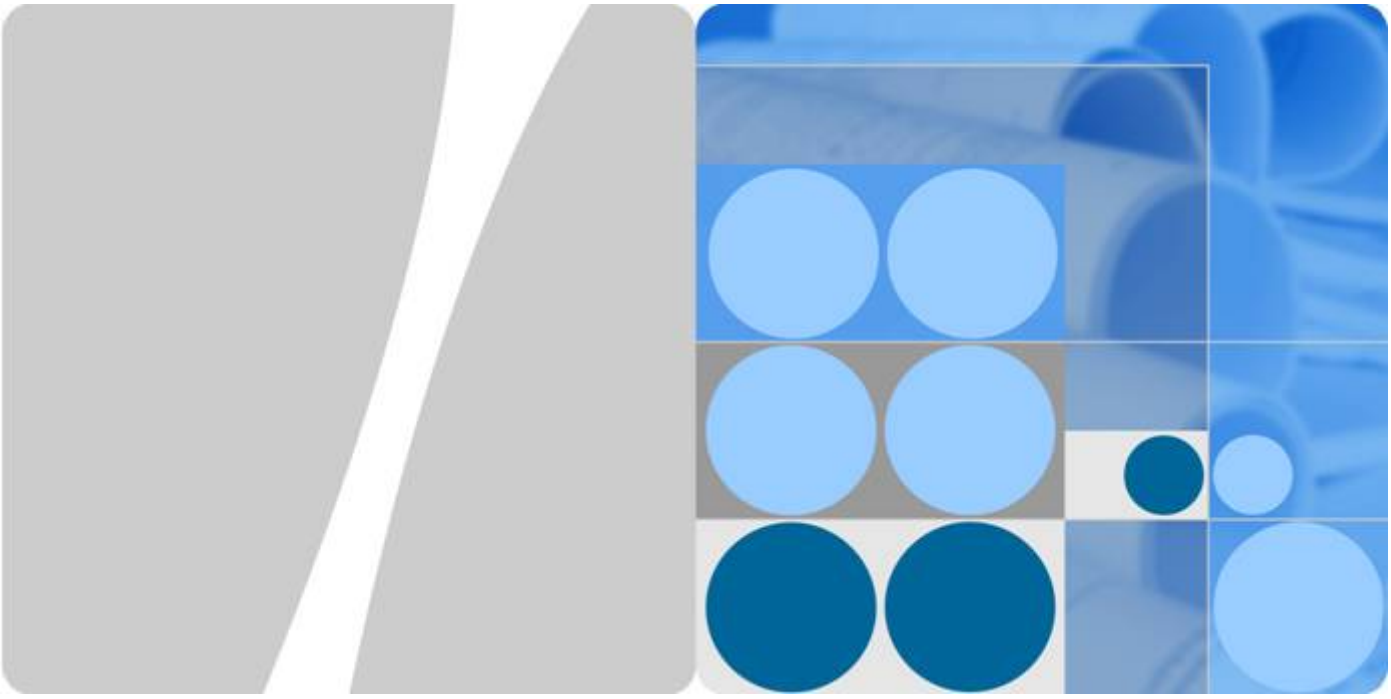

Virtual Private Network

User Guide

Date **2025-12-18**

1 Overview

1.1 What Is VPN?

Overview

Virtual Private Network (VPN) establishes secure, reliable, and cost-effective encrypted connections between your on-premises network or data center and a virtual network on the cloud.

NOTE

VPN supports only non-inter-border connections.

Site-to-Cloud VPN (S2C VPN) uses the Internet Protocol Security (IPsec) protocol and has two editions: Classic VPN and Enterprise Edition VPN.

S2C VPN involves three key components: VPN gateway, customer gateway, and VPN connection.

- A VPN gateway provides an Internet egress for a Virtual Private Cloud (VPC) to connect to a customer gateway in your on-premises data center.
- A VPN connection connects a VPN gateway to a customer gateway through encrypted tunnels, enabling communication between a VPC and your on-premises data center. This helps quickly establish a secure hybrid cloud environment.

Figure 1-1 shows the VPN networking.

Figure 1-1 VPN networking

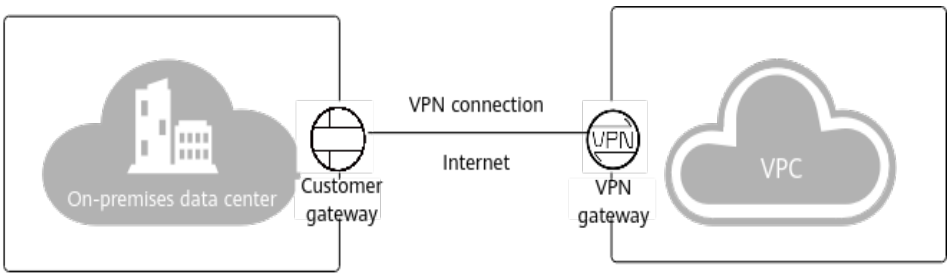

Components

- **VPN gateway**: a virtual gateway of VPN on the cloud. It establishes secure private connections with a customer gateway in your on-premises network or data center.
- **Customer gateway**: a resource that provides information to the cloud about your customer gateway device. It can be a physical device or software application in your on-premises data center.
- **VPN connection**: a secure channel between a VPN gateway and a customer gateway. VPN connections use the Internet Key Exchange (IKE) and IPsec protocols to encrypt the transmitted data.

1.2 Product Advantages

Enterprise Edition VPN has the following advantages:

- **High security**
  - Data is encrypted using IKE/IPsec or SSL, ensuring high data security.
  - A VPN gateway is exclusive to a tenant, isolating tenants from each other.
- **High availability**
  - A VPN gateway provides two IP addresses to establish dual independent VPN connections with a customer gateway. If one VPN connection fails, traffic can be quickly switched to the other VPN connection.
  - Active-active gateways are deployed in different availability zones (AZs) to ensure AZ-level high availability.
  - Active/Standby mode: In normal cases, a VPN gateway communicates with a customer gateway through the active connection. If the active connection fails, traffic is automatically switched to the standby VPN connection. After the fault is rectified, traffic is switched back to the active VPN connection.
  - High availability (HA) mode: S2C VPN supports active/standby and active-active modes.
- **Cost-effectiveness**
  - IPsec connections over the Internet provide a cost-effective alternative to Direct Connect.
  - A VPN gateway can be bound to elastic IP addresses (EIPs) that share bandwidth, reducing bandwidth costs.
  - The bandwidth can be adjusted when an EIP instance is created.
- **Easy to use**
  - A VPN gateway supports multiple connection modes, including policy-based, static routing, and BGP routing, to meet different access requirements of customer gateways.
  - A VPN gateway on the cloud can function as a VPN hub, enabling on-premises branch sites to access each other.
  - A VPN connection can be created in a few simple steps on the VPN device in an on-premises data center and on the VPN console, and is ready to use immediately after being created.

- VPN can be used together with the enterprise router service, allowing enterprises to build more flexible cloud-based networks.
- Private VPN gateways are supported to encrypt traffic transmitted over Direct Connect connections, improving data transmission security.
- In S2C VPN, you can flexibly enable or disable branch interconnection to implement interconnection or isolation between customer gateways, respectively.

1.3 Application Scenarios

Hybrid Cloud Deployment

You can use a VPN to connect your on-premises data center to a VPC and use the elastic and fast scaling capabilities of the cloud to expand application computing capabilities. **Figure 1-2** shows the hybrid cloud deployment.

Figure 1-2 Hybrid cloud deployment

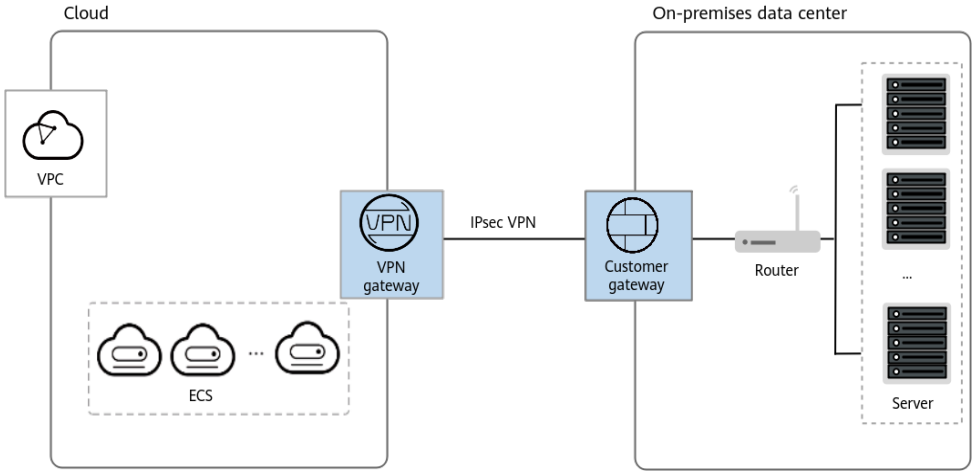

Inter-Region Interconnection Between VPCs

With VPNs, you can connect VPCs in different regions to enable connectivity between user services in these regions, as shown in **Figure 1-3**.

Figure 1-3 Inter-region interconnection between VPCs

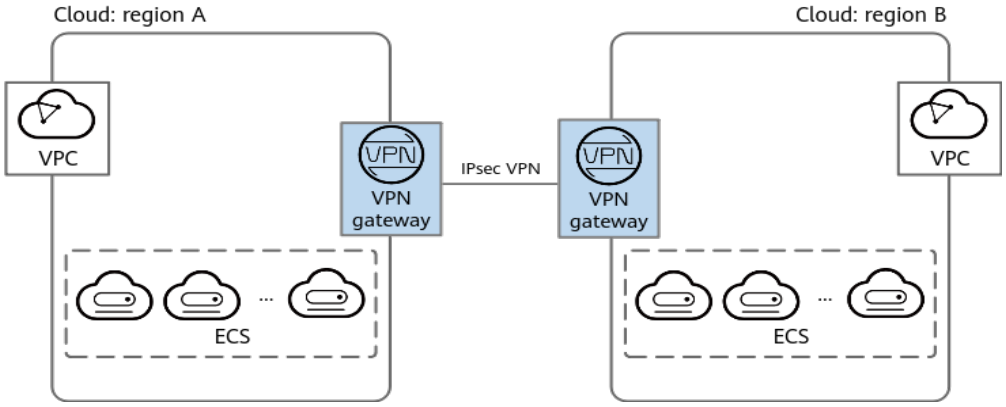

Enterprise Branch Interconnection

A VPN gateway functions as a VPN hub to connect enterprise branches, as shown in **Figure 1-4**. This eliminates the need to configure VPN connections between every two branches.

Figure 1-4 Enterprise branch interconnection

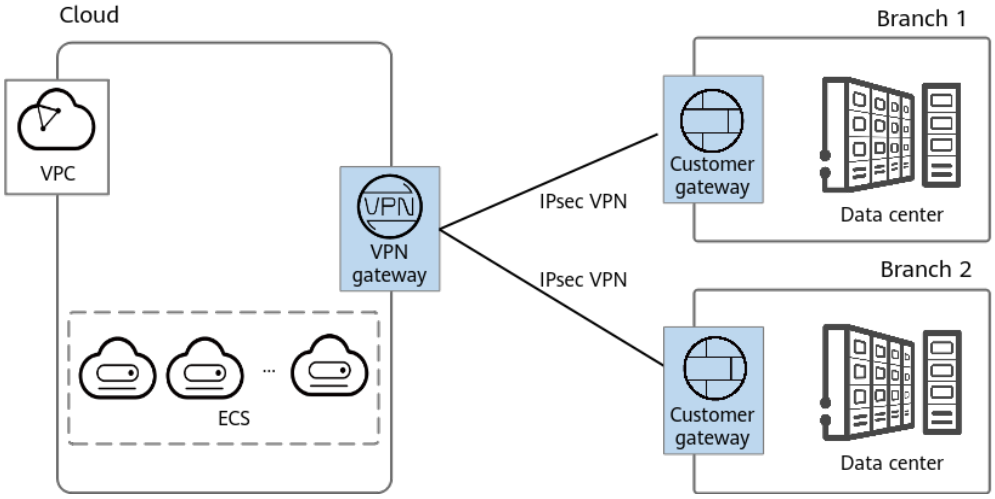

1.4 Product Specifications

1.4.1 S2C VPN

NOTE

- The maximum forwarding bandwidth provided in this section is measured under the following conditions. The actual forwarding bandwidth may vary according to various factors such as the conditions of the customer network and the Internet.

  In the IKE policy, the IKE version is IKEv2, the authentication algorithm is MD5, the encryption algorithm is AES128-GCM, the DH algorithm is group 15, and the local and customer IDs are IP addresses.

Table 1-1 S2C VPN specifications

Item	Professional 1	Professional 2
Exclusive gateway resources	Supported	Supported
Dual connections	Supported	Supported
Active-active gateways	Supported	Supported
Active/Standby gateways	Supported	Supported
Policy-based mode	Supported	Supported
Routing mode: static routing	Supported	Supported

Item	Professional 1	Professional 2
Routing mode: BGP routing	Supported	Supported
Policy template mode	Not supported	Not supported
Maximum forwarding bandwidth	300 Mbit/s	1 Gbit/s
Maximum number of VPN connection groups	100	100
Interconnection with an enterprise router	Supported	Supported
Interconnection with a VPC	Supported	Supported
Access through a public IP address	Supported	Supported
Private network	Supported	Supported
Supported regions	Subject to the regions available on the management console	Subject to the regions available on the management console

1.5 Quotas and Constraints

1.5.1 S2C VPN

VPN Gateway

Table 1-2 Constraints on VPN gateways

VPN Type	Resource	Default Quota
Enterprise Edition VPN	VPN gateways per tenant in each region	50 • If you have only one VPC, you can create a maximum of 50 VPN gateways for the VPC. • If you have multiple VPCs, you can create a maximum of 50 VPN gateways for all these VPCs.
	VPN connection groups per VPN gateway	100
	Local subnets per VPN gateway	50

VPN Type	Resource	Default Quota
	Number of BGP routes that can be accepted by VPN gateways of different specifications	100 • VPN gateway of the Professional 1 specification: 200 • VPN gateway of the Professional 2 specification: 300
	Maximum number of routes supported by a VPN gateway	10000
	Maximum number of ACL rules supported by VPN gateways of different specifications	300
Classic VPN	VPN gateways per tenant in each region	2 Only one VPN gateway can be created for a VPC.

- By default, the maximum length of TCP packets supported by a VPN gateway is 1300 bytes.

Customer Gateway

Table 1-3 Constraints on customer gateways

VPN Type	Resource	Default Quota
Enterprise Edition VPN	Customer gateways per tenant in each region	100

- Enable NAT traversal on the customer gateway based on the networking.
  - If the customer gateway is connected to the Internet through a NAT device, enable NAT traversal on the customer gateway.
  - If the customer gateway is directly connected to the Internet, you do not need to enabled NAT traversal on the customer gateway.
- Dead Peer Detection (DPD) must be enabled on a customer gateway.
- A customer gateway must support IPsec tunnel interfaces and be configured with a corresponding security policy.
- When Network Quality Analysis (NQA) is enabled for a connection in static routing mode, the IPsec tunnel interface of a customer gateway must have an IP address and be able to respond to ICMP requests.
- It is recommended that the maximum segment size (MSS) of TCP packets be set to a value less than 1399 on a customer gateway, so as to prevent fragmentation caused by addition of an IPsec header.

VPN Connection

Table 1-4 Constraints on VPN connections

VPN Type	Resource	Default Quota	How to Increase Quota
Enterprise Edition VPN	Policy rules per VPN connection	5	The quotas cannot be increased.
	Customer subnets per VPN connection	50	
Classic VPN	VPN connections per tenant in each region	12	This quota cannot be increased.

- In multi-subnet scenarios, you are advised to use VPN connections in routing mode. For a VPN connection in policy-based or policy template mode, a VPN gateway creates a communications tunnel for each pair of the local and customer subnets by default. If there are multiple local or customer subnets for a VPN connection in policy-based or policy template mode, multiple communications tunnels are created.

  For a VPN gateway of the Professional 1 or Professional 2 specification, a single IP address can be used to establish a maximum of 300 communications tunnels with customer gateways..

  – In routing mode, each VPN connection occupies only one communications tunnel of the corresponding VPN gateway IP address.
  – In policy-based or policy template mode, each VPN connection occupies *M* x *N* communications tunnels of the corresponding VPN gateway IP address. *M* indicates the number of local subnets, and *N* indicates the number of customer subnets.

    For a VPN gateway of the Professional 1 or Professional 2 specification, a maximum of 300 source and destination subnet pairs are supported for each IP address used to establish connections with customer gateways.

  If the number of communications tunnels occupied by all VPN connections in different modes established by a single gateway IP address exceeds 100, excess VPN connections will fail to be created.

  If the number of communications tunnels occupied by VPN connections in different modes established by a single IP address of a VPN gateway of the Professional 1 or Professional 2 specification exceeds 300, excess VPN connections will fail to be created.
- When creating a VPN connection in policy-based mode and adding multiple policy rules, ensure that the source and destination CIDR blocks in different policy rules do not overlap. Otherwise, data flows may be incorrectly matched or IPsec tunnels may flap.

