a Customer Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a customer subnet and a customer gateway, respectively.

A customer gateway's IP address is a public IP address.

1.5.11 How Do I Disable PFS When Creating a VPN Connection?

Cloud side

In the VPN connection configuration, set **PFS** in the IPsec policy to **Disable**. By default, PFS is enabled on the cloud side.

Customer gateway in your on-premises data center

By default, PFS is disabled on some vendors' devices. For details about how to disable PFS, see the corresponding product documentation.

□ NOTE

Ensure that the PFS settings on the cloud side and the customer gateway are consistent. Otherwise, the negotiation will fail.

For security purposes, you are advised to enable PFS on both the cloud side and the customer gateway.

1.5.12 How Many Local and Customer Subnets Can I Add to a VPN?

- You can configure a maximum of 50 local subnets for each VPN gateway.
- You can configure a maximum of 50 customer subnets for each VPN connection.
- You can configure a maximum of five policy rules for each VPN connection.
 You can configure 1 source CIDR block and 50 destination CIDR blocks in each policy rule.

1.5.13 What Are the Precautions for Configuring the Local and Customer Subnets for a VPN Connection?

- The number of local subnets and the number of customer subnets are limited. If the number of local or customer subnets exceeds the upper limit, aggregate the subnets.
 - Maximum number of local subnets for each VPN gateway: 50
 - Maximum number of customer subnets for each VPN connection: 50
- A local subnet cannot include the CIDR block of a customer subnet, whereas a customer subnet can include the CIDR block of a local subnet.
- There are routes pointing to the local subnets in the VPC where the VPN gateway resides.
- If there are two connections (connection A and connection B) created for a VPN gateway, and the customer subnet of connection A is within that of connection B, when the destination network to be accessed belongs to the overlapped CIDR block, the connection created first is matched first, regardless of the connection status. (Mask length match is not used for VPN connections created in policy-based mode.)

1.5.14 Why Is a VPN Connection in Not Connected State on the Management Console When It Is Already Available?

There is a certain delay in updating the VPN connection state on the management console.

If the service access is normal, the VPN connection has been established. The state of the VPN connection will be updated to **Connected** after several minutes.

1.5.15 What Can I Do If a Message Is Displayed Indicating That the VPN Connection Does Not Exist After Negotiation Policies Are Modified?

This problem is caused by the page refresh interval.

When you modify advanced policy settings, the system deletes the VPN connection and then creates one. If the page temporarily displays a message indicating that the connection is being deleted or created, do not create the same connection with the same local subnet, customer subnet, and customer gateway again.

If the page remains in the connection deleting or creating state for a long time, submit a service ticket.

1.5.16 What Is the Maximum Bandwidth Supported by a VPN Gateway?

The maximum bandwidth supported by a VPN gateway is 1 Gbit/s.

1.5.17 Which IKE Version Should I Select When I Create a VPN Connection?

IKEv2 is recommended because IKEv1 is not secure. In addition, IKEv2 outperforms IKEv1 in connection negotiation and establishment, authentication methods, dead peer detection (DPD) timeout processing, and security association (SA) timeout processing.

IKEv2 will be widely used, and IKEv1 will gradually phase out.

Introduction to IKEv1 and IKEv2

- As a hybrid protocol, IKEv1 brings some security and performance defects due to its complexity. As such, it has become a bottleneck in the IPsec system.
- IKEv2 addresses the issues of IKEv1 while retaining basic functions of IKEv1.
 IKEv2 is more simplified, efficient, secure, and robust than IKEv1. Additionally, IKEv2 is defined by RFC 4306 in a single document, whereas IKEv1 are defined in multiple documents. By minimizing core functions and default password algorithms, IKEv2 greatly improves interoperability between different IPsec VPNs.

Security Risks of IKEv1

- The cryptographic algorithms supported by IKEv1 have not been updated for more than 10 years. In addition, IKEv1 does not support strong cryptographic algorithms such as AES-GCM and ChaCha20-Poly1305. For IKEv1, the E (Encryption) bit in the ISALMP header specifies that the payloads following the ISALMP header are encrypted, but any data integrity verification of those payloads is handled by a separate hash payload. This separation of encryption from data integrity protection prevents the use of authenticated encryption (AES-GCM) with IKEv1.
- IKEv1 is vulnerable to DoS amplification attacks and half-open connection attacks. After responding to spoofed packets, the responder maintains initiator-responder relationships, consuming a large number of system resources.

This defect is inherent to IKEv1 and is addressed in IKEv2.

The aggressive mode of IKEv1 is not secure. In this mode, information packets are not encrypted, posing risks of information leakage. There are also bruteforce attacks targeting at the aggressive mode, such as man-in-the-middle attacks.

Differences Between IKEv1 and IKEv2

Negotiation process

- IKEv1 is complex and consumes a large amount of bandwidth. IKEv1 SA negotiation consists of two phases. In IKEv1 phase 1, an IKE SA is established in either main mode or aggressive mode. Main mode requires three exchanges between peers totaling six ISAKMP messages, whereas aggressive mode requires two exchanges totaling three ISAKMP messages. Aggressive mode is faster, but does not provide identity protection for peers as key exchange and identity authentication are performed simultaneously. In IKEv1 phase 2, IPsec SAs are established through three ISAKMP messages in quick mode.
- Compared with IKEv1, IKEv2 simplifies the SA negotiation process. IKEv2 requires only two exchanges, totaling four messages, to establish an IKE SA and a pair of IPsec SAs. To create multiple pairs of IPsec SAs, only one additional exchange is needed for each additional pair of SAs.

For IKEv1 negotiation, its main mode involves nine (6+3) messages, and its aggressive mode involves six (3+3) messages. In contrast, IKEv2 negotiation requires only four (2+2) messages.

Authentication methods

- Only IKEv1 (requiring an encryption card) supports digital envelope authentication (HSS-DE).
- IKEv2 supports Extensible Authentication Protocol (EAP) authentication.
 IKEv2 can use an AAA server to remotely authenticate mobile and PC users and assign private IP addresses to these users. IKEv1 does not provide this function and must use L2TP to assign private IP addresses.
- Only IKEv2 supports IKE SA integrity algorithms.

DPD timeout processing

- Only IKEv1 supports the retry-interval parameter. If a device sends a DPD packet but receives no reply within the specified retry-interval, the device records a DPD failure event. When the number of DPD failure events reaches 5, both the IKE SA and IPsec SAs are deleted. IKE SA negotiation will start again only when there is traffic to be transmitted over the IPsec tunnel.
- In IKEv2, the retransmission interval increases from 1, 2, 4, 8, 16, 32 to 64, in seconds. If no reply is received within eight consecutive transmissions, the peer end is considered dead, and the IKE SA and IPsec SAs are deleted.

• IKE SA timeout processing and IPsec SA timeout processing

In IKEv2, the IKE SA soft lifetime is 9/10 of the IKE SA hard lifetime plus or minus a random number. This reduces the likelihood that two ends initiate renegotiation simultaneously. Therefore, you do not manually set the soft lifetime in IKEv2.

Advantages of IKEv2 over IKEv1

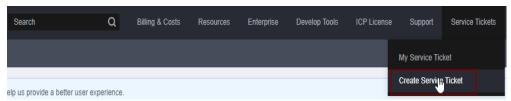
- Simplifies the SA negotiation process, improving efficiency.
- Fixes many cryptographic security vulnerabilities, improving security.
- Supports EAP authentication, improving authentication flexibility and scalability.

- EAP is an authentication protocol that supports multiple authentication methods. The biggest advantage of EAP is its scalability. That is, new authentication methods can be added without changing the original authentication system. EAP authentication has been widely used in dial-up access networks.
- Employs an Encrypted Payload on basis of ESP. This payload contains both an encryption algorithm and a data integrity algorithm. AES-GCM ensures confidentiality, integrity, and authentication, and works well with IKEv2.

1.5.18 What Types of VPN Service Tickets Are There? How Do I Create a VPN Service Ticket?

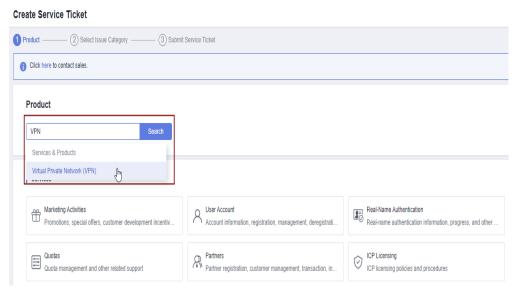
- 1. Log in to the management console.
- 2. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket**.

Figure 1-6 Create Service Ticket



3. Search for VPN and select Virtual Private Network (VPN).

Figure 1-7 Selecting Virtual Private Network (VPN)



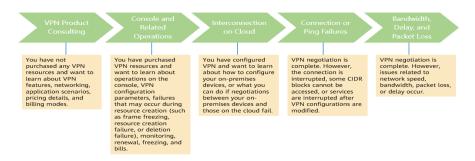
4. Select an issue category.

Figure 1-8 Select Issue Category



When you **submit a service ticket**, select an issue category to facilitate problem handling.

Figure 1-9 Issue category and classification basis



1.5.19 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

□ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

1.5.20 What VPN Resources Can Be Monitored?

VPN gateway

The following bandwidth information of a VPN gateway IP address can be monitored: inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view the monitoring information, click $\stackrel{\triangle}{=}$ in the **Gateway IP Address** column in the VPN gateway list.

VPN connection

The following information about a VPN connection can be monitored: VPN connection status, average link round-trip time (RTT), maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate.

To monitor average link RTT, maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate, click the VPN connection name and click **Add** in the **Health Check** area on the **Summary** tab page to add health check items. Private network metrics can be configured only when the VPN connection uses the static routing mode and the NQA function is enabled.

To view the status of a VPN connection, click $\stackrel{\triangle}{=}$ in the **Monitoring** column of the VPN connection.

1.5.21 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

To view the status of a VPN connection, click $\stackrel{\triangle}{=}$ in the **Monitoring** column of the VPN connection.

1.6 VPN Negotiation and Interconnection

1.6.1 What Devices Can Be Connected to Huawei Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

■ NOTE

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- The following products can connect to the cloud through VPNs:
 - Devices: Huawei firewalls and access routers (ARs), Hillstone firewalls, and Check Point firewalls
 - Cloud services: Alibaba Cloud, Tencent Cloud, Amazon Web Services (AWS), and Microsoft Azure
 - Software: strongSwan
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.
 - Most enterprise-class routers and firewalls support the IPsec protocol.
- Some devices support IPsec VPN only after you purchase required software licenses.
 Your on-premises data center administrator can check with the device vendor whether a license is required based on the device model.

1.6.2 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 1-4 VPN negotiation parameters

Protocol	Parameter	Value
IKE	Authentication Algorithm	 MD5(Insecure. Not recommended.) SHA1(Insecure. Not recommended.) SHA2-256 SHA2-384 SHA2-512 The default value is SHA2-256.
	Encryption Algorithm	 3DES (Insecure. Not recommended.) AES-128(Insecure. Not recommended.) AES-192(Insecure. Not recommended.) AES-256(Insecure. Not recommended.) AES-128-GCM-16 AES-256-GCM-16 The default value is AES-128.

Protocol	Parameter	Value
	DH Algorithm	 Group 1(Insecure. Not recommended.) Group 2(Insecure. Not recommended.) Group 5(Insecure. Not recommended.) Group 14(Insecure. Not recommended.) Group 15 Group 16 Group 19 Group 20 Group 21 The default value is Group 15.
	Version	 v1 (For security reasons, IKEv1 is not recommended. If your devices support IKEv2, select IKEv2. For VPN connections set up using SM series cryptographic algorithms, only IKEv1 is supported.) v2 The default value is v2.
	Lifetime (s)	86400 (default value) Unit: second Value range: 60 to 604800
	Local ID	 IP Address The local IP address is automatically displayed as the EIP of the VPN gateway, removing the need to manually configure it. FQDN By default, the local ID type is IP address and the local ID value is the EIP of the VPN gateway.
	Customer ID	 IP Address FQDN By default, the customer ID type is IP address and the customer ID value is the public IP address of the customer gateway.

Protocol	Parameter	Value
IPsec	Authentication Algorithm	 SHA1(Insecure. Not recommended.) MD5(Insecure. Not recommended.) SHA2-256 SHA2-384 SHA2-512 The default value is SHA2-256.
	Encryption Algorithm	 3DES (Insecure. Not recommended.) AES-128(Insecure. Not recommended.) AES-192(Insecure. Not recommended.) AES-256(Insecure. Not recommended.) AES-128-GCM-16 AES-256-GCM-16 The default value is AES-128.
	PFS	 Disable(Insecure. Not recommended.) DH group 1(Insecure. Not recommended.) DH group 2(Insecure. Not recommended.) DH group 5(Insecure. Not recommended.) DH group 14(Insecure. Not recommended.) DH group 15 DH group 16 DH group 19 DH group 20 DH group 21 The default value is Group 15.
	Transfer Protocol	ESP (default value)
	Lifetime (s)	3600 (default value) Unit: second Value range: 30 to 604800

• Perfect Forward Secrecy (PFS) is a security feature.

IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. When PFS is enabled, an additional DH exchange will be performed during IPsec SA negotiation to generate a new IPsec SA key, improving IPsec SA security.

- For security purposes, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the gateway device in your on-premises data center and the PFS settings on both ends are the same. Otherwise, the negotiation will fail.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB on the cloud side and cannot be changed for the VPN. This parameter is not involved in negotiation and has no impact on the establishment of an IPsec SA.

1.6.3 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

1.6.4 How Do I Configure a VPN on an On-premises Device? (Example of Configuring VPN on a Huawei USG6600 Series Firewall)

VPN settings on the device in your on-premises data center must be consistent with those on the cloud. Otherwise, the VPN cannot be established.

To set up a VPN, you also need to configure an IPsec VPN tunnel on the router or firewall in your on-premises data center. The configuration method varies according to your network device in use. For details, see the configuration guide of your network device.

The following uses a Huawei USG6600 series firewall running V100R001C30SPC300 as an example to describe how to configure a VPN on an on-premises device.

Assume that the subnets of an on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the public IP address of the IPsec tunnel egress in the on-premises data center is 1.1.1.2. The subnets of a VPC are 192.168.1.0/24 and 192.168.2.0/24, and the public IP address of the IPsec tunnel egress in the VPC is 1.1.1.1.

Procedure

- 1. Log in to the command line interface (CLI) of the firewall.
- 2. Check firewall version information.

display version 17:20:502017/03/09 Huawei Versatile Security Platform Software Software Version: USG6600 V100R001C30SPC300 (VRP (R) Software, Version 5.30)

3. Create an ACL.

acl number 3065 vpn-instance vpn64 rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255

rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 a

4. Create an IKE proposal.

ike proposal 64 dh group5 authentication-algorithm sha1 integrity-algorithm hmac-sha2-256 sa duration 3600

5. Create an IKE peer and bind it to the created IKE proposal. The peer IP address is 1.1.1.1.

ike peer vpnikepeer_64
pre-shared-key ******** (******** indicates a pre-shared key.)
ike-proposal 64
undo version 2
remote-address vpn-instance vpn64 1.1.1.1
sa binding vpn-instance vpn64

6. Configure an IPsec proposal.

IPsec proposal IPsecpro64 encapsulation-mode tunnel esp authentication-algorithm sha1 q

7. Configure an IPsec policy and bind the IPsec proposal to it.

IPsec policy vpnIPsec64 1 isakmp security acl 3065 pfs dh-group5 ike-peer vpnikepeer_64 proposal IPsecpro64 local-address 1.1.1.2

8. Apply the IPsec policy to the corresponding sub-interface.

interface GigabitEthernet0/0/2.64 IPsec policy vpnIPsec64 q

9. Test connectivity.

Test the connectivity between your ECS on the cloud and a host in your onpremises data center, as shown in **Figure 1-10**.

Figure 1-10 Connectivity test

1.6.5 Does VPN Support Interconnection with a Customer Gateway Through a Domain Name?

No. VPN supports interconnection with a customer gateway only through the public IP address of the customer gateway.

1.6.6 How Many Tunnels Does My VPN Connection Have?

Number of tunnels in a VPN connection = Number of local subnets x Number of customer subnets

- An IPsec tunnel is in Active state when data traffic is transmitted between two subnets at the two ends of the IPsec tunnel.
- A VPN connection is in Connected state as long as one of its tunnels is in Active state.

1.6.7 How Do I Allow Specific Hosts to Access a VPC Subnet Through a Created VPN Connection?

Restrictions in the on-premises data center:

- Access control policies on the VPN device
- ACL rules on the router or switch

Restrictions at the cloud side:

- Security group rules that permit access only from specified IP addresses
- ACL rules configurations

You are advised not to change the local or customer subnet to control access.

1.6.8 Do VPNs Have the DPD Function Enabled?

Yes.

By default, the dead peer detection (DPD) function is enabled for VPNs to detect the state of the IKE process in an on-premises data center.

After three consecutive detection failures, the IKE process in the on-premises data center is considered abnormal, and the tunnel on the cloud is automatically deleted.

The DPD protocol does not require that the peer end also be configured with DPD, but it requires that the peer end be able to respond to DPD detections. To ensure consistent tunnel states at the two ends, it is recommended that you enable DPD on your on-premises gateway to detect the IKE process state of the VPN service.

Deleting the tunnel in the case of DPD detection failures will not affect service stability.

1.6.9 How Can I Use Security Groups to Prevent VPN Access to Some ECSs in a VPC to Implement Security Isolation?

You can configure security groups to allow access only to specific CIDR blocks or ECSs in a VPC through a VPN.

Configuration example: Prevent the customer subnet 192.168.1.0/24 from accessing ECSs in the VPC subnet 10.1.0.0/24.

Procedure:

- 1. Create security groups 1 and 2.
- 2. Configure security group 1 to deny access from subnet 192.168.1.0/24.
- 3. Configure security group 2 to permit access from subnet 192.168.1.0/24.
- 4. Associate ECSs in subnet 10.1.0.0/24 with security group 1 and associate other ECSs in the VPC with security group 2.

1.6.10 Will a VPN Connection Be Re-established After Its Configuration Is Modified?

A VPN connection consists of local subnets, customer subnets, customer gateway, pre-shared keys (PSKs), IKE negotiation policy, and IPsec negotiation policy. A VPN connection is modified if any of the following happens:

- If the local and customer subnets are modified, the connection ID will remain unchanged. If not all subnets are updated, the established tunnel between subnets will not be re-established.
- If the IP address of the customer gateway is changed, the connection ID will remain unchanged, but the VPN connection will be re-established.

- If only the PSKs are changed, the connection ID and status will remain unchanged. The PSK will be checked again during renegotiation. If the PSKs do not match, the renegotiation fails.
- If a negotiation policy is modified (PSK verification is required), the connection ID will be changed and the connection needs to be re-established.

1.6.11 Why Cannot I Initiate Negotiation from Amazon Web Services to Huawei Cloud After They Are Interconnected?

After a VPN connection is established between Amazon Web Services (AWS) and Huawei Cloud, AWS works in Response mode and does not initiate negotiation. As such, SA establishment will not be triggered when an AWS EC2 accesses a Huawei Cloud ECS.

According to the AWS documentation, the customer side (the cloud connected to AWS) initiates negotiation by default, and you can also enable the AWS side to initiate negotiation.

1.6.12 How Do I Configure DPD for Interconnection with the Cloud?

By default, DPD is enabled on the cloud side and cannot be disabled.

You can configure DPD as follows:

• DPD-type: on-demand

• DPD idle-time: 30s

• DPD retransmit-interval: 15s

DPD retry-limit: 3

DPD msg: seq-hash-notify

The **DPD msg** format at both ends of the VPN connection must be the same, but the DPD type, idle time, retransmission interval, and retry limit can be different.

1.6.13 What Should I Do If My Firewall Cannot Receive Response Packets from the VPN Gateway in IKE Phase 1?

- 1. Check whether the public IP addresses of the two ends can communicate with each other by running the ping command. By default, the VPN gateway EIPs can be pinged.
- 2. Verify that the on-premises gateway (firewall) and VPN gateway can exchange packets with UDP ports 500 and 4500.
- 3. Verify that the source port number is not translated when the on-premises gateway connects to the VPN gateway. In a NAT traversal scenario, ensure that the source port number is not changed after NAT traversal.
- 4. Verify that IKE negotiation parameter settings are consistent at the two ends of the VPN.

In a NAT traversal scenario, set the customer ID type to IP address and the value to the post-NAT public IP address of the on-premises gateway.

1.6.14 What Should I Do If My Firewall Cannot Receive Response Packets from a VPN Subnet?

- 1. Check the routes, security policies, NAT configuration, interesting traffic, and negotiation policies for phase 2 negotiation on the on-premises gateway device.
 - Route configurations: Route the data for accessing cloud subnets to tunnels.
 - Security policies: Allow traffic from on-premises subnets to cloud subnets.
 - NAT policies: Do not perform source NAT on the traffic originated from on-premises subnets to cloud subnets.
 - Interesting traffic: The interesting traffic configurations at both ends are reversed at the two ends of a VPN connection. The address object name cannot be used for the interesting traffic configured using IKEv2.
 - Negotiation policies: Ensure the negotiations policies, especially PFS, at both ends are the same.
- 2. After confirming that both phase 1 and phase 2 negotiations are normal, ensure that the security groups on the cloud permit ICMP packets originated from on-premises subnets to cloud subnets.

1.6.15 How Many Bits Do the DH Groups Used by VPN Have?

Diffie-Hellman (DH) groups determine the strength of the key used in the key exchange process. Higher DH group numbers are usually more secure, but more time is required to calculate the key.

Table 1-5 lists the number of bits corresponding to the DH groups used by VPN.

Table 1-5 Number of bits corresponding to each DH group

DH Group	Modulus
1	768 bits
2	1024 bits
5	1536 bits
14	2048 bits
15	3072 bits
16	4096 bits
19	256 bits (ECP)
20	384 bits (ECP)
21	521 bits (ECP)

The following DH algorithms have security risks and are not recommended: DH group 1, DH group 2, and DH group 5.

1.7 Connection or Ping Failure

1.7.1 Why Is a VPN Connection Always in Not Connected State After Its Configuration Is Complete?

The configuration may be incorrect.

- At the two ends (cloud and on-premises data center) of the VPN connection, ensure that the pre-shared keys (PSKs) and negotiation information are consistent, the local and remote subnets are reversed, and the local and remote gateways are also reversed.
- 2. Ensure that routes, NAT, and security policies are correctly configured on the device in your on-premises data center.

1.7.2 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes of disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- Dead Peer Detection (DPD) is not configured on the device in your onpremises data center.
- Configuration is modified when the VPN connection is in use.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- At the two ends of the VPN connection, the local and customer subnet configurations are reversed.
- SA lifetime settings at both ends of the VPN connection are the same.
- DPD is enabled on the on-premises gateway device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.
- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway device is large enough for the VPN connection.

 VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway device.