1.6 Reference Standards and Protocols

The following standards and protocols are associated with VPN:

- RFC 2403: The Use of HMAC-MD5-96 within ESP and AH
- RFC 2404: The Use of HMAC-SHA-1-96 within ESP and AH
- RFC 2409: The Internet Key Exchange (IKE)
- RFC 2451: The ESP CBC-Mode Cipher Algorithms
- RFC 3526: More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)
- RFC 3566: The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec
- RFC 3602: The AES-CBC Cipher Algorithm and Its Use with IPsec
- RFC 3664: The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)
- RFC 4106: The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)
- RFC 4109: Algorithms for Internet Key Exchange version 1 (IKEv1)
- RFC 4434: The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)
- RFC 4868: Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPsec
- RFC 4301: Security Architecture for the Internet Protocol
- RFC 4302: IP Authentication Header
- RFC 4303: IP Encapsulating Security Payload (ESP)
- RFC 4305: Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)
- RFC 4306: Internet Key Exchange (IKEv2)Protocol
- RFC 4307: Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)
- RFC 4308: Cryptographic Suites for IPsec
- RFC 5282: Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange version 2 (IKEv2) Protocol
- RFC 6989: Additional Diffie-Hellman Tests for the Internet Key Exchange Protocol Version 2 (IKEv2)
- RFC 7296: Internet Key Exchange Protocol Version 2 (IKEv2)
- RFC 7321: Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)
- RFC 8247: Algorithm Implementation Requirements and Usage Guidance for the Internet Key Exchange Protocol Version 2 (IKEv2)
- RFC 3947: Negotiation of NAT-Traversal in the IKE
- RFC 3948: UDP Encapsulation of IPsec ESP Packets
- RFC 3706: A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers
- RFC 4271: A Border Gateway Protocol 4 (BGP-4)

1.7 Differences between S2C Enterprise Edition VPN and Classic VPN

Table 1-5 Differences between Enterprise Edition VPN and Classic VPN

Category	Item	Enterprise Edition VPN	Classic VPN
Tenant isolation	Tenant-exclusive gateway	Supported	Not supported
Features	Policy-based mode	Supported	Supported
	Routing mode	Static routing and BGP routing	Not supported
	VPN hub	Supported	Not supported
	Enterprise router	Supported	Not supported
	Network type	Public network and private network	Public network
Capacity	Number of subnets	• Route-based mode: 50 • Policy-based mode: 5	Policy-based mode: 5
	For more information, see **Table 1-1**.		
Reliability	Gateway protection mode	Active/Standby or active-active	-
	Cross-AZ gateway deployment	Supported	Not supported
	Active-active VPN connections	Supported	Not supported
	Backup with Direct Connect	Supported	Not supported

1.8 Security

1.8.1 Identity Authentication and Access Control

An S2C VPN connection supports authentication of a customer gateway using a pre-shared key (PSK).

The identity authentication succeeds and the VPN connection can be set up only when the PSK configured on the customer gateway is the same as that configured for the VPN connection.

Reference link:

PSK

1.8.2 Data Protection Technologies

- S2C VPN is a tunneling technology that provides IP-layer security using the IKE/IPsec protocol suite. It ensures confidentiality and integrity of IP data packets and prevents them from being intercepted, disclosed, or tampered with on insecure networks (such as the Internet).
- When creating an S2C VPN connection, you can configure data encryption and authentication algorithms in a policy.

  Table 1-6 lists the algorithms recommended for S2C VPN in descending order of security.

  Table 1-6 Parameters for configuring an S2C VPN policy

Parameter		Description
IKE Policy	Version	• v2 • v1 (IKEv1 has low security. If the device supports IKEv2, IKEv2 is recommended.) The default value is **v2**.
	Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: • SHA2-512 • SHA2-384 • SHA2-256 • MD5(Insecure. Not recommended.) • SHA1(Insecure. Not recommended.) By default, the SHA2-256 algorithm is used.

Parameter		Description
	Encryption Algorithm	The following encryption algorithms are supported: ● AES-256-GCM-16 (supported only by Enterprise Edition VPN) ● AES-128-GCM-16 (supported only by Enterprise Edition VPN) ● AES-256(Insecure. Not recommended.) ● AES-192(Insecure. Not recommended.) ● AES-128(Insecure. Not recommended.) ● 3DES(Insecure. Not recommended.) The default value is **AES-128**.
	DH Algorithm	The following algorithms are supported: ● Group 21 ● Group 20 ● Group 19 ● Group 16 ● Group 15 ● Group 14(Insecure. Not recommended.) ● Group 5(Insecure. Not recommended.) ● Group 2(Insecure. Not recommended.) ● Group 1(Insecure. Not recommended.) By default, Group 15 is used.
IPsec Policy	Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: ● SHA2-512 ● SHA2-384 ● SHA2-256 ● MD5(Insecure. Not recommended.) ● SHA1(Insecure. Not recommended.) By default, the SHA2-256 algorithm is used.

Parameter		Description
	Encryption Algorithm	The following encryption algorithms are supported: ● AES-256-GCM-16 ● AES-128-GCM-16 ● AES-256(Insecure. Not recommended.) ● AES-192(Insecure. Not recommended.) ● AES-128(Insecure. Not recommended.) ● 3DES(Insecure. Not recommended.) The default value is **AES-128**.

PFS

Perfect Forward Secrecy (PFS) ensures that the compromise of the keys of an IPsec tunnel does not affect the security of other tunnels by leveraging that the keys of these tunnels are irrelevant to each other. By default, the PFS function is enabled for S2C VPN.

Each IPsec VPN connection consists of at least one IPsec tunnel, each of which uses an independent set of keys to protect user traffic.

S2C VPN supports the following algorithms:

- DH group 1 (This algorithm is insecure. Exercise caution when using it.)
- DH group 2 (This algorithm is insecure. Exercise caution when using it.)
- DH group 5 (This algorithm is insecure. Exercise caution when using it.)
- DH group 14
- DH group 15
- DH group 16
- DH group 19
- DH group 20
- DH group 21

Anti-replay

Anti-replay uses sequence numbers to protect IPsec encrypted packets against replay attacks, which are initiated by repeatedly sending intercepted data packets. By default, the anti-replay function is enabled for the VPN service.

Resource Isolation

A VPN gateway is exclusive to a tenant. As such, tenants are isolated from each, ensuring tenant data security.

1.8.3 Audit and Logs

VPN records the create, delete, and modify operations performed on all resources initiated by your account, and sends the records to Cloud Trace Service (CTS) in log files for query, audit, and source tracing.

Figure 1-5 Audit and logs

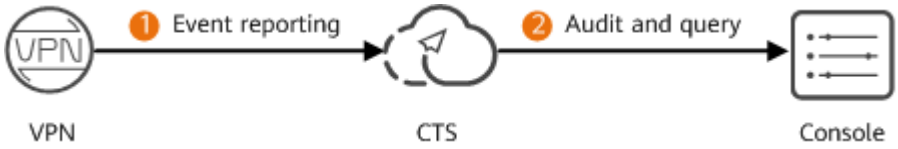

Reference link:

For details about how to view audit logs, see the *Cloud Trace Service User Guide*.

1.8.4 Service Resilience

VPN provides the dual-AZ disaster recovery function. You can create a VPN gateway in two AZs in the same region, and create a VPN connection between the customer gateway and each AZ.

Dual-AZ disaster recovery is supported only by Enterprise Edition VPN, but not Classic VPN.

Figure 1-6 Scenario where services are running properly

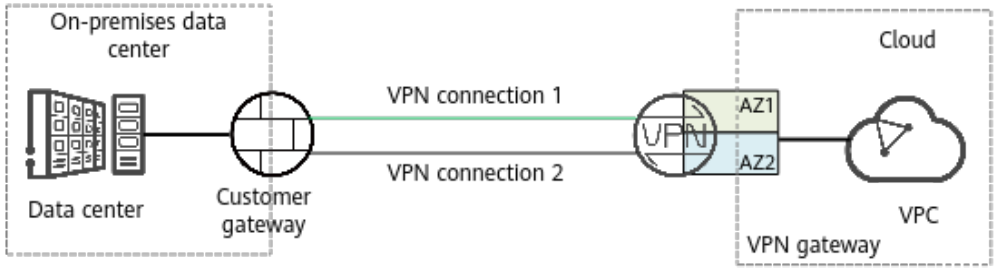

If the VPN gateway or VPN connection in an AZ is faulty, traffic is automatically switched to the other VPN connection, ensuring normal service running.

Figure 1-7 Failover scenario

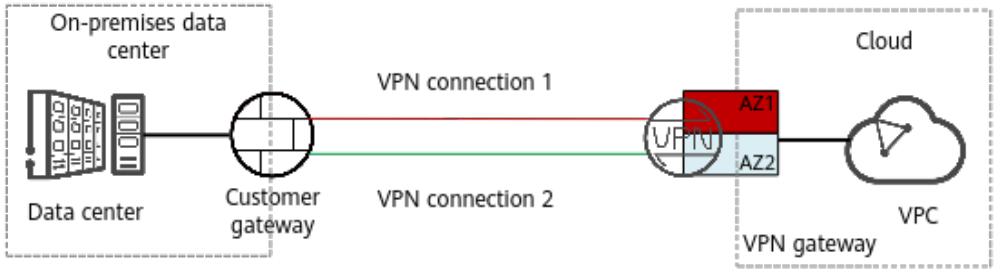

1.9 Permissions Management

1.9.1 IAM-based Permissions Management

If you need to assign different permissions to personnel in your enterprise to access your VPN resources created on the cloud service platform, Identity and Access Management (IAM) is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you securely manage access to your resources.

With IAM, you can use your account to create IAM users, and assign permissions to the users to control their access to specific cloud resources. For example, some software developers in your enterprise need to use VPN resources but should not be allowed to delete them or perform any high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the permissions required for using VPN resources.

If your account does not need individual IAM users for permissions management, skip this section, which has no impact on using functions of VPN.

IAM is a free service. You only pay for the resources in your account.

For more information about IAM, see **IAM Service Overview**.

VPN Permissions

New IAM users do not have any permissions assigned by default. You need to first add them to one or more groups and attach policies or roles to these groups. The users then inherit permissions from the groups and can perform specified operations on cloud services based on the permissions they have been assigned.

VPN is a project-level service deployed for specific regions. When you set **Scope** to **Region-specific projects** and select the specified projects in the specified regions, the users only have permissions for VPN in the selected projects. If you set **Scope** to **All resources**, users have permissions for VPN in all region-specific projects. When accessing VPN, the users need to switch to the authorized region.

You can grant permissions by using roles or policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. There are only a limited number of roles for granting permissions to users. Some roles depend other roles to take effect. When you assign such roles to users, remember to assign the roles they depend on. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- Policies: a type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under certain conditions. This mechanism allows for more flexible policy-based authorization, meeting requirements for secure access control. For example, administrators can grant IAM users only permissions to manage VPN resources of a certain type.

Table 1-7 lists all system-defined permissions for VPN.

Table 1-7 System-defined permissions for VPN

System Role/ Policy Name	Description	Dependency
VPN Administrator (not recommended)	Administrator permissions for VPN. Users with these permissions can perform all operations on VPN. This role must be used together with the **Tenant Guest** and **VPC Administrator** roles in the same project. • **VPC Administrator**: project-level policy, which is selected in the same project as **VPN Administrator**. • **Tenant Guest**: project-level policy, which is selected in the same project as **VPN Administrator**.	-
VPN FullAccess (recommended)	Full permissions for VPN. **NOTE** All actions that are used to query list information do not support authorization based on enterprise projects. You need to configure actions in the IAM view separately.	The actions of global services and the region-level actions cannot be configured in the same policy. As such, the following global actions are added: • "tms:predefineTags:list" • "scm:cert:list" • "scm:cert:get" • "scm:cert:download"
VPN ReadOnlyAccess	Read-only permissions on VPN resources. Users who have these permissions can only view information about VPN resources. **NOTE** All actions that are used to query list information do not support authorization based on enterprise projects. You need to configure actions in the IAM view separately.	The actions of global services and the region-level actions cannot be configured in the same policy. As such, the following global actions are added: • "tms:predefineTags:list" • "scm:cert:list" • "scm:cert:get"

Table 1-8 lists the common operations supported by system-defined permissions for S2C VPN.

Table 1-8 Common operations supported by system-defined permissions for S2C VPN

Operation	VPN Administrator (Not Recommended)	VPN FullAccess (Recommended)	VPN ReadOnlyAccess
Creating a VPN gateway	√	• Enterprise Edition VPN: √ • Classic VPN: ×	×
Viewing a VPN gateway	√	√	√
Querying the VPN gateway list	√	√	√
Updating a VPN gateway	√	• Enterprise Edition VPN: √ • Classic VPN: ×	×
Deleting a VPN gateway	√	• Enterprise Edition VPN: √ • Classic VPN: ×	×
Creating a VPN connection	√	• Enterprise Edition VPN: √ • Classic VPN: √	×
Viewing a VPN connection	√	√	√
Querying the VPN connection list	√	√	√
Updating a VPN connection	√	• Enterprise Edition VPN: √ • Classic VPN: √	×
Deleting a VPN connection	√	• Enterprise Edition VPN: √ • Classic VPN: √	×
Creating a customer gateway	√	• Enterprise Edition VPN: √ • Classic VPN: N/A	×
Viewing a customer gateway	√	• Enterprise Edition VPN: √ • Classic VPN: N/A	√

Operation	VPN Administrator (Not Recommended)	VPN FullAccess (Recommended)	VPN ReadOnlyAccess
Querying the customer gateway list	√	• Enterprise Edition VPN: √ • Classic VPN: N/A	√
Updating a customer gateway	√	• Enterprise Edition VPN: √ • Classic VPN: N/A	×
Deleting a customer gateway	√	• Enterprise Edition VPN: √ • Classic VPN: N/A	×
Creating a VPN connection monitor	√	• Enterprise Edition VPN: √ • Classic VPN: ×	×
Querying a VPN connection monitor	√	• Enterprise Edition VPN: √ • Classic VPN: ×	√
Querying the VPN connection monitor list	√	• Enterprise Edition VPN: √ • Classic VPN: ×	√
Deleting a VPN connection monitor	√	• Enterprise Edition VPN: √ • Classic VPN: ×	×
Querying VPN connection logs	√	• Enterprise Edition VPN: √ • Classic VPN: ×	√

Helpful Links

- **IAM Service Overview**
- **Creating a User Group and User and Granting VPN Permissions**

1.10 Basic Concepts

1.10.1 IPsec VPN

Internet Protocol Security (IPsec) VPN uses a secure network protocol suite that authenticates and encrypts the packets of data to provide secure encrypted

communication between different networks. S2C VPN is an IPsec VPN technology on the cloud.

In the example shown in **Figure 1-8**, assume that you have created a VPC with two subnets (192.168.1.0/24 and 192.168.2.0/24) on the cloud, and the router in your on-premises data center also has two subnets (192.168.3.0/24 and 192.168.4.0/24). In this case, you can create a VPN to connect the VPC subnets and the data center subnets.

Figure 1-8 IPsec VPN

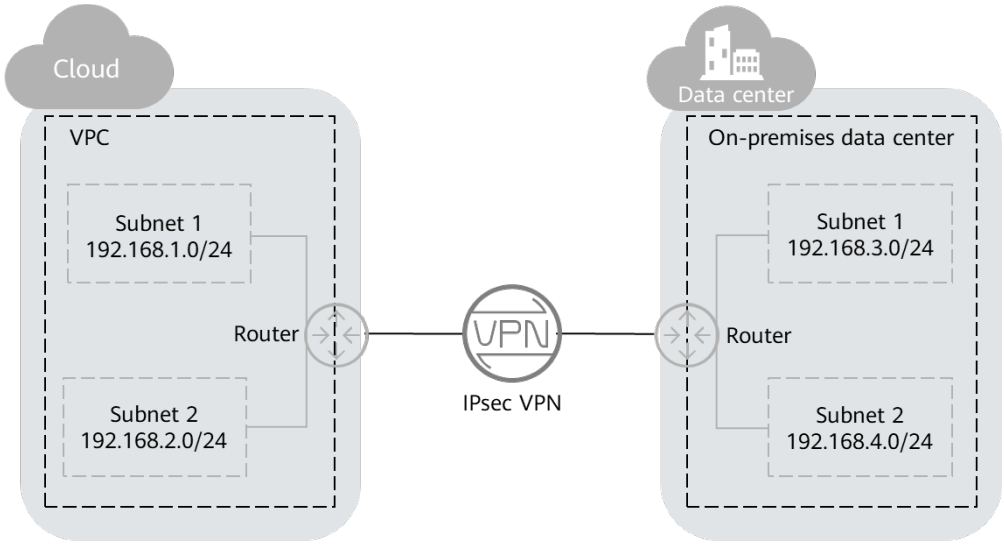

Site-to-site VPN is supported to enable communication between VPC subnets and on-premises data center subnets.

Reference link:

S2C VPN

1.10.2 VPN Gateway

A VPN gateway is a virtual gateway of VPN on the cloud. It establishes secure private connections with a customer gateway in your on-premises network or data center. A VPN gateway needs to work with a customer gateway in your on-premises data center.

Reference link:

- **Creating an S2C VPN Gateway**
- **Creating a Classic VPN Gateway**

1.10.3 VPN Connection

A VPN connection is a secure channel between a VPN gateway and a customer gateway. VPN connections use the IKE and IPsec protocols to encrypt the transmitted data, ensuring data security and reliability.

Reference link:

- **Creating Enterprise Edition VPN Connections**
- **Creating a Classic VPN Connection**
- **What Are a VPC, a VPN Gateway, and a VPN Connection?**

1.10.4 VPN Gateway Bandwidth

The bandwidth you purchased for a VPN gateway refers to outbound bandwidth, that is, bandwidth for traffic sent from a VPC on the cloud to a customer gateway in an on-premises data center.

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

If your VPN gateway is billed by traffic on a pay-per-use basis, the bandwidth size of the VPN gateway does not affect the total price. But it is recommended that you set the bandwidth size based on actual requirements to prevent a large amount of traffic caused by program errors or malicious access.