1.7.3 How Do I Quickly Restore an Interrupted IPsec VPN Connection?

- If negotiation cannot be triggered, check connectivity between the public IP addresses of gateways at both ends of the IPsec VPN connection. For example, you can run the ping command to check the connectivity. By default, a VPN gateway responds to ICMP packets.
- 2. If connectivity is normal, check whether link switching occurs between outbound interfaces. That is, check whether the traffic for access to the VPN gateway is forwarded out from a non-negotiated interface.
- 3. If traffic is forwarded through the correct link, change the PSKs at both ends of the IPsec VPN connection to trigger re-negotiation.
- 4. If re-negotiation fails, check whether the negotiation policies configured at both ends are consistent and whether the interesting traffic configurations at both ends are reversed (same number of configurations and same subnets).
- 5. If the negotiation policies and interesting traffic configurations are correct, disable the VPN connection on the on-premises device. After the VPN connection state changes to **Not connected**, enable the VPN connection on the on-premises device and trigger a data flow.
- 6. If negotiation still fails, perform the following operations:
 - a. Record the negotiation policies, PSK, local subnets, customer gateway, and customer subnets of the VPN connection.
 - b. Use the existing VPN gateway to create another VPN connection. The negotiation policies, PSK, and local subnets are the same as those of the original VPN connection. The customer gateway and customer subnets can be configured randomly.
 - c. After the new VPN connection is created, delete the original VPN connection, and change the customer gateway and customer subnets of the new VPN connection to be the same as those of the original VPN connection.
 - d. Trigger the negotiation again.

If the fault persists, **submit a service ticket** to customer service personnel.

1.7.4 What Will Happen If Traffic Exceeds the Bandwidth of a VPN Gateway?

The VPN gateway bandwidth applies to traffic in the outbound direction of a VPC. If outbound traffic in the VPC exceeds the bandwidth, network congestion will occur, some subnets cannot be accessed, or even the VPN connection will be interrupted due to VPN detection timeout.

In this case, you are advised to increase the VPN gateway bandwidth.

□ NOTE

The maximum bandwidth of a VPN gateway is 1 Gbit/s.

1.7.5 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

1.7.6 Why Cannot ECSs at the Two Ends of a Normal Cross-Region VPN Connection Ping Each Other?

By default, a security group permits outbound traffic with any port number. To allow inbound traffic, add inbound rules to the security group. Ensure that the security group associated with the ECS that needs to receive ping packets allows inbound ICMP requests.

1.7.7 Why Cannot Subnets at the Two Ends of a Normal VPN Connection Access Each Other?

The VPN connection is normal, indicating that the negotiation parameters at both ends of the VPN connection are correct. You need to perform the following operations:

- Verify that routes to the VPN device in your on-premises data center are correctly configured.
- Verify that inter-subnet data exchange is allowed on the VPN device.
- Verify that NAT is not performed on the on-premises subnets that need to access the cloud.
- Verify that mutual access between the public IP addresses of the VPN gateway and customer gateway is permitted.

1.7.8 What Do I Do If a VPN Connection Is Interrupted and a Message Indicating Data Flow Mismatch Is Displayed?

This is usually caused by ACL configuration mismatch between the local and customer gateways.

- 1. Verify that at the two ends of the VPN connection, the local and remote subnets are reversed and the ACL configurations are also reversed.
- Use the subnet/mask format when you configure interesting traffic in your on-premises data center. Do not use the address object mode since it may cause incompatibility issues.

1.7.9 What Do I Do If a VPN Connection Is Interrupted and a Message Indicating DPD Timeout Is Displayed?

This happens because there is no data exchange over the VPN connection. When the SA lifetime ends, the VPN connection is deleted as the peer end does not respond to the dead peer detection (DPD).

Solution

1. Enable DPD on the on-premises gateway device, and verify that data flows from both ends can trigger connection establishment.

2. Deploy a ping shell script on the servers at both ends. Alternatively, configure a keepalive function (for example, NQA) on the on-premises gateway device to keep the connection alive.

1.7.10 Why Is a VPN Connection in Not Connected State on the Management Console When It Is Already Available?

There is a certain delay in updating the VPN connection state on the management console.

If the service access is normal, the VPN connection has been established.

1.7.11 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

To view the status of a VPN connection, click in the **Monitoring** column of the VPN connection.

1.7.12 What Do I Do If a VPN Connection Fails to Be Established?

- Log in to the management console, and choose Virtual Private Network > Enterprise - VPN Connections.
- 2. In the VPN connection list, locate the target VPN connection, and choose **More** > **Modify Policy Settings** on the right to view IKE and IPsec policies of the VPN connection.
- 3. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.
 If the IKE SA has been set up in phase 1 but no IPsec SA has been established in phase 2, the IPsec policies at both ends of the VPN connection may be inconsistent.
- 4. Check whether the ACL configurations are correct.
 - If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

5. Ping the two ends of the VPN connection from each other to check whether the VPN connection is normal.

1.7.13 What Should I Do If I Cannot Access the ECSs on the Cloud from My On-premises Data Center or LAN After the VPN Connection Has Been Set Up?

The security group denies access from all sources by default. If you want to access your ECSs, configure security group rules to permit access from your on-premises subnets.

1.7.14 Why Is the State of a Successfully Created VPN Connection Displayed as Not Connected?

There is a delay in updating the state of a VPN connection on the management console. Please refresh the page in about 2 minutes.

1.7.15 Do VPNs Have the DPD Function Enabled?

Yes.

By default, the dead peer detection (DPD) function is enabled for VPNs to detect the state of the IKE process in an on-premises data center.

After three consecutive detection failures, the IKE process in the on-premises data center is considered abnormal, and the tunnel on the cloud is automatically deleted.

The DPD protocol does not require that the peer end also be configured with DPD, but it requires that the peer end be able to respond to DPD detections. To ensure consistent tunnel states at the two ends, it is recommended that you enable DPD on your on-premises gateway to detect the IKE process state of the VPN service.

□ NOTE

Deleting the tunnel in the case of DPD detection failures will not affect service stability. DPD can detect exceptions in the IKE process at the peer end in time and reset the tunnel to ensure tunnel synchronization between the two ends. After a tunnel is deleted, if there is traffic transmitted over the tunnel, the tunnel can be re-established through negotiation.

1.8 Public Addresses

1.8.1 Can the EIP of a VPN Gateway Be Retained After the VPN Gateway Is Deleted?

If a pay-per-use EIP is bound to a pay-per-use VPN gateway, deleting the VPN gateway will also delete the bound EIP.

To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway.

1.8.2 Can EIPs Be Used as VPN Gateway IP Addresses?

Yes.

When creating a VPN gateway, you can bind EIPs as the gateway IP addresses.

1.8.3 Do I Need to Purchase EIPs for Hosts to Communicate with Each Other Through a VPN?

If your on-premises hosts need to access an ECS on the cloud through a VPN, you do not need to purchase any EIPs for the ECS.

If an ECS needs to provide services accessible from the Internet, you need to purchase an EIP for the ECS.

1.8.4 Why Does an ECS Have EIP Access Information After I Enable a VPN?

A possible cause is that the ECS has an EIP bound before the VPN is used. In this scenario, you can access the ECS through both the VPN and the EIP.

To allow only hosts on the VPN to access the ECS, unbind the EIP from the ECS after the VPN connection is established.

1.8.5 Can My On-premises Gateway Have a Non-fixed Public IP Address?

Yes.

If the VPN gateway that you purchased supports access via non-fixed IP addresses, your customer gateway device in the on-premises data center can use a non-fixed IP address to connect to the cloud.

■ NOTE

Whether a VPN gateway supports access via non-fixed IP addresses depends on the region selected on the management console.

1.9 Route Configurations

1.9.1 What Are a Customer Gateway and a Customer Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a customer subnet and a customer gateway, respectively.

A customer gateway's IP address is a public IP address.

1.9.2 Where Can I Add Routes to Customer Subnets on the VPN Console?

When a VPN connection is created, routes to customer subnets are automatically delivered.

1.9.3 Do I Need to Add a Route for an ECS with Multiple NICs to Reach the On-premises Network?

- If the primary network interface card (NIC) is used to establish a VPN connection with the on-premises network, no route needs to be added.
- If a non-primary NIC is used to establish a VPN connection with the onpremises network, add a route to the on-premises network with the gateway address of the non-primary NIC as the next hop.

1.9.4 Huawei Cloud VPN NQA

What Is NQA?

Network Quality Analysis (NQA) is a technology to measure network performance and collect statistics on network indicators such as delay, jitter, and packet loss rate. It helps administrators learn network service quality in real time and effectively diagnose and locate network faults.

NQA Fundamentals

Figure 1-11 NQA test



In an NQA test, the source is called an NQA client, and the destination is called an NQA server. To enable an NQA client to initiate an NQA test, you need to create a test instance of a specific type on the NQA client. The NQA client then constructs packets that comply with the corresponding protocol, adds timestamps to the packets, and sends the packets to the server.

An NQA server listens to the NQA test packets with the specified IP address and port number and responds to the test accordingly. The client then calculates performance indicators, such as the connectivity, delay, and packet loss rate, based on statistics about the sent and received packets.

Processing Mechanism of NQA Tests

In an ICMP test, ICMP packets are sent to check reachability of the destination and calculate the network response time and packet loss rate.

A source constructs an ICMP Echo Request packet and sends it to a destination. When receiving the packet, the destination returns an ICMP Echo Reply packet to the source.

Upon receipt of the ICMP Echo Reply packet, the source calculates the time between when it sends the ICMP Echo Request packet and when it receives the ICMP Echo Reply packet. The test result reflects network performance and connectivity.

The NQA detection interval is 10s, and three ICMP requests are sent within 10s.

Why Do We Need NQA?

As value-added services develop, users and carriers demand higher quality of service (QoS). Especially after voice and video services are provisioned on conventional IP networks, carriers and users reach service level agreements (SLAs) to implement QoS guaranteed services.

To provide committed bandwidth for users, carriers need to collect statistics about network indicators such as the delay, jitter, and packet loss rate, and analyze the statistics to obtain network performance. Conventional network performance analysis methods (such as ping and tracert) cannot meet carriers' requirements for real-time monitoring on diverse services. Against this backdrop, NQA can be deployed to accurately test the network running status and export statistics. NQA can measure the performance of various protocols running on the network. This facilitates real-time collection of different network performance indicators, such as the total HTTP connection delay, TCP connection delay, DNS resolution delay, file transfer rate, FTP connection delay, and DNS resolution error rate. Network carriers control these indicators to provide users with network services of various grades. In addition, NQA is an effective tool to diagnose and locate faults on the network.

NQA for Static Routes

- Static routes do not have a dedicated detection mechanism. If an indirect link fails, a network administrator must manually delete the corresponding static route from the IP routing table. This process delays link switchover and causes service interruption for a significant amount of time.
- When creating VPN connections in static routing mode, you can enable NQA
 to detect faults in links for static routes. This prevents the preceding problems
 and ensures stability of VPN connections. When using NQA, ensure that the
 customer gateway device supports ICMP and is correctly configured with the
 customer tunnel interface IP addresses of the VPN connections. Otherwise,
 traffic will fail to be forwarded.
- If NQA detection fails for a VPN connection in static routing mode, the corresponding route is withdrawn. The customer gateway needs to permit ICMP traffic from the local tunnel interface address to the remote tunnel interface address of the VPN connection.
- The NQA detection results of VPN connections in health check are reported only to Cloud Eye. There is no impact if the detection fails. The customer gateway needs to permit ICMP traffic from the public IP addresses of the VPN gateway to the public IP address of the customer gateway.

1.10 Subnet Configurations

1.10.1 What Are the Precautions for Configuring the Local and Customer Subnets for a VPN Connection?

- The number of local subnets and the number of customer subnets are limited. If the number of local or customer subnets exceeds the upper limit, aggregate the subnets.
 - Maximum number of local subnets for each VPN gateway: 50
 - Maximum number of customer subnets for each VPN connection: 50
- A local subnet cannot include the CIDR block of a customer subnet, whereas a customer subnet can include the CIDR block of a local subnet.
- There are routes pointing to the local subnets in the VPC where the VPN gateway resides.
- If there are two connections (connection A and connection B) created for a VPN gateway, and the customer subnet of connection A is within that of connection B, when the destination network to be accessed belongs to the overlapped CIDR block, the connection created first is matched first, regardless of the connection status. (Mask length match is not used for VPN connections created in policy-based mode.)

1.10.2 How Many Local and Customer Subnets Can I Add to a VPN?

- You can configure a maximum of 50 local subnets for each VPN gateway.
- You can configure a maximum of 50 customer subnets for each VPN connection.

1.10.3 What Do I Do If an Exception Occurs When I Add a Customer Subnet During VPN Connection Creation?

Check whether this customer subnet is involved in a route of a VPC peering, Cloud Connect, or Direct Connect connection. If so, a route conflict occurs and you need to delete the route and create a new one to prevent the conflict.

1.10.4 Can the EIP of a VPN Gateway Be Retained After the VPN Gateway Is Deleted?

If a pay-per-use EIP is bound to a pay-per-use VPN gateway, deleting the VPN gateway will also delete the bound EIP.

To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway.

1.10.5 How Do I Plan CIDR Blocks for Access to a VPC Through a VPN Connection?

- The CIDR blocks of a VPC cannot conflict with on-premises CIDR blocks.
- To avoid conflicts with cloud service addresses, do not use 127.0.0.0/8, 169.254.0.0/16, 224.0.0.0/3, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 for your on-premises network.

If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.

1.10.6 How Is a VPN Gateway IP Address Allocated?

VPN gateway IP addresses are a group of IP addresses planned before VPN gateways are purchased. These IP addresses are preset with VPN configurations.

When you buy a VPN gateway, the system randomly assigns an IP address and binds it to the VPC you selected. This IP address can be bound to only one VPC.

You cannot change the IP address of a VPN gateway as this IP address has preset configurations. When a VPN gateway is deleted, the binding relationship between the gateway IP address and the gateway VPC is released. When a new VPN gateway is purchased, the system randomly allocates a new gateway IP address.

1.11 VPN Interesting Traffic

1.11.1 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to configure policy rules (ACL rules) for a VPN connection on the management console only when **VPN Type** is set to **Policy-based**.

1.11.2 How Do I Configure and Modify the Interesting Traffic of a VPN on the Cloud?

The number of rules that specify interesting traffic is the product of the number of local subnets and the number of customer subnets. For example, when there are local subnets A and B and customer subnets C, D, and E, the following six ACL rules need to be configured to specify interesting traffic:

```
rule 1 permit ip source A destination C
rule 2 permit ip source A destination D
rule 3 permit ip source A destination E
rule 4 permit ip source B destination C
rule 5 permit ip source B destination D
rule 6 permit ip source B destination E
```

If you modify the local or customer subnets on the management console, the interesting traffic configuration is automatically updated. That is, ACL rules on the cloud are modified.

1.12 Keeping VPN Connections Alive

1.12.1 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes of disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- Dead Peer Detection (DPD) is not configured on the device in your onpremises data center.
- Configuration is modified when the VPN connection is in use.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- At the two ends of the VPN connection, the local and customer subnet configurations are reversed.
- SA lifetime settings at both ends of the VPN connection are the same.
- DPD is enabled on the on-premises gateway device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.
- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway device is large enough for the VPN connection.
- VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway device.

1.13 Monitoring

1.13.1 What VPN Resources Can Be Monitored?

VPN gateway

The following bandwidth information of a VPN gateway IP address can be monitored: inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view the monitoring information, click $\stackrel{\triangle}{=}$ in the **Gateway IP Address** column in the VPN gateway list.

VPN connection

The following information about a VPN connection can be monitored: VPN connection status, average link round-trip time (RTT), maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate.

To monitor average link RTT, maximum link RTT, link packet loss rate, average tunnel RTT, maximum tunnel RTT, and tunnel packet loss rate, click the VPN connection name and click **Add** in the **Health Check** area on the **Summary** tab page to add health check items. Private network metrics can be configured only when the VPN connection uses the static routing mode and the NQA function is enabled.

To view the status of a VPN connection, click $\stackrel{\square}{=}$ in the **Monitoring** column of the VPN connection.

1.13.2 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

To view the status of a VPN connection, click in the **Monitoring** column of the VPN connection.

1.13.3 Can I View the Traffic of Each VPN Connection?

No. VPN traffic is monitored on a per-VPN gateway basis. You can view the inbound and outbound traffic as well as the inbound and outbound bandwidths of a VPN gateway, but cannot view the traffic statistics of a specific VPN connection.

1.13.4 Will I Be Notified of Abnormal VPN Monitoring Results?

Yes.

You can configure, on the Simple Message Notification (SMN) and Cloud Eye consoles, to receive notifications if abnormal VPN monitoring results occur.

1.14 Bandwidth and Network Speed

1.14.1 How Is the Network Speed of a VPN Connection Tested?

Test environment: A VPN connection has been created. ECSs have been created on the local subnets of VPCs at the two ends of the VPN connection. The ECSs can ping each other.

When the bandwidth of a purchased VPN gateway is 200 Mbit/s:

1.	When the ECSs at the two ends of the VPN connection run Windows, iPerf3
	and FileZilla (a free FTP application for file upload and download) are used to
	test the network speed. The test result is 180 Mbit/s, meeting requirements.

The TCP-based FTP protocol has a congestion control mechanism, and the IPsec protocol adds new headers to original packets. As such, it is normal in the industry, to have a network speed deviation of about 10%.

Figure 1-12 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 client.

C:\Windows>iperf3 -c 10.1.0.180 -i 1 -w 1M Connecting to host 10.1.0.180, port 5201 41 local 192.168.0.75 port 49212 connected to 10.1.0.180 port 5201 Transfer Bandwidth IDl Interval 0.00-1.01 sec 17.1 MBytes 142 Mbits/sec 253 Mbits/sec 1.01-2.00 sec 30.0 MBytes 2.00-3.01 19.8 MBytes 165 Mbits/sec 194 Mbits/sec sec 23.2 MBytes 3.01-4.01 4.01-5.00 18.9 MBytes 5.00-6.01 sec 26.2 MBytes 219 Mbits/sec 18.4 MBytes 153 Mhits/sec 6.01-7.01 sec sec 23.2 MBytes sec 21.1 MBytes 7.01-8.01 195 Mbits/sec 180 Mbits/sec 8.01-9.00 sec 9.00-10.01 sec 21.0 MButes 174 Mbits/sec ID] Interval Iransfer Bandwidth 0.00-10.01 sec 0.00-10.01 sec 219 MBytes 183 Mbits/sec sender 183 Mbits/sec 219 MBytes receiver perf Done.

Figure 1-12 Test result for 200 Mbit/s bandwidth (iPerf3 client)

Figure 1-13 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 server.

Figure 1-13 Test result for 200 Mbit/s bandwidth (iPerf3 server)

- 2. When the ECSs at the two ends of the VPN connection run CentOS 7, iPerf3 is used to test the network speed. The test result is 180 Mbit/s, meeting requirements.
- 3. When the ECS functioning as a server runs CentOS 7 and the ECS functioning as a client runs Windows, iPerf3 and FileZilla are used to test the network speed. The test result is 20 Mbit/s, failing to meet requirements.

This is because TCP implementations on Windows and Linux are different.

Figure 1-14 shows the result of using iPerf3 to test the network speed between two ECSs running different operating systems.

Figure 1-14 Test result on iPerf3

```
:\Windows>iperf3 -c 10.1.0.182 -i 1 -w 1M
onnecting to host 10.1.0.182, port 5201
 41 local 192.168.0.75 port 49426 connected to 10.1.0.182 port 5201
ID] Interval
                         Transfer
                                       Bandwidth
                         4.38 MBytes 36.7 Mbits/sec
4.50 MBytes 37.7 Mbits/sec
      0.00-1.00
 4]
                         4.38 MBytes
                   sec
       1.00-2.00
                                      43.0 Mbits/sec
                         5.12 MBytes
                         1.75
                              MBytes
                                       14.7
                                            Mbits/sec
                   sec
                         2.12
                                       17.8
                              MBytes
                                            Mbits/sec
                   sec
                         3.25
                                       27.3
                    sec
                              MBytes
                                            Mbits/sec
                         2.12 MBytes
                                       17.8
                                            Mbits/sec
       7.00-8.00
                    sec
                         1.25
                              MBytes
                                       10.5 Mbits/sec
                         2.25 MBytes
                                       18.9 Mbits/sec
       9.00-10.00
                         2.38 MBytes
                                       19.9 Mbits/sec
IDI
     Interval
                         Iransfer
                                       Bandwidth
       0.00-10.00 sec
                         29.1 MBytes 24.4 Mbits/sec
                                                                         sender
       0.00-10.00
                   sec
                         28.2 MBytes
                                       23.6 Mbits/sec
                                                                         receiver
```

When the bandwidth of a purchased VPN gateway is 1000 Mbit/s:

Ⅲ NOTE

Some regions support only 300 Mbit/s bandwidth by default. If higher bandwidth is required, select 300 Mbit/s bandwidth and then **submit a service ticket** for capacity expansion.

The VPN gateway bandwidth is shared by all of its VPN connections. To fully use the large bandwidth of 1000 Mbit/s, deploy multiple ECSs with high specifications as the forwarding performance of a single ECS is limited. ECSs with their NICs supporting the bandwidth of 2 Gbit/s or higher are recommended.

Conclusions: Based on the preceding test results, bandwidths of VPN gateways meet requirements. To fully use your purchased bandwidth, you are advised to use servers running the same operating system and using NICs meeting certain requirements at the two ends of a VPN connection.

1.14.2 In Which Direction Is the VPN Bandwidth Limited? What Is the Unit of Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited as follows:

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is greater than 10 Mbit/s, the bandwidth in the inbound direction is the same as the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

1.14.3 How Do I Change the VPN Bandwidth?

- 1. In the VPN gateway list, click the name of a VPN gateway. The gateway details page is displayed.
- In the EIP area, click Change next to Bandwidth.
- Change the EIP bandwidth.

When changing the EIP bandwidth, follow the EIP bandwidth configuration requirements. For details, see **Modifying an EIP Bandwidth**.

The EIP bandwidth cannot exceed the VPN bandwidth.

1.14.4 What Will Happen If Traffic Exceeds the Bandwidth of a VPN Gateway?

The VPN gateway bandwidth applies to traffic in the outbound direction of a VPC. If outbound traffic in the VPC exceeds the bandwidth, network congestion will occur, some subnets cannot be accessed, or even the VPN connection will be interrupted due to VPN detection timeout.

In this case, you are advised to increase the VPN gateway bandwidth.

□ NOTE

The maximum bandwidth of a VPN gateway is 1000 Mbit/s.

1.14.5 Why Does the VPN Bandwidth Change Not Take Effect?

There is a delay for the VPN bandwidth change to take effect.

Test the bandwidth 5 minutes after you change the bandwidth.

◯ NOTE

Changing the VPN bandwidth will not interrupt services on networks.

1.14.6 What Are the Differences Between the Bandwidth of a VPN Connection and That of a Direct Connect Connection?

Concepts

- The bandwidth of a Direct Connect connection is the bandwidth of the physical connection created by a user.
- The bandwidth of a VPN connection applies to the outbound direction of the cloud.

Maximum Bandwidth

- By default, the maximum bandwidth of a Direct Connect connection is 1000 Mbit/s. When you create a connection on the management console and set
 Port Type to 10GE single-mode optical port, the maximum bandwidth is 10 Gbit/s.
- The maximum bandwidth of a VPN is 1000 Mbit/s.

Network Quality

- A Direct Connect user has a dedicated connection with high network quality.
- VPN connections share the bandwidth of their VPN gateway. That is, the total bandwidth of VPN connections cannot exceed the bandwidth of the

corresponding VPN gateway. The network quality will be affected by the Internet quality.