1.10.5 Local Subnet

Local subnets are VPC subnets that need to communicate with an on-premises network through VPN. When you buy a VPN gateway, you can set **Local Subnet** to either of the following options:

- **Select subnet**: Select subnets from the drop-down list. This is recommended if all subnets that require VPN communication are in the VPC.
- **Enter CIDR block**: Enter a subnet using CIDR notation (example: 192.168.0.0/16). If multiple subnets are specified, separate them by a comma (,). This is recommended if the CIDR blocks requiring VPN communication are not in the VPC to which the VPN gateway belongs. For example, CIDR blocks (such as 0.0.0.0/0) that are connected using a VPC peering are not in the VPC to which the VPN gateway belongs.

1.10.6 Customer Gateway

A customer gateway can be a physical device or software application in your on-premises data center. A customer gateway is a resource that provides information on the management console about your customer gateway device.

Reference link:

Creating a Customer Gateway

1.10.7 Customer Subnet

Customer subnets are subnets in an on-premises data center that access a VPC on the cloud through a VPN.

- You need to enter subnets using CIDR notation (example: 192.168.0.0/16), and with each entry separated by a comma.
- After configuring a customer subnet, you do not need to add a route for it. The VPN service will automatically deliver routes pointing to the customer subnet.

NOTE

A customer subnet cannot be set to a Class D or Class E IP address or an IP address starting with 127.

1.10.8 PSK

A pre-shared key (PSK) is a key configured for a VPN connection on the cloud. It is used for IKE negotiation between VPN devices at both ends of a VPN connection. Ensure that the PSK configurations at both ends of the VPN connection are the same. Otherwise, the IKE negotiation will fail.

Reference link:

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

2 Getting Started

2.1 Configuring Enterprise Edition VPN to Connect an On-premises Data Center to a VPC

2.1.1 Overview

Scenario

To meet business development requirements, enterprise A needs to implement communication between its on-premises data center and its VPC. In this case, enterprise A can use the VPN service to create connections between the on-premises data center and the VPC.

- If the on-premises data center has only one customer gateway and this gateway can be configured with only one IP address, it is recommended that the VPN gateway use the active-active mode. **Figure 2-1** shows the networking.

  In active-active mode, if connection 1 is faulty, traffic is automatically switched to connection 2 for transmission, without affecting enterprise services.

  Figure 2-1 Active-active mode

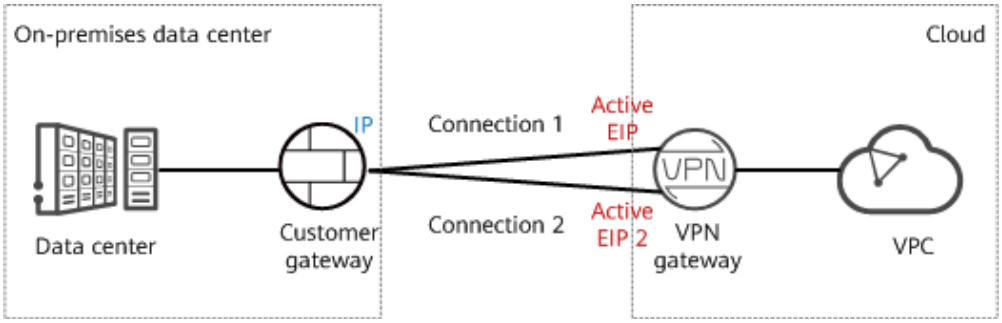

- If the on-premises data center has two customer gateways or has only one customer gateway that can be configured with two IP addresses, it is

recommended that the VPN gateway use the active/standby mode. **Figure 2-2** shows the networking.

In active/standby mode, connection 1 is the active link and connection 2 is the standby link. By default, traffic is transmitted only through the active link. If the active link fails, traffic is automatically switched to the standby link, without affecting enterprise services. After the active link recovers, traffic is switched back to the active link.

Figure 2-2 Active/Standby mode

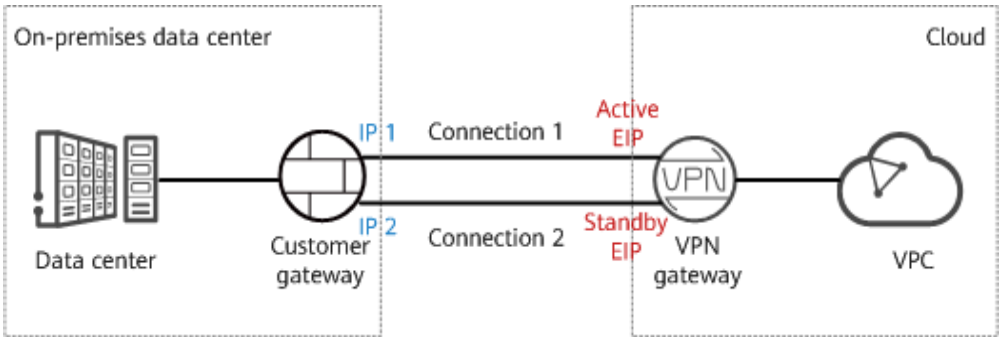

Limitations and Constraints

- The customer gateway device must support standard IKE and IPsec protocols.
- The interconnection subnets of the on-premises data center cannot overlap with those of the VPC and cannot contain reserved CIDR blocks such as 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console.
  - If the VPC uses Direct Cloud or Cloud Connect connections to communicate with other VPCs, the on-premises data center subnets cannot overlap with those of these VPCs.

Data Plan

Table 2-1 Data plan

Category	Item	Data
VPC	Subnet that needs to access the on-premises data center	192.168.0.0/16
VPN gateway	Interconnection subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses. 192.168.2.0/24
	HA mode	Active-active

Category	Item	Data
	EIP	EIPs are automatically generated when you buy them. By default, a VPN gateway uses two EIPs. In this example, the EIPs are as follows: • Active EIP: 11.xx.xx.11 • Active EIP 2: 11.xx.xx.12
VPN connection	Tunnel interface addresses under **Connection 1's Configuration**	The IP addresses are used to establish an IPsec tunnel between a VPN gateway and a customer gateway. At the two ends of the IPsec tunnel, the configured local and remote tunnel interface addresses must be reversed. • Local tunnel interface address: 169.254.70.1/30 • Customer tunnel interface address: 169.254.70.2/30
	Tunnel interface addresses under **Connection 2's Configuration**	• Local tunnel interface address: 169.254.71.1/30 • Customer tunnel interface address: 169.254.71.2/30
On-premises data center	Subnet that needs to access the VPC	172.16.0.0/16
Customer gateway	Gateway IP address	The gateway IP address is assigned by a carrier. In this example, the gateway IP address is: 22.xx.xx.22

Operation Process

Figure 2-3 shows the process of using the VPN service to enable communication between an on-premises data center and a VPC.

Figure 2-3 Operation process

Table 2-2 Operation process description

No.	Step	Description
1	**Step 1: Creating a VPN Gateway**	Bind two EIPs to the VPN gateway. If you have purchased EIPs, you can directly bind them to the VPN gateway.
2	**Step 2: Creating a Customer Gateway**	Configure the VPN device in the on-premises data center as the customer gateway.
3	**Step 3: Creating VPN Connections**	• Create two VPN connections between the VPN gateway (active EIP and active EIP 2) and the customer gateway. • The connection mode, PSK, IKE policy, and IPsec policy settings of connection 2 must be the same as those of connection 1.
4	**Step 4: Configuring the Customer Gateway Device**	• The local and remote tunnel interface addresses configured on the customer gateway device must be the same as the customer and local tunnel interface addresses of the VPN connections, respectively. • The connection mode, PSK, IKE policy, and IPsec policy settings on the customer gateway device must be same as those of the VPN connections.
5	**Step 5: Verifying Network Connectivity**	Log in to an ECS and run the **ping** command to verify the network connectivity.

2.1.2 Step 1: Creating a VPN Gateway

Prerequisites

- A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
- Security group rules have been configured for ECSs in the VPC, and allow the customer gateway in the on-premises data center to access VPC resources. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner, and choose > **Virtual Private Network**.

Step 3 In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.

Step 4 Click **Buy S2C VPN Gateway**.

Step 5 Set parameters as prompted, click **Buy Now**, and complete the payment.

Step 6 The following describes only key parameters.

Table 2-3 Key VPN gateway parameters

Parameter	Description	Example Value
Billing Mode	**Pay-per-use** is supported.	Pay-per-use
Region	Select the region nearest to you.	-
Name	Enter the name of a VPN gateway.	vpngw-001
Network Type	• **Public network**: A VPN gateway communicates with a customer gateway in an on-premises data center through the Internet. • **Private network**: A VPN gateway communicates with a customer gateway in an on-premises data center through a private network.	Public network
Associate With	• **VPC**: Through a VPC, the VPN gateway sends messages to the customer gateway or servers in the local subnet. • **Enterprise Router**: Through an enterprise router, the VPN gateway sends messages to the customer gateway or servers in the subnets of all VPCs connected to the enterprise router.	VPC
VPC	Select the VPC that needs to access the on-premises data center.	vpc-001(192.168.0.0/16)
Interconnection Subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	192.168.2.0/24
Local Subnet	Specify the VPC subnet that needs to access the on-premises data center. You can manually enter a CIDR block or select a subnet from the drop-down list box.	192.168.0.0/24
Specification	Select **Professional 1**.	Professional 1

Parameter	Description	Example Value
HA Mode	Select **Active-active**.	Active-active
Active EIP	You can buy a new EIP or use an existing EIP.	11.xx.xx.11
Active EIP 2		11.xx.xx.12

----End

Verification

Check the created VPN gateway on the **VPN Gateways** page. The initial state of the VPN gateway is **Creating**. When the VPN gateway state changes to **Normal**, the VPN gateway is successfully created.

2.1.3 Step 2: Creating a Customer Gateway

Procedure

Step 1 In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.

Step 2 On the **Customer Gateway** page, click **Create Customer Gateway**.

Step 3 Set parameters as prompted and click **Create Now**.

The following describes only key parameters.

Table 2-4 Customer gateway parameters

Parameter	Description	Example Value
Name	Name a customer gateway.	cgw-001
Identifier	Enter the IP address of the customer gateway.	IP Address 22.xx.xx.22
BGP ASN	Enter the ASN of your on-premises data center or private network. The BGP ASN of the customer gateway must be different from that of the VPN gateway.	65000

----End

Verification

Check the created customer gateway on the **Customer Gateway** page.

2.1.4 Step 3: Creating VPN Connections

Procedure

Step 1 In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.

Step 2 On the **VPN Connection** page, click **Create VPN Connection**.

Step 3 Set VPN connection parameters as prompted and click **Buy Now**.

The following describes only key parameters.

Table 2-5 Description of VPN connection parameters

Parameter	Description	Example Value
Name	Enter the name of VPN connection 1.	vpn-001
VPN Gateway	Select the VPN gateway created in **Step 1: Creating a VPN Gateway**.	vpngw-001
VPN Gateway IP of Connection 1	Select the active EIP of the VPN gateway.	11.xx.xx.11
Customer Gateway of Connection 1	Select the customer gateway of connection 1.	cgw-001
VPN Gateway IP of Connection 2	Select active EIP 2 of the VPN gateway.	11.xx.xx.12
Customer Gateway of Connection 2	Select the customer gateway of connection 2.	cgw-001
VPN Type	Select **Static routing**.	Static routing

Parameter	Description	Example Value
Customer Subnet	Enter the subnet of the on-premises data center that needs to access the VPC. **NOTE** • The customer subnet can overlap with the local subnet but cannot be the same as the local subnet. • A customer subnet cannot be included in the existing subnets of the VPC associated with the VPN gateway. It also cannot be the destination address in the route table of the VPC associated with the VPN gateway. • Customer subnets cannot be the reserved CIDR blocks of VPCs, for example, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. • If the interconnection subnet is associated with an ACL rule, ensure that the ACL rule permits the TCP port for traffic between all local and customer subnets. • Address groups cannot be used to configure the source and destination subnets in a policy on customer gateway devices. • When **Associate With** is set to **Enterprise Router** and **VPN Type** is set to **BGP routing** or **Policy-based**, you do not need to configure customer subnets.	172.16.0.0/16
Connection 1's Configuration	Configure the IP address assignment mode of gateway interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, and policies for connection 1.	Set parameters based on the site requirements.
Interface IP Address Assignment	The options include **Manually specify** and **Automatically assign**.	Manually specify

Parameter	Description	Example Value
Local Tunnel Interface Address	Specify the tunnel interface address of the VPN gateway. **NOTE** The local and remote interface addresses configured on the customer gateway device must be the same as the values of **Customer Tunnel Interface IP Address** and **Local Tunnel Interface IP Address**, respectively.	169.254.70.2/30
Customer Tunnel Interface Address	Specify the tunnel interface address of the customer gateway device.	169.254.70.1/30
Link Detection	This function is used for route reliability detection in multi-link scenarios. **NOTE** When enabling this function, ensure that the customer gateway supports ICMP and is correctly configured with the customer interface IP address of the VPN connection. Otherwise, VPN traffic will fail to be forwarded.	**NQA** enabled
PSK, Confirm PSK	Specify the negotiation key of the VPN connection. The PSKs configured on the VPN console and the customer gateway device must be the same.	Test@123
Policy Settings	Configure the IKE and IPsec policies, which define the encryption algorithms used by the VPN tunnel. The policy settings on the VPN console and the customer gateway device must be the same.	Default
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**.	Disabled
Local Tunnel Interface Address	Specify the tunnel interface address of the VPN gateway.	169.254.71.2/30
Customer Tunnel Interface Address	Specify the tunnel interface address of the customer gateway device.	169.254.71.1/30

----End

Verification

Check the created VPN connection on the **VPN Connection** page. The initial state of the VPN connection is **Creating**. As the customer gateway device has not been configured, no VPN connection can be established. After about 2 minutes, the VPN connection state changes to **Not connected**.

2.1.5 Step 4: Configuring the Customer Gateway Device

Procedure

NOTE

In this example, the customer gateway device is an AR router.

Step 1 Log in to the AR router.

Step 2 Enter the system view.

```
<AR651>system-view
```

Step 3 Configure an IP address for the WAN interface. In this example, the WAN interface of the AR router is GigabitEthernet 0/0/8.

```
[AR651]interface GigabitEthernet 0/0/8
[AR651-GigabitEthernet0/0/8]ip address 22.xx.xx.22 255.255.255.0
[AR651-GigabitEthernet0/0/8]quit
```

Step 4 Configure a default route.

```
[AR651]ip route-static 0.0.0.0 0.0.0.0 22.xx.xx.1
```

In this command, *22.xx.xx.1* is the gateway address of the AR router's public IP address. Replace it with the actual gateway address.

Step 5 Enable the SHA-2 algorithm to be compatible with the standard RFC algorithms.

```
[AR651]IPsec authentication sha2 compatible enable
```

Step 6 Configure an IPsec proposal.

```
[AR651]IPsec proposal hwproposal1
[AR651-IPsec-proposal-hwproposal1]esp authentication-algorithm sha2-256
[AR651-IPsec-proposal-hwproposal1]esp encryption-algorithm aes-128
[AR651-IPsec-proposal-hwproposal1]quit
```

Step 7 Configure an IKE proposal.

```
[AR651]ike proposal 2
[AR651-ike-proposal-2]encryption-algorithm aes-128
[AR651-ike-proposal-2]dh group14
[AR651-ike-proposal-2]authentication-algorithm sha2-256
[AR651-ike-proposal-2]authentication-method pre-share
[AR651-ike-proposal-2]integrity-algorithm hmac-sha2-256
[AR651-ike-proposal-2]prf hmac-sha2-256
[AR651-ike-proposal-2]quit
```

Step 8 Configure IKE peers.

```
[AR651]ike peer hwpeer1
[AR651-ike-peer-hwpeer1]undo version 1
[AR651-ike-peer-hwpeer1]pre-shared-key cipher Test@123
[AR651-ike-peer-hwpeer1]ike-proposal 2
[AR651-ike-peer-hwpeer1]local-address 22.xx.xx.22
[AR651-ike-peer-hwpeer1]remote-address 11.xx.xx.11
[AR651-ike-peer-hwpeer1]rsa encryption-padding oaep
[AR651-ike-peer-hwpeer1]rsa signature-padding pss
[AR651-ike-peer-hwpeer1]ikev2 authentication sign-hash sha2-256
[AR651-ike-peer-hwpeer1]quit
```

```
[AR651]ike peer hwpeer2
[AR651-ike-peer-hwpeer2]undo version 1
[AR651-ike-peer-hwpeer2]pre-shared-key cipher Test@123
[AR651-ike-peer-hwpeer2]ike-proposal 2
[AR651-ike-peer-hwpeer2]local-address 22.xx.xx.22
[AR651-ike-peer-hwpeer2]remote-address 11.xx.xx.12
[AR651-ike-peer-hwpeer2]rsa encryption-padding oaep
[AR651-ike-peer-hwpeer2]rsa signature-padding pss
[AR651-ike-peer-hwpeer2]ikev2 authentication sign-hash sha2-256
[AR651-ike-peer-hwpeer2]quit
```

The commands are described as follows:

- **pre-shared-key cipher**: configures a PSK, which must be the same as that configured on the VPN console.
- **local-address**: specifies the public IP address of the AR router.
- **remote-address**: specifies the active EIP or active EIP 2 of the VPN gateway.

Step 9 Configure an IPsec profile.

```
[AR651]IPsec profile hwpro1
[AR651-IPsec-profile-hwpro1]ike-peer hwpeer1
[AR651-IPsec-profile-hwpro1]proposal hwproposal1
[AR651-IPsec-profile-hwpro1]pfs dh-group14
[AR651-IPsec-profile-hwpro1]quit
[AR651]IPsec profile hwpro2
[AR651-IPsec-profile-hwpro2]ike-peer hwpeer2
[AR651-IPsec-profile-hwpro2]proposal hwproposal1
[AR651-IPsec-profile-hwpro2]pfs dh-group14
[AR651-IPsec-profile-hwpro2]quit
```

Step 10 Configure virtual tunnel interfaces.

```
[AR651]interface Tunnel0/0/1
[AR651-Tunnel0/0/1]mtu 1400
[AR651-Tunnel0/0/1]ip address 169.254.70.1 255.255.255.252
[AR651-Tunnel0/0/1]tunnel-protocol IPsec
[AR651-Tunnel0/0/1]source 22.xx.xx.22
[AR651-Tunnel0/0/1]destination 11.xx.xx.11
[AR651-Tunnel0/0/1]IPsec profile hwpro1
[AR651-Tunnel0/0/1]quit
[AR651]interface Tunnel0/0/2
[AR651-Tunnel0/0/2]mtu 1400
[AR651-Tunnel0/0/2]ip address 169.254.71.1 255.255.255.252
[AR651-Tunnel0/0/2]tunnel-protocol IPsec
[AR651-Tunnel0/0/2]source 22.xx.xx.22
[AR651-Tunnel0/0/2]destination 11.xx.xx.12
[AR651-Tunnel0/0/2]IPsec profile hwpro2
[AR651-Tunnel0/0/2]quit
```

The commands are described as follows:

- **interface Tunnel0/0/1** and **interface Tunnel0/0/2**: indicate the tunnel interfaces corresponding to the two VPN connections.

  In this example, Tunnel0/0/1 establishes a VPN connection with the active EIP of the VPN gateway, and Tunnel0/0/2 establishes a VPN connection with active EIP 2 of the VPN gateway.
- **ip address**: configures an IP address for a tunnel interface on the AR router.
- **source**: specifies the public IP address of the AR router.
- **destination**: specifies the active EIP or active EIP 2 of the VPN gateway.

Step 11 Configure NQA.

```
[AR651]nqa test-instance IPsec_nqa1 IPsec_nqa1
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]test-type icmp
```

```
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]destination-address ipv4 169.254.70.2
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]source-address ipv4 169.254.70.1
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]frequency 15
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]ttl 255
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]start now
[AR651-nqa-IPsec_nqa1-IPsec_nqa1]quit
[AR651]nqa test-instance IPsec_nqa2 IPsec_nqa2
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]test-type icmp
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]destination-address ipv4 169.254.71.2
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]source-address ipv4 169.254.71.1
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]frequency 15
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]ttl 255
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]start now
[AR651-nqa-IPsec_nqa2-IPsec_nqa2]quit
```

The commands are described as follows:

- **nqa test-instance IPsec_nqa1 IPsec_nqa1** and **nqa test-instance IPsec_nqa2 IPsec_nqa2**: configure two NQA test instances named **IPsec_nqa1** and **IPsec_nqa2**.

  In this example, the test instance **IPsec_nqa1** is created for the VPN connection to which the active EIP of the VPN gateway belongs; the test instance **IPsec_nqa2** is created for the VPN connection to which active EIP 2 of the VPN gateway belongs.
- **destination-address**: specifies the tunnel interface address of the VPN connection.
- **source-address**: specifies the tunnel interface address of the AR router.

Step 12 Configure association between the static route and NQA.

```
[AR651]ip route-static 192.168.0.0 255.255.255.0 Tunnel0/0/1 track nqa IPsec_nqa1 IPsec_nqa1
[AR651]ip route-static 192.168.0.0 255.255.255.0 Tunnel0/0/2 track nqa IPsec_nqa2 IPsec_nqa2
```

The parameters are described as follows:

- **192.168.0.0** indicates the local subnet of the VPC.
- **Tunnel***x* and **IPsec_nqa***x* in the same command correspond to the same VPN connection.

----End

Verification

Step 1 Log in to the management console.

Step 2 Click in the upper left corner, and choose > **Virtual Private Network**.

Step 3 In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.

Verify that the states of the two VPN connections are both **Normal**.

----End

2.1.6 Step 5: Verifying Network Connectivity

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 **Computing** > **Elastic Cloud Server**.

Step 4 Log in to an ECS.

Login using VNC on the management console is used as an example.

Step 5 Run the following command on the ECS:

ping 172.16.0.100

172.16.0.100 is the IP address of a server in the on-premises data center. Replace it with an actual server IP address.

If information similar to the following is displayed, the VPC on the cloud and the on-premises data center can communicate with each other.

```
Reply from xx.xx.xx.xx: bytes=32 time=28ms TTL=245
Reply from xx.xx.xx.xx: bytes=32 time=28ms TTL=245
Reply from xx.xx.xx.xx: bytes=32 time=28ms TTL=245
Reply from xx.xx.xx.xx: bytes=32 time=27ms TTL=245
```

----End

2.2 Configuring S2C Classic VPN to Connect an On-premises Data Center to a VPC

2.2.1 Buying a VPN Gateway

Scenarios

To connect your on-premises data center or private network to your ECSs in a VPC, buy a VPN gateway first.

Prerequisites

- A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
- Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.

4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

5. On the **VPN Gateways** page, click **Create VPN Gateway**.
6. Configure parameters based on **Table 2-6**, and click **Create Now**.

Table 2-6 Description of VPN gateway parameters

Parameter	Description	Example Value
Billing Mode	Billing mode of a VPN gateway, which can be pay-per-use **Pay-per-use**: When you buy a pay-per-use VPN gateway, you must buy a VPN connection together with the VPN gateway.	Pay-per-use
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across regions. For low network latency and fast resource access, select the region nearest to your target users.	-
Name	Name of a VPN gateway.	vpngw-001
VPC	Name of the VPC to which the VPN gateway connects.	vpc-001
Type	VPN type. **IPsec** is selected by default.	IPsec

Table 2-7 Description of VPN connection parameters

Parameter	Description	Example Value
Name	Name of a VPN connection.	vpn-001
VPN Gateway	Name of the VPN gateway for which the VPN connection is created.	vpcgw-001

Parameter	Description	Example Value
Local Subnet	VPC subnets that will access your on-premises network through a VPN. You can set the local subnet using either of the following methods: • **Select subnet**: Select the subnets that need to access your on-premises data center or private network. • **Specify CIDR block**: Enter the CIDR blocks that need to access your on-premises data center or private network.	192.168.1.0/24, 192.168.2.0/24
Remote Gateway	The public IP address of the gateway in your data center or on the private network. This IP address is used for communicating with your VPC.	N/A
Remote Subnet	The subnets of your on-premises network that will access a VPC through a VPN. The remote and local subnets cannot overlap with each other. The remote subnet cannot overlap with CIDR blocks involved in existing VPC peering, Direct Connect, or Cloud Connect connections created for the local VPC.	192.168.3.0/24, 192.168.4.0/24
PSK	PSKs configured at both ends of a VPN connection must be the same. The PSK: • Contains 6 to 128 characters. • Can contain only: – Digits – Letters – Special characters: ~ ` ! @ # $ % ^ ( ) - _ + = [ ] { } | \ , . / : ;	Test@123
Confirm PSK	Enter the PSK again.	Test@123
Advanced Settings	• **Default**: Use default IKE and IPsec policies. • **Custom**: Use custom IKE and IPsec policies. For details, see **Table 2-8** and **Table 2-9**.	Custom

Table 2-8 IKE policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: • MD5(Insecure. Not recommended.) • SHA1(Insecure. Not recommended.) • SHA2-256 • SHA2-384 • SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: • AES-128 • AES-192 • AES-256 • 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128
DH Algorithm	• Group 1(Insecure. Not recommended.) • Group 2(Insecure. Not recommended.) • Group 5(Insecure. Not recommended.) • Group 14 • Group 15 • Group 16 • Group 19 • Group 20 • Group 21 The default value is **Group 14**. DH algorithms configured at both ends of a VPN connection must be the same. Otherwise, the negotiation will fail.	Group 14
Version	Version of the IKE protocol. The value can be one of the following: • v1 (not recommended due to security risks) • v2 The default value is **v2**.	v2

Parameter	Description	Example Value
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **86400**.	86400

Table 2-9 IPsec policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: ● SHA1(Insecure. Not recommended.) ● MD5(Insecure. Not recommended.) ● SHA2-256 ● SHA2-384 ● SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: ● AES-128 ● AES-192 ● AES-256 ● 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128

Parameter	Description	Example Value
PFS	Algorithm used by the Perfect forward secrecy (PFS) function. ● DH group 1(Insecure. Not recommended.) ● DH group 2(Insecure. Not recommended.) ● DH group 5(Insecure. Not recommended.) ● DH group 14 ● DH group 15 ● DH group 16 ● DH group 19 ● DH group 20 ● DH group 21 The default algorithm is **DH group 14**.	DH group 14
Transfer Protocol	Security protocol used in IPsec to transmit and encapsulate user data. The following protocols are supported: ● ESP ● AH ● AH-ESP The default protocol is **ESP**.	ESP
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **3600**.	3600

7. Confirm the VPN gateway information and click **Submit**.

2.2.2 Buying a VPN Connection

Scenarios

To connect your on-premises data center or private network to your ECSs in a VPC, you need to create a VPN connection after a VPN gateway is obtained.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.

3. Click ☰ in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.

4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Connections**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

5. On the **VPN Connections** page, click **Create VPN Connection**.

6. Configure the parameters as prompted and click **Pay Now**. **Table 2-10** describes the VPN connection parameters.

Table 2-10 Description of VPN connection parameters

Parameter	Description	Example Value
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across regions. For low network latency and fast resource access, select the region nearest to your target users.	N/A
Name	Name of a VPN connection.	vpn-001
VPN Gateway	Name of the VPN gateway for which the VPN connection is created.	vpcgw-001
Local Subnet	VPC subnets that will access your on-premises network through a VPN. You can set the local subnet using either of the following methods: • **Select subnet**: Select the subnets that need to access your on-premises data center or private network. • **Specify CIDR block**: Enter the CIDR blocks that need to access your on-premises data center or private network.	192.168.1.0/24, 192.168.2.0/24
Remote Gateway	The public IP address of the gateway in your data center or on the private network. This IP address is used for communicating with your VPC.	N/A

Parameter	Description	Example Value
Remote Subnet	The subnets of your on-premises network that will access a VPC through a VPN. The remote and local subnets cannot overlap with each other. The remote subnet cannot overlap with CIDR blocks involved in existing VPC peering, Direct Connect, or Cloud Connect connections created for the local VPC.	192.168.3.0/24, 192.168.4.0/24
PSK	Private key shared by two ends of a VPN connection for negotiation. PSKs configured at both ends of the VPN connection must be the same. The PSK: • Contains 6 to 128 characters. • Can contain only: – Digits – Letters – Special characters: ~ ` ! @ # $ % ^ ( ) - _ + = [ ] { } | \ , . / : ;	Test@123
Confirm PSK	Enter the PSK again.	Test@123
Advanced Settings	• **Default**: Use default IKE and IPsec policies. • **Existing**: Use existing IKE and IPsec policies. • **Custom**: including **IKE Policy** and **IPsec Policy**, which specifies the encryption and authentication algorithms of a VPN tunnel. For details, see **Table 2-11** and **Table 2-12**.	Custom

Table 2-11 IKE policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: • MD5(Insecure. Not recommended.) • SHA1(Insecure. Not recommended.) • SHA2-256 • SHA2-384 • SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: • AES-128 • AES-192 • AES-256 • 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128
DH Algorithm	• Group 1(Insecure. Not recommended.) • Group 2(Insecure. Not recommended.) • Group 5(Insecure. Not recommended.) • Group 14 • Group 15 • Group 16 • Group 19 • Group 20 • Group 21 The default algorithm is **Group 14**.	Group 14
Version	Version of the IKE protocol. The value can be one of the following: • v1 (not recommended due to security risks) • v2 The default value is **v2**.	v2

Parameter	Description	Example Value
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **86400**.	86400

Table 2-12 IPsec policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: ● SHA1(Insecure. Not recommended.) ● MD5(Insecure. Not recommended.) ● SHA2-256 ● SHA2-384 ● SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: ● AES-128 ● AES-192 ● AES-256 ● 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128

Parameter	Description	Example Value
PFS	Algorithm used by the Perfect forward secrecy (PFS) function. ● DH group 1(Insecure. Not recommended.) ● DH group 2(Insecure. Not recommended.) ● DH group 5(Insecure. Not recommended.) ● DH group 14 ● DH group 15 ● DH group 16 ● DH group 19 ● DH group 20 ● DH group 21 The default algorithm is **DH group 14**.	DH group 14
Transfer Protocol	Security protocol used in IPsec to transmit and encapsulate user data. The following protocols are supported: ● AH ● ESP ● AH-ESP The default protocol is **ESP**.	ESP
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **3600**.	3600

NOTE

An IKE policy specifies the encryption and authentication algorithms to be used in the negotiation phase of an IPsec tunnel. An IPsec policy specifies the protocol, encryption algorithm, and authentication algorithm to be used in the data transmission phase of an IPsec tunnel. The IKE and IPsec policies must be the same at both ends of a VPN connection. If they are different, the VPN connection cannot be set up.

3 Management

3.1 S2C Enterprise Edition VPN

3.1.1 Enterprise Edition VPN Gateway Management

3.1.1.1 Creating a VPN Gateway

Scenario

To connect your on-premises data center or private network to your ECSs in a VPC, you need to create a VPN gateway before creating a VPN connection.

Context

The recommended networking varies according to the number of customer gateway IP addresses, as described in **Table 3-1**.

Table 3-1 Networking

Number of Customer Gateway IP Addresses	Recommended Networking	Description
1	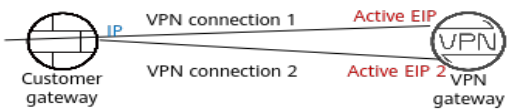	It is recommended that the VPN gateway uses the active-active mode. In this case, one VPN connection group is used.

Number of Customer Gateway IP Addresses	Recommended Networking	Description
2	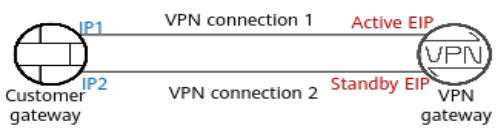	It is recommended that the VPN gateway uses the active/standby mode. In this case, two VPN connection groups are used.

- If your on-premises data center has only one customer gateway configured with only one IP address, it is recommended that the VPN gateway uses the active-active mode. In this mode, you need to create a VPN connection between each of the active EIP and active EIP 2 of the VPN gateway and the IP address of the customer gateway. In this scenario, only one VPN connection group is used.
- If your on-premises data center has two customer gateways or one customer gateway configured with two IP addresses, it is recommended that the VPN gateway uses the active/standby mode. In this mode, you need to create a VPN connection with each of the customer gateway IP addresses using the active and standby EIPs of the VPN gateway. In this scenario, two VPN connection groups are used.

Prerequisites

- A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
- Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.
- An enterprise router has been created if you want to use it to connect to a VPN gateway. For details, see the enterprise router documentation.

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 > **Virtual Private Network**.

Step 4 In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.

Step 5 On the **S2C VPN Gateways** page, click **Create S2C VPN Gateway**.

Step 6 Set parameters as prompted and click **Create Now**.

Table 3-2 lists the VPN gateway parameters.