1.14.7 How Do I Determine My VPN Bandwidth?

Consider the following when you determine the bandwidth:

- Amount of data transmitted over a VPN tunnel in a period of time (Reserve enough bandwidth to prevent link congestion.)
- Egress bandwidths at the two ends of a VPN connection: The egress bandwidth at the cloud side must be less than that at the on-premises side.

1.15 Quotas

1.15.1 What Quotas Does a VPN Have?

What Is a Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

Resource Types

VPN resources include VPN gateways, VPN connections, and customer gateways. The total quota of each resource type varies according to regions.

How Do I View My Quotas?

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- Choose Resources > My Quotas in the upper right corner of the page.
 The Service Quota page is displayed.
- 4. View the used and total quota of each type of resources on the displayed page.

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

How Do I Apply for a Higher Quota?

- 1. Log in to the management console.
- Choose Resources > My Quotas in the upper right corner of the page.
 The Service Quota page is displayed.
- 3. Click **Increase Quota** in the upper right corner of the page.
- On the Create Service Ticket page, configure parameters as required.
 In the Problem Description area, enter the required quota and the reason for the quota adjustment.

5. Select the agreement and click **Submit**.

1.15.2 How Many VPN Gateways and VPN Connections Can I Create By Default?

By default, each user can create a maximum of 50 VPN gateways and 100 customer gateways. Each VPN gateway can have a maximum of 100 connection groups. When two EIPs of a VPN gateway are connected to the same public IP address of a customer gateway, one VPN connection group is used. When two EIPs of a VPN gateway are connected to two customer gateways or two public IP addresses of the same customer gateway, two VPN connection groups are used.

Before purchasing VPN gateways, check your available quota. If the quota is insufficient, **submit a service ticket** to increase the quota.

1.15.3 How Do I Change My VPN Gateway and Connection Quotas?

- 1. Log in to the management console, and choose **Service Tickets** > **Create Service Ticket** in the menu bar.
- 2. On the Create Service Ticket page, click Quotas in the Services area.
- 3. Click **Quota Application** under **Issue Categories**.
- 4. Click **Create Now**.

Enter required information and click Submit.

1.15.4 How Many IPsec VPNs Can I Have?

By default, each user can create a maximum of 50 VPN gateways and 100 customer gateways. Each VPN gateway can have a maximum of 100 connection groups. When two EIPs of a VPN gateway are connected to the same public IP address of a customer gateway, one VPN connection group is used. When two EIPs of a VPN gateway are connected to two customer gateways or two public IP addresses of the same customer gateway, two VPN connection groups are used.

Before purchasing VPN gateways, check your available quota. If the quota is insufficient, **submit a service ticket** to increase the quota.

1.16 Account Permissions

1.16.1 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

Ⅲ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

1.16.2 What Should I Do If the System Displays a Message Indicating that I Do Not Have the Permissions to Create a VPN?

- Check whether your account is an IAM account.
- Ensure that your IAM account has the VPN FullAccess permission. For details, see Creating a User Group and Assigning Permissions and Adding Users to or Removing Users from a User Group.

1.16.3 How Do I Determine that a VPN Cannot Be Created in My Account Due to Insufficient Permissions?

If the system displays a message indicating that you do not have the permissions to create a VPN gateway or VPN connection, add the required permissions.

For details about the permissions required for using the VPN service, see **IAM-based Permissions Management**.

Page - S2C Classic VPN

2.1 General Questions

2.1.1 What Devices Can Be Connected to the Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

Ⅲ NOTE

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- Devices that can interconnect with the VPN service are usually from the following:
 - Vendors such as Huawei (routers and firewalls), H3C (routers and firewalls), Cisco (routers and firewalls), Ruijie (routers and firewalls), ZTE, Sangfor, Fortinet, 360, Topsec, Hillstone, NetentSec, NSFOCUS, DELL, ZyXEL, and Juniper
 - Cloud service providers such as Alibaba Cloud, Tencent Cloud, and Amazon Web Services
 - Software vendors such as Openswan, strongSwan, and TheGreenBow
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.
 - Most enterprise-class routers and firewalls support the IPsec protocol.
- However, some devices support IPsec VPN only after you purchase required software licenses.

Contact your on-premises data center administrator to confirm the device model with the vendor.

2.1.2 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 2-1 VPN negotiation parameters

Policy	Parameter	Value
IKE	Authentication Algorithm	 MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-256 AES-192 AES-128 (default value)
	DH Algorithm	 Group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 14 (default value) Group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 15 Group 16 Group 19 Group 20 Group 21 NOTE In some regions, only Group 14, Group 2, and Group 5 are available.
	Version	v1 (not recommended due to security risks)v2 (default value)

Policy	Parameter	Value
	Lifetime (s)	86400 (default) Unit: second Value range: 60 to 604800
IPsec	Authentication Algorithm	 SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 AES-128 (default value) AES-192 AES-256 3DES (This algorithm is insecure. Exercise caution when using this algorithm.)
	PFS	 DH group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 14 (default value) DH group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 15 DH group 16 DH group 19 DH group 20 DH group 21 Disable NOTE In some regions, only DH group 14, DH group 2, and DH group 5 are available.
	Transfer Protocol	ESP (default value)AHAH-ESP

Policy	Parameter	Value
	Lifetime (s)	3600 (default)
		Unit: second
		Value range: 480 to 604800

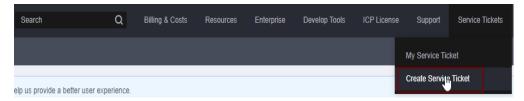
◯ NOTE

- Perfect Forward Secrecy (PFS) is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. After PFS is configured, an additional DH exchange will be performed during IPsec SA negotiation, and a new IPsec SA key will be generated, improving IPsec SA security.
- To ensure security, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the on-premises gateway. Otherwise, the negotiation will fail.
- To enable PFS, ensure that the configurations at both ends of a VPN are the same.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB and cannot be changed for the VPN. This lifetime does not affect the establishment of an IPsec SA.

2.1.3 What Are the Categories of VPN Service Tickets? How Do I Create a VPN Service Ticket?

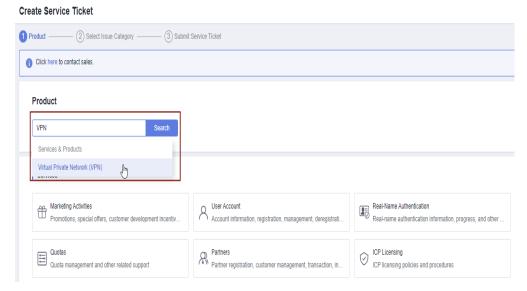
- 1. Log in to the management console.
- 2. Choose **Service Tickets** > **Create Service Ticket** in the upper right corner.

Figure 2-1 Create Service Ticket



3. Search for "VPN" and choose Virtual Private Network (VPN).

Figure 2-2 Selecting Virtual Private Network (VPN)



4. Select an issue category.

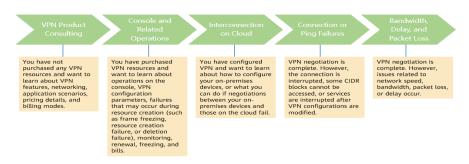
Figure 2-3 Selecting an issue category



□ NOTE

When you **submit a service ticket**, select an issue category to facilitate problem handling.

Figure 2-4 Issue category and classification basis



2.1.4 Can I Deploy Applications on the Cloud, Databases in an On-premises Data Center, and Then Connect Them Through a VPN?

VPN connects a VPC and an on-premises network.

After a VPN is set up, service traffic can be transmitted between the VPC and on-premises data center. For an application server on the cloud, access to an on-premises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check whether the network latency and packet loss adversely affect service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

2.1.5 Can I Visit Websites Across International Borders Using a VPN?

No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

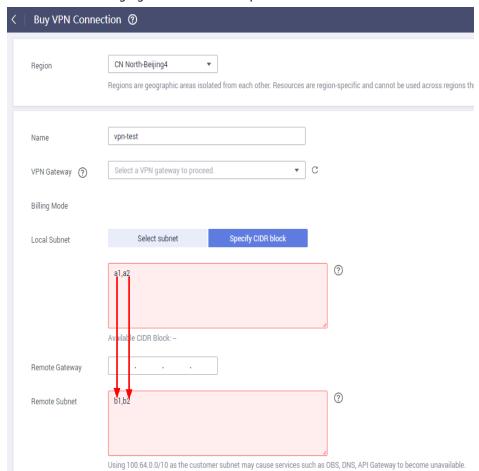
2.1.6 What Is a VPN Connection? How Do I Set the Number of VPN Connections When Buying a VPN Gateway?

A VPN connection is an IPsec connection established between a VPN gateway and an independent public IP address of an on-premises data center. You can configure multiple local subnets (VPC subnets) and remote subnets (on-premises subnets) for one VPN connection.

The number of VPN connections to be created is determined by the number of onpremises data centers. Each VPN connection can connect a VPC to only one onpremises data center.



In the preceding figure, if subnets a1 and a2 on the cloud need to communicate with subnets b1 and b2 on the on-premises network, you only need to create one VPN connection, with source CIDR blocks set to a1 and a2 and destination CIDR blocks set to b1 and b2. The following figure shows an example.



2.1.7 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

After a VPN connection is created, you can locate the row that contains the VPN connection and choose **Operation** > **View Metric** to view the VPN connection status.

2.1.8 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

Ⅲ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

2.1.9 What Are the Differences Between IPsec VPN and SSL VPN in Application Scenarios and Connection Modes?

Scenarios

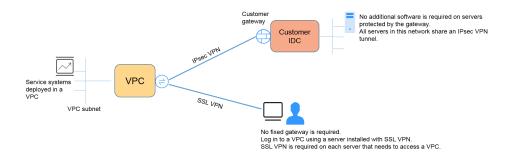
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to complete IPsec VPN negotiation.

SSL VPN needs to install a specified client software on the server, then the server connects to the SSL device through the username and password.



IPsec VPN and SSL VPN are supported.

2.1.10 Will an IPsec VPN Connection Be Established Automatically?

After you complete configurations at both ends of an IPsec VPN connection, the VPN connection will not be automatically established only after data flows between the two ends of the connection. If no data flows between the cloud and the on-premises data center, the VPN connection will always be in the down state. Any data generated by accessing servers or pinging between servers can trigger the establishment of a VPN connection.

The establishment of a VPN connection can be triggered in either of the following two conditions: The VPN gateway and the remote gateway automatically trigger negotiation. The cloud and on-premises servers access each other via the VPN connection to be established.

However, automatic establishment of a VPN connection cannot be triggered by a VPN gateway on the cloud. Verify that the establishment of your VPN connection can be triggered by the data flows between the two ends of the VPN connection. That is, check whether a VPN connection can be established after you ping a cloud server from an on-premises server, and whether a VPN connection can be established after you disconnect the connection and ping an on-premises server from a cloud server.

The source and destination addresses of the ping packets must be protected by the VPN. Before a VPN connection is established, the gateway IP addresses at both ends can be pinged. However, pinging the gateway IP addresses does not trigger the establishment of the VPN connection.

2.1.11 What Will I Be Charged for Creating a VPN? Will I Be Charged for VPN Gateway IP Addresses?

VPN support the yearly/monthly and pay-per-use billing modes, and are billed by the gateway bandwidth and VPN connections. The billing modes available for a region are subject to those displayed on the page.

VPN gateways can be billed by traffic or bandwidth.

- 1. A yearly/monthly VPN gateway can only be billed by bandwidth. The price of a yearly/monthly VPN gateway includes the price of the VPN connections that can be created for the gateway and the bandwidth price.
- 2. The billing cycle of the pay-per-use billing mode is 1 hour. If you choose a pay-per-use VPN gateway, a VPN connection must be purchased together with the VPN gateway. The price includes the VPN gateway bandwidth or traffic price and the price of the VPN connection created together with the gateway. If you create another connection for the gateway, you will be charged for this connection.

- The IP address of a VPN gateway is free of charge. Only the bandwidth of the VPN gateway is billed.
- A VPN gateway cannot share bandwidth with an EIP bound to an ECS.

2.1.12 Can a VPN Gateway IP Address Be Retained After the VPN Gateway Is Deleted?

No. The VPN gateway IP address will be released after the VPN gateway is deleted.

Deleting a VPN gateway will also delete the resources associated with the gateway, such as public IP addresses.

NOTICE

Deleting the last connection of a pay-per-use VPN gateway will also delete the gateway. If you want to retain the IP address, do not delete the last VPN connection.

2.1.13 Which VPN Resources Can Be Monitored?

VPN Gateway

Bandwidth information that can be monitored includes inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view VPN gateway metrics, locate the target VPN gateway and click **View Metric** in the **Operation** column.

VPN Connection

The VPN connection status can be monitored.

Value 1 indicates that the connection is normal.

Value **0** indicates that the connection is not connected.

To view the VPN connection status, locate the target VPN connection and click **View Metric** in the **Operation** column.

2.1.14 Which Direction of the Bandwidth Is Limited and What Is the Unit of the Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited.

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is more than 10 Mbit/s, the bandwidth in the inbound direction is the same as that of the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

2.1.15 What Is the Actual Network Speed of a VPN Connection?

A VPN connection has been created. Two ECSs have been created with one at the local end and the other at the remote end. The two ECSs can ping each other.

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 200 Mbit/s:

 If the ECSs at the two ends of the VPN run Windows, use iPerf3 and FileZilla (a free FTP application for file uploading and downloading) to test the network speed.

The test shows that the average network speed of the VPN is 180 Mbit/s, and there is about 10% network speed deviation. The TCP and FTP protocols have the congestion control mechanism, and the IPsec protocol adds a new IP header. Therefore, about 10% network speed deviation is normal for the VPN network.

Figure 2-5 shows the test result.

Figure 2-5 Test result for 200 Mbit/s bandwidth (iPerf3 client)

```
C:\Windows>iperf3 -c 10.1.0.180 -i 1 -w 1M
Connecting to host 10.1.0.180, port 5201
   41 local 192.168.0.75 port 49212 connected to 10.1.0.180 port 5201
                          Transfer
                                       Bandwidth
 ID1 Interval
       0.00-1.01
                    sec
                         17.1 MBytes
                                        142 Mbits/sec
                                        253 Mbits/sec
       1.01-2.00
                    sec 30.0 MBytes
       2.00-3.01
                         19.8 MBytes
                                        165 Mbits/sec
                                        194 Mbits/sec
                         23.2 MBytes
       3.01-4.01
                    sec
       4.01-5.00
                         18.9 MBytes
       5.00-6.01
                    sec 26.2 MBytes
                                        219 Mbits/sec
                         18.4 MBytes
                                        153 Mhits/sec
       6.01-7.01
                    sec
                    sec 23.2 MBytes
sec 21.1 MBytes
                                        195 Mbits/sec
180 Mbits/sec
       7.01 - 8.01
       8.01-9.00
       9.00-10.01 sec
                         21.0 MBytes
                                        174 Mbits/sec
 ID] Interval
                          Iransfer
                                       Bandwidth
       0.00-10.01 sec
                          219 MBytes
                                        183 Mbits/sec
                                                                         sender
                          219 MBytes
                                        183 Mbits/sec
        0.00-10.01 sec
                                                                         receiver
 perf Done.
```

Figure 2-6 shows the test result.

Figure 2-6 Test result for 200 Mbit/s bandwidth (iPerf3 server)

- 2. If the ECSs at the two ends of the VPN run CentOS 7, use iPerf3 to test the network speed. The network speed can reach 180 Mbit/s.
- 3. If the ECS functioning as the server runs CentOS 7, and the ECS functioning as the client runs Windows, use iPerf3 and FileZilla to test the network speed.

The network speed is about 20 Mbit/s, a slow network speed. That is because TCP implementations on Windows and that on Linux are different. Therefore, if the ECSs at the two ends of the VPN run different OSs, the VPN network speed does not meet the bandwidth requirements.

Figure 2-7 shows the test result.

Figure 2-7 Test result when ECSs at the two ends run different OSs (iPerf3)

```
::\Windows>iperf3 -c 10.1.0.182 -i 1 -w 1M
Connecting to host 10.1.0.182, port 5201
41 local 192.168.0.75 port 49426 connected to 10.1.0.182 port 5201
 ID] Interval
                          Transfer
                                        Bandwidth
                                        36.7 Mbits/sec
       0.00-1.00
                          4.38 MBytes
  41
                    26
                          4.50 MBytes 37.7 Mbits/sec
       1.00-2.00
                                        43.0 Mbits/sec
                          5.12 MBytes
                    sec
                               MBytes
                                        14.7
                                              Mbits/sec
                          1.75
                    sec
                                              Mbits/sec
                          2.12
                               MBytes
                    sec
                          3.25
                               MBytes
                                              Mbits/sec
                    sec
                    sec
                          2.12
                               MBytes
                                              Mbits/sec
        7.00-8.00
                          1.25
                               MBytes
                                        10.5
                                              Mbits/sec
                    sec
                          2.25
                               MBytes
                                              Mbits/sec
                    sec
                                        19.9 Mbits/sec
       9.00-10.00
                          2.38 MBytes
       0.00-10.00 sec
                          29.1 MBytes
                                        24.4 Mbits/sec
                                                                            sender
       0.00-10.00
                          28.2 MBytes
                                        23.6 Mbits/sec
                                                                            receiver
perf Done
```

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 1,000 Mbit/s:

The VPN gateway bandwidth is shared by all of its VPN connections. If the bandwidth size is large, multiple ECSs are required to test the VPN gateway bandwidth because the forwarding performance of an ECS is limited. This scenario has high requirements on ECS specifications. The ECSs must have NICs that support the bandwidth of 2 Gbit/s or higher.

Conclusions: Based on the preceding test results, bandwidths of VPN gateways meet requirements. To fully use your purchased bandwidth, you are advised to use servers running the same operating system and using NICs meeting certain requirements at the two ends of a VPN connection.

2.1.16 Can a VPN Billed by Traffic Use a Shared Data Package?

No.

The VPN service is billed independently and cannot use shared data packages.

2.1.17 How Do I Change the Billing Mode of a VPN Gateway from Pay-Per-Use to Yearly/Monthly?

Prerequisites

- A pay-per-use VPN gateway is billed by bandwidth.
 - To change the billing mode of a VPN gateway billed by traffic from pay-peruse to yearly/monthly, first change the VPN gateway from being billed by traffic to being billed by bandwidth and then from pay-per-use to yearly/ monthly.
- The number of created VPN connections is less than 10.
- At least 10 more VPN connections can be created in your account.

Procedure

Perform the following operations:

- 1. Log in to the management console.
- 2. Click in the upper left corner and select the desired region and project.
- 3. Click **Service List** in the upper left corner. Under **Networking**, select **Virtual Private Network**.
- 4. In the navigation pane on the left, choose Virtual Private Network > Classic.
- 5. On the **VPN Gateways** page, locate the row that contains the target VPN gateway, choose **More** > **Change Billing Mode** in the **Operation** column.
- 6. In the displayed **Change Billing Mode** dialog box, click **OK**.

- The billing mode of a VPN gateway cannot be changed from yearly/monthly to pay-per-use.
- Resource quotas of a yearly/monthly VPN gateway can be decreased upon a renewal.
- In the yearly/monthly billing mode, Required VPN Connections indicates the total number of VPN connections that can be created for the VPN gateway free of charge.
- After you change the billing mode of a VPN gateway from pay-per-use to yearly/monthly, the number of VPN connections that can be created for the VPN gateway is 10 by default.
- 7. Confirm the VPN gateway information and set a renewal duration. Click **Pay**.
- 8. On the payment page, confirm the order information, select coupon or discount, and select the payment method.

Click Pay.

■ NOTE

Changing the billing mode of a VPN gateway from pay-per-use to yearly/monthly will not affect your services.

2.1.18 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

- VPC
 - A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
 - You can use the VPN service to securely access ECSs in a VPC.
- VPN gateway
 - A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have only one VPN gateway, whereas one VPN gateway can have multiple VPN connections.
 - With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.

VPN connection

A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

The number of VPN connections is irrelevant to the number of local subnets or the number of remote subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

2.1.19 What Is a Remote Gateway and Remote Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a remote subnet and a remote gateway, respectively.

A remote gateway's IP address is a public IP address.

2.1.20 How Many VPN Connections Do I Need to Connect to Multiple On-premises Servers?

VPN uses the IPsec technology to connect your on-premises data center to a VPC on the cloud. As such, the number of VPN connections is related to the number of data centers where the servers to be connected to the cloud are located, but not to the number of servers.

In most cases, one on-premises data center has one public gateway. All servers connect to the Internet through this gateway. Therefore, you only need to configure one VPN connection to allow communications between the VPC and your on-premises data center.

2.1.21 Does a VPN Allow for Communications Between Two VPCs?

- If the two VPCs are deployed in the same region, use a VPC peering connection to connect them.
- If the two VPCs are deployed in different regions, use a VPN connection to connect them. The detailed operations are as follows:
 - a. Create a VPN gateway for each VPC and create VPN connections for the two VPN gateways.
 - b. Set the remote gateway address of each VPN connection to the gateway IP address of the peer side.
 - c. Set the remote subnets of each VPN connection to the subnets of the peer VPC.
 - d. The pre-shared keys and algorithm parameters of the two VPN connections must be the same.

2.1.22 What Are the Impacts of a VPN on an On-premises Network? What Are the Changes to the Route for Accessing an ECS?

When you configure a VPN, perform the following operations on the on-premises gateway:

- 1. Configure IKE and IPsec policies.
- 2. Specify interesting traffic (ACL rules).
- 3. Check the route configuration on the gateway to ensure that traffic destined for a VPC can be routed to the correct outbound interface (interface having an IPsec policy bound).

After the VPN configuration is complete, only the traffic matching the ACL rules enters the VPN tunnel.

For example, before a VPN is created, on-premises users access the ECS through the EIP bound to the ECS. After a VPN is created, data flows matching the ACL rules access the private IP address of the ECS through the VPN tunnel.

2.1.23 Can I Use a Network with Two Egresses to Establish Two VPN Connections with the Same VPC?

No.

When creating a VPN on the cloud, a local subnet is a VPC subnet, and a remote subnet is an on-premises subnet. If the two connections use the same local subnet and remote subnet, the VPN connections will fail.

2.1.24 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes for disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- DPD is not configured in your on-premises data center.
- Configuration is modified when the VPN connection is in use.
- Packets are fragmented because the data size exceeds the MTU.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- Local and remote subnets are matched pairs.
- SA lifetime settings at both ends of the VPN connection are the same.

- DPD is enabled on the on-premises gateway device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.
- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway is large enough to be used by the VPN connection.
- VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway.
- Ping the subnets at both ends continuously. The script is as follows:

```
#!/bin/sh
host=$1
if [ -z $host ]; then
    echo "Usage: `basename $0` [HOST]"
    exit 1
fi
log_name=$host".log"

while :; do
    result=`ping -W 1 -c 1 $host | grep 'bytes from '`
    if [ $? -gt 0 ]; then
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is down"| tee -a $log_name
    else
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is ok -`echo $result | cut -d ':' -f 2`"| tee -a
$log_name
    fi
sleep 5 # avoid ping rain
done
#./ping.sh x.x.x.x >>/dev/null &
```

◯ NOTE

- 1. Use the vi editor to copy the preceding script to the **ping.sh** file.
- 2. Run the **chmod 777 ping.sh** command to grant permissions to the file.
- 3. Run the ping command:

```
./ping.sh x.x.x.x >> /dev/null &
```

x.x.x.x indicates the IP address to be pinged.