Table 3-2 Description of VPN gateway parameters

Parameter	Description	Example Value
Region	For low network latency and fast resource access, select the region nearest to your target users. Resources cannot be shared across regions.	Select a region as required.
Name	Name of a VPN gateway. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	vpngw-001
Associate With	• VPC Through a VPC, the VPN gateway sends messages to the customer gateway or servers in the local subnet. • Enterprise Router Through an enterprise router, the VPN gateway sends messages to the customer gateway or servers in the subnets of all VPCs connected to the enterprise router. **NOTE** In this scenario, pay attention to the upper limit of entries in the routing table of the enterprise router. If the number of routes advertised by the customer gateway and VPN gateway exceeds this upper limit, the enterprise router cannot learn the excess routes. As a result, traffic will fail to be forwarded between the VPN gateway and the customer gateway.	VPC
VPC	Select a VPC.	vpc-001(192.168.0.0/16)
Enterprise Router	Select an enterprise router.	er-001
Access VPC	This parameter is available only when **Associate With** is set to **Enterprise Router**. If a VPN gateway needs to connect to different VPCs in the southbound and northbound directions, set the VPC in the northbound direction as the access VPC.	vpc-001(192.168.0.0/16)
Access Subnet	This parameter is available only when **Associate With** is set to **Enterprise Router**. An access subnet is used by the VPN gateway to connect to the Internet.	subnet-001(192.168.0.0/24)

Parameter	Description	Example Value
Gateway IP Address	This parameter is available only when **Associate With** is set to **Enterprise Router** and **Network Type** is set to **Private network**. • Auto-assigned IP address (default) An IP address on the access subnet will be automatically assigned to the VPN gateway. You can view the automatically assigned IP address on the **VPN Gateway** page. • Manually-specified IP address Manually configure IP addresses on the access subnet for the VPN gateway.	Auto-assigned IP address
Interconnection Subnet	This parameter is available only when **Associate With** is set to **VPC**. This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	192.168.66.0/24
Local Subnet	This parameter is available only when **Associate With** is set to **VPC**. Specify the VPC subnets with which your on-premises data center needs to communicate through the customer gateway. • Select subnet Select subnets of the local VPC. • Enter CIDR block Enter subnets of the local VPC or subnets of the VPC that establishes a peering connection with the local VPC.	192.168.1.0/24,192.168.2.0/24
BGP ASN	BGP ASN of the VPN gateway, which must be different from that of the customer gateway. The BGP ASN ranges from 1 to 4294967295.	64512

Parameter	Description	Example Value
HA Mode	• Active-active – When **Associate With** is set to **VPC**, the outgoing traffic from the VPN gateway to the customer subnet is preferentially forwarded through the first VPN connection (VPN connection 1) set up between the customer subnet and an EIP. If VPN connection 1 fails, the outgoing traffic is automatically switched to the other VPN connection (VPN connection 2) set up with the customer subnet. After VPN connection 1 recovers, the outgoing traffic is still transmitted through VPN connection 2 and will not be switched back to VPN connection 1. – 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. • Active/Standby The outgoing traffic from the VPN gateway to the customer subnet is preferentially transmitted through the VPN connection (VPN connection 1) set up between the customer subnet and the active EIP. If VPN connection 1 fails, the outgoing traffic is automatically switched to the other VPN connection (VPN connection 2) set up between the customer subnet and the standby EIP. After VPN connection 1 recovers, the outgoing traffic is automatically switched back to VPN connection 1.	Active-active
Specification	Two options are available: **Professional 1** and **Professional 2**.	Professional 1
Bandwidth Name	Specify the name of the EIP bandwidth.	Vpngw-bandwidth2
Active EIP	EIP used by the VPN gateway to communicate with a customer gateway. • **Create now**: Create an EIP. • **Use existing**: Use an existing EIP.	Create Now

Parameter	Description	Example Value
Bandwidth (Mbit/s)	Bandwidth of the EIP, in Mbit/s. ● All VPN connections created using the EIP share the bandwidth of the EIP. The total bandwidth consumed by all the VPN connections cannot exceed the bandwidth of the EIP. If network traffic exceeds the bandwidth of the EIP, network congestion may occur and VPN connections may be interrupted. As such, ensure that you configure enough bandwidth. ● You can configure alarm rules on Cloud Eye to monitor the bandwidth. ● You can customize the bandwidth within the allowed range.	10 Mbit/s
Active EIP 2	A VPN gateway needs to be bound to a group of EIPs (active EIP and active EIP 2). You can plan the bandwidth for each EIP. The EIPs can share bandwidth with the EIPs of other network services.	Create Now
Standby EIP	A VPN gateway needs to be bound to a group of EIPs (active EIP and standby EIP). You can plan the bandwidth for each EIP. The EIPs can share bandwidth with the EIPs of other network services. **NOTE** When **Billing Mode** of the VPN gateway is **Pay-per-use** and the backup EIP is billed by traffic, you are advised to configure alarm rules on Cloud Eye to monitor the backup EIP. This prevents traffic fee overrun caused by VPN connection switching due to a fault of the active VPN connection. For details about how to configure alarm rules on Cloud Eye, see the *Elastic IP User Guide*.	Create Now
Enterprise Project	Enterprise project to which the VPN belongs. An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.	default
Access VPC	If a VPN gateway needs to connect to different VPCs in the southbound and northbound directions, set the VPC in the northbound direction as the access VPC. The VPC in the southbound direction is the VPC associated with the VPN gateway.	Same as the associated VPC

Parameter	Description	Example Value
Access Subnet	By default, a VPN gateway uses the interconnection subnet to connect to the associated VPC. Set this parameter when another subnet needs to be used.	Same as the interconnection subnet

Step 7 Confirm the VPN gateway information and click **Create Now**

----End

3.1.1.2 Viewing a VPN Gateway

Scenario

After creating a VPN gateway, you can view its details.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** tab page, view the VPN gateway list.
6. Click the name of a VPN gateway to view its details.
   - For VPN gateways of the public network type, you can view their basic information, EIPs, and tags.
   - For VPN gateways of the private network type, you can view their basic information and advanced settings.

3.1.1.3 Modifying a VPN Gateway

Scenario

You can modify basic information about a VPN gateway, including the name and local subnet.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** page, click **Modify Basic Information** in the **Operation** column of the target VPN gateway.

To modify only the name of a VPN gateway, you can also click on the right of the VPN gateway name.

6. Modify the name and local subnet of the VPN gateway as prompted.
7. Click **OK**.

Table 3-3 describes the parameters for modifying the VPN gateway.

Table 3-3 Parameters for modifying the VPN gateway

Parameter	Description	Modifiable or Not
Name	Name of a VPN gateway. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	Y
Local Subnet	VPC subnets with which your on-premises data center needs to communicate through the customer gateway.	Y
Billing Mode	The value is **Pay-per-use**.	Y
Region	For low network latency and fast resource access, select the region nearest to your target users. Resources cannot be shared across regions.	N
Specification	Two options are available: **Professional 1** and **Professional 2**.	The supported specifications are subject to those displayed on the management console.
Associate With	The value can be **VPC** or **Enterprise Router**.	N
Enterprise Router	The associated enterprise router needs to be specified only when **Associate With** is set to **Enterprise Router**.	N
VPC	VPC that the on-premises data center needs to access.	N
Interconnection Subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	N

Parameter	Description	Modifiable or Not
BGP ASN	BGP AS number.	N
AZ	An AZ is a geographic location with independent power supply and network facilities in a region. AZs in the same VPC are interconnected through private networks and are physically isolated.	N

3.1.1.4 Changing the Specification of a VPN Gateway

Scenario

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 Professional 1 and Professional 2.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** page, choose **More** > **Change Specification** or click **Change Specification** in the **Operation** column of the target VPN gateway.
6. Modify the gateway specification as prompted.

3.1.1.5 Binding an EIP to a VPN Gateway

Scenario

You can bind EIPs to a VPN gateway that has been created.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.

5. On the **S2C VPN Gateways** page, click **Bind EIP** in the **Operation** column of the target VPN gateway.
   - If the VPN gateway uses the active-active mode, the VPN gateway can have an active EIP and active EIP 2 bound.
   - If the VPN gateway uses the active/standby mode, the VPN gateway can have an active EIP and a standby EIP bound.
6. Select the desired EIP and click **OK**.

3.1.1.6 Unbinding an EIP from a VPN Gateway

Scenario

After a VPN gateway is created, you can unbind an EIP from it.

Notes and Constraints

An EIP that is in use by a VPN connection cannot be unbound from a VPN gateway.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** page, choose **More** > **Unbind EIP** in the **Operation** column of the target VPN gateway.
   - If the VPN gateway uses the active-active mode, the active EIP and active EIP 2 can be unbound from the VPN gateway.
   - If the VPN gateway uses the active/standby mode, the active EIP and standby EIP can be unbound from the VPN gateway.
6. Click **OK**.

3.1.1.7 Deleting a VPN Gateway

Scenario

You can delete a VPN gateway that is no longer required.

Notes and Constraints

- The delete operation is not supported for a VPN gateway that is being created, updated, or deleted.
- An EIP bound to a VPN gateway will be released when the VPN gateway is deleted.

  To retain such a pay-per-use EIP, unbind it before deleting the VPN gateway. For details about how to unbind an EIP, see **3.1.1.6 Unbinding an EIP from a VPN Gateway**.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** page, choose **More** > **Delete** or click **Delete** in the **Operation** column of the target VPN gateway.
6. In the **Delete VPN Gateway** dialog box, click **Auto Enter**.
7. Click **OK**.

   NOTE

   The impact of shared bandwidth freezing on EIPs is subject to the EIP documentation.

3.1.1.8 Searching for VPN Gateways by Tag

Scenario

When using the VPN service, you can classify VPN resources based on specific rules to facilitate resource management and fee calculation.

With the Tag Management Service (TMS), you can add tags to your VPN resources to classify them. Additionally, you can quickly search for VPN resources by tag on the management console.

Prerequisites

You have added tags to VPN resources.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Gateways**.
5. On the **S2C VPN Gateways** page, click in the text box for selecting a property or entering a keyword, choose a tag key under **Resource Tag**, and select a tag value.
   - You can only select existing keys and values from the drop-down list.
   - You can select a maximum of 20 tags to search for VPN resources. If you select multiple tags, the relationship between them is AND.
   - You can use tags together with other types of filter criteria. The relationship between them is AND.

3.1.1.9 Upgrading a Gateway

Overview

You can determine whether a VPN gateway can be upgraded by checking whether the upgrade button is available in the **Operation** column of the VPN gateway.

- If no upgrade button is available, the VPN gateway cannot be upgraded.
- If the upgrade button is available, the VPN gateway can be upgraded.

You can determine whether to perform a rollback when the gateway status is **Upgrade to be committed**.

Upgrade Impact

- VPN connections will be interrupted for about 10 minutes during the upgrade.
- You cannot perform operations on a VPN gateway or its VPN connections during the upgrade.

Rollback

After the upgrade, you need to check whether services are normal. If there are any exceptions, you can roll back the upgrade. If services are normal, you can commit the upgrade, after which a rollback is not supported.

Procedure

Figure 3-1 Upgrade flowchart

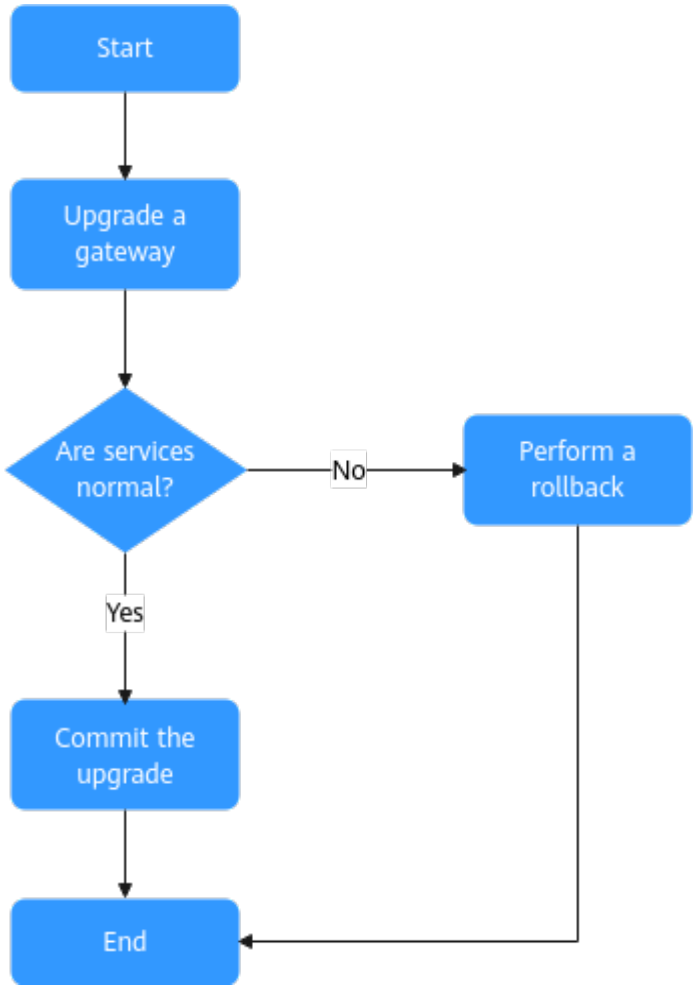

Step 1 Upgrade a gateway.

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. On the **S2C VPN Gateways** page, locate the target VPN gateway, and click **Upgrade Gateway** in the **Operation** column.
5. In the dialog box that is displayed, read the upgrade impact and rollback information, select **I understand the above information**, and click **OK**.
6. Check the upgrade status. During the upgrade, you can click **View task** in the **Status** column of the VPN gateway to view the upgrade progress.
   - If the upgrade is successful, the gateway status changes to **Upgrade to be committed**. Go to 2.
   - If the upgrade fails, a rollback is automatically performed. You can view the failure information in the upper right corner of the VPN gateway list.

Step 2 Verify services.

1. If services are normal, click **Commit Upgrade** in the **Operation** column to commit the upgrade.

   NOTICE

   After you commit the upgrade, a rollback is not supported. Exercise caution when performing this operation.

2. If services are abnormal, click **Roll Back** in the **Operation** column, and submit a service ticket to contact technical support.

----End

3.1.2 Customer Gateway Management of Enterprise Edition VPN

3.1.2.1 Creating a Customer Gateway

Scenario

To connect your on-premises data center or private network to your ECSs in a VPC, you need to create a customer gateway before creating a VPN connection.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.
5. On the **Customer Gateway** page, click **Create Customer Gateway**.
6. Set parameters as prompted and click **Create Now**.

   Table 3-4 lists the customer gateway parameters.

Table 3-4 Description of customer gateway parameters

Parameter	Description	Example Value
Name	Name of a customer gateway. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	cgw-001
BGP ASN	Enter the ASN of your on-premises data center or private network. The BGP ASN of the customer gateway must be different from that of the VPN gateway.	65000

Parameter	Description	Example Value
Advanced Settings > Tags	Tag of a VPN resource. The value consists of a key and a value. A maximum of 20 tags can be added. You can select predefined tags or customize tags. To view predefined tags, click **View predefined tags**.	-

7. (Optional) If there are two customer gateways, repeat the preceding operations to configure the other customer gateway with a different identifier.

Related Operations

You need to configure an IPsec VPN tunnel on the router or firewall in your on-premises data center.

3.1.2.2 Viewing a Customer Gateway

Scenario

After creating a customer gateway, you can view its details.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.
5. On the **Customer Gateway** page, view the customer gateway list.
6. Click the name of a customer gateway to view its details.
   - In the **Basic Information** area, you can view the name, identifier, ID, BGP ASN, and VPN connections of the customer gateway.

3.1.2.3 Modifying a Customer Gateway

Scenario

After creating a customer gateway, you can modify its name.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.

3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.
5. On the **Customer Gateway** page, click next to the name of a customer gateway.
6. Enter a new name for the customer gateway and click **OK**.

   Table 3-5 describes the parameters related to customer gateway modification.

Table 3-5 Parameters related to customer gateway modification

Parameter	Description	Modifiable or Not
Name	Name of a VPN connection. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	Y
Advanced Settings > Tags	Resource tag of the VPN service, including a key and a value.	Y
BGP ASN	BGP AS number.	N
Identifier	IP address used by the customer gateway to communicate with the VPN gateway. The value must be a static address.	N

3.1.2.4 Deleting a Customer Gateway

Scenario

You can delete a customer gateway that you have created.

Notes and Constraints

Before deleting a customer gateway associated with a VPN connection, remove the customer gateway from the VPN connection.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.

5. On the **Customer Gateway** page, locate the customer gateway to delete, and click **Delete** in the **Operation** column.
6. Click **OK**.

3.1.2.5 Searching for Customer Gateways by Tag

Scenario

When using the VPN service, you can classify VPN resources based on specific rules to facilitate resource management and fee calculation.

With the Tag Management Service (TMS), you can add tags to your VPN resources to classify them. Additionally, you can quickly search for VPN resources by tag on the management console.

Prerequisites

You have added tags to VPN resources.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – Customer Gateways**.
5. Click in the text box for selecting a property or entering a keyword, choose a tag key under **Resource Tag**, and select a tag value.
   - You can only select existing keys and values from the drop-down list.
   - You can select a maximum of 20 tags to search for VPN resources. If you select multiple tags, the relationship between them is AND.
   - You can use tags together with other types of filter criteria. The relationship between them is AND.

3.1.3 Enterprise Edition VPN Connection Management

3.1.3.1 Creating VPN Connections

Scenario

To connect your on-premises data center or private network to your ECSs in a VPC, you need to create VPN connections after creating a VPN gateway and a customer gateway.

Notes and Constraints

- When creating a VPN connection in static routing mode, ensure that the customer gateway supports ICMP and is correctly configured with the customer interface IP address of the VPN connection before enabling NQA. Otherwise, traffic will fail to be forwarded.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, click **Create VPN Connection**.

   NOTE

   A VPN gateway can establish two VPN connections with a customer gateway using EIPs, improving reliability.
6. Set parameters as prompted and click **Buy Now**.

   Table 3-6 lists the VPN connection parameters.

Table 3-6 Description of VPN connection parameters

Parameter	Description	Example Value
Name	VPN connection name. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	vpn-001
VPN Gateway	Name of the VPN gateway for which VPN connections are created. You can also click **Create VPN Gateway** to create a VPN gateway. For details about related parameters, see Table 3-2.	vpngw-001
VPN Gateway IP of Connection 1	• When **Network Type** is set to **Public network**, the value is the active EIP of the VPN gateway. • When **Network Type** is set to **Private network**, the value is the active IP address of the VPN gateway. The same address of a VPN gateway cannot be repeatedly selected when you create VPN connections between the VPN gateway and the same customer gateway.	11.xx.xx.11

Parameter	Description	Example Value
Customer Gateway of Connection 1	Select the customer gateway of connection 1. You can also click **Create Customer Gateway** to create a customer gateway. For details about related parameters, see **Table 3-4**. **NOTE** If a customer gateway connects to multiple VPN gateways, the BGP ASNs and VPN types of the VPN gateways must be the same.	cgw-001
VPN Gateway IP of Connection 2	• When **Network Type** is set to **Public network** and **HA Mode** is set to **Active-active**, the value is active EIP 2 of the VPN gateway. • When **Network Type** is set to **Private network** and **HA Mode** is set to **Active-active**, the value is active IP address 2 of the VPN gateway. • When **Network Type** is set to **Public network** and **HA Mode** is set to **Active/Standby**, the value is the standby EIP of the VPN gateway. • When **Network Type** is set to **Private network** and **HA Mode** is set to **Active/Standby**, the value is the standby IP address of the VPN gateway. The VPN gateway IP address must be unique for each connection with a customer gateway.	11.xx.xx.12
Customer Gateway of Connection 2	Select the customer gateway of connection 2. You can also click **Create Customer Gateway** to create a customer gateway. For details about related parameters, see **Table 3-4**. **NOTE** If a customer gateway connects to multiple VPN gateways, the BGP ASNs and VPN types of the VPN gateways must be the same.	cgw-001

Parameter	Description	Example Value
VPN Type	IPsec connection mode, which can be route-based or policy-based. ● Static routing Determines the data that enters the IPsec VPN tunnel based on the route configuration (local subnet and customer subnet). **Application scenario: Communication between customer gateways** ● BGP routing Determines the traffic that can enter the IPsec VPN tunnel based on BGP routes. **Application scenario: Communication between customer gateways, many or frequently changing interconnection subnets, or backup between VPN and Direct Connect** ● Policy-based Determines the data that enters the IPsec VPN tunnel based on the policy (between the customer network and VPC). Policy rules can be defined based on the source and destination CIDR blocks. **Application scenario: Isolation between customer gateways** **NOTE** By default, the VPN type, customer subnet, branch interconnection setting (BGP routing mode), and policy rules (policy-based mode) of the two connections are the same.	Static routing

Parameter	Description	Example Value
Customer Subnet	Customer-side subnet that needs to access the VPC on the cloud through VPN connections. If there are multiple customer subnets, separate them with commas (,). **NOTE** ● The customer subnet can overlap with the local subnet but cannot be the same as the local subnet. ● A customer subnet cannot be included in the existing subnets of the VPC associated with the VPN gateway. It also cannot be the destination address in the route table of the VPC associated with the VPN gateway. ● Customer subnets cannot be the reserved CIDR blocks of VPCs, for example, 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. ● If the interconnection subnet is associated with an ACL rule, ensure that the ACL rule permits the TCP port for traffic between all local and customer subnets. ● Address groups cannot be used to configure the source and destination subnets in a policy on customer gateway devices. ● When **Associate With** is set to **Enterprise Router** and **VPN Type** is set to **BGP routing** or **Policy-based**, you do not need to configure customer subnets.	172.16.1.0/24,172.16.2.0/24
Branch Interconnection	This parameter is available only when **VPN Type** is set to **BGP routing**. ● Enabled ● Disabled This function is disabled by default. **NOTE** When this function is disabled, only local subnet routes are advertised.	Disabled

Parameter	Description	Example Value
Policy	This parameter is available only when **VPN Type** is set to **Policy-based**. Defines the data flow that enters the encrypted VPN connections between the local and customer subnets. You need to configure the source and destination CIDR blocks in each policy rule. By default, a maximum of five policy rules can be configured. • Source CIDR Block The source CIDR block must contain some CIDR blocks of the local subnets. **0.0.0.0/0** indicates any IP address. A maximum of five source CIDR blocks can be configured for a VPN connection. • Destination CIDR Block The destination CIDR block must contain all the CIDR blocks of the customer subnets. A policy rule supports a maximum of 50 destination CIDR blocks, which are separated by commas (,).	• Source CIDR block 1: 192.168.1.0/24 • Destination CIDR block 1: 172.16.1.0/24,172.16.2.0/24 • Source CIDR block 2: 192.168.2.0/24 • Destination CIDR block 2: 172.16.1.0/24,172.16.2.0/24
Connection 1's Configuration	Configure the IP address assignment mode of tunnel interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, policies, and advanced settings for connection 1.	Set parameters based on the site requirements.

Parameter	Description	Example Value
Interface IP Address Assignment	This parameter is available only when **VPN Type** is set to **Static routing** or **BGP routing**. **NOTE** • Set interface IP addresses to the tunnel interface IP addresses used by the VPN gateway and customer gateway to communicate with each other. • If the tunnel interface address of the customer gateway is fixed, select **Manually specify**, and set the tunnel interface address of the VPN gateway based on the tunnel interface address of the customer gateway. • Manually specify – Set **Local Tunnel Interface Address** to the tunnel interface address of the VPN gateway, which can reside only on the CIDR block 169.254.*x.x*/30 (except 169.254.195.*x*/30). Then, the system automatically sets **Customer Tunnel Interface Address** based on the value of **Local Tunnel Interface Address**. For example, when you set **Local Tunnel Interface Address** to **169.254.1.6/30**, the system automatically sets **Customer Tunnel Interface Address** to **169.254.1.5/30**. – When you set **VPN Type** to **BGP routing** and configure tunnel interface addresses in **Manually specify** mode, ensure that the local and remote tunnel interface addresses configured on the customer gateway device (the other end of the VPN connection) are the same as the values of **Customer Tunnel Interface Address** and **Local Tunnel Interface Address**, respectively. • Automatically assign	Automatically assign

Parameter	Description	Example Value
	– By default, an IP address on the CIDR block 169.254.*x.x*/30 is assigned to the tunnel interface of the VPN gateway. – To view the automatically assigned local and customer interface IP addresses, click **Modify VPN Connection** on the **VPN Connection** page. – When you set **VPN Type** to **BGP routing** and select **Automatically assign**, check the automatically assigned local and customer tunnel interface addresses after the VPN connection is created. Ensure that the local and remote tunnel interface addresses configured on the customer gateway device (the other end of the VPN connection) are the reverse of the settings on the cloud side.	
Local Tunnel Interface Address	This parameter is available only when **Interface IP Address Assignment** is set to **Manually specify**. Tunnel interface IP address of the VPN gateway.	N/A
Customer Tunnel Interface Address	This parameter is available only when **Interface IP Address Assignment** is set to **Manually specify**. Tunnel interface IP address of the customer gateway device.	N/A

Parameter	Description	Example Value
Link Detection	This parameter is available only when **VPN Type** is set to **Static routing**. **NOTE** When enabling this function, ensure that the customer gateway supports ICMP and is correctly configured with the customer interface IP address of the VPN connection. Otherwise, traffic will fail to be forwarded. After this function is enabled, the VPN gateway automatically performs Network Quality Analysis (NQA) on the customer interface IP address of the customer gateway.	Selected
PSK	The PSKs configured for the VPN gateway and customer gateway must be the same. The PSK: • Contains 8 to 128 characters. • Can contain only three or more types of the following characters: – Digits – Uppercase letters – Lowercase letters – Special characters: ~ ! @ # $ % ^ ( ) - _ + = { } , . / : ;	Test@123
Confirm PSK	Enter the PSK again.	Test@123
Policy Settings	• **Default**: Use default IKE and IPsec policies. • **Custom**: Use custom IKE and IPsec policies. For details about the policies, see **Table 3-7** and **Table 3-8**. **NOTE** When **Local ID** and **Customer ID** are set to **IP Address**, you can specify specific IP addresses as the local and customer IDs, which must be different.	Custom

Parameter	Description	Example Value
Tags	• Tag of a VPN resource. The value consists of a key and a value. A maximum of 20 tags can be added. • You can select predefined tags or customize tags. • To view predefined tags, click **View predefined tags**.	-
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**. • Enabled • Disabled	Enabled

Table 3-7 IKE policy

Parameter	Description	Example Value
Version	Version of the IKE protocol. The value can be one of the following: • v1 (v1 has low security. If the device supports v2, v2 is recommended.) • v2 The default value is **v2**.	v2
Negotiation Mode	This parameter is available only when **Version** is **v1**. • Main • Aggressive	Main
Authentication Algorithm	Hash algorithm used for authentication. The following options are available: • SHA1(Insecure. Not recommended.) • MD5(Insecure. Not recommended.) • SHA2-256 • SHA2-384 • SHA2-512 The default value is **SHA2-256**.	SHA2-256

Parameter	Description	Example Value
Encryption Algorithm	Encryption algorithm. The following options are available: ● 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 When this encryption algorithm is used, the IKE version can only be **v2**. The default value is **AES-128**.	AES-128
DH Algorithm	The following algorithms are supported: ● 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**.	Group 15
Lifetime (s)	Lifetime of a security association (SA). An SA will be renegotiated when its lifetime expires. ● Unit: second ● The value ranges from **60** to **604800**. ● The default value is **86400**.	86400

Parameter	Description	Example Value
Local ID	Authentication identifier of the VPN gateway used in IPsec negotiation. The peer ID configured on the customer gateway must be the same as the local ID configured here. Otherwise, IPsec negotiation fails. ● IP Address (default value) – The system automatically sets this parameter to the IP address of the VPN gateway. – You can configure a specific IP address as the local ID, which must be different from the customer ID. ● FQDN Set the FQDN to a string of 1 to 128 case-sensitive characters that can contain letters, digits, and special characters (excluding &, <, >, [,], \, ?, and spaces).	IP Address
Customer ID	Authentication identifier of the customer gateway used in IPsec negotiation. The local ID configured on the customer gateway must be the same as the customer ID configured here. Otherwise, IPsec negotiation fails. ● **IP Address** (default) – The system automatically sets this parameter to the IP address of the customer gateway. – You can configure a specific IP address as the customer ID, which must be different from the local ID. ● **FQDN** Set the FQDN to a string of 1 to 128 case-sensitive characters that can contain letters, digits, and special characters (excluding &, <, >, [,], \, ?, and spaces).	IP Address

Table 3-8 IPsec policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following options are available: ● SHA1(Insecure. Not recommended.) ● MD5(Insecure. Not recommended.) ● SHA2-256 ● SHA2-384 ● SHA2-512 The default value is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following options are available: ● 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**.	AES-128

Parameter	Description	Example Value
PFS	Algorithm used by the Perfect forward secrecy (PFS) function. PFS supports the following algorithms: ● 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 **DH group 15**.	DH group 15
Transfer Protocol	Security protocol used in IPsec to transmit and encapsulate user data. The following protocol is supported: ESP The default value is **ESP**.	ESP
Lifetime (s)	Lifetime of an SA. An SA will be renegotiated when its lifetime expires. ● Unit: second ● The value ranges from **30** to **604800**. ● The default value is **3600**.	3600

NOTE

An IKE policy specifies the encryption and authentication algorithms to use in the negotiation phase of an IPsec tunnel. An IPsec policy specifies the protocol, encryption algorithm, and authentication algorithm to use in the data transmission phase of an IPsec tunnel. The policy settings for VPN connections must be the same at the VPC and on-premises data center sides. If they are different, VPN negotiation will fail, causing the failure to establish VPN connections.

The following algorithms are not recommended because they are not secure enough:

- Authentication algorithms: SHA1 and MD5
- Encryption algorithms: 3DES, AES-128, AES-192, and AES-256

  Because some customer devices do not support secure encryption algorithms, the default encryption algorithm of VPN connections is still AES-128. You are advised to use a more secure encryption algorithm if customer devices support secure encryption algorithms.
- DH algorithms: Group 1, Group 2, Group 5, and Group 14

7. Confirm the VPN connection configuration and click **Create Now**.

3.1.3.2 Configuring Health Check

Scenario

After VPN connections are created, you can configure health check to enable the VPN gateway to send probe packets to the customer gateway to collect statistics about the round-trip time and packet loss rate of physical links. The statistics help you learn about the VPN connection quality. Enabling health check does not affect tunnels. The Cloud Eye service monitors the round-trip time and packet loss rate of VPN links. For details, see **Metrics (S2C Enterprise Edition VPN)**.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, click the name of the target VPN connection. On the **Summary** tab page, click **Add** in the **Health Check** area.
6. In the **Add Health Check** dialog box, click **OK**.

3.1.3.3 Viewing a VPN Connection

Scenario

After creating a VPN connection, you can view its details.

Procedure

1. Log in to the management console.

2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, view the VPN connection list.
6. Click the name of a VPN connection to view its basic information, policy configuration, and tags.
   - When **VPN Type** is **Static routing**, the basic information includes the VPN connection information and health check information.
   - When **VPN Type** is **BGP routing**, the basic information includes the VPN connection information and health check information.
   - When **VPN Type** is **Policy-based**, the basic information includes the VPN connection information, policy rule information, and health check information.

NOTE

- 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.
- In the VPN connection list, you can locate the target VPN connection and click **View Metric** to view monitoring information about the VPN connection.

  Check the value of **VPN Connection Status**. If the value is **0**, the VPN connection is not connected. If the value is **1**, the VPN connection is connected. If the value is **2**, the VPN connection status is unknown.
- In the VPN connection list, dual connections to the same customer gateway are identified by . If such dual connections are displayed on different pages, and are also displayed on different pages.

  The dual-connection identifier will be unavailable if you sort VPN connections by any field in the VPN connection list. The identifier will be restored after you cancel field-based sorting.
- In the VPN connection list, you can click **View Logs** corresponding to the target VPN connection to view its IPsec negotiation logs.

  If a VPN connection is in **Not connected** state, you can determine the cause of the disconnection based on the VPN connection log details.
- On the **VPN Connection** page, the **Export** and setting buttons are available above the gateway list.
  - You can click **Export** in the upper left corner and select the data to be exported from the drop-down list.
  - You can click in the upper right corner and set the columns to be displayed as required.

3.1.3.4 Modifying a VPN Connection

Scenario

A VPN connection is an encrypted communications channel established between a VPN gateway in a VPC and a customer gateway in your on-premises data center. You can modify a VPN connection when required.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, locate the VPN connection to modify, and click **Modify VPN Connection** or **Modify Policy Settings**.
6. Modify VPN connection parameters as prompted.
   - For a VPN connection in BGP routing mode, you can enable or disable branch Interconnection on the **Modify VPN Connection** page.
7. Click **OK**.

CAUTION

If you change the PSK or modify the IKE or IPsec policy of a VPN connection, ensure that the new configurations are consistent with those on the customer gateway. Otherwise, the VPN connection will be interrupted.

Only some of the parameters take effect immediately after being modified, as described in **Table 3-9**.

Table 3-9 Time when new parameter settings take effect

Item	Parameter	When New Settings Take Effect	How to Modify
-	PSK	• When IKEv1 is used, the new setting takes effect in the next negotiation period. • When IKEv2 is used, the new setting takes effect after the VPN connection is re-established.	• When IKEv1 is used: Locate the VPN connection to modify, choose **More** > **Reset PSK** on the right, and change the PSK as prompted. • When IKEv2 is used: 1. Delete the current VPN connection. 2. Create a new VPN connection.
IKEv1 policy	Encryption Algorithm	The new settings take effect in the next negotiation period.	Locate the VPN connection to modify, and click **Modify VPN Configuration**.

Item	Parameter	When New Settings Take Effect	How to Modify
	Authentication Algorithm		
	DH Algorithm		
	Negotiation Mode		
	Local ID		
	Customer ID		
	Lifetime (s)		
	Version	The new settings take effect immediately.	
IKEv2 policy	Encryption Algorithm	The new settings take effect in the next negotiation period.	Locate the VPN connection to modify, and click **Modify VPN Configuration**.
	Authentication Algorithm		
	DH Algorithm		
	Lifetime (s)		
	Version	The new settings take effect immediately.	
	Local ID	The new settings take effect after the VPN connection is re-established.	1. Delete the current VPN connection. 2. Create a new VPN connection.
	Customer ID		

Item	Parameter	When New Settings Take Effect	How to Modify
IPsec policy	Encryption Algorithm	The new settings take effect in the next negotiation period.	Locate the VPN connection to modify, and click **Modify VPN Configuration**.
	Authentication Algorithm		
	PFS		
	Lifetime (s)		
	Transfer Protocol	This parameter cannot be modified on the management console.	

Table 3-10 describes the parameters related to VPN connection modification.

Table 3-10 Parameters related to VPN connection modification

Parameter	Description	Modifiable or Not
Name	VPN connection name. The value can contain only letters, digits, underscores (_), hyphens (-), and periods (.).	Y
Customer Gateway	Gateway used for communicating with a VPC through VPN.	Y
Customer Subnet	Subnet in the on-premises data center that needs to access the VPC.	Y
Policy Settings	There are IKE and IPsec policies.	Y
PSK	The PSKs configured for the VPN gateway and customer gateway must be the same.	Y
Local Tunnel Interface Address	Tunnel interface IP address configured on the VPN gateway.	Y

Parameter	Description	Modifiable or Not
Customer Tunnel Interface Address	Tunnel interface IP address configured on the customer gateway device.	Y
Branch Interconnection	This parameter is available only when **VPN Type** is set to **BGP routing**.	Y
EIP	This parameter is available only when **Network Type** is set to **Public network**.	N
Private IP address	This parameter is available only when **Network Type** is set to **Private network**.	N
VPN Gateway	VPN gateway that has been created.	N
Identifier	IP address used by the customer gateway to communicate with the VPN gateway. The value must be a static address. Ensure that UDP port 4500 is permitted in a firewall rule on the customer gateway in your on-premises data center or private network.	N
Interface IP Address Assignment	Mode in which IP addresses of the local and customer interfaces are assigned. The options include **Manually specify** and **Automatically assign**.	N

3.1.3.5 Deleting a VPN Connection

Scenario

If a VPN connection is no longer required, you can delete it to release network resources.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.

4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, choose **More** > **Delete** in the **Operation** column of a VPN connection.
6. In the **Delete VPN Connection** dialog box, click **Auto Enter**.
7. Click **OK**.

3.1.3.6 Resetting a VPN Connection

Scenario

You can reset a VPN connection if an exception occurs when you use the VPN connection.

Notes and Constraints

- You can reset a VPN connection only when its status is **Not connected**, **Normal**, or **Unknown**.
- Whether VPN connections can be reset is subject to the actual GUI on the management console.

CAUTION

Resetting a VPN connection will interrupt the VPN connection. Exercise caution when performing this operation.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. Locate the target VPN connection, , and click **OK**.
6. Check the VPN connection status.