4. Run the following command:

```
tail -f x.x.x.x.log
```

You can view the ping result in real time.

2.1.25 Why Is Not Connected Displayed as the Status for a Successfully Created VPN Connection?

After a VPN connection is created, its status changes to **Normal** only after servers at both ends of the VPN connection communicate with each other.

• IKE v1:

If no traffic goes through the VPN connection for a period of time, the VPN connection needs to be renegotiated. The negotiation time depends on the **Lifecycle (s)** value in the IPsec policy. Generally, **Lifecycle (s)** is set to **3600** (1 hour), indicating that the negotiation will be initiated in the fifty-fourth minute. If the negotiation succeeds, the connection remains to the next round of negotiation. If the negotiation fails, the VPN connection status changes to **Not Connected** within one hour. The connection can be restored only after

the two ends of the VPN connection communicate with each other. The disconnection can be avoided by using a network monitoring tool, such as IP SLA, to generate packets.

• IKE v2: If no traffic goes through the VPN connection for a period of time, the VPN connection remains in the connected status.

2.1.26 What Can I Do If VPN Connection Setup Fails?

- 1. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.
 - a. If the IKE policy has been set up during phase one and the IPsec policy has not been enabled in phase two, the IPsec policies at both ends of the VPN connection may be inconsistent.
 - b. If you use a Cisco physical device in your on-premises data center, it is recommended that you use MD5, and set **Authentication Mode** to **MD5** when configuring the IPsec policy for the VPN connection on the cloud.
- 2. Check whether the ACL rules are correct.

If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

3. Check whether the VPN connection is normal by pinging the local end from the remote end and pinging the remote end from the local end.

2.1.27 Can an EIP Be Used as a VPN Gateway IP Address?

EIPs can be used as VPN gateway IP addresses only in Enterprise Edition VPN, but not in Classic VPN.

2.1.28 Why Is the VPN Connection Always in the Not Connected State Even After Its Configuration Is Complete?

Ensure that the pre-shared keys and negotiation information at both ends are consistent. The local subnets and VPN gateway on the cloud are the remote subnets and remote gateway in the on-premises data center. The remote gateway and remote subnets on the cloud are the local gateway and local subnets in the on-premises data center.

Ensure that routes, NAT, and security policy rules are correctly configured on your on-premises gateway device. Then, ping the servers in subnets at both ends.

◯ NOTE

VPN is triggered based on data flows. After you configure VPN, ping a device in the peer subnet. Before running the ping command, disable the firewall function on the device, and allow inbound ICMP packets in the security group on the cloud.

Pinging the gateway IP address cannot trigger VPN negotiation. Ping the server in the subnet protected by the gateway.

2.1.29 Which Remote VPN Devices Are Supported?

Most devices that meet IPsec VPN standards and reference protocol requirements can be used as remote VPN devices, for example, Cisco ASA firewalls, Huawei USG6000 series firewalls, USG9000 series firewalls, Hillstone firewalls, and Cisco ISR routers. **Table 2-2** lists the supported Huawei USG6000 and USG9000 firewalls.

Table 2-2 Huawei VPN devices

Supported Remote VPN Device	Description
Huawei USG6000 series	USG6320/6310/6510-SJJ USG6306/6308/6330/6350/6360/6370/6380/6390/6507/ 6530/6550/6570: 2048 USG6620/6630/6650/6660/6670/6680
Huawei USG9000 series	USG9520/USG9560/USG9580

Other devices that meet the requirements outlined in **Reference Standards and Protocols** can also be deployed. However, some devices may not be supported because of the inconsistent protocol implementation methods of these devices.

2.1.30 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to create ACL rules on your on-premises gateway device. The ACL rules will be referenced by IPsec policies.

When you configure VPN on the cloud, ACL rules will be automatically generated based on the local and remote subnets entered on the management console and then delivered to the VPN gateway.

Number of ACL rules = Number of local subnets x Number of remote subnets

2.2 Product Consultation

2.2.1 What Are the Typical Scenarios of IPsec VPN?

A VPN is a point-to-point connection that implements private network access between two points.

- Applicable scenarios:
 - A VPN is created between different regions to enable cross-region VPC communications.
 - A VPN is created between VPCs of Huawei Cloud and another cloud, for example, Alibaba Cloud.

- A VPN is created between a VPC of Huawei Cloud and an on-premises data center
- A VPN hub is used together with VPC peering connections and Cloud Connect connections to enable communications between an on-premises data center and multiple VPCs on the cloud.
- A VPN is used together with source NAT to enable access to specific IP addresses across clouds.
- A VPN can be used between the cloud and your home network that uses PPPoE dial-up.
- A VPN can be used between the cloud and 4G/5G routers.
- A VPN can be used between the cloud and your personal terminals.
- Not applicable scenarios:
 - A VPN cannot be used to connect VPCs in the same region. It is recommended that you use VPC peering connections to enable communications between VPCs in the same region.

2.2.2 What Are a VPC, a VPN Gateway, and a VPN Connection?

VPC enables you to create private, isolated virtual networks. You can use VPN to securely access ECSs in VPCs.

A VPN gateway is an egress gateway for a VPC. With a VPN gateway, you can create a secure, reliable, and encrypted connection between a VPC and an onpremises data center or between two VPCs in different regions.

A VPN connection is a secure and reliable IPsec encrypted communications tunnel established between a VPN gateway and the remote gateway in an on-premises data center.

To create a VPN on the cloud, perform the following operations:

- 1. Create a VPN gateway. You need to specify the VPC to be connected, as well as the bandwidth and EIPs of the VPN gateway.
- 2. Create a VPN connection. You need to specify the gateway EIP used to connect to the remote gateway, subnets, and negotiation policies.

2.2.3 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

- VPC
 - A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
 - You can use the VPN service to securely access ECSs in a VPC.
- VPN gateway

- A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have only one VPN gateway, whereas one VPN gateway can have multiple VPN connections.
- With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.

VPN connection

A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

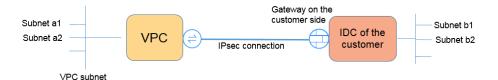
MOTE

The number of VPN connections is irrelevant to the number of local subnets or the number of remote subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

2.2.4 What Is a VPN Connection? How Do I Set the Number of VPN Connections When Buying a VPN Gateway?

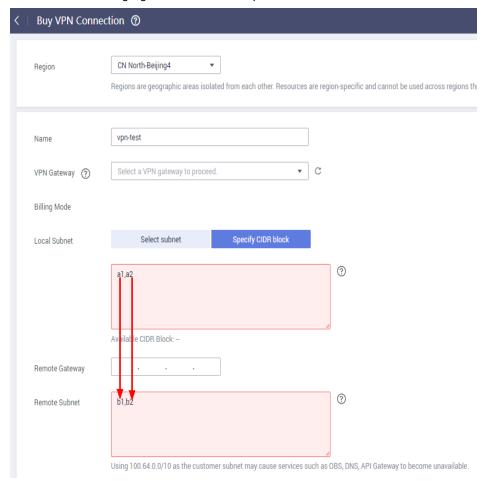
A VPN connection is an IPsec connection established between a VPN gateway on Huawei Cloud and an independent public IP address of an on-premises data center. You can configure multiple local subnets (VPC subnets) and remote subnets (on-premises subnets) for one VPN connection.

The number of VPN connections to be created is determined by the number of onpremises data centers. Each VPN connection can connect a VPC to only one onpremises data center.



Ⅲ NOTE

In the preceding figure, if subnets a1 and a2 on the cloud need to communicate with subnets b1 and b2 on the on-premises network, you only need to create one VPN connection, with source CIDR blocks set to a1 and a2 and destination CIDR blocks set to b1 and b2. The following figure shows an example.



2.2.5 What Is a Remote Gateway and Remote Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a remote subnet and a remote gateway, respectively.

A remote gateway's IP address is a public IP address.

2.2.6 How Do I Plan the CIDR Block of a VPC Accessed over a VPN Connection?

- The VPC CIDR block cannot conflict with the on-premises CIDR block.
- To avoid conflicts with cloud service addresses, do not use 127.0.0.0/8, 169.254.0.0/16, 224.0.0.0/3, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 for your on-premises network.

If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.

2.2.7 Will an IPsec VPN Connection Be Established Automatically?

After you complete configurations at both ends of an IPsec VPN connection, the VPN connection will not be automatically established only after data flows between the two ends of the connection. If no data flows between the cloud and the on-premises data center, the VPN connection will always be in the down state. Any data generated by accessing servers or pinging between servers can trigger the establishment of a VPN connection.

The establishment of a VPN connection can be triggered in either of the following two conditions: The VPN gateway and the remote gateway automatically trigger negotiation. The cloud and on-premises servers access each other via the VPN connection to be established.

However, automatic establishment of a VPN connection cannot be triggered by a VPN gateway on the cloud. Verify that the establishment of your VPN connection can be triggered by the data flows between the two ends of the VPN connection. That is, check whether a VPN connection can be established after you ping a cloud server from an on-premises server, and whether a VPN connection can be established after you disconnect the connection and ping an on-premises server from a cloud server.

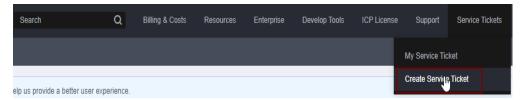
□ NOTE

The source and destination addresses of the ping packets must be protected by the VPN. Before a VPN connection is established, the gateway IP addresses at both ends can be pinged. However, pinging the gateway IP addresses does not trigger the establishment of the VPN connection.

2.2.8 What Are the Categories of VPN Service Tickets? How Do I Create a VPN Service Ticket?

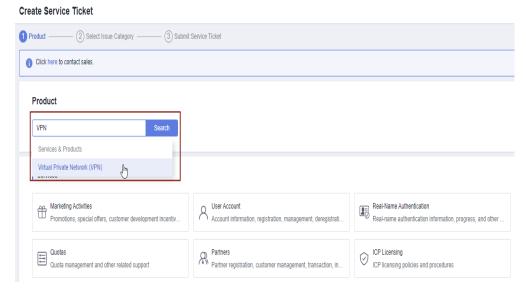
- 1. Log in to the management console.
- 2. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket**.

Figure 2-8 Create Service Ticket



Search for VPN and select Virtual Private Network (VPN).

Figure 2-9 Selecting Virtual Private Network (VPN)



4. Select an issue category.

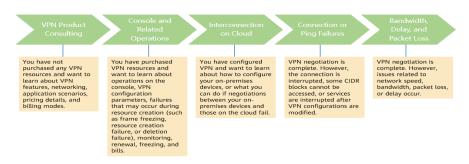
Figure 2-10 Select Issue Category



Ⅲ NOTE

When you **submit a service ticket**, select an issue category to facilitate problem handling.

Figure 2-11 Issue category and classification basis



2.2.9 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 2-3 VPN negotiation parameters

Protocol	Parameter	Value
IKE	Authentication Algorithm	 MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 3DES (This algorithm is insecure. Exercise caution when using this algorithm.) AES-256 AES-192 AES-128 (default value)
	DH Algorithm	 Group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 14 (default value) Group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) Group 15 Group 16 Group 19 Group 20 Group 21 NOTE In some regions, only Group 14, Group 2, and Group 5 are available.
	Version	v1 (not recommended due to security risks)v2 (default value)

Protocol	Parameter	Value
	Lifetime (s)	86400 (default value) Unit: second Value range: 60 to 604800
IPsec	Authentication Algorithm	 SHA1 (This algorithm is insecure. Exercise caution when using this algorithm.) MD5 (This algorithm is insecure. Exercise caution when using this algorithm.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 AES-128 (default value) AES-192 AES-256 3DES (This algorithm is insecure. Exercise caution when using this algorithm.)
	PFS	 DH group 5 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 2 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 14 (default value) DH group 1 (This algorithm is insecure. Exercise caution when using this algorithm.) DH group 15 DH group 16 DH group 19 DH group 20 DH group 21 Disable NOTE In some regions, only DH group 14, DH group 2, and DH group 5 are available.
	Transfer Protocol	ESP (default value)AHAH-ESP

Protocol	Parameter	Value
	Lifetime (s)	3600 (default value)
		Unit: second
		Value range: 480 to 604800

◯ NOTE

- Perfect Forward Secrecy (PFS) is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. After PFS is configured, an additional DH exchange will be performed during IPsec SA negotiation, and a new IPsec SA key will be generated, improving IPsec SA security.
- To ensure security, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the on-premises gateway. Otherwise, the negotiation will fail.
- To enable PFS, ensure that the configurations at both ends of a VPN are the same.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB and cannot be changed for the VPN. This lifetime does not affect the establishment of an IPsec SA.

2.2.10 What Devices Can Be Connected to the Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- 1. Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- The following products can connect to the cloud through VPNs:
 - Devices: Huawei firewalls and access routers (ARs), Hillstone firewalls, and Check Point firewalls
 - Cloud services: Alibaba Cloud, Tencent Cloud, Amazon Web Services (AWS), and Microsoft Azure
 - Software: strongSwan
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.

Most enterprise-class routers and firewalls support the IPsec protocol.

 However, some devices support IPsec VPN only after you purchase required software licenses.

Contact your on-premises data center administrator to confirm the device model with the vendor.

2.2.11 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

2.2.12 How Do I Allow Specific Servers to Access a VPC Subnet Through a Created VPN Connection?

Configurations in the on-premises data center

- Configure deny rules on VPN devices.
- Configure ACL rules on the router or switch.

Configurations on the cloud

- Configure security group rules to deny access from specific IP addresses.
- Configure ACL rules.

□ NOTE

All rules must be added to the device before the VPN tunnel is established. Do not change the local subnet and the remote subnet to restrict the access.

2.2.13 Which VPN Resources Can Be Monitored?

VPN Gateway

Bandwidth information that can be monitored includes inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view VPN gateway metrics, locate the target VPN gateway and click **View Metric** in the **Operation** column.

VPN Connection

The VPN connection status can be monitored.

Value 1 indicates that the connection is normal.

Value **0** indicates that the connection is not connected.

To view the VPN connection status, locate the target VPN connection and click **View Metric** in the **Operation** column.

2.2.14 Can an EIP Be Used as a VPN Gateway IP Address?

Nο

The IP address of a VPN gateway is of preset configurations and is automatically assigned when the VPN gateway is created. An EIP cannot be used by a VPN gateway.

2.2.15 Do I Need to Purchase EIPs for Servers That Communicate with Each Other Through a VPN?

If your on-premises server needs to access an ECS on the cloud through a VPN, you do not need to purchase an EIP.

If the ECS needs to provide services accessible from the Internet, an EIP is required.

2.2.16 Are SSL VPNs Supported?

Currently, SSL VPNs are supported.

2.2.17 How Long Does It Take for Delivered VPN Configurations to Take Effect?

It takes 1 to 5 minutes for the VPN configurations to take effect.

□ NOTE

After VPN configurations take effect, configure your gateway device on your on-premises network to complete tunnel negotiation with the VPN gateway.

2.2.18 What Should I Do If I Cannot Create Connections for a VPN Gateway That Has No Bandwidth Information?

If a Classic VPN gateway has no bandwidth information, the VPN is of the earlier edition. This type of VPN cannot be created on the cloud anymore.

- Only one VPN connection can be created for each VPN gateway of the earlier edition and its bandwidth is not guaranteed. You can delete the gateway and create one of the new edition. But services will be affected.
- By default, the bandwidth of a VPN gateway changed to the new edition is 10 Mbit/s. You can adjust the bandwidth as required.

2.2.19 Does VPN Support IPv6?

Nο

Currently, only IPv4 is supported.

2.2.20 How Do I Determine My VPN Bandwidth Size?

Consider the following when you determine the bandwidth:

- Amount of data transmitted over a VPN tunnel in a period of time (Reserve enough bandwidth to prevent link congestion.)
- The egress bandwidth at the end of the VPN connection on the cloud must be less than that at the end of the VPN connection in the on-premises data center.

2.2.21 Which IKE Version Should I Select When I Create a VPN Connection?

IKEv2 is recommended because IKEv1 is not secure. In addition, IKEv2 outperforms IKEv1 in connection negotiation and establishment, authentication methods, dead peer detection (DPD) timeout processing, and security association (SA) timeout processing.

IKEv2 will be widely used, and IKEv1 will gradually phase out.

Introduction to IKEv1 and IKEv2

- The complexity of IKEV1, a hybrid protocol, inevitably brings some security and performance defects. This has become the bottleneck for the current IPsec system.
- The IKEv2 protocol reserves basic functions of IKEv1 and overcomes some problems brought by IKEv1. Moreover, for simplicity, efficiency, security, and robustness, RFC 4306, a document describes version 2 of IKE, combines the contents of what were previously separate IKEv1 documents. By minimizing core functions and default password algorithms, IKEv2 greatly improves the interoperation capability among different IPsec VPNs.

IKEv1 Security Vulnerabilities

• The cryptographic algorithms supported by IKEv1 have not been updated for more than 10 years. Also, IKEv1 does not support strong cryptographic algorithms such as AES-GCM and ChaCha20-Poly1305. For IKEv1, the E (Encryption) bit in the ISALMP header specifies that the payloads following the ISALMP header are encrypted, but any data integrity verification of those payloads is handled by a separate hash payload. This separation of encryption from data integrity protection prevents the use of authenticated encryption (AES-GCM) with IKEv1.

- IKEv1 protocol is vulnerable to DoS amplification attacks. IKEv1 is vulnerable to half-open connections.
 - IKEv2 can defend against DoS attacks.
- The IKEv1 aggressive mode is not secure enough. In aggressive mode, information packets are not encrypted. There are also brute-force attacks targeting at the aggressive mode, such as man-in-the-middle attacks.

Differences Between IKEv1 and IKEv2

Negotiation process

- IKEv1 SA negotiation consists of two phases. IKEv1 is complex and occupies a large amount of bandwidth. IKEv1 phase 1 negotiation aims to establish the IKE SA. This process supports the main mode and aggressive mode. Main mode uses six ISAKMP messages to establish the IKE SA, but aggressive mode uses only three. Therefore, aggressive mode is faster in IKE SA establishment. However, aggressive mode does not provide peer identity protection because key exchange and identity authentication are performed at the same time. IKEv1 phase 2 negotiation aims to set up the IPsec SA for data transmission. This process uses the fast exchange mode (3 ISAKMP messages) to complete the negotiation.
- Compared with IKEv1, IKEv2 simplifies the SA negotiation process. IKEv2 uses two exchanges (a total of 4 messages) to create an IKE SA and a pair of IPsec SAs. To create multiple pairs of IPsec SAs, only one additional exchange is needed for each additional pair of SAs.

For IKEv1 negotiation, its main mode requires nine (6+3) packets in total and its aggressive mode requires 6 (3+3) packets. IKEv2 negotiation requires only 4 (2+2) packets.

Authentication methods

- Only IKEv1 (requiring an encryption card) supports digital envelope authentication (HSS-DE).
- IKEv2 supports Extensible Authentication Protocol (EAP) authentication.
 IKEv2 can use an AAA server to remotely authenticate mobile and PC users and assign private IP addresses to these users. IKEv1 does not provide this function and must use L2TP to assign private IP addresses.
- Only IKEv2 supports IKE SA integrity algorithms.

• DPD timeout

- Only IKEv1 supports the retry-interval parameter. If a device sends a
 DPD packet but receives no reply within the specified retry-interval, the
 device records a DPD failure event. When the number of failure events
 reaches five, both the IKE SA and IPsec SA are deleted. The IKE SA
 negotiation will be started again when the device has IPsec traffic to
 handle.
- In IKEv2 mode, the retransmission interval increases from 1, 2, 4, 8, 16, 32 to 64 seconds. If no reply is received within eight consecutive transmissions, the peer end is considered dead, and the IKE SA and IPsec SA will be deleted.
- IKE SA timeout processing and IPsec SA timeout processing

In IKEv2, the IKE SA soft lifetime is 9/10 of the IKE SA hard lifetime plus or minus a random value to reduce the likelihood that two ends initiate renegotiation at the same time. Therefore, soft lifetime does not require manual settings in IKEv2.

Advantages of IKEv2 Over IKEv1

- Simplified SA negotiation process and improved negotiation efficiency.
- Closed many cryptographic loopholes, improving security.
- Supports EAP authentication, improving authentication flexibility and scalability.
- EAP is an authentication protocol that supports multiple authentication methods. The biggest advantage of EAP is scalability. That is, new authentication modes can be added without changing the original authentication system. EAP authentication has been widely used in dial-up access networks.
- IKEv2 employs an encrypted payload that is based on the design of ESP. The IKEv2 encrypted payload associates encryption and data integrity protection in a fashion that makes it possible to use authenticated encryption algorithms. AES-GCM ensures confidentiality, integrity, and authentication.

2.2.22 How Many Bits Do the DH Groups Used by VPN Have?

The Diffie-Hellman (DH) groups determine the strength of the key used in the key exchange process. Higher DH group numbers are usually more secure, but extra time is required to calculate the key.

Table 2-4 lists the bits corresponding to the DH groups used by VPN.

Table 2-4 Bit corresponding to each DH group

DH Group	Modulus
1	768 bits
2	1024 bits
5	1536 bits
14	2048 bits
15	3072 bits
16	4096 bits
19	ecp256 bits
20	ecp384 bits
21	ecp521 bits

□ NOTE

The following DH algorithms have security risks and are not recommended: DH group 1, DH group 2, and DH group 5.

2.2.23 Can I Visit Websites Across International Borders Using a VPN?

No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

2.2.24 Can I Deploy Applications on the Cloud, Databases in an On-premises Data Center, and Then Connect Them Through a VPN?

VPN connects a VPC and an on-premises network.

After a VPN is set up, service traffic can be transmitted between the VPC and onpremises data center. For an application server on the cloud, access to an onpremises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check whether the network latency and packet loss adversely affect service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

2.2.25 What Are the Differences Between the Application Scenarios and Connection Modes of IPsec and SSL VPNs?

Scenarios

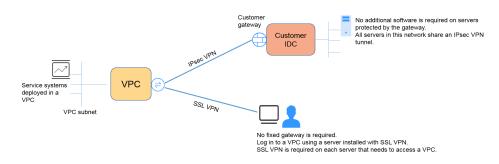
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to complete IPsec VPN negotiation.

SSL VPN needs to install a specified client software on the server, then the server connects to the SSL device through the username and password.



Ⅲ NOTE

IPsec VPN and SSL VPN are supported.

2.2.26 What Will I Be Charged for Creating a VPN? Will I Be Charged for VPN Gateway IP Addresses?

VPNs are billed on a yearly/monthly or pay-per-use basis. You need to pay for both the VPN gateway bandwidth/traffic and VPN connections. The billing modes available for a region are subject to those displayed on the page.

VPN gateways can be billed by traffic or bandwidth.

- 1. A yearly/monthly VPN gateway can only be billed by bandwidth. The price of a yearly/monthly VPN gateway includes the price of the VPN connections that can be created for the gateway and the bandwidth price.
- 2. The billing cycle of the pay-per-use billing mode is one hour. If you choose a pay-per-use VPN gateway, a VPN connection must be purchased together with the VPN gateway. The price includes the VPN gateway bandwidth or traffic price and the price of the VPN connection created together with the gateway. If you create another connection for the gateway, you will be charged for the additional connection.