   On the **VPN Connection** page, the VPN connection status is **Resetting** or **Not connected**.

CAUTION

If the VPN connection status is **Not connected** one minute after the VPN connection is reset, you need to manually refresh the page.

If the VPN connection status is still **Not connected** after you refresh the page, the VPN connection negotiation fails. For details, see **The State of a VPN Connection Is Not connected**.

3.1.3.7 Viewing VPN Connection Logs

Scenario

Logs are generated when VPN connection negotiation succeeds or fails. You can view connection logs to locate VPN connection faults.

Notes and Constraints

You can view the latest 200 VPN connection logs.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. On the **VPN Connection** page, locate the target VPN connection, and click **View Logs** to view connection logs.

   On the **View Logs** page, you can export the log data showing the time and information.

3.1.3.8 Searching for VPN Connections by Tag

Scenario

When using the VPN service, you can classify VPN resources based on specific rules to facilitate resource management and fee calculation.

With the Tag Management Service (TMS), you can add tags to your VPN resources to classify them. Additionally, you can quickly search for VPN resources by tag on the management console.

Prerequisites

You have added tags to VPN resources.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, and choose > **Virtual Private Network**.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.
5. Click in the text box for selecting a property or entering a keyword, choose a tag key under **Resource Tag**, and select a tag value.
   - You can only select existing keys and values from the drop-down list.

- You can select a maximum of 20 tags to search for VPN resources. If you select multiple tags, the relationship between them is AND.
- You can use tags together with other types of filter criteria. The relationship between them is AND.

3.2 S2C Classic VPN

3.2.1 Classic VPN Gateway Management

3.2.1.1 Buying a VPN Gateway

Scenarios

To connect your on-premises data center or private network to your ECSs in a VPC, buy a VPN gateway first.

Prerequisites

- A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
- Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.
5. On the **VPN Gateways** page, click **Create VPN Gateway**.
6. Configure parameters based on **Table 3-11**, and click **Create Now**.

Table 3-11 Description of VPN gateway parameters

Parameter	Description	Example Value
Billing Mode	Billing mode of a VPN gateway, which can be pay-per-use **Pay-per-use**: When you buy a pay-per-use VPN gateway, you must buy a VPN connection together with the VPN gateway.	Pay-per-use
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across regions. For low network latency and fast resource access, select the region nearest to your target users.	-
Name	Name of a VPN gateway.	vpngw-001
VPC	Name of the VPC to which the VPN gateway connects.	vpc-001
Type	VPN type. **IPsec** is selected by default.	IPsec

Table 3-12 Description of VPN connection parameters

Parameter	Description	Example Value
Name	Name of a VPN connection.	vpn-001
VPN Gateway	Name of the VPN gateway for which the VPN connection is created.	vpcgw-001
Local Subnet	VPC subnets that will access your on-premises network through a VPN. You can set the local subnet using either of the following methods: • **Select subnet**: Select the subnets that need to access your on-premises data center or private network. • **Specify CIDR block**: Enter the CIDR blocks that need to access your on-premises data center or private network.	192.168.1.0/24, 192.168.2.0/24

Parameter	Description	Example Value
Remote Gateway	The public IP address of the gateway in your data center or on the private network. This IP address is used for communicating with your VPC.	N/A
Remote Subnet	The subnets of your on-premises network that will access a VPC through a VPN. The remote and local subnets cannot overlap with each other. The remote subnet cannot overlap with CIDR blocks involved in existing VPC peering, Direct Connect, or Cloud Connect connections created for the local VPC.	192.168.3.0/24, 192.168.4.0/24
PSK	PSKs configured at both ends of a VPN connection must be the same. The PSK: • Contains 6 to 128 characters. • Can contain only: – Digits – Letters – Special characters: ~ ` ! @ # $ % ^ ( ) - _ + = [ ] { } | \ , . / : ;	Test@123
Confirm PSK	Enter the PSK again.	Test@123
Advanced Settings	• **Default**: Use default IKE and IPsec policies. • **Custom**: Use custom IKE and IPsec policies. For details, see **Table 3-13** and **Table 3-14**.	Custom

Table 3-13 IKE policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: • MD5(Insecure. Not recommended.) • SHA1(Insecure. Not recommended.) • SHA2-256 • SHA2-384 • SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: • AES-128 • AES-192 • AES-256 • 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128
DH Algorithm	• Group 1(Insecure. Not recommended.) • Group 2(Insecure. Not recommended.) • Group 5(Insecure. Not recommended.) • Group 14 • Group 15 • Group 16 • Group 19 • Group 20 • Group 21 The default value is **Group 14**. DH algorithms configured at both ends of a VPN connection must be the same. Otherwise, the negotiation will fail.	Group 14
Version	Version of the IKE protocol. The value can be one of the following: • v1 (not recommended due to security risks) • v2 The default value is **v2**.	v2

Parameter	Description	Example Value
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **86400**.	86400

Table 3-14 IPsec policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: ● SHA1(Insecure. Not recommended.) ● MD5(Insecure. Not recommended.) ● SHA2-256 ● SHA2-384 ● SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: ● AES-128 ● AES-192 ● AES-256 ● 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128

Parameter	Description	Example Value
PFS	Algorithm used by the Perfect forward secrecy (PFS) function. ● DH group 1(Insecure. Not recommended.) ● DH group 2(Insecure. Not recommended.) ● DH group 5(Insecure. Not recommended.) ● DH group 14 ● DH group 15 ● DH group 16 ● DH group 19 ● DH group 20 ● DH group 21 The default algorithm is **DH group 14**.	DH group 14
Transfer Protocol	Security protocol used in IPsec to transmit and encapsulate user data. The following protocols are supported: ● ESP ● AH ● AH-ESP The default protocol is **ESP**.	ESP
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **3600**.	3600

7. Confirm the VPN gateway information and click **Submit**.

3.2.1.2 Viewing a VPN Gateway

Scenarios

After creating a VPN gateway, you can view its details.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

The VPN console is displayed.

4. In the navigation pane on the left, choose **Virtual Private Network**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

5. View VPN gateway information.

3.2.1.3 Modifying a VPN Gateway

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 **Network** > **Virtual Private Network**.
4. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.

5. In the navigation pane on the left, choose **Virtual Private Network**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

6. On the Classic page, click the **VPN Gateways** tab.
   - Locate the row that contains the target VPN gateway, and choose **More** > **Modify Bandwidth** in the **Operation** column.
   - Locate the row that contains the target VPN gateway, and choose **More** > **Modify Basic Information** in the **Operation** column.
   - Locate the row that contains the target VPN gateway, and choose **More** > **Modify Specifications** in the **Operation** column.
7. Modify the VPN gateway bandwidth, name, or description as required.
8. Click **OK**.

Modifying Basic Information About a VPN Gateway

Scenario

You can modify the name and description of a VPN gateway.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.

4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic – VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

5. Locate the row that contains the VPN gateway that you want to modify, and choose **More** > **Modify Basic Information** in the **Operation** column.

6. Modify the VPN gateway name or description as required.

   NOTE

   The name of a VPN gateway can contain only letters, digits, underscores (_), hyphens (-), and periods (.).

7. Click **OK**.

Modifying VPN Gateway Bandwidth

Scenario

When the bandwidth of a VPN gateway cannot meet your service requirements, you can modify the VPN gateway bandwidth.

Procedure

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. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
5. In the navigation pane on the left, choose **Virtual Private Network** > **Classic – VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.
6. On the **VPN Gateways** page, locate the row that contains the target VPN gateway and click **Modify Bandwidth** in the **Operation** column.
7. Modify the bandwidth as required.
8. Click **Submit**.

Modifying VPN Gateway Specifications

Scenario

When the number of VPN connections associated with a VPN gateway does not meet your requirements, you can modify the VPN gateway specifications.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic – VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.

5. Locate the row that contains the target VPN gateway, and choose **More** > **Modify Specifications** in the **Operation** column.
6. Enter the number of required VPN connections on the displayed page.
7. Click **Submit**.

   NOTE

   The number of required VPN connections can only be increased.

3.2.1.4 Deleting a VPN Gateway

Scenarios

If a VPN gateway is no longer required, you can delete it to release network resources as long as it has no VPN connections configured.

If it has any connections configured, delete the connections first.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Gateways**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.
5. Locate the row that contains the target VPN gateway, and choose **More** > **Delete** in the **Operation** column.

   If Enterprise Edition VPN is available for the selected region, locate the row that contains the target VPN gateway, and choose **More** > **Delete**.
6. In the displayed dialog box, click **Yes**.

3.2.2 Classic VPN Connection Management

3.2.2.1 Buying a VPN Connection

Scenarios

To connect your on-premises data center or private network to your ECSs in a VPC, you need to create a VPN connection after a VPN gateway is obtained.

Procedure

1. Log in to the management console.

2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Connections**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**.
5. On the **VPN Connections** page, click **Create VPN Connection**.
6. Configure the parameters as prompted and click **Pay Now**. **Table 3-15** describes the VPN connection parameters.

Table 3-15 Description of VPN connection parameters

Parameter	Description	Example Value
Region	Regions are geographic areas that are physically isolated from each other. The networks inside different regions are not connected to each other, so resources cannot be shared across regions. For low network latency and fast resource access, select the region nearest to your target users.	N/A
Name	Name of a VPN connection.	vpn-001
VPN Gateway	Name of the VPN gateway for which the VPN connection is created.	vpcgw-001
Local Subnet	VPC subnets that will access your on-premises network through a VPN. You can set the local subnet using either of the following methods: • **Select subnet**: Select the subnets that need to access your on-premises data center or private network. • **Specify CIDR block**: Enter the CIDR blocks that need to access your on-premises data center or private network.	192.168.1.0/24, 192.168.2.0/24
Remote Gateway	The public IP address of the gateway in your data center or on the private network. This IP address is used for communicating with your VPC.	N/A

Parameter	Description	Example Value
Remote Subnet	The subnets of your on-premises network that will access a VPC through a VPN. The remote and local subnets cannot overlap with each other. The remote subnet cannot overlap with CIDR blocks involved in existing VPC peering, Direct Connect, or Cloud Connect connections created for the local VPC.	192.168.3.0/24, 192.168.4.0/24
PSK	Private key shared by two ends of a VPN connection for negotiation. PSKs configured at both ends of the VPN connection must be the same. The PSK: • Contains 6 to 128 characters. • Can contain only: – Digits – Letters – Special characters: ~ ` ! @ # $ % ^ ( ) - _ + = [ ] { } | \ , . / : ;	Test@123
Confirm PSK	Enter the PSK again.	Test@123
Advanced Settings	• **Default**: Use default IKE and IPsec policies. • **Existing**: Use existing IKE and IPsec policies. • **Custom**: including **IKE Policy** and **IPsec Policy**, which specifies the encryption and authentication algorithms of a VPN tunnel. For details, see **Table 3-16** and **Table 3-17**.	Custom

Table 3-16 IKE policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: • MD5(Insecure. Not recommended.) • SHA1(Insecure. Not recommended.) • SHA2-256 • SHA2-384 • SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: • AES-128 • AES-192 • AES-256 • 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128
DH Algorithm	• Group 1(Insecure. Not recommended.) • Group 2(Insecure. Not recommended.) • Group 5(Insecure. Not recommended.) • Group 14 • Group 15 • Group 16 • Group 19 • Group 20 • Group 21 The default algorithm is **Group 14**.	Group 14
Version	Version of the IKE protocol. The value can be one of the following: • v1 (not recommended due to security risks) • v2 The default value is **v2**.	v2

Parameter	Description	Example Value
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **86400**.	86400

Table 3-17 IPsec policy

Parameter	Description	Example Value
Authentication Algorithm	Hash algorithm used for authentication. The following algorithms are supported: ● SHA1(Insecure. Not recommended.) ● MD5(Insecure. Not recommended.) ● SHA2-256 ● SHA2-384 ● SHA2-512 The default algorithm is **SHA2-256**.	SHA2-256
Encryption Algorithm	Encryption algorithm. The following algorithms are supported: ● AES-128 ● AES-192 ● AES-256 ● 3DES(Insecure. Not recommended.) The default algorithm is **AES-128**.	AES-128

Parameter	Description	Example Value
PFS	Algorithm used by the Perfect forward secrecy (PFS) function. • DH group 1(Insecure. Not recommended.) • DH group 2(Insecure. Not recommended.) • DH group 5(Insecure. Not recommended.) • DH group 14 • DH group 15 • DH group 16 • DH group 19 • DH group 20 • DH group 21 The default algorithm is **DH group 14**.	DH group 14
Transfer Protocol	Security protocol used in IPsec to transmit and encapsulate user data. The following protocols are supported: • AH • ESP • AH-ESP The default protocol is **ESP**.	ESP
Lifetime (s)	Lifetime of an SA, in seconds An SA will be renegotiated when its lifetime expires. The default value is **3600**.	3600

NOTE

An IKE policy specifies the encryption and authentication algorithms to be used in the negotiation phase of an IPsec tunnel. An IPsec policy specifies the protocol, encryption algorithm, and authentication algorithm to be used in the data transmission phase of an IPsec tunnel. The IKE and IPsec policies must be the same at both ends of a VPN connection. If they are different, the VPN connection cannot be set up.

3.2.2.2 Viewing a VPN Connection

Scenarios

After creating a VPN connection, you can view its details.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Connections**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**. Then, click the **VPN Connections** tab.
5. View the VPN connection information. You can also locate the row that contains the target VPN connection, and click **View Policy** in the **Operation** column to view IKE and IPsec policy details of the VPN connection.

3.2.2.3 Modifying a VPN Connection

Scenarios

A VPN connection is an encrypted communications channel established between the VPN gateway in your VPC and that in an on-premises data center. The VPN connection can be modified after creation.

⚠ CAUTION

If you modify the advanced settings, network communications may be interrupted. Exercise caution when performing this operation.

Changing the PSK only will not delete the current VPN connection. The new PSK takes effect during IKE renegotiation after the IKE lifetime expires.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Connections**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**. Then, click the **VPN Connections** tab.
5. Locate the row that contains the target VPN connection, and click **Modify** in the **Operation** column.
6. In the displayed **Modify VPN Connection** dialog box, modify parameters as required.

NOTE

The name of a VPN gateway can contain only letters, digits, underscores (_), hyphens (-), and periods (.).

7. Click **OK**.

3.2.2.4 Deleting a VPN Connection

Scenarios

If a VPN connection is no longer required, you can delete it to release network resources.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. Click in the upper left corner, select a region, and choose **Network** > **Virtual Private Network**.

   The VPN console is displayed.
4. In the navigation pane on the left, choose **Virtual Private Network** > **Classic** > **VPN Connections**.

   If Enterprise Edition VPN is available for the selected region, choose **Virtual Private Network** > **Classic**. Then, click the **VPN Connections** tab.
5. Locate the row that contains the target VPN connection, and choose **More** > **Delete** in the **Operation** column.
6. In the displayed dialog box, click **Yes**.

3.3 Monitoring

3.3.1 Monitoring VPN

Monitoring is the key to ensuring VPN performance, reliability, and availability. You can determine VPN resource usage based on monitoring data. The cloud platform provides Cloud Eye to help you obtain the running statuses of your VPNs. You can use Cloud Eye to automatically monitor VPNs in real time and manage alarms and notifications, so that you can know VPN performance metrics in a timely manner.

Reference link:

For more monitored information, see the *Cloud Eye User Guide*.

3.3.2 Metrics (S2C Enterprise Edition VPN)

Description

This section describes monitored metrics reported by VPN to Cloud Eye as well as their namespaces and dimensions. You can use the Cloud Eye management

console to query the metrics of the monitored objects and alarms generated for VPN.