- The IP address of the VPN gateway will not be billed.
- A VPN gateway cannot share a bandwidth with an EIP bound to an ECS.

2.2.27 What Is the Difference Between Billing a VPN Gateway by Bandwidth and by Traffic?

A pay-per-use VPN gateway can be billed by bandwidth or by traffic.

Their differences are as follows:

 Billed by bandwidth: The billing cycle is one hour. The generated fee depends on the bandwidth size. • Billed by traffic: The traffic fees generated each hour will be collected. The billing is based on the generated traffic going out of a VPC. The bandwidth size does not affect the price of the public traffic per GB.

2.2.28 Can a VPN Billed by Traffic Use a Shared Data Package?

No.

The VPN service is billed independently and cannot use shared data packages.

2.2.29 Can a VPN Gateway IP Address Be Retained After the VPN Gateway Is Deleted?

No. The VPN gateway IP address will be released after the VPN gateway is deleted.

Deleting a VPN gateway will also delete the resources associated with the gateway, such as public IP addresses.

NOTICE

Deleting the last connection of a pay-per-use VPN gateway will also delete the gateway. If you want to retain the IP address, do not delete the last VPN connection.

2.2.30 Do I Need to Purchase EIPs for Servers That Communicate with Each Other Through a VPN?

If your on-premises server needs to access an ECS on the cloud through a VPN, you do not need to purchase an EIP.

If the ECS needs to provide services accessible from the Internet, an EIP is required.

2.2.31 Where Can I Add Routes on the VPN Console to Reach the Remote Subnets?

When a VPN connection is created, routes are automatically delivered to reach the remote subnets.

2.2.32 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

After a VPN connection is created, you can locate the row that contains the VPN connection and choose **Operation** > **View Metric** to view the VPN connection status.

2.2.33 What Can I Do If VPN Connection Setup Fails?

- 1. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.
 - a. If the IKE policy has been set up during phase one and the IPsec policy has not been enabled in phase two, the IPsec policies at both ends of the VPN connection may be inconsistent.
 - b. If you use a Cisco physical device in your on-premises data center, it is recommended that you use MD5, and set **Authentication Mode** to **MD5** when configuring the IPsec policy for the VPN connection on the cloud.
- 2. Check whether the ACL rules are correct.

If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

3. Check whether the VPN connection is normal by pinging the local end from the remote end and pinging the remote end from the local end.

2.2.34 Which Direction of the Bandwidth Is Limited and What Is the Unit of the Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited.

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is more than 10 Mbit/s, the bandwidth in the inbound direction is the same as that of the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

2.3 Networking and Application Scenarios

2.3.1 Can I Visit Websites Across International Borders Using a VPN?

No.

VPN enables site-to-site connections by connecting the network of an on-premises data center to a VPC on the cloud.

2.3.2 Can I Deploy Applications on the Cloud, Databases in an On-premises Data Center, and Then Connect Them Through a VPN?

VPN connects a VPC and an on-premises network.

After a VPN is set up, service traffic can be transmitted between the VPC and onpremises data center. For an application server on the cloud, access to an onpremises database is logically the same as access to other hosts in the same LAN. Given this, it is feasible to use a VPN to connect an application on the cloud to a database in an on-premises data center.

This is a typical IPsec VPN scenario.

Additionally, there are no limitations on the service initiator. That is, service requests can be initiated from the cloud or the on-premises data center.

NOTICE

- After a VPN is set up, check whether the network latency and packet loss adversely affect service running.
- It is recommended that you run the ping command to check the packet loss and network latency details.

2.3.3 How Many VPN Connections Do I Need to Connect to Multiple On-premises Servers?

VPN uses the IPsec technology to connect your on-premises data center to a VPC on the cloud. As such, the number of VPN connections is related to the number of data centers where the servers to be connected to the cloud are located, but not to the number of servers.

In most cases, one on-premises data center has one public gateway. All servers connect to the Internet through this gateway. Therefore, you only need to configure one VPN connection to allow communications between the VPC and your on-premises data center.

2.3.4 Do I Need to Install IPsec Software on Each Server That Needs to Access an ECS to Establish a VPN Connection?

No.

VPN connects two LANs. Multiple servers in the on-premises data center use the same public IP address to access the cloud. If you install IPsec software on the on-premises servers, the VPN gateway on the cloud will receive negotiation packets from different servers and then the system receives a large amount of repeated negotiation information, which causes connection exceptions or even connection unavailability.

It is recommended that you use the egress firewall to configure a VPN to connect to the cloud. When creating a VPN, you can specify multiple CIDR blocks. You

should only allow servers of developers to access the ECS on the cloud based on the security group rules on the cloud or the security rules of the on-premises data center.

2.3.5 What Are the Differences Between the Application Scenarios and Connection Modes of IPsec and SSL VPNs?

Scenarios

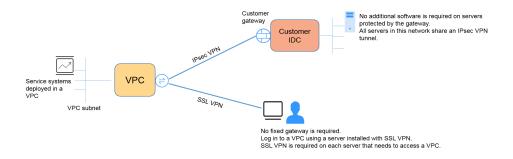
IPsec VPN connects two LANs, such as a branch and its headquarters (or a VPC), or an on-premises data center and a VPC.

SSL VPN connects a client to a LAN. For example, the portable computer of an employee on a business trip accesses the internal network of the company.

Connection Modes

IPsec VPN requires fixed gateways, such as firewalls or routers, at both ends. The administrator needs to configure gateways at both ends to complete IPsec VPN negotiation.

SSL VPN needs to install a specified client software on the server, then the server connects to the SSL device through the username and password.



MOTE

IPsec VPN and SSL VPN are supported.

2.3.6 Does a VPN Allow for Communications Between Two VPCs?

- If the two VPCs are deployed in the same region, use a VPC peering connection to connect them.
- If the two VPCs are deployed in different regions, use a VPN connection to connect them. The detailed operations are as follows:
 - a. Create a VPN gateway for each VPC and create VPN connections for the two VPN gateways.
 - b. Set the remote gateway address of each VPN connection to the gateway IP address of the peer side.
 - c. Set the remote subnets of each VPN connection to the subnets of the peer VPC.

d. The pre-shared keys and algorithm parameters of the two VPN connections must be the same.

2.3.7 What Are the Impacts of a VPN on an On-premises Network? What Are the Changes to the Route for Accessing an ECS?

When you configure a VPN, perform the following operations on the on-premises gateway:

- 1. Configure IKE and IPsec policies.
- 2. Specify the to-be-protected traffic (ACL rules).
- 3. Check the route configuration on the gateway to ensure that traffic destined for a VPC can be routed to the correct outbound interface (interface having an IPsec policy bound).

After the VPN configuration is complete, only the traffic matching the ACL rules enters the VPN tunnel.

For example, before a VPN is created, on-premises users access the ECS through the EIP bound to the ECS. After a VPN is created, data flows matching the ACL rules access the private IP address of the ECS through the VPN tunnel.

2.3.8 What Configurations Are Required at Both Ends of a VPN that Connects an On-premises Data Center to a VPC?

To implement the VPN interconnection, create a VPN on the cloud and configure the VPN device in the on-premises data center.

- Creating a VPN on the cloud: Buy a VPN gateway by selecting the billing mode, VPC, and bandwidth. Buy a VPN connection by specifying the gateway IP addresses, subnets, and negotiation policies at both ends.
- Configuring the on-premises VPN device: Select the public IP address on the on-premises data center side, complete configurations of IPsec negotiation phases 1 and 2 on the device that supports IPsec VPN, and then configure network routes, NAT, and security rules.

2.3.9 Can I Use a Network with Two Egresses to Establish Two VPN Connections with the Same VPC?

No.

When creating a VPN on the cloud, a local subnet is a VPC subnet, and a remote subnet is an on-premises subnet. If the two connections use the same local subnet and remote subnet, the VPN connections will fail.

2.3.10 Can I Connect Two VPCs in the Same Region Through a VPN?

No.

For two VPCs in the same region, you can use a VPC peering or Cloud Connect connection to connect them.

2.3.11 How Can I Connect Two VPCs in the Same Region?

Two VPCs in the same region can be connected using a VPC peering or Cloud Connect connection. VPC Peering can only connect VPCs in the same region, and Cloud Connect can also connect VPCs in different regions.

2.3.12 How Do I Replace a Direct Connect Connection with a VPN?

- 1. Ensure that the on-premises gateway supports IPsec VPN.
- 2. Create a VPN gateway and a VPN connection on the cloud. Select the VPC to which the Direct Connect connection uses for the VPN gateway.

NOTICE

When creating a VPN connection, configure its remote subnet as follows to avoid routing conflicts.

- Delete the virtual interface of the Direct Connect connection first and then configure the VPN connection.
- Divide the remote subnet into two subnets and configure the VPN connection. After the Direct Connect connection is deleted, configure the VPN connection again.

2.3.13 How Do I Enable Communications Between Two VPCs and an On-premises Network?

Network Topology

IDC-VPC 1-VPC 2



□ NOTE

IDC indicates the on-premises data center. A VPN connection is established between VPC 1 and the IDC.

Procedure

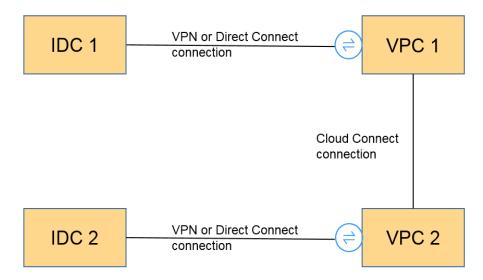
- 1. Check whether the two VPCs are in the same region.
 - If the two VPCs are in the same region, use a VPC peering or Cloud Connect connection (free of charge) to connect them.
 - If the two VPCs are in different regions, use a Cloud Connect connection.
 (You need to pay for the bandwidth fee.)
- 2. Establish a VPN connection between the on-premises and a VPC. In the onpremises data center, set the remote subnet to the subnets of VPC 1 and VPC

2. The local subnet of VPC 1 must contain the subnet connected through a VPC peering or Cloud Connect connection. The subnet route of the VPC peering or Cloud Connect connection should destine for the on-premises subnet.

2.3.14 How Do I Connect Four Subnets?

Figure 2-12 shows the network topology.

Figure 2-12 Network topology



- 1. Use a VPN or Direct Connect connection to connect IDC 1 to VPC 1.
- 2. Use a Cloud Connect connection to connect VPC 1 to VPC 2. (You can also use a VPC peering connection to enable communications between VPC 1 and VPC 2 if they are in the same region.)
- 3. Use a VPN or Direct Connect connection to connect IDC 2 to VPC 2.
- 4. Configure routes to allow traffic to go and from the four subnets involved in the VPN, Cloud Connect, and Direct Connect connections.

2.3.15 Do I Need Two VPN Connections to Connect Four Subnets of Two Regions If Each Region Has Two Subnets?

No.

Only one VPN connection is required between two regions. The subnets can all be added to the VPN connection.

In this scenario, if you attempt to create a second VPN connection, the management console displays a message indicating that a conflict occurs because the two connections have the same remote gateway address.

2.3.16 Can I Access OBS Through a VPN?

Yes.

With the help of the VPC endpoint service, you can access OBS through a VPN. You need to create two VPC endpoints for the private DNS server and OBS, respectively.

Then, configure the private DNS server and routes in your on-premises data center

2.3.17 How Do I Connect My Personal Computer to the Cloud Through a VPN?

Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.

To use VPN, on-premises devices must support the standard IPsec protocol.

2.3.18 How Do I Access ECSs at Home When My Enterprise Network Has Been Connected to the Cloud Through a VPN?

A VPN is an IPsec VPN that connects a VPC on the cloud and an on-premises local area network (LAN).

The home network is not a part of the LAN of your enterprise and cannot be directly connected to the VPC on the cloud.

If your host at home needs to access VPC resources on the cloud, your host can directly access the EIP of the cloud service or connect to the LAN of your enterprise through SSL VPN (if your enterprise supports SSL access) and then access VPC resources on the cloud through the LAN.

2.3.19 How Do I Establish a VPN Connection Temporarily If No IPsec-Capable On-Premises Device Is Available After I Purchase a VPN Gateway and VPN Connection?

To establish a VPN connection with the cloud, you must have an on-premises device that supports the standard IPsec protocol and have a fixed public IP address.

If the preceding requirements are not met, you can install third-party IPsec software on a host to temporarily connect to the cloud.

Recommended third-party IPsec software includes strongSwan, Openswan, and TheGreenBow. For details about the interconnection, see **Administrator Guide**.

2.3.20 How Do I Select a Proper Region on the Cloud When I Am Buying a VPN Gateway?

You can select a VPC in any region when you buy a VPN gateway.

But it is recommended that you select the region where your on-premises data center locates for lower network latency.

- For multiple VPCs in the same region, you only need to create one VPN gateway because the VPCs can be connected using VPC peering connections (free of charge).
- To connect to multiple VPCs in different regions, you can use VPN and Cloud Connect.

2.4 Billing and Payments

2.4.1 What Is the Difference Between Billing a VPN Gateway by Bandwidth and by Traffic?

A pay-per-use VPN gateway can be billed by bandwidth or by traffic. Their differences are as follows:

- Billed by bandwidth: The billing cycle is one hour. The generated fee depends on the bandwidth size.
- Billed by traffic: The traffic fees generated each hour will be collected. The billing is based on the generated traffic going out of a VPC. The bandwidth size does not affect the price of the public traffic per GB.

2.4.2 Can a VPN Billed by Traffic Use a Shared Data Package?

Nο.

The VPN service is billed independently and cannot use shared data packages.

2.4.3 For How Many VPN Connections Will I Be Charged to Connect VPCs in Different Regions of Huawei Cloud?

VPNs can be used to connect VPCs in different regions. The VPN bandwidth and connections of each region will be billed independently. Example:

In Region A, you establish one VPN connection with Region B and another VPN connection with Region C, then

- The VPN gateway of Region A has two connections.
- The VPN gateway of Region B has one connection.
- The VPN gateway of Region C has one connection.

In this case, you will be charged for four VPN connections.

2.4.4 How Do I Change the Billing Mode of a VPN Gateway from Pay-Per-Use to Yearly/Monthly?

Prerequisites

A pay-per-use VPN gateway is billed by bandwidth.
 To change the billing mode of a VPN gateway billed by traffic from pay-per-use to yearly/monthly, first change the VPN gateway from being billed by traffic to being billed by bandwidth and then from pay-per-use to yearly/monthly.

- The number of created VPN connections is less than 10.
- At least 10 more VPN connections can be created in your account.

Procedure

Perform the following operations:

- 1. Log in to the management console.
- 2. Click in the upper left corner and select the desired region and project.
- 3. Click **Service List** in the upper left corner. Under **Networking**, select **Virtual Private Network**.
- 4. In the navigation pane on the left, choose **Virtual Private Network > Classic**.
- 5. On the **VPN Gateways** page, locate the row that contains the target VPN gateway, choose **More** > **Change Billing Mode** in the **Operation** column.
- 6. In the displayed **Change Billing Mode** dialog box, click **OK**.

\cap	1	N		т	E
		IN	U	"	ᆮ

- The billing mode of a VPN gateway cannot be changed from yearly/monthly to pay-per-use.
- Resource quotas of a yearly/monthly VPN gateway can be decreased upon a renewal.
- In the yearly/monthly billing mode, **Required VPN Connections** indicates the total number of VPN connections that can be created for the VPN gateway free of charge.
- After you change the billing mode of a VPN gateway from pay-per-use to yearly/ monthly, the number of VPN connections that can be created for the VPN gateway is 10 by default.
- 7. On the **Change Subscription** page that is displayed, confirm the information about the VPN gateway and configure the renewal duration.
- 8. On the payment page, confirm the order information, select coupon or discount, and select the payment method.

Click Pay.

Changing the billing mode of a VPN gateway from pay-per-use to yearly/monthly will not affect your services.

2.4.5 Will a Yearly/Monthly VPN Gateway Be Automatically Renewed?

Yes.

Renewal fees will be automatically collected from your balance.

A yearly/monthly VPN gateway needs to be prepaid. To ensure that your connection is normal, top up your account if your balance is not enough.

2.4.6 Can I Unsubscribe from a Yearly/Monthly VPN Gateway?

Yes.

On the **VPN Gateways** page, locate the row that contains the VPN gateway you want to unsubscribe and click **Unsubscribe** in the **Operation** column. After you unsubscribe from a yearly/monthly VPN gateway, all VPN connections created for the gateway will also be deleted and cannot be recovered.

After the unsubscription, the remaining prepaid fees will be refunded.

2.4.7 When Will My VPN Resources Be Frozen? How Can I Unfreeze the VPN Resources?

- If pay-per-use VPN resources are in arrears, the resources enter the grace period, during which you can still access and use the resources. If the grace period ends and you have not paid off the arrears, the resources enter the retention period, during which the resources are frozen. Frozen resources are unavailable and cannot be modified or deleted. If the retention period ends and you still have not topped up your account and paid off the arrears, the resources will be released and cannot be restored. To ensure that resources are available, top up your account and pay off the arrears before the resources expire.
- Frozen VPN resources will become available after you renew them or top up your account. If a VPN connection is in the not connected state, initiate data flows to trigger the VPN connection and let it be in the normal state. For example, you can ping hosts on different subnets to trigger data flows.

2.5 Related Operations on the Console

2.5.1 What Are the Relationships Between a VPC, a VPN Gateway, and a VPN Connection?

When your on-premises data center needs to access ECSs in a VPC, you can create a VPN gateway to establish a VPN connection between your on-premises data center and the VPC.

VPC

- A VPC is a private network on the cloud. Multiple VPCs can be created in the same region while they are isolated from each other. A VPC can be divided into multiple subnets.
- You can use the VPN service to securely access ECSs in a VPC.

VPN gateway

- A VPN gateway is created in a VPC and is the access point of a VPN connection. One VPC can have only one VPN gateway, whereas one VPN gateway can have multiple VPN connections.
- With a VPN gateway, a secure, reliable, and encrypted connection can be established between a VPC and an on-premises data center or between VPCs in different regions.

VPN connection

A VPN connection is created for a VPN gateway and connects a VPC to an onpremises data center (or a VPC in another region).

The number of VPN connections is irrelevant to the number of local subnets or the number of remote subnets. It is only related to the number of on-premises data centers (or VPCs in other regions) to be connected to your VPC. The created VPN connections are displayed in the VPN connection list. You can also view the number of VPN connections created for each VPN gateway.

2.5.2 How Long Does It Take for Delivered VPN Configurations to Take Effect?

It takes 1 to 5 minutes for the VPN configurations to take effect.

After VPN configurations take effect, configure your gateway device on your on-premises network to complete tunnel negotiation with the VPN gateway.

2.5.3 Why Is the VPN Connection Always in the Not Connected State Even After Its Configuration Is Complete?

Ensure that the pre-shared keys and negotiation information at both ends are consistent. The local subnets and VPN gateway on the cloud are the remote subnets and remote gateway in the on-premises data center. The remote gateway and remote subnets on the cloud are the local gateway and local subnets in the on-premises data center.

Ensure that routes, NAT, and security policy rules are correctly configured on your on-premises gateway device. Then, ping the servers in subnets at both ends.

VPN is triggered based on data flows. After you configure VPN, ping a device in the peer subnet. Before running the ping command, disable the firewall function on the device, and allow inbound ICMP packets in the security group on the cloud.

Pinging the gateway IP address cannot trigger VPN negotiation. Ping the server in the subnet protected by the gateway.

2.5.4 Can a VPN Gateway IP Address Be Retained After the VPN Gateway Is Deleted?

No. The VPN gateway IP address will be released after the VPN gateway is deleted.

Deleting a VPN gateway will also delete the resources associated with the gateway, such as public IP addresses.

NOTICE

Deleting the last connection of a pay-per-use VPN gateway will also delete the gateway. If you want to retain the IP address, do not delete the last VPN connection.

2.5.5 Do I Need to Create a VPN Gateway or a VPN Connection for Creating a VPN? Which Information About a Created VPN Can Be Modified?

Prerequisites for creating a VPN

Create a VPC and VPC subnets. The VPC subnets cannot conflict with the on-premises subnets.

To create a VPN, you need to:

- Create a VPN gateway. Set Billing Mode, Region, Name, VPC, Billed By, and Bandwidth (Mbit/s). An IP address will be assigned to the VPN gateway after the gateway is created. Only configurations for Name and Bandwidth can be modified after the VPN gateway is created.
- Create a VPN connection. Specify the connection name, associated VPN gateway, local subnets, remote gateway, remote subnets, PSK, and negotiation policies. The connection name, local subnets, PSK, remote gateway, remote subnets, and negotiation policies can be modified after the VPN connection is created.

2.5.6 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to create ACL rules on your on-premises gateway device. The ACL rules will be referenced by IPsec policies.

When you configure VPN on the cloud, ACL rules will be automatically generated based on the local and remote subnets entered on the management console and then delivered to the VPN gateway.

Number of ACL rules = Number of local subnets x Number of remote subnets

2.5.7 What Do I Do If an Exception Occurs When I Add a Remote Subnet During VPN Connection Creation?

Check whether this remote subnet has been used as the destination of a VPC peering, Cloud Connect, or Direct Connect connection route, which causes route conflicts. If yes, delete the route and create a new one.

2.5.8 Where Can I Add Routes on the VPN Console to Reach the Remote Subnets?

When a VPN connection is created, routes are automatically delivered to reach the remote subnets.

2.5.9 Can I Call APIs to Manage Huawei Cloud VPN Resources?

VPN requires complex configurations. Currently, VPN resources cannot be created, queried, or modified through APIs. You can only manage VPN resources on the VPN console.

2.5.10 What Is a Remote Gateway and Remote Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a remote subnet and a remote gateway, respectively.

A remote gateway's IP address is a public IP address.

2.5.11 How Do I Disable PFS When Creating a VPN Connection?

The Perfect Forward Secrecy (PFS) function can be disabled for some regions. You are advised to enable PFS in your on-premises data center, because it improves IKE negotiation security in phase 2.

By default, PFS is disabled on some vendors' devices. Check the device configuration manual to ensure that PFS is enabled.

□ NOTE

- PFS is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. After PFS is configured, an additional DH exchange will be performed during IPsec SA negotiation, and a new IPsec SA key will be generated, improving IPsec SA security.
- To ensure security, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the on-premises gateway. Otherwise, the negotiation will fail.

2.5.12 How Many Local and Remote Subnets Can I Add to a VPN? Why Is an Error Message Displayed When I Update the Local Subnet by Specifying a CIDR Block?

- You can configure up to 5 local subnets. The product of the number of local subnets and the number of remote subnets cannot exceed 225.
- A VPC delivers VPC subnet routes based on the remote subnets of the VPN connection, remote subnets of the Direct Connect connection, and subnets of the VPC peering connection. Each subnet has one subnet route.
- The number of VPC subnet routes cannot exceed 200. That is, the total number of remote subnets of the VPN connection, remote subnets of the Direct Connect connection, subnets of the VPC peering connection, and custom routes in a VPC cannot exceed 200.

2.5.13 What Are the Precautions for Configuring the Local and Remote Subnets of a VPN Connection?

• The number of local subnets and the number of customer subnets are limited. If the number of local or customer subnets exceeds the upper limit, aggregate the subnets.

- The local subnet cannot include the CIDR block of the remote subnet. The remote subnet can include the CIDR block of the local subnet.
- There are routes pointing to the local subnets in the VPC where the VPN gateway resides.
- If there are two connections (connection A and connection B) created for a VPN gateway, and the remote subnet of connection A is within that of connection B, when the destination network to be accessed belongs to the overlapped CIDR block, the connection created first is matched first, regardless of the connection status. (Mask length match is not used for the policy-based VPN.)

2.5.14 Why the Status of a VPN Connection Is Not Connected on the Management Console When It Is Already Available?

There is a latency to display the latest VPN connection status on the management console.

If the service access is normal, the VPN connection is established. After several minutes, the VPN connection status will be **Connected**.

2.5.15 What Can I Do If a Message Is Displayed Indicating That the VPN Connection Does Not Exist After Negotiation Policies Are Modified?

This problem is caused by the page refresh interval.

When you modify the advanced settings, the system first deletes the VPN connection and then creates one. If the page contains the message indicating that the connection is being deleted or created displays for a short period of time, do not create the same connection (with the same local subnet, remote subnet, and remote gateway) again.

If the page remains in the connection deleting or creating state for a long time, submit a service ticket.

2.5.16 What Should I Do If I Cannot Create Connections for a VPN Gateway That Has No Bandwidth Information?

If a Classic VPN gateway has no bandwidth information, the VPN is of the earlier edition. This type of VPN cannot be created on the cloud anymore.

- Only one VPN connection can be created for each VPN gateway of the earlier edition and its bandwidth is not guaranteed. You can delete the gateway and create one of the new edition. But services will be affected.
- You can **submit a service ticket** to upgrade the gateway to the new edition. This will not affect your services.

By default, the bandwidth of a VPN gateway changed to the new edition is 10 Mbit/s. You can adjust the bandwidth as required.

2.5.17 How Do I Reset a VPN Connection?

- Disable the VPN connection on the on-premises device. After the status of the VPN connection on the cloud changes to **Not connected**, enable the VPN connection on the on-premises device.
- Change the remote gateway IP address of the VPN connection on the cloud to any other IP address. After the status of the connection in the on-premises data center changes to inactive, change the remote gateway IP address on the cloud to the previous IP address.

2.5.18 What Is the Maximum Bandwidth Supported by a VPN Gateway?

The maximum bandwidth supported by a VPN gateway is 300 Mbit/s.

2.5.19 Which IKE Version Should I Select When I Create a VPN Connection?

IKEv2 is recommended because IKEv1 is not secure. In addition, IKEv2 outperforms IKEv1 in connection negotiation and establishment, authentication methods, dead peer detection (DPD) timeout processing, and security association (SA) timeout processing.

IKEv2 will be widely used, and IKEv1 will gradually phase out.

Introduction to IKEv1 and IKEv2

- The complexity of IKEV1, a hybrid protocol, inevitably brings some security and performance defects. This has become the bottleneck for the current IPsec system.
- The IKEv2 protocol reserves basic functions of IKEv1 and overcomes some problems brought by IKEv1. Moreover, for simplicity, efficiency, security, and robustness, RFC 4306, a document describes version 2 of IKE, combines the contents of what were previously separate IKEv1 documents. By minimizing core functions and default password algorithms, IKEv2 greatly improves the interoperation capability among different IPsec VPNs.

IKEv1 Security Vulnerabilities

- The cryptographic algorithms supported by IKEv1 have not been updated for more than 10 years. Also, IKEv1 does not support strong cryptographic algorithms such as AES-GCM and ChaCha20-Poly1305. For IKEv1, the E (Encryption) bit in the ISALMP header specifies that the payloads following the ISALMP header are encrypted, but any data integrity verification of those payloads is handled by a separate hash payload. This separation of encryption from data integrity protection prevents the use of authenticated encryption (AES-GCM) with IKEv1.
- IKEv1 protocol is vulnerable to DoS amplification attacks. IKEv1 is vulnerable to half-open connections.

IKEv2 can defend against DoS attacks.

• The IKEv1 aggressive mode is not secure enough. In aggressive mode, information packets are not encrypted. There are also brute-force attacks targeting at the aggressive mode, such as man-in-the-middle attacks.

Differences Between IKEv1 and IKEv2

Negotiation process

- IKEv1 SA negotiation consists of two phases. IKEv1 is complex and occupies a large amount of bandwidth. IKEv1 phase 1 negotiation aims to establish the IKE SA. This process supports the main mode and aggressive mode. Main mode uses six ISAKMP messages to establish the IKE SA, but aggressive mode uses only three. Therefore, aggressive mode is faster in IKE SA establishment. However, aggressive mode does not provide peer identity protection because key exchange and identity authentication are performed at the same time. IKEv1 phase 2 negotiation aims to set up the IPsec SA for data transmission. This process uses the fast exchange mode (3 ISAKMP messages) to complete the negotiation.
- Compared with IKEv1, IKEv2 simplifies the SA negotiation process. IKEv2 uses two exchanges (a total of 4 messages) to create an IKE SA and a pair of IPsec SAs. To create multiple pairs of IPsec SAs, only one additional exchange is needed for each additional pair of SAs.

○ NOTE

For IKEv1 negotiation, its main mode requires nine (6+3) packets in total and its aggressive mode requires 6 (3+3) packets. IKEv2 negotiation requires only 4 (2+2) packets.

Authentication methods

- Only IKEv1 (requiring an encryption card) supports digital envelope authentication (HSS-DE).
- IKEv2 supports Extensible Authentication Protocol (EAP) authentication.
 IKEv2 can use an AAA server to remotely authenticate mobile and PC users and assign private IP addresses to these users. IKEv1 does not provide this function and must use L2TP to assign private IP addresses.
- Only IKEv2 supports IKE SA integrity algorithms.

DPD timeout

- Only IKEv1 supports the retry-interval parameter. If a device sends a DPD packet but receives no reply within the specified retry-interval, the device records a DPD failure event. When the number of failure events reaches five, both the IKE SA and IPsec SA are deleted. The IKE SA negotiation will be started again when the device has IPsec traffic to handle.
- In IKEv2 mode, the retransmission interval increases from 1, 2, 4, 8, 16, 32 to 64 seconds. If no reply is received within eight consecutive transmissions, the peer end is considered dead, and the IKE SA and IPsec SA will be deleted.

• IKE SA timeout processing and IPsec SA timeout processing

In IKEv2, the IKE SA soft lifetime is 9/10 of the IKE SA hard lifetime plus or minus a random value to reduce the likelihood that two ends initiate renegotiation at the same time. Therefore, soft lifetime does not require manual settings in IKEv2.

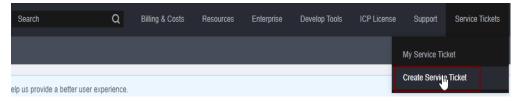
Advantages of IKEv2 Over IKEv1

- Simplified SA negotiation process and improved negotiation efficiency.
- Closed many cryptographic loopholes, improving security.
- Supports EAP authentication, improving authentication flexibility and scalability.
- EAP is an authentication protocol that supports multiple authentication methods. The biggest advantage of EAP is scalability. That is, new authentication modes can be added without changing the original authentication system. EAP authentication has been widely used in dial-up access networks.
- IKEv2 employs an encrypted payload that is based on the design of ESP. The IKEv2 encrypted payload associates encryption and data integrity protection in a fashion that makes it possible to use authenticated encryption algorithms. AES-GCM ensures confidentiality, integrity, and authentication.

2.5.20 What Are the Categories of VPN Service Tickets? How Do I Create a VPN Service Ticket?

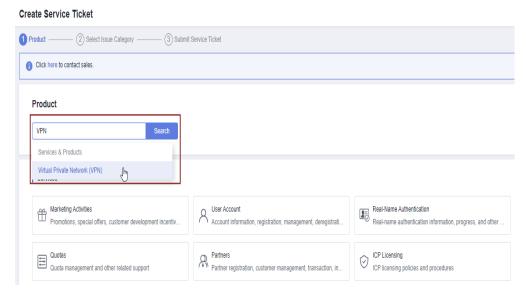
- 1. Log in to the management console.
- 2. In the upper right corner of the management console, choose **Service Tickets** > **Create Service Ticket**.

Figure 2-13 Create Service Ticket



Search for VPN and select Virtual Private Network (VPN).

Figure 2-14 Selecting Virtual Private Network (VPN)



4. Select an issue category.

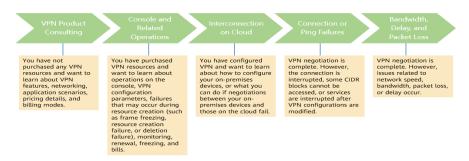
Figure 2-15 Select Issue Category



□ NOTE

When you **submit a service ticket**, select an issue category to facilitate problem handling.

Figure 2-16 Issue category and classification basis



2.5.21 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

◯ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

2.5.22 Which VPN Resources Can Be Monitored?

VPN Gateway

Bandwidth information that can be monitored includes inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view VPN gateway metrics, locate the target VPN gateway and click **View Metric** in the **Operation** column.

VPN Connection

The VPN connection status can be monitored.

Value 1 indicates that the connection is normal.

Value **0** indicates that the connection is not connected.

To view the VPN connection status, locate the target VPN connection and click **View Metric** in the **Operation** column.

2.5.23 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

After a VPN connection is created, you can locate the row that contains the VPN connection and choose **Operation** > **View Metric** to view the VPN connection status.

2.6 VPN Negotiation and Interconnection

2.6.1 What Devices Can Be Connected to the Cloud Through a VPN?

VPN supports the standard Internet Protocol Security (IPsec) protocol. A device in your on-premises data center can connect to the cloud if the device meets the following requirements:

- 1. Supports IPsec VPN.
- 2. Has a fixed public IP address, which can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

Most devices are routers and firewalls. For details about the interconnection configuration, see **Administrator Guide**.

■ NOTE

- Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.
- The following products can connect to the cloud through VPNs:
 - Devices: Huawei firewalls and access routers (ARs), Hillstone firewalls, and Check Point firewalls
 - Cloud services: Alibaba Cloud, Tencent Cloud, Amazon Web Services (AWS), and Microsoft Azure
 - Software: strongSwan
- The IPsec protocol is a standard IETF protocol. Devices that support IPsec can interconnect with the cloud through a VPN.

Most enterprise-class routers and firewalls support the IPsec protocol.

 However, some devices support IPsec VPN only after you purchase required software licenses.

Contact your on-premises data center administrator to confirm the device model with the vendor.

2.6.2 What Are VPN Negotiation Parameters? What Are Their Default Values?

Table 2-5 VPN negotiation parameters

Policy	Parameter	Value
IKE	Authentication Algorithm	 MD5(Insecure. Not recommended.) SHA1(Insecure. Not recommended.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 3DES (Insecure. Not recommended.) AES-256 AES-192 AES-128 (default value)

Policy	Parameter	Value
	DH Algorithm	Group 5(Insecure. Not recommended.)
		Group 2(Insecure. Not recommended.)
		Group 14 (default value)
		Group 1(Insecure. Not recommended.)
		Group 15
		Group 16
		Group 19
		Group 20
		Group 21
		NOTE In some regions, only Group 14, Group 2, and Group 5 are available.
	Version	• v1
		• v2 (default value)
	Lifetime (s)	86400 (default)
		Unit: second
		Value range: 60 to 604800
IPsec	Authentication Algorithm	 SHA1 (Insecure. Not recommended.) MD5 (Insecure. Not recommended.) SHA2-256 (default value) SHA2-384 SHA2-512
	Encryption Algorithm	 AES-128 (default value) AES-192 AES-256 3DES (Insecure. Not recommended.)

Policy	Parameter	Value	
	PFS	DH group 5(Insecure. Not recommended.)	
		DH group 2(Insecure. Not recommended.)	
		DH group 14 (default value)	
		DH group 1(Insecure. Not recommended.)	
		DH group 15	
		DH group 16	
		DH group 19	
		DH group 20	
		DH group 21	
		Disable	
		NOTE In some regions, only DH group 14, DH group 2, and DH group 5 are available.	
	Transfer Protocol	ESP (default value)	
		• AH	
		AH-ESP	
	Lifetime (s)	3600 (default)	
		Unit: second	
		Value range: 480 to 604800	

Ⅲ NOTE

- Perfect Forward Secrecy (PFS) is a security feature.
 - IKE negotiation has two phases, phase one and phase two. The key of phase two (IPsec SA) is derived from the key generated in phase one. Once the key in phase one is disclosed, the security of the IPsec VPN may be adversely affected. To improve the key security, IKE provides PFS. After PFS is configured, an additional DH exchange will be performed during IPsec SA negotiation, and a new IPsec SA key will be generated, improving IPsec SA security.
- To ensure security, PFS is enabled on the cloud side by default. Ensure that PFS is also enabled on the on-premises gateway. Otherwise, the negotiation will fail.
- To enable PFS, ensure that the configurations at both ends of a VPN are the same.
- The default traffic-based lifetime of an IPsec SA is 1,843,200 KB and cannot be changed for the VPN. This lifetime does not affect the establishment of an IPsec SA.

2.6.3 Will an IPsec VPN Connection Be Established Automatically?

After you complete configurations at both ends of an IPsec VPN connection, the VPN connection will not be automatically established only after data flows between the two ends of the connection. If no data flows between the cloud and

the on-premises data center, the VPN connection will always be in the down state. Any data generated by accessing servers or pinging between servers can trigger the establishment of a VPN connection.

The establishment of a VPN connection can be triggered in either of the following two conditions: The VPN gateway and the remote gateway automatically trigger negotiation. The cloud and on-premises servers access each other via the VPN connection to be established.

However, automatic establishment of a VPN connection cannot be triggered by a VPN gateway on the cloud. Verify that the establishment of your VPN connection can be triggered by the data flows between the two ends of the VPN connection. That is, check whether a VPN connection can be established after you ping a cloud server from an on-premises server, and whether a VPN connection can be established after you disconnect the connection and ping an on-premises server from a cloud server.

The source and destination addresses of the ping packets must be protected by the VPN. Before a VPN connection is established, the gateway IP addresses at both ends can be pinged. However, pinging the gateway IP addresses does not trigger the establishment of the VPN connection.

2.6.4 How Should I Configure an On-premises Gateway When I Use a VPN to Connect to the Cloud?

Determine on-premises subnets, VPC subnets, and gateway IP addresses at both ends.

Configure IPsec policies on the on-premises gateway according to the IPsec policies configured on the cloud. Add rules to the security group associated with the VPC to allow ICMP packets in both the inbound and outbound directions.

- Route setting: Add routes starting from the on-premises gateway and destining for the VPN gateway. The next hop of the route on the VPN gateway is the public gateway IP address in the outbound direction.
- NAT setting: On the on-premises gateway, disable NAT for the on-premises subnets that will access the VPC subnets. Add security group rules to allow mutual access between the on-premises subnets and the VPC subnets, and allow the UDP 500, UDP 4500, ESP (IP protocol 50), and AH (IP protocol 51) packets both from and to IP addresses of the VPN gateway on the cloud and the on-premises gateway.

2.6.5 Does VPN Support Interconnection with a Remote Gateway Through a Domain Name?

No. A VPN connection can only connect to a remote gateway through the gateway public IP address.

2.6.6 How Many Tunnels Does My VPN Connection Have?

The number of tunnels in a VPN connection = The number of local subnets x The number of remote subnets of the VPN connection

The status of a VPN connection is normal as long as one of the tunnels is in the active state. If you need each tunnel to be in the active state, data flows need to be triggered between every two subnets.

2.6.7 How Do I Allow Specific Servers to Access a VPC Subnet Through a Created VPN Connection?

Configurations in the on-premises data center

- Configure deny rules on VPN devices.
- Configure ACL rules on the router or switch.

Configurations on the cloud

- Configure security group rules to deny access from specific IP addresses.
- Configure ACL rules.

□ NOTE

All rules must be added to the device before the VPN tunnel is established. Do not change the local subnet and the remote subnet to restrict the access.

2.6.8 Do VPNs Have the DPD Function Enabled?

Yes.

By default, the dead peer detection (DPD) function is enabled for VPNs to detect the state of the IKE process in an on-premises data center.

After three consecutive detection failures, the IKE process in the on-premises data center is considered abnormal, and the tunnel on the cloud is automatically deleted.

The DPD protocol does not require that the peer end also be configured with DPD, but it requires that the peer end be able to respond to DPD detections. To ensure consistent tunnel states at the two ends, it is recommended that you enable DPD on your on-premises gateway to detect the IKE process state of the VPN service.

Ⅲ NOTE

After DPD fails, the tunnel will be deleted without affecting service stability.

DPD can detect exceptions in the IKE process at the peer end in time and reset the tunnel to ensure tunnel synchronization between the two ends. After a tunnel is deleted, if there is traffic transmitted over the tunnel, the tunnel can be re-established through negotiation.

2.6.9 How Can I Use Security Groups to Prevent ECSs in a VPC From Being Accessed Through a VPN to Implement Security Isolation?

You can configure security groups to allow access only to specific CIDR blocks or ECSs in a VPC through a VPN.

Configuration example: Prevent ECSs in VPC subnet 10.1.0.0/24 from accessing on-premises subnet 192.168.1.0/24.

Procedure:

- 1. Create security groups 1 and 2.
- 2. Security group 1 denies access from subnet 192.168.1.0/24.
- 3. Security group 2 allows access from subnet 192.168.1.0/24.
- 4. Associate ECSs in subnet 10.1.0.0/24 with security group 1 and associate other ECSs in the VPC with security group 2.

2.6.10 Will a VPN Connection Be Reestablished After Its Configuration Is Modified?

A VPN connection consists of local subnets, remote subnets, remote gateway, preshared key, IKE negotiation policies, and IPsec negotiation policies. A VPN connection is modified if any of the following happens:

- If the local and remote subnets are modified, the connection ID will remain unchanged. If not all subnets are updated, the established tunnel between subnets will not be re-established.
- If the IP address of the remote gateway is changed, even the connection ID will remain unchanged, the remote end has changed. Therefore, the VPN connection needs to be reestablished.
- If only the pre-shared keys of the connection are changed, the connection ID and status will remain unchanged. The keys will be checked again during renegotiation. If the keys do not match, the renegotiation fails.
- If the negotiation policy is modified (pre-shared key authentication is required), the connection ID will be changed and the connection needs to be re-established.

2.6.11 Why Cannot I Initiate Negotiation from Amazon Web Services to Huawei Cloud After They Are Interconnected Through a VPN?

After a VPN connection is established between Amazon Web Services (AWS) and Huawei Cloud, AWS works in Response mode and does not initiate negotiation. As such, SA establishment will not be triggered when an AWS EC2 accesses a Huawei Cloud ECS.

According to the AWS documentation, the customer side (the cloud connected to AWS) initiates negotiation by default, and you can also enable the AWS side to initiate negotiation.

2.6.12 Why Does the Negotiation Fail When Multiple Subnets Are Interconnected with Alibaba Cloud?

According to the configuration requirements of Alibaba Cloud, only IKEv2 can be used when multiple subnets are interconnected with the Alibaba cloud.

The number of connections you need to configure on Huawei Cloud is determined by multiplying the number of subnets on the Huawei Cloud side and that on the Alibaba Cloud side. For example, if there are two subnets on the Huawei Cloud side and three subnets on the Alibaba Cloud side, configure six connections on Huawei Cloud. Each connection protects only one flow.

2.6.13 How Do I Configure DPD for Interconnection with the Cloud?

By default, DPD is enabled on the cloud side and cannot be disabled.

Configure DPD as follows:

DPD-type: on-demand

DPD idle-time: 30s

• DPD retransmit-interval: 15s

• DPD retry-limit: 3

DPD msg: seq-hash-notify

The **DPD msg** format at both ends of the VPN connection must be the same, but the DPD type, idle time, retransmission interval, and retry limit can be different.

2.6.14 What Should I Do If My Firewall Cannot Receive Response Packets from the VPN Gateway in IKE Phase 1?

- Check whether the public IP addresses of the two ends can communicate with each other by running the ping command. By default, the cloud-side gateway IP address can be pinged.
- 2. Verify that the on-premises gateway (firewall) and cloud-side gateway can exchange packets through UDP ports 500 and 4500.
- 3. Verify that the source port number is not translated when the on-premises gateway connects to the cloud-side gateway. In a NAT traversal scenario, ensure that the source port number is not changed after NAT traversal.
- 4. The IKE negotiation settings at both ends must be the same. In the NAT traversal scenario, set the ID type in the on-premises data center to IP and the local ID on the cloud side to the post-NAT public IP address.

2.6.15 What Should I Do If My Firewall Cannot Receive Response Packets from a VPN Subnet?

- 1. Check the routes, security policies, NAT configuration, interesting traffic, and negotiation policies for phase 2 negotiation on the on-premises gateway device.
 - Route configurations: Route the data for accessing cloud subnets to tunnels.
 - Security policies: Allow traffic from on-premises subnets to cloud subnets.
 - NAT policies: Do not perform NAT when on-premises subnets access cloud subnets.
 - Interesting traffic: The interesting traffic configurations at both ends of a VPN connection are reversed. The address object name cannot be used for the interesting traffic configured using IKEv2.
 - Negotiation policies: Ensure the negotiations policies, especially PFS, at both ends are the same.

2. After confirming that both phase 1 and phase 2 negotiations are normal, ensure that the security group rules on the cloud allow the on-premises subnets to access the cloud subnets using ICMP.

2.6.16 How Many Bits Do the DH Groups Used by VPN Have?

The Diffie-Hellman (DH) groups determine the strength of the key used in the key exchange process. Higher DH group numbers are usually more secure, but extra time is required to calculate the key.

Table 2-6 lists the bits corresponding to the DH groups used by VPN.

Table 2-6 Bit corresponding to each DH group

DH Group	Modulus
1	768 bits
2	1024 bits
5	1536 bits
14	2048 bits
15	3072 bits
16	4096 bits
19	ecp256 bits
20	ecp384 bits
21	ecp521 bits

□ NOTE

The following DH algorithms have security risks and are not recommended: DH group 1, DH group 2, and DH group 5.

2.6.17 Which Remote VPN Devices Are Supported?

Most devices that meet IPsec VPN standards and reference protocol requirements can be used as remote VPN devices, for example, Cisco ASA firewalls, Huawei USG6000 series firewalls, USG9000 series firewalls, Hillstone firewalls, and Cisco ISR routers. **Table 2-7** lists the supported Huawei USG6000 and USG9000 firewalls.

Table 2-7 Huawei VPN devices

Supported Remote VPN Device	Description
Huawei USG6000 series	USG6320/6310/6510-SJJ USG6306/6308/6330/6350/6360/6370/6380/6390/6507/ 6530/6550/6570: 2048 USG6620/6630/6650/6660/6670/6680
Huawei USG9000 series	USG9520/USG9560/USG9580

Other devices that meet the requirements outlined in **Reference Standards and Protocols** can also be deployed. However, some devices may not be supported because of the inconsistent protocol implementation methods of these devices. If connections cannot be set up, rectify the fault by referring to **1.1.24 What Do I Do If a VPN Connection Fails to Be Established?** or contact technical support.

2.7 Connection or Ping Failure

2.7.1 Why Is the VPN Connection Always in the Not Connected State Even After Its Configuration Is Complete?

Ensure that the pre-shared keys and negotiation information at both ends are consistent. The local subnets and VPN gateway on the cloud are the remote subnets and remote gateway in the on-premises data center. The remote gateway and remote subnets on the cloud are the local gateway and local subnets in the on-premises data center.

Ensure that routes, NAT, and security policy rules are correctly configured on your on-premises gateway device. Then, ping the servers in subnets at both ends.

∩ NOTE

VPN is triggered based on data flows. After you configure VPN, ping a device in the peer subnet. Before running the ping command, disable the firewall function on the device, and allow inbound ICMP packets in the security group on the cloud.

Pinging the gateway IP address cannot trigger VPN negotiation. Ping the server in the subnet protected by the gateway.

2.7.2 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes for disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- DPD is not configured in your on-premises data center.
- Configuration is modified when the VPN connection is in use.
- Packets are fragmented because the data size exceeds the MTU.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- Local and remote subnets are matched pairs.
- SA lifetime settings at both ends of the VPN connection are the same.
- DPD is enabled on the on-premises gateway device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.
- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway is large enough to be used by the VPN connection.
- VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway.
- Ping the subnets at both ends continuously. The script is as follows:

```
#!/bin/sh
host=$1
if [ -z $host ]; then
    echo "Usage: `basename $0` [HOST]"
    exit 1
fi
log_name=$host".log"

while :; do
    result=`ping -W 1 -c 1 $host | grep 'bytes from '`
    if [ $? -gt 0 ]; then
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is down"| tee -a $log_name
    else
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is ok -`echo $result | cut -d ':' -f 2`"| tee -a $log_name
    fi
sleep 5 # avoid ping rain
done
#./ping.sh x.x.x.x >>/dev/null &
```

■ NOTE

- 1. Use the vi editor to copy the preceding script to the **ping.sh** file.
- 2. Run the **chmod 777 ping.sh** command to grant permissions to the file.
- 3. Run the ping command:

./ping.sh x.x.x.x >> /dev/null &

x.x.x.x indicates the IP address to be pinged.

4. Run the following command:

```
tail -f x.x.x.x.log
```

You can view the ping result in real time.

2.7.3 How Do I Quickly Restore an Interrupted IPsec VPN Connection?

- 1. Trigger IPsec negotiation by private network data flows. For example, two private networks at both ends of the VPN connection ping each other. If traffic can be properly triggered, deploy a continuous ping script. For details, see 2.12.1 How Can I Prevent VPN Disconnections?
- 2. If the negotiation cannot be triggered, check the Internet connectivity by pinging the VPN gateway IP address and the remote gateway IP address. By default, a VPN gateway responds to ICMP packets.
- 3. If connectivity is normal, check whether link switching occurs between outbound interfaces. That is, check whether the traffic for access to the VPN gateway is forwarded out from a non-negotiated interface.
- 4. If there are no multiple ports or the port path is normal, change the PSKs at both ends of the tunnel to trigger negotiation again.
- 5. If the negotiation fails, check whether the negotiation policies configured at both ends are consistent and whether the interesting traffic at both ends is mutually mirrored.
- 6. If the negotiation policies and interesting traffic configurations are correct, disable the VPN connection on the on-premises device. After the VPN connection state changes to **Not connected**, enable the VPN connection on the on-premises device and trigger a data flow.
- 7. If the negotiation still cannot be triggered, perform the following operations:
 - a. Record the negotiation policies, PSK, local subnets, remote gateway, and remote subnets of the VPN connection.
 - b. Use the existing VPN gateway to create another VPN connection. The negotiation policy, PSK, and local subnets are the same as those of the original VPN connection. Randomly configure the remote gateway and remote subnets.
 - c. After the new VPN connection is created, delete the original VPN connection and change the remote gateway and remote subnets of the new VPN connection to the recorded information.
 - d. Trigger the negotiation again.

If the fault persists, submit a service ticket to customer service personnel.

2.7.4 What Happens If the Bandwidth of a VPN Gateway Exceeds the Size I Specified When I Create the Gateway?

The VPN gateway bandwidth is used in the outbound direction of a VPC. If the bandwidth exceeds the size specified, network congestion will occur, some subnets cannot be accessed, or even the VPN connection will be interrupted, because the VPN detection packets may not be received.

In this case, you are advised to increase the VPN gateway bandwidth.

∩ NOTE

The maximum bandwidth of a VPN connection is 300 Mbit/s.

2.7.5 Will an IPsec VPN Connection Be Established Automatically?

After you complete configurations at both ends of an IPsec VPN connection, the VPN connection will not be automatically established only after data flows between the two ends of the connection. If no data flows between the cloud and the on-premises data center, the VPN connection will always be in the down state. Any data generated by accessing servers or pinging between servers can trigger the establishment of a VPN connection.

The establishment of a VPN connection can be triggered in either of the following two conditions: The VPN gateway and the remote gateway automatically trigger negotiation. The cloud and on-premises servers access each other via the VPN connection to be established.

However, automatic establishment of a VPN connection cannot be triggered by a VPN gateway on the cloud. Verify that the establishment of your VPN connection can be triggered by the data flows between the two ends of the VPN connection. That is, check whether a VPN connection can be established after you ping a cloud server from an on-premises server, and whether a VPN connection can be established after you disconnect the connection and ping an on-premises server from a cloud server.

■ NOTE

The source and destination addresses of the ping packets must be protected by the VPN.

Before a VPN connection is established, the gateway IP addresses at both ends can be pinged. However, pinging the gateway IP addresses does not trigger the establishment of the VPN connection.

2.7.6 Why ECSs at Both Ends of a Normal Cross-Region VPN Connection Cannot Access Each Other?

By default, a security group allows all outbound traffic. To allow inbound traffic, add inbound rules to the security group. Ensure that the security group associated with the ECS that needs to receive ping packets allows inbound ICMP requests.

2.7.7 Why Subnets at Both Ends of a Normal VPN Connection Cannot Access Each Other?

The VPN connection is normal, indicating that the negotiation parameters at both ends of the VPN connection are correct. You need to perform the following operations:

- Verify that routes to the VPN device in your on-premises data center are correctly configured.
- Verify that inter-subnet data exchange is allowed on the VPN device.
- Verify that NAT is not performed on the on-premises subnets that need to access the cloud.
- Verify that mutual access between the public IP addresses of the VPN gateway and customer gateway is permitted.

2.7.8 What Do I Do If a VPN Connection In Use Is Interrupted and a Message Is Displayed Indicating That Traffic from IP Addresses Not Whitelisted Generates?

This is usually caused by ACL configuration mismatch between the local and remote gateways.

- 1. Verify that at the two ends of the VPN connection, the local and remote subnets are reversed and the ACL configurations are also reversed.
- 2. Use the subnet/mask format when you configure interesting traffic in your on-premises data center. Do not use the address object mode since it may cause incompatibility issues.

2.7.9 What Do I Do If a VPN Connection Is Interrupted and a Message Is Displayed Indicating That the DPD Times Out?

This happens because there is no data exchange over the VPN connection. When the SA lifetime ends, the VPN connection is deleted as the peer end does not respond to the dead peer detection (DPD).

Solution

- 1. Enable DPD on the on-premises gateway device and test whether data flows from both ends can trigger connection establishment.
- 2. Deploy a ping shell script on the servers at both ends. Alternatively, configure a keepalive function (for example, NQA on Huawei devices and IP SLA on Cisco devices) on the on-premises gateway device to keep the connection alive.

2.7.10 Why the Status of a VPN Connection Is Not Connected on the Management Console When It Is Already Available?

There is a latency to display the latest VPN connection status on the management console.

If the service access is normal, the VPN connection is established.

2.7.11 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

After a VPN connection is created, you can locate the row that contains the VPN connection and choose **Operation** > **View Metric** to view the VPN connection status.

2.7.12 What Can I Do If VPN Connection Setup Fails?

1. Check the IKE and IPsec policies to see whether the negotiation modes and encryption algorithms at both ends of the VPN connection are the same.

- a. If the IKE policy has been set up during phase one and the IPsec policy has not been enabled in phase two, the IPsec policies at both ends of the VPN connection may be inconsistent.
- b. If you use a Cisco physical device in your on-premises data center, it is recommended that you use MD5, and set **Authentication Mode** to **MD5** when configuring the IPsec policy for the VPN connection on the cloud.
- 2. Check whether the ACL rules are correct.

If the subnets of your on-premises data center are 192.168.3.0/24 and 192.168.4.0/24, and the VPC subnets are 192.168.1.0/24 and 192.168.2.0/24, configure the ACL rules for each on-premises subnet to allow communication with the VPC subnets. The following provides an example of ACL configurations:

```
rule 1 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 2 permit ip source 192.168.3.0 0.0.0.255 destination 192.168.2.0 0.0.0.255 rule 3 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.1.0 0.0.0.255 rule 4 permit ip source 192.168.4.0 0.0.0.255 destination 192.168.2.0 0.0.0.255
```

3. Check whether the VPN connection is normal by pinging the local end from the remote end and pinging the remote end from the local end.

2.7.13 What Should I Do If I Cannot Access the ECSs on the Cloud from My On-premises Data Center or LAN After the VPN Connection Has Been Set Up?

The security group denies access from all sources by default. If you want to access your ECSs, modify the security group rules and allow the access from the onpremises subnets.

2.7.14 Why Is Not Connected Displayed as the Status for a Successfully Created VPN Connection?

After a VPN connection is created, its status changes to **Normal** only after servers at both ends of the VPN connection communicate with each other.

• IKE v1:

If no traffic goes through the VPN connection for a period of time, the VPN connection needs to be renegotiated. The negotiation time depends on the value of **Lifetime (s)** in the IPsec policy. Generally, **Lifetime (s)** is set to **3600** (1 hour), indicating that the negotiation will be initiated in the fifty-fourth minute. If the negotiation succeeds, the connection remains to the next round of negotiation. If the negotiation fails, the VPN connection status changes to **Not Connected** within one hour. The connection can be restored only after the two ends of the VPN connection communicate with each other. The disconnection can be avoided by using a network monitoring tool, such as IP SLA, to generate packets.

• IKE v2: If no traffic goes through the VPN connection for a period of time, the VPN connection remains in the connected status.

2.7.15 Do VPNs Have the DPD Function Enabled?

Yes.

By default, the dead peer detection (DPD) function is enabled for VPNs to detect the state of the IKE process in an on-premises data center.

After three consecutive detection failures, the IKE process in the on-premises data center is considered abnormal, and the tunnel on the cloud is automatically deleted.

The DPD protocol does not require that the peer end also be configured with DPD, but it requires that the peer end be able to respond to DPD detections. To ensure consistent tunnel states at the two ends, it is recommended that you enable DPD on your on-premises gateway to detect the IKE process state of the VPN service.

After DPD fails, the tunnel will be deleted without affecting service stability.

DPD can detect exceptions in the IKE process at the peer end in time and reset the tunnel to ensure tunnel synchronization between the two ends. After a tunnel is deleted, if there is traffic transmitted over the tunnel, the tunnel can be re-established through negotiation.

2.8 EIPs

2.8.1 Can a VPN Gateway IP Address Be Retained After the VPN Gateway Is Deleted?

No. The VPN gateway IP address will be released after the VPN gateway is deleted.

Deleting a VPN gateway will also delete the resources associated with the gateway, such as public IP addresses.

NOTICE

Deleting the last connection of a pay-per-use VPN gateway will also delete the gateway. If you want to retain the IP address, do not delete the last VPN connection.

2.8.2 Can an EIP Be Used as a VPN Gateway IP Address?

No.

The IP address of a VPN gateway is of preset configurations and is automatically assigned when the VPN gateway is created. An EIP cannot be used by a VPN gateway.

2.8.3 Do I Need to Purchase EIPs for Servers That Communicate with Each Other Through a VPN?

If your on-premises server needs to access an ECS on the cloud through a VPN, you do not need to purchase an EIP.

If the ECS needs to provide services accessible from the Internet, an EIP is required.

2.8.4 Why Does an ECS Have EIP Access Information After I Enable a VPN?

This occurs because the ECS has an EIP bound before the VPN is used. That is, you can access the ECS through the VPN or the EIP.

After the VPN is established, traffic from servers meeting ACL rules can enter the tunnel to access ECSs.

- If an EIP is bound to an ECS, devices on a non-VPN network can access the ECS using the EIP.
- If the ECS can be accessed only through a VPN, unbind the EIP from the ECS
 after the VPN connection is active. When an ECS needs an EIP, you can use
 ACL rules to specify the traffic that can access the ECS through the EIP.

Retaining an EIP or not depends on your services. If an ECS is used to access an on-premises data center through a VPN, and is also used to provide services accessible from the Internet, its EIP needs to be retained.

2.8.5 Can My On-premises Gateway Have No Fixed Public IP Address?

No.

To connect your on-premises data center to the cloud through a VPN, your onpremises gateway must have a fixed public IP address. This public IP address can be statically configured, or translated through NAT in NAT traversal scenarios (your device is deployed behind a NAT gateway).

□ NOTE

Common home broadband routers, Windows hosts providing VPN services (such as L2TP), and personal mobile terminals cannot connect to the cloud through a VPN.

2.9 Route Configurations

2.9.1 What Is a Remote Gateway and Remote Subnet in a VPN Connection?

A VPN connection is created on the cloud. As such, a subnet of a VPC is a local subnet, and a VPN gateway created on the cloud is a local gateway. The subnet and gateway in an on-premises data center connected to the VPC are a remote subnet and a remote gateway, respectively.

A remote gateway's IP address is a public IP address.

2.9.2 Where Can I Add Routes on the VPN Console to Reach the Remote Subnets?

When a VPN connection is created, routes are automatically delivered to reach the remote subnets.

2.9.3 Do I Need to Add a Route for an ECS with Multiple NICs to Reach the On-premises Network?

- If a primary NIC is used to establish a VPN with the on-premises network, no route needs to be added.
- If a non-primary NIC is used to establish a VPN with the on-premises network, add a route to reach the gateway of the non-primary NIC.

2.10 Subnet Setting

2.10.1 What Are the Precautions for Configuring the Local and Remote Subnets of a VPN Connection?

- The number of local subnets and the number of customer subnets are limited. If the number of local or customer subnets exceeds the upper limit, aggregate the subnets.
- The local subnet cannot include the CIDR block of the remote subnet. The remote subnet can include the CIDR block of the local subnet.
- There are routes pointing to the local subnets in the VPC where the VPN gateway resides.
- If there are two connections (connection A and connection B) created for a VPN gateway, and the remote subnet of connection A is within that of connection B, when the destination network to be accessed belongs to the overlapped CIDR block, the connection created first is matched first, regardless of the connection status. (Mask length match is not used for the policy-based VPN.)

2.10.2 How Many Local and Remote Subnets Can I Add to a VPN? Why Is an Error Message Displayed When I Update the Local Subnet by Specifying a CIDR Block?

- You can configure up to 5 local subnets. The product of the number of local subnets and the number of remote subnets cannot exceed 225.
- A VPC delivers VPC subnet routes based on remote subnets of a VPN connection, remote subnets of a Direct Connect connection, subnets of a VPC peering connection, and subnets of a Cloud Connect connection. Each subnet has one route.
- The number of VPC subnet routes cannot exceed 200. That is, in a VPC, the total number of remote subnets of a VPN connection, remote subnets of a Direct Connect connection, subnets of a VPC peering connection, and subnets of a Cloud Connect connection, and custom routes cannot exceed 200.

2.10.3 What Do I Do If an Exception Occurs When I Add a Remote Subnet During VPN Connection Creation?

Check whether this remote subnet has been used as the destination of a VPC peering, Cloud Connect, or Direct Connect connection route, which causes route conflicts. If yes, delete the route and create a new one.

2.10.4 Can a VPN Gateway IP Address Be Retained After the VPN Gateway Is Deleted?

No. The VPN gateway IP address will be released after the VPN gateway is deleted.

Deleting a VPN gateway will also delete the resources associated with the gateway, such as public IP addresses.

NOTICE

Deleting the last connection of a pay-per-use VPN gateway will also delete the gateway. If you want to retain the IP address, do not delete the last VPN connection.

2.10.5 How Do I Plan the CIDR Block of a VPC Accessed over a VPN Connection?

- The VPC CIDR block cannot conflict with the on-premises CIDR block.
- To avoid conflicts with cloud service addresses, do not use 127.0.0.0/8, 169.254.0.0/16, 224.0.0.0/3, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 for your on-premises network.

If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.

2.10.6 How Is a VPN Gateway IP Address Allocated?

VPN gateway IP addresses are a group of IP addresses planned before VPN gateways are purchased. These IP addresses are preset with VPN configurations.

When you buy a VPN gateway, the system randomly assigns an IP address and binds it to the VPC you selected. This IP address can be bound to only one VPC.

The IP address of the VPN gateway has preset data. Therefore, it is not interchangeable with an EIP, and you cannot specify an EIP as the VPN gateway IP address when you are buying the VPN gateway. The VPN gateway IP address can only be assigned randomly from the preset VPN IP address pool. When a VPN gateway is deleted, the binding relationship between the gateway IP address and the gateway VPC is released. When a new VPN gateway is purchased, the system randomly allocates a new gateway IP address.

2.10.7 What Is the Limitation on the Number of Local and Remote Subnets of a VPN?

- A maximum of five local subnets are supported. The product of the number of local subnets and that of remote subnets in a VPN cannot exceed 225.
- A VPC delivers VPC subnet routes based on the remote subnets of VPN connections, remote subnets of Direct Connect connections, and subnets of VPC peering connections. Each subnet has one subnet route.
- The number of VPC subnet routes cannot exceed 200. That is, the total number of remote subnets of VPN connections, remote subnets of Direct Connect connections, subnets of VPC peering connections, and custom routes in the same VPC cannot exceed 200.

2.11 VPN Interesting Traffic

2.11.1 Do I Need to Configure ACL Rules on the Console After I Configure ACL Rules on the On-premises Gateway Device?

You need to create ACL rules on your on-premises gateway device. The ACL rules will be referenced by IPsec policies.

When you configure VPN on the cloud, ACL rules will be automatically generated based on the local and remote subnets entered on the management console and then delivered to the VPN gateway.

Number of ACL rules = Number of local subnets x Number of remote subnets

2.11.2 How Do I Configure and Modify the Interesting Traffic of a VPN on the Cloud?

The number of rules that specify interesting traffic is the product of the number of local subnets and the number of remote subnets. For example, when there are local subnets A and B and remote subnets C, D, and E, the following six ACL rules need to be configured to specify interesting traffic:

```
rule 1 permit ip source A destination C
rule 2 permit ip source A destination D
rule 3 permit ip source A destination E
rule 4 permit ip source B destination C
rule 5 permit ip source B destination D
rule 6 permit ip source B destination E
```

If you modify the local or remote subnets on the management console, the interesting traffic configuration is automatically updated. That is, ACL rules on the cloud are modified.

2.12 Keeping VPN Connection Alive

2.12.1 How Can I Prevent VPN Disconnections?

VPN connections are renegotiated when the IPsec SA lifetime is about to expire or the data transmitted through a VPN connection exceeds 20 GB. Usually, renegotiation does not interrupt VPN connections.

Most disconnections are caused by incorrect configurations at the two ends of the VPN connection or renegotiation failures due to Internet exceptions.

Common causes for disconnections are as follows:

- ACLs at both ends of the VPN connection do not match.
- SA lifetime settings at both ends of the VPN connection are different.
- DPD is not configured in your on-premises data center.
- Configuration is modified when the VPN connection is in use.
- Packets are fragmented because the data size exceeds the MTU.
- Jitter occurs on the carrier's network.

As such, ensure that the following VPN configurations are correct to keep VPN connections alive:

- Local and remote subnets are matched pairs.
- SA lifetime settings at both ends of the VPN connection are the same.
- DPD is enabled on the on-premises device, and the number of detection times is 3 or more.
- Parameters are modified at both ends of the VPN connection during the use of the VPN connection.
- Set TCP MAX-MSS to 1300 for the on-premises gateway device.
- The bandwidth of the on-premises gateway is large enough to be used by the VPN connection.
- VPN connection negotiation can be triggered by both ends and active negotiation has been enabled on the on-premises gateway.
- Ping the subnets at both ends continuously. The script is as follows:

```
#!/bin/sh
host=$1
if [ -z $host ]; then
    echo "Usage: `basename $0` [HOST]"
    exit 1
fi
log_name=$host".log"

while :; do
    result=`ping -W 1 -c 1 $host | grep 'bytes from '`
    if [ $? -gt 0 ]; then
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is down"| tee -a $log_name
    else
        echo -e "`date +'%Y/%m/%d %H:%M:%S'` - host $host is ok -`echo $result | cut -d ':' -f 2`"| tee -a $log_name
    fi
sleep 5 # avoid ping rain
done
#./ping.sh x.x.x.x >>/dev/null &
```


- 1. Use the vi editor to copy the preceding script to the ping.sh file.
- 2. Run the **chmod 777 ping.sh** command to grant permissions to the file.
- 3. Run the ping command:

./ping.sh x.x.x.x >>/dev/null &

x.x.x.x indicates the IP address to be pinged.

4. After the ping command is executed, the *x.x.x.x.***log** file is generated. Run the following command:

tail -f x.x.x.x.log

You can view the ping result in real time.

2.13 Monitoring

2.13.1 Which VPN Resources Can Be Monitored?

VPN Gateway

Bandwidth information that can be monitored includes inbound traffic, inbound bandwidth, outbound traffic, outbound bandwidth, and outbound bandwidth usage.

To view VPN gateway metrics, locate the target VPN gateway and click **View Metric** in the **Operation** column.

VPN Connection

The VPN connection status can be monitored.

Value 1 indicates that the connection is normal.

Value **0** indicates that the connection is not connected.

To view the VPN connection status, locate the target VPN connection and click **View Metric** in the **Operation** column.

2.13.2 Will I Be Notified If a VPN Connection Is Interrupted?

The VPN connection status can be monitored. After a VPN connection is created, the VPN service reports the connection status information to Cloud Eye, but does not send alarm notifications to you. To receive notifications, create alarm rules and enable **Alarm Notification** on the Cloud Eye console.

After a VPN connection is created, you can locate the row that contains the VPN connection and choose **Operation** > **View Metric** to view the VPN connection status.

2.13.3 Can I View the Traffic of Each VPN Connection?

No. VPN traffic monitoring is about the VPN gateway. You can view the inbound and outbound traffic as well as the inbound and outbound bandwidths of a VPN gateway, but cannot view the traffic usage of a specific VPN connection.

2.13.4 Will I Be Notified When the VPN Monitoring Result Is Abnormal?

Yes.

You can configure, on the Simple Message Notification (SMN) and Cloud Eye consoles, to receive notifications if abnormal VPN monitoring results occur.

Creating Topics and Adding Subscriptions on the SMN Console

- Log in to the management console.
 Under Management & Governance, select Simple Message Notification.
- Choose Topic Management > Topics and click Create Topic to create a topic, for example, VPN-huaweicloud.
- Choose Topic Management > Subscriptions and click Add Subscription.
 Select a topic, set Protocol to Email, and enter an email address for receiving alarm notifications in the Endpoint box.

-	\neg	B 1	_	
		NI	"	

After the subscription is added, the system will send a confirmation email to your email address. Confirm the subscription in your email.

Creating VPN Alarm Rules on the Cloud Eye Console

- Log in to the management console.
 Under Management & Governance, select Cloud Eye.
- 2. Create an alarm rule for monitoring the bandwidth usage of a VPN gateway. Enter an alarm rule name, select Elastic IP and Bandwidth for Resource Type, set Dimension to Bandwidths, Monitoring Scope to Specific resources and select the target VPN gateway, set Method to Create manually, and Alarm Policy to Outbound Bandwidth Usage, 5 consecutive periods, >, and 90. Set Notification Object to an SMN topic and use the default settings for other parameters.
- 3. Create an alarm rule for monitoring the VPN connection status.

 The creation process is similar to that of the bandwidth. Select Virtual Private Network for Resource Type, set Dimension to VPN connections, Monitoring Scope to Specific resources and select the target VPN connection, set Method to Create manually, and Alarm Policy to VPN Connection Status, <, and 1. Set Notification Object to an SMN topic and use the default settings for other parameters.
- 4. Create an alarm rule for monitoring your on-premises links.
 - Create a website monitoring task, set **Type** to **PING**, **URL** to the gateway IP address of your on-premises data center, and retain the default settings for other parameters. Create an alarm rule, select **Website Monitoring** for **Resource Type**, set **Monitoring Scope** to **Specific resources** and select the target website monitoring task, set **Method** to **Create manually**, and **Alarm Policy** to **Available Monitoring Location Count**, and configure other parameter as required. Set **Notification Object** to an SMN topic and use the default settings for other parameters.

2.14 Bandwidth and Network Speed

2.14.1 What Is the Actual Network Speed of a VPN Connection?

A VPN connection has been created. Two ECSs have been created with one at the local end and the other at the remote end. The two ECSs can ping each other.

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 200 Mbit/s:

 If the ECSs at the two ends of the VPN run Windows, use iPerf3 and FileZilla (a free FTP application for file uploading and downloading) to test the network speed.

□ NOTE

The test shows that the average network speed of the VPN is 180 Mbit/s, and there is about 10% network speed deviation. The TCP and FTP protocols have the congestion control mechanism, and the IPsec protocol adds a new IP header. Therefore, about 10% network speed deviation is normal for the VPN network.

Figure 2-17 shows the test result.

Figure 2-17 Test result for 200 Mbit/s bandwidth (iPerf3 client)

```
C:\Windows>iperf3 -c 10.1.0.180 -i 1 -w 1M
Connecting to host 10.1.0.180, port 5201
4] local 192.168.0.75 port 49212 connected to 10.1.0.180 port 5201
 ID] Interval
                            Iransfer
                                          Bandwidth
        0.00-1.01
                     sec 17.1 MBytes
sec 30.0 MBytes
                                           142 Mbits/sec
        1.01-2.00
                                           253 Mbits/sec
                           19.8 MBytes
                                           165 Mbits/sec
        2.00-3.01
                     sec
                                           194 Mbits/sec
        3.01-4.01
                           23.2 MBytes
                     sec
        4.01-5.00
                           18.9 MBytes
                                           161 Mbits/sec
                     sec
                           26.2 MBytes
                                           219 Mbits/sec
                     sec
        6.01-7.01
                           18.4 MBytes
                                           153 Mbits/sec
                     sec
                                            195 Mbits/sec
                           23.2 MBytes
                           21.1 MBytes
                                            180 Mbits/sec
        9.00-10.01
                           21.0 MBytes
                                            174 Mbits/sec
 ID] Interval
                            Transfer
                                          Bandwidth
                            219 MBytes
219 MBytes
        0.00-10.01 sec
                                           183 Mbits/sec
183 Mbits/sec
                                                                                sender
        0.00-10.01 sec
                                                                                receiver
perf Done.
```

Figure 2-18 shows the test result.

erver listening on 5201 ted connection from 192.168.0.75, port 49211 local 10.1.0.180 port 5201 connected to 192.168.0.75 port 49212 Transfer Bandwidth
15.1 MBytes 127 Mbits
30.2 MBytes 252 Mbits
19.7 MBytes 166 Mbits 0.00-1.00 1.00-2.01 2.01-3.00 3.00-4.01 4.01-5.01 127 Mbits/sec 252 Mbits/sec 19.7 MBytes 23.6 MBytes 18.6 MBytes 26.3 MBytes 18.4 MBytes 23.4 MBytes 21.5 MBytes 20.4 MBytes 1.32 MBytes Mbits/sec Mbits/sec Mbits/sec Mbits/sec Mbits/sec Mbits/sec 173 Mbits/sec 162 Mbits/sec Interval Transfer B sec 0.00 Bytes 0. B≱ndwidth 0.00 bits/sec 0.00-10.07 sender

Figure 2-18 Test result for 200 Mbit/s bandwidth (iPerf3 server)

2. If the ECSs at the two ends of the VPN run CentOS 7, use iPerf3 to test the network speed. The network speed can reach 180 Mbit/s.

182 Mbits/sec

219 MBytes

3. If the ECS functioning as the server runs CentOS 7, and the ECS functioning as the client runs Windows, use iPerf3 and FileZilla to test the network speed. The network speed is about 20 Mbit/s, a slow network speed. That is because TCP implementations on Windows and that on Linux are different. Therefore, if the ECSs at the two ends of the VPN run different OSs, the VPN network speed does not meet the bandwidth requirements.

Figure 2-19 shows the test result.

Figure 2-19 Test result when ECSs at the two ends run different OSs (iPerf3)

```
::\Windows>iperf3 -c 10.1.0.182 -i 1 -w 1M
Connecting to host 10.1.0.182, port 5201
41 local 192.168.0.75 port 49426 connected to 10.1.0.182 port 5201
 ID] Interval
                            Transfer
                                          Bandwidth
                           4.38 MBytes 36.7 Mbits/sec
4.50 MBytes 37.7 Mbits/sec
       0.00 - 1.00
        1.00-2.00
                      នមន
        2.00-3.00
                      sec
                           5.12 MBytes 43.0 Mbits/sec
        3.00-4.00
                      sec
                           1.75 MBytes 14.7 Mbits/sec
        4.00-5.00
                           2.12 MBytes
                                          17.8 Mbits/sec
                      sec
        5.00-6.00
                           3.25 MBytes 27.3 Mbits/sec
                      sec
                           2.12 MBytes
  41
        6.00-7.00
                                          17.8 Mbits/sec
                      sec
                           1.25 MBytes
        7.00-8.00
                                          10.5 Mbits/sec
  41
                      sec
                                          18.9 Mbits/sec
19.9 Mbits/sec
   4]
        8.00-9.00
                           2.25 MBytes
                      sec
  4]
        9.00-10.00
                           2.38 MBytes
                     sec
 ID1 Interval
                            Iransfer
                                           Bandwidth
        0.00-10.00 sec
                           29.1 MBytes
                                          24.4 Mbits/sec
                                                                               sender
        0.00-10.00
                           28.2 MBytes
                                          23.6 Mbits/sec
                                                                               receiver
iperf Done
```

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 1,000 Mbit/s:

The VPN gateway bandwidth is shared by all of its VPN connections. If the bandwidth size is large, multiple ECSs are required to test the VPN gateway bandwidth because the forwarding performance of an ECS is limited. This scenario has high requirements on ECS specifications. The ECSs must have NICs that support the bandwidth of 2 Gbit/s or higher.

Conclusions: Based on the preceding test results, bandwidths of VPN gateways meet requirements. To fully use your purchased bandwidth, you are advised to use servers running the same operating system and using NICs meeting certain requirements at the two ends of a VPN connection.

2.14.2 Which Direction of the Bandwidth Is Limited and What Is the Unit of the Bandwidth?

Your purchased VPN gateway bandwidth applies to the outbound direction of the cloud. To achieve a balance between bandwidths in the inbound and outbound directions, the bandwidth in the inbound direction is limited as follows:

- If the purchased bandwidth is 10 Mbit/s or less, the bandwidth in the inbound direction is limited to 10 Mbit/s.
- If the purchased bandwidth is more than 10 Mbit/s, the bandwidth in the inbound direction is the same as that of the purchased bandwidth.

The unit of bandwidth is Mbit/s and that of traffic is GB.

2.14.3 How Do I Change the VPN Bandwidth Size?

- 1. On the **VPN Gateways** page, locate the row that contains the target VPN gateway and choose **More** > **Modify Bandwidth** in the **Operation** column.
- 2. On the Modify Bandwidth page, select your required bandwidth size.
- Click Submit.

2.14.4 What Happens If the Bandwidth of a VPN Gateway Exceeds the Size I Specified?

The VPN gateway bandwidth is used in the outbound direction of a VPC. If the bandwidth exceeds the size specified, network congestion will occur, some subnets cannot be accessed, or even the VPN connection will be interrupted, because the VPN detection packets may not be received.

In this case, you are advised to increase the VPN gateway bandwidth.

The maximum bandwidth of a VPN connection is 300 Mbit/s.

2.14.5 Why Does the VPN Bandwidth Change Not Take Effect?

There is a latency for the VPN bandwidth change to take effect.

Test the bandwidth 5 minutes after you change the bandwidth.

Changing the VPN bandwidth will not interrupt workload running and networks.

2.14.6 Can a VPN Share Bandwidth with an EIP?

No.

Currently, a public IP address is automatically generated and its bandwidth is set when you create a VPN gateway. The VPN cannot share bandwidth with an EIP.

2.14.7 What Are the Differences Between the Bandwidth of a VPN Connection and that of a Direct Connect Connection?

Concepts

- The bandwidth of a Direct Connect connection is the bandwidth of the connection created by a user.
- The bandwidth of a VPN connection applies to the outbound direction of the cloud.

Bandwidth Size

- The default maximum bandwidth of a Direct Connect connection is 1,000 Mbit/s. When you create a connection on the management console and set Port Type to 10GE single-mode optical port, the maximum bandwidth is 10 Gbit/s.
- The maximum bandwidth of a VPN connection is 300 Mbit/s.

Network Quality

- A Direct Connect user has a dedicated connection with high network quality.
- VPN connections shared the bandwidth of their VPN gateway. The total bandwidth of VPN connections cannot exceed the bandwidth of their gateway. The network quality will be affected by the Internet quality.

2.14.8 How Do I Determine My VPN Bandwidth Size?

Consider the following when you determine the bandwidth:

- Amount of data transmitted over a VPN tunnel in a period of time (Reserve enough bandwidth to prevent link congestion.)
- The egress bandwidth at the end of the VPN connection on the cloud must be less than that at the end of the VPN connection in the on-premises data center.

2.15 Quotas

2.15.1 What Is the VPN Quota?

What Is a Quota?

Quotas can limit the number or amount of resources available to users, such as the maximum number of ECSs or EVS disks that can be created.

If the existing resource quota cannot meet your service requirements, you can apply for a higher quota.

How Do I View My Quota?

1. Log in to the management console.

- 2. Click in the upper left corner and select the desired region and project.
- Choose Resources > My Quotas in the upper right corner of the page.
 The Service Quota page is displayed.
- 4. View the used and total quota of each type of resources on the displayed page.

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

How Do I Apply for a Higher Quota?

- 1. Log in to the management console.
- Choose Resources > My Quotas in the upper right corner of the page.
 The Service Quota page is displayed.
- 3. Click **Increase Quota** in the upper right corner of the page.
- On the Create Service Ticket page, configure parameters as required.
 In the Problem Description area, enter the required quota and the reason for the quota adjustment.
- 5. Select the agreement and click **Submit**.

2.15.2 How Many VPN Gateways and VPN Connections Can I Create By Default?

- VPN: By default, each user can create up to 50 VPN gateways and 100 remote gateways. Before purchasing VPN gateways, check your remaining quota. If the quota has been reached, submit a service ticket to request for quota increase.
- Classic VPN: By default, each user can create two VPN gateways and 12 VPN connections. Before purchasing VPN gateways, check your remaining quota. If the quota has been reached, submit a service ticket to request for quota increase.

2.15.3 How Do I Change My VPN Gateway and Connection Quotas?

- 1. Log in to the management console. In the upper right corner of the page, choose **Service Tickets** > **Create Service Ticket**.
- 2. On the **Create Service Ticket** page, click **Quotas** in the **Services** area.
- 3. Choose **Quota Application** under **Issue Categories**.
- 4. Click Create Now.

Enter required information and click **Submit**.

2.15.4 How Many IPsec VPNs Can I Have?

By default, a user can have a maximum of five IPsec VPNs. If the quota cannot fulfill your service requirements, request for quota increase.

2.16 Account Permissions

2.16.1 Are a Username and Password Required for Creating an IPsec VPN Connection?

No. IPsec VPN uses a pre-shared key (PSK) for authentication. The PSK is configured on a VPN gateway, and a connection will be established after VPN negotiation is complete. Therefore, no username or password is required for creating an IPsec VPN connection. Generally, SSL, PPTP, and L2TP VPNs use usernames and passwords for authentication.

■ NOTE

IPsec XAUTH provides extended authentication for IPsec VPN. It requires users to enter their usernames and passwords during VPN negotiation.

Currently, VPN does not support IPsec XAUTH.

2.16.2 What Should I Do If the System Displays a Message Indicating that I Do Not Have the Permissions to Create a VPN?

Check whether your account is an IAM user account. If yes, perform operations on the IAM console as the Huawei Cloud account user to authorize you the VPC operation permissions. Ensure that your account has the VPC Administrator, Tenant Guest, and VPN Administrator permissions.

2.16.3 How Do I Determine that a VPN Cannot Be Created in My Account Due to Insufficient Permissions?

- The VPN gateways and connections created by a Huawei Cloud account are invisible to IAM user accounts.
- A message will be displayed indicating that the system is busy if you create a VPN gateway or connection using an IAM user account.

For details about the permissions required for creating a VPN connection, see 2.16.2 What Should I Do If the System Displays a Message Indicating that I Do Not Have the Permissions to Create a VPN?

3 FAQs - P2C VPN

3.1 How Do I Test the Bandwidth of a P2C VPN Gateway?

You are advised to use the iPerf3 tool to test the bandwidth of P2C VPN gateways. Using commands such as FTP and SCP commands for file transfer is not recommended, because the file transfer rate is affected by the disk read/write speed and the bandwidth test result is not accurate.

Prerequisites

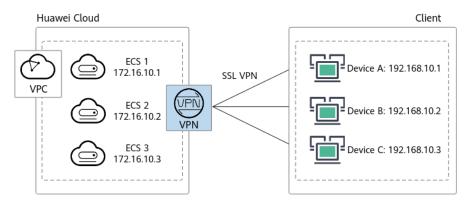
- The VPN gateway and server have been configured, and clients can connect to the VPN gateway.
 - In this example, the specification of the VPN gateway is Professional 1 (maximum forwarding bandwidth: 300 Mbit/s).
- Three ECS instances have been deployed in the VPC where the VPN gateway is located to simulate resource nodes on the cloud.
 - In this example, the flavor of each ECS is c6.large.2 (2 vCPUs, 4 GB memory, CentOS 8.0 64-bit system image).
- Three on-premises devices have been prepared to simulate clients.

 In this example, device A and device B are Linux servers (4U8G, running the ubuntu-20.04.6-live-server-amd64 operating system), and device C is a PC (i7 processor, running the Windows 10 operating system).
- The bandwidth capabilities of on-premises devices, ECS interfaces, and networks meet the following requirements: The upstream and downstream bandwidths are greater than or equal to 100 Mbit/s.
- The quality of the network between on-premises devices and the VPN gateway is good.

Networking Scenario

This section describes how to use the iPerf3 tool to test the VPN gateway bandwidth. Figure 3-1 shows the networking diagram.

Figure 3-1 Networking diagram



Installing iPerf3

The following describes how to install iPerf3 on the on-premises devices used in this test.

Installing iPerf3 on Linux

- 1. Open the CLI.
- 2. Run the following command to install iPerf3:

yum install -y iperf3

 Run the following command to check whether the installation is successful: iperf3 -v

If the iPerf version is displayed, the installation is successful.

Installing iPerf3 on Windows

Download the iPerf3 software package based on the operating system version from the iPerf3 official website.

Using iPerf3 to Test the Bandwidth of a VPN Gateway

Overview of iPerf3

Table 3-1 describes the key parameters of iPerf3.

Table 3-1 Key parameters of iPerf3

Param eter	Description
-S	Specifies that iPerf3 runs in server mode.
-С	Specifies that iPerf3 runs in client mode.
-р	Specifies the listening port of the server, that is, the server port to which a client needs to connect. (The settings on the server and client must be the same.)
-i	Specifies the interval for sending data, in seconds.

Param eter	Description
-l	Specifies the length of the buffer to read or write. You are advised to set this parameter to 1300 to simulate service data whose payload length is 1300 bytes.
-P	Specifies the number of threads. If this parameter is not specified, a single thread is used by default.

On-premises Devices Functioning as Servers

- 1. Run the following commands on on-premises devices to start the iPerf3 process in server mode, with different listening ports specified. The following is an example:
 - Device A (Linux)

```
iperf3 -s -p 20001
```

Device B (Linux)

iperf3 -s -p 20002

Device C (Windows)

```
iperf3.exe -s -p 20003
```

2. Run the following command on the three ECSs to start the iPerf3 process in client mode, with the server listening port of the on-premises device specified.

```
iperf3 -c server-ip -p server-port -l 1300 -P 10
```

```
The following is an example:
```

```
iperf3 -c 192.168.10.1 -p 20001 -l 1300 -P 10 iperf3 -c 192.168.10.2 -p 20002 -l 1300 -P 10 iperf3 -c 192.168.10.3 -p 20003 -l 1300 -P 10
```

On-premises Devices Functioning as Clients

1. Run the following command on the three ECSs to start the iPerf3 process in server mode, with different listening ports specified.

```
iperf3 -s -p server-port
```

The following is an example:

```
iperf3 -s -p 20001
iperf3 -s -p 20002
iperf3 -s -p 20003
```

- 2. Run the following commands on on-premises devices to start the iPerf3 process in client mode, with different server listening ports specified. The following is an example:
 - Device A

```
iperf3 -c 172.16.10.1 -p 20001 -l 1300 -P 10
```

Device B

```
iperf3 -c 172.16.10.2 -p 20002 -l 1300 -P 10
```

- Device C

iperf3.exe -c 172.16.10.3 -p 20003 -l 1300 -P 10

Test Result

After the iPerf3 process is executed, the following information is displayed:

```
Connecting to host 172.16.10.1, port 20001
[ 4] local 192.168.10.1 port 20001 connected to 172.16.10.1 port 20001
[ ID] Interval
                     Transfer Bandwidth
[ 4] 0.00-1.00 sec 8.62 MBytes 72.1 Mbits/sec
[ 4] 1.00-2.01 sec 9.88 MBytes 82.2 Mbits/sec [ 4] 2.01-3.01 sec 9.88 MBytes 82.9 Mbits/sec
[ 4] 3.01-4.00 sec 9.50 MBytes 80.4 Mbits/sec
[ 4] 4.00-5.01 sec 9.88 MBytes 82.1 Mbits/sec [ 4] 5.01-6.01 sec 9.62 MBytes 81.2 Mbits/sec
[ 4] 6.01-7.00 sec 9.12 MBytes 77.0 Mbits/sec
[ 4] 7.00-8.01 sec 10.0 MBytes 83.2 Mbits/sec [ 4] 8.01-9.01 sec 9.50 MBytes 79.9 Mbits/sec
[ 4] 9.01-10.01 sec 8.62 MBytes 72.4 Mbits/sec
[ ID] Interval
                     Transfer Bandwidth
[ 4] 0.00-10.01 sec 94.6 MBytes 79.3 Mbits/sec
                                                                       sender
[ 4] 0.00-10.01 sec 94.6 MBytes 79.3 Mbits/sec
```

According to the preceding iPerf3 test result, the transmission rate from 192.168.10.1 to 172.16.10.1 is about 79.3 Mbit/s. The test lasted for 10 seconds, during which 94.6 MB data is sent.

3.2 Does a P2C VPN Gateway Support Domain Name Access?

A P2C VPN gateway supports domain name access. This means users can use domain names to access cloud services.

3.3 Failed to Upload Certificates to CCM

3.3.1 Error Message "The certificate chain length must be greater than 1." Is Displayed During Certificate Upload

Symptom

When a certificate is uploaded to CCM, the error message "The certificate chain length must be greater than 1." is displayed.

Possible Causes

The upper-level CA certificate is not uploaded together with the current certificate.

Constraints

Before uploading certificates to CCM, ensure that the server certificate, CA certificate, and server private key have been generated. For details about how to generate certificates, see **Using Easy-RSA to Issue Certificates (Server and Client Sharing a CA Certificate)**.

Procedure

- **Step 1** Log in to the management console.
- **Step 2** Click in the upper left corner and select the desired region and project.
- Step 3 Click in the upper left corner, and choose Networking > Virtual Private
 Network
- **Step 4** In the navigation pane on the left, choose **Virtual Private Network > Enterprise VPN Gateways**.
- **Step 5** Click the **P2C VPN Gateways** tab. In the P2C VPN gateway list, locate the target P2C VPN gateway, and click **View Server** in the **Operation** column.
- **Step 6** On the **Server** tab page, set **Server Certificate** to **Existing certificate**, and click **Upload** in the drop-down list box. The **Cloud Certificate Manager** page is displayed.
- **Step 7** On the **SSL Certificate Manager** page, click the **Hosted Certificates** tab and then click **Upload Certificate**.
- **Step 8** Use Notepad or Notepad++ to open the server certificate.
- **Step 9** Copy the certificate content to the **Certificate File** text box.

The certificate file format is as follows:

```
-----BEGIN CERTIFICATE-----
Server certificate content
-----END CERTIFICATE-----
```

- **Step 10** Use Notepad or Notepad++ to open the CA certificate used to issue the server certificate.
- **Step 11** Copy the certificate content to the **Certificate File** text box.

The certificate file format is as follows:

```
----BEGIN CERTIFICATE----
Server certificate content
----END CERTIFICATE----
----BEGIN CERTIFICATE----
CA certificate content
----END CERTIFICATE----
```

- **Step 12** Use Notepad or Notepad++ to open the private key of the server certificate.
- **Step 13** Copy the private key content to the **Private Key** text box.

The private key format is as follows:

```
-----BEGIN PRIVATE KEY-----
Server private key content
-----END PRIVATE KEY-----
```

Step 14 Click **Submit**. The certificates are uploaded.

----End

3.3.2 Error Message "Incorrect format of the domain bound to the certificate to be uploaded." Is Displayed During Certificate Upload

Symptom

When a certificate is uploaded to CCM, the error message "Incorrect format of the domain bound to the certificate to be uploaded." is displayed.

Possible Causes

The domain name format of the certificate is incorrect. The correct format is xxx.com or xxx.cn.

Procedure

- Step 1 Re-generate certificates with the correct domain name format. For details, see Using Easy-RSA to Issue Certificates (Server and Client Sharing a CA Certificate).
- **Step 2** Log in to the management console.
- **Step 3** Go to the CCM console and upload the certificates.

For details, see Using the CCM to Manage a Server Certificate.

----End

3.4 When a Client Fails to Connect to a VPN Gateway, No Error Information Is Displayed and the Client Is Always in Connecting State.

Symptom

When a client fails to connect to a vpn gateway, no error information is displayed and the client is always in connecting state.

Possible Causes

- The client device cannot access the Internet.
- The client version does not meet the requirements.

Procedure

- 1. Verify that the client can ping the EIPs of the VPN gateway.
- 2. If the client runs on the Windows operating system, disable the firewall function and try again.
- 3. Use OpenVPN 2.5 or later.

If the problem persists, **submit a service ticket** to contact Huawei technical support.