Namespace

SYS.VPN

Metrics

Table 3-18 Metrics supported for Enterprise Edition VPN gateways

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
gateway_send_pkt_rate	Outbound Packet Rate	Average number of data packets leaving the cloud per second.	≥ 0	pps	N/A	Gateway	1 minute
gateway_recv_pkt_rate	Inbound Packet Rate	Average number of data packets entering the cloud per second.	≥ 0	pps	N/A	Gateway	1 minute
gateway_send_rate	Outbound Bandwidth	Average volume of traffic leaving the cloud per second.	0-1	bps	1024(IEC)	Gateway	1 minute
gateway_recv_rate	Inbound Bandwidth	Average volume of traffic entering the cloud per second.	0-1	bps	1024(IEC)	Gateway	1 minute
gateway_send_rate_usage	Outbound Bandwidth Usage	Bandwidth utilization for traffic leaving the cloud.	0-100	percentage(%)	N/A	Gateway	1 minute

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
gateway_recv_rate_usage	Inbound Bandwidth Usage	Bandwidth utilization for traffic entering the cloud.	0-100	percentage(%)	N/A	Gateway	1 minute
gateway_connection_num	Number of Connections	Number of VPN connections.	≥ 0	count	N/A	Gateway	1 minute

Table 3-19 Enterprise Edition VPN connection metrics

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
tunnel_average_latency	Average Tunnel RTT	Average round-trip time on the tunnel between the VPN gateway and customer gateway.	0–5000	ms	N/A	VPN connection	15s
tunnel_max_latency	Maximum Tunnel RTT	Maximum round-trip time on the tunnel between the VPN gateway and customer gateway.	0–5000	ms	N/A	VPN connection	15s
tunnel_packet_loss_rate	Tunnel Packet Loss Rate	Packet loss rate on the tunnel between the VPN gateway and customer gateway.	0–100	percentage(%)	N/A	VPN connection	15s

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
link_average_latency	Average Link RTT	Average round-trip time on the physical link between the VPN gateway and customer gateway.	0–5000	ms	N/A	VPN connection	15s
link_max_latency	Maximum Link RTT	Maximum round-trip time on the physical link between the VPN gateway and customer gateway.	0–5000	ms	N/A	VPN connection	15s
link_packet_loss_rate	Link Packet Loss Rate	Packet loss rate on the physical link between the VPN gateway and customer gateway.	0–100	percentage(%)	N/A	VPN connection	15s
connection_status	VPN Connection Status	Status of a VPN connection: • **0**: not connected • **1**: connected • **2**: unknown	0, 1, or 2	N/A	N/A	VPN connection	1 minute
recv_pkt_rate	Packet Receive Rate	Average number of data packets received per second.	≥ 0	pps	N/A	VPN connection	1 minute
send_pkt_rate	Packet Send Rate	Average number of data packets sent per second.	≥ 0	pps	N/A	VPN connection	1 minute
recv_rate	Traffic Receive Rate	Average volume of traffic received per second.	0-1	bps	1024(IEC)	VPN connection	1 minute
send_rate	Traffic Send Rate	Average volume of traffic sent per second.	0-1	bps	1024(IEC)	VPN connection	1 minute

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
sa_send_pkt_rate	SA Packet Send Rate	Average number of data packets sent over an SA per second.	≥ 0	pps	N/A	SA of a VPN connection	1 minute
sa_recv_pkt_rate	SA Packet Receive Rate	Average number of data packets received over an SA per second.	≥ 0	pps	N/A	SA of a VPN connection	1 minute
sa_recv_rate	SA Traffic Receive Rate	Average volume of traffic received over an SA per second.	0-1	bps	1024(IEC)	SA of a VPN connection	1 minute
sa_send_rate	SA Traffic Send Rate	Average volume of traffic sent over an SA per second.	0-1	bps	1024(IEC)	SA of a VPN connection	1 minute

Dimensions

key	Value
evpn_connection_id	Enterprise Edition S2C VPN connection
evpn_sa_id	SAs of an Enterprise Edition S2C VPN connection
evpn_gateway_id	Enterprise Edition S2C VPN gateway

3.3.3 Metrics (S2C Classic VPN)

Description

This section describes monitored metrics reported by VPN to Cloud Eye as well as their namespaces and dimensions. You can use the Cloud Eye management console to query the metrics of the monitored objects and alarms generated for VPN.

Namespace

SYS.VPC

Metrics

Table 3-20 Metrics supported for Classic VPN bandwidth

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
upstream_bandwidth	Outbound Bandwidth	Network rate of outbound traffic (previously called "Upstream Bandwidth").	≥ 0	bit/s	1000(SI)	Bandwidth or EIP	1 minute
downstream_bandwidth	Inbound Bandwidth	Network rate of inbound traffic (previously called "Downstream Bandwidth").	≥ 0	bit/s	1000(SI)	Bandwidth or EIP	1 minute
upstream_bandwidth_usage	Outbound Bandwidth Usage	Usage of outbound bandwidth. Outbound bandwidth usage = Outbound bandwidth/ Purchased bandwidth	0-100	%	N/A	Bandwidth or EIP	1 minute

Metric ID	Metric Name	Description	Value Range	Unit	Conversion Rule	Monitored Object (Dimension)	Monitoring Interval (Raw Data)
downstream_bandwidth_usage	Inbound Bandwidth Usage	Usage of inbound bandwidth. Inbound bandwidth usage = Inbound bandwidth/ Purchased bandwidth **NOTE** • Up to 10 Mbit/s inbound bandwidth is provided for some sites that purchase an inbound bandwidth of less than 10 Mbit/s. As such, the inbound bandwidth usage may be greater than 100%. • If you change the bandwidth of an EIP in use, there is a delay of 5–10 minutes for the metrics to update for the new bandwidth.	0-100	%	N/A	Bandwidth or EIP	1 minute
up_stream	Outbound Traffic	Outbound network traffic (previously called "Upstream Traffic")	≥ 0	Byte	1000(SI)	Bandwidth or EIP	1 minute
down_stream	Inbound Traffic	Inbound network traffic (previously called "Downstream Traffic")	≥ 0	Byte	1000(SI)	Bandwidth or EIP	1 minute

Dimensions

key	Value
publicip_id	EIP ID

key	Value
bandwidth_id	Bandwidth ID

3.3.4 Viewing Metrics

Scenarios

View the VPN connection status and usages of bandwidth and EIP. You can view data of the last 15 minutes, last 30 minutes, last 1 hour, last 2 hours, last 3 hours, last 12 hours, last 24 hours, last 7 days, last 30 days, or a custom time range.

Support for Metrics

Table 3-21 Support for metrics

Metric Name	Support	Enabled by Default?
VPN Connection Status	Supported by both Enterprise Edition VPN and Classic VPN	Yes
• Average Link RTT • Maximum Link RTT • Link Packet Loss Rate • Packet Receive Rate • Packet Send Rate • Traffic Receive Rate • Traffic Send Rate • SA Packet Receive Rate • SA Packet Send Rate • SA Traffic Receive Rate • SA Traffic Send Rate	Supported only by Enterprise Edition VPN	No You can click the name of a VPN connection and add a health check item on the **Summary** tab page.

Metric Name	Support	Enabled by Default?
• Average Tunnel RTT • Maximum Tunnel RTT • Tunnel Packet Loss Rate	Supported only by Enterprise Edition VPN	Yes Private network monitoring metrics are supported only when a VPN connection uses the static routing mode and has NQA enabled.

Viewing VPN Gateway Metrics

- Viewing metrics on the Cloud Eye console
  a. Log in to the management console.
  b. Click in the upper left corner and select the desired region and project.
  c. In the upper left corner of the page, click and choose **Management & Governance** > **Cloud Eye**.
  d. Choose **Cloud Service Monitoring** > **Virtual Private Network**.
  e. Click **View Metric** in the **Operation** column of the target VPN gateway to view its status.
     You can view data of the last 15 minutes, last 30 minutes, last 1 hour, last 2 hours, last 3 hours, last 12 hours, last 24 hours, last 7 days, last 30 days, or a custom time range.

3.3.5 Creating a Monitoring Alarm Rule

Scenarios

You can create monitoring alarm rules to customize monitored objects and notification policies, so that you can be well-informed of the VPN service status.

Procedure

1. Log in to the management console.
2. Click in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click and choose **Management & Governance** > **Cloud Eye**.
4. Choose **Cloud Service Monitoring** > **Virtual Private Network VPN**, and configure alarm rules for different types of alarms as required.
   – Alarms related to VPN gateways in S2C Enterprise Edition VPN: Click the **S2C VPN Gateway** tab, and click **Create Alarm Rule** in the **Operation** column of a VPN gateway.
   – Alarms related to VPN connections in S2C Enterprise Edition VPN: Click the **S2C VPN Connection** tab, and click **Create Alarm Rule** in the **Operation** column of a VPN connection.

5. Click **Create**.

   After the VPN monitoring alarm rule is configured, if you have enabled alarm notifications and configured related parameters, you will receive notifications once an alarm is triggered.

   NOTE

   For more information about VPN alarm rules, see the *Cloud Eye User Guide*.

3.4 Audit

3.4.1 Key Operations That Can Be Recorded by CTS

Table 3-22 Operations related to S2C Enterprise Edition VPN that can be recorded by CTS

Operation	Resource Type	Trace Name
Creating a customer gateway	customer-gateway	createCgw
Updating a customer gateway	customer-gateway	updateCgw
Deleting a customer gateway	customer-gateway	deleteCgw
Creating a VPN gateway	vpn-gateway	createVgw
Updating a VPN gateway	vpn-gateway	updateVgw
Deleting a VPN gateway	vpn-gateway	deleteVgw
Updating the specification of a pay-per-use VPN gateway	vpn-gateway	updatePostpaidVgwSpecification
Creating a VPN connection	vpn-connection	createVpnConnection
Updating a VPN connection	vpn-connection	updateVpnConnection
Deleting a VPN connection	vpn-connection	deleteVpnConnection
Creating a resource tag	instance	batchCreateResourceTags

Operation	Resource Type	Trace Name
Deleting a resource tag	instance	batchDeleteResourceTags
Querying the customer gateway list	customer-gateway	listCgws
Querying a customer gateway	customer-gateway	showCgw
Querying resource tags	instance	showResourceTags
Querying project tags	instance	listProjectTags
Querying resource instances by tag	instance	listResourcesByTags
Querying the number of resource instances by tag	instance	countResourcesByTags
Querying a VPN gateway	vpn-gateway	showVgw
Querying the AZs of VPN gateways	availability_zone	listAvailabilityZones
Querying the VPN connection list	vpn-connection	listVpnConnections
Querying a VPN connection	vpn-connection	showVpnConnection
Querying the VPN gateway list	vpn-gateway	listVgws
Querying a VPN connection monitor	connection-monitor	showConnectionMonitor
Querying the VPN connection monitor list	connection-monitor	listConnectionMonitors
Querying quotas of a specified tenant	quota	showQuotasInfo
Querying VPN connection logs	vpn-connection	showVpnConnectionLog

Operation	Resource Type	Trace Name
Creating VPN connections in batches	vpn-connection	batchCreateVpnConnection
Resetting a VPN connection	vpn-connection	resetVpnConnection
Upgrading an S2C VPN gateway	vpn-gateway	upgradeVpnGateway
Querying the S2C VPN gateway task list	job	listVpnGatewayJobs
Deleting an S2C VPN gateway task	job	deleteVpnGatewayJob
Creating a VPN connection monitor	connection-monitor	createConnectionMonitor
Deleting a VPN connection monitor	connection-monitor	deleteConnectionMonitor

3.4.2 Querying CTS Traces

After you enable CTS and the management tracker is created, CTS starts recording operations performed on VPN resources. You can view the operation records in the last seven days on the CTS console.

Reference link:

For details about how to view audit logs, see the *Cloud Trace Service User Guide*.

3.5 Permissions Management

3.5.1 Creating a User and Granting VPN Permissions

Use the **Identity and Access Management (IAM)** service to implement fine-grained permissions control over your VPN resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing VPN resources.
- Grant users only the permissions required to perform a given task based on their job responsibilities.
- Grant the permission to perform professional and efficient O&M on your VPN resources to other accounts or cloud services.

If your account meets your permissions requirements, you can skip this section.

This section describes the procedure for granting permissions (see **Figure 3-2**).

Prerequisites

You have learned about the permissions supported by VPN, and determined the permissions to be granted to a user group. Before granting permissions of other services, learn about all **system-defined permissions** supported by IAM.

Process Flow

Figure 3-2 Process of granting VPN permissions

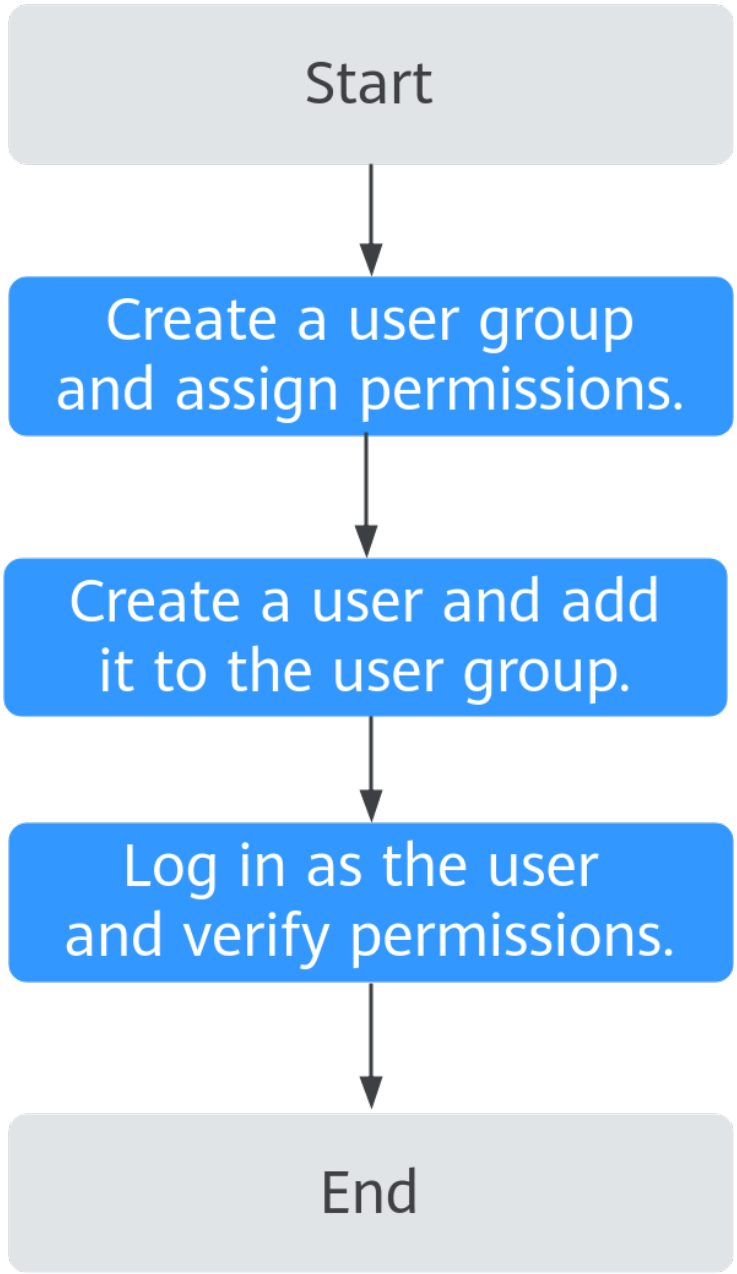

1. **Create a user group and assign permissions** to it.

   Create a user group on the IAM console, and attach the **VPN FullAccess** policy to the group.

2. **Create a user and add it to the user group**.

   Create a user on the IAM console and add the user to the group created in **1**.

3. **Log in** and verify permissions.

   Log in to the management console as the created user. Switch to the authorized region and verify the permissions.

   - Click ☰ in the upper left corner, and choose **Network** > **Virtual Private Network** > **Enterprise – VPN Gateways**. On the **S2C VPN Gateways** tab page, click **Create S2C VPN Gateway** in the upper right corner to create a VPN gateway. If the VPN gateway is successfully created, the **VPN FullAccess** policy has already taken effect.
   - Click ☰ in the upper left corner, and choose **Network** > **Virtual Private Network** > **Classic**. Click **Buy VPN Gateway** to create a VPN gateway. If the VPN gateway is successfully created, the **VPN FullAccess** policy has already taken effect.
   - Click ☰ in the upper left corner, and select any service except the VPN service. Assume that the current policy contains only **VPN FullAccess**. If a message appears indicating that you have insufficient permissions to access the service, the **VPN FullAccess** policy has already taken effect.

3.5.2 VPN Custom Policies

Custom policies can be created to supplement the system-defined policies of VPN.

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**. The following section contains examples of common VPN custom policies.

Example VPN custom policy

- Example 1: Grant permission to delete VPN gateways.

  You need to add the following dependent actions. If they are not added, an exception may occur.

```
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "vpn:vpnGateways:delete"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "vpc:subNetworkInterfaces:update",
```

```
            "vpc:routeTables:update",
            "vpc:subnets:delete",
            "vpc:publicIps:list",
            "vpc:publicIps:delete",
            "vpc:vpcs:get",
            "vpc:routeTables:get",
            "vpc:ports:get",
            "vpc:publicIps:update",
            "vpc:subnets:get",
            "vpc:bandwidths:list",
            "vpc:publicIps:get",
            "vpc:vpcs:list"
         ]
      },
      {
         "Effect": "Allow",
         "Action": [
            "er:instances:get",
            "er:instances:list"
         ]
      }
   ]
}
```

- Example 2: Deny VPN connection deletion.

  A policy with only "Deny" permissions must be used together with other policies. If the permissions granted to an IAM user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

  The following method can be used if you need to assign permissions of the **VPN FullAccess** policy to a user but also forbid the user from deleting VPN connections. Create a custom policy for denying VPN connection deletion, and assign both policies to the group the user belongs to. Then the user can perform all operations on VPN except deleting VPN connections. The following is an example of a deny policy:

  ```
  {
     "Version": "1.1",
     "Statement": [
        {
           "Effect": "Deny",
           "Action": [
              "vpn:vpnConnections:delete"
           ]
        }
     ]
  }
  ```

- Example 3: defining multiple actions in a policy

  A custom policy can contain the actions of one or multiple services that are of the same type (global or project-level). The following is an example policy containing multiple actions.

  ```
  {
     "Version": "1.1",
     "Statement": [
        {
           "Effect": "Allow",
           "Action": [
              "vpn:vpnGateways:create",
              "vpn:vpnConnections:create",
              "vpn:customerGateways:create"
           ]
        },
        {
           "Effect": "Deny",
           "Action": [
              "vpn:vpnGateways:delete",
  ```

```
            "vpn:vpnConnections:delete",
            "vpn:customerGateways:create"
        ]
    },
    {
        "Effect": "Allow",
        "Action": [
            "vpc:vpcs:list",
            "vpc:subnets:get"
        ]
    }
  ]
}
```

4 Best Practice

4.1 S2C Enterprise Edition VPN

4.1.1 Connecting an On-premises Data Center to a VPC on the Cloud Through VPN (Active-Active Mode)

4.1.1.1 Overview

Scenario

VPN can be used to enable communication between an on-premises data center and ECSs in a VPC.

Networking

In this example, a group of VPN connections are set up between an on-premises data center and a VPC to ensure network reliability. If one VPN connection fails, traffic is automatically switched to the other VPN connection, ensuring service continuity.

Figure 4-1 Networking diagram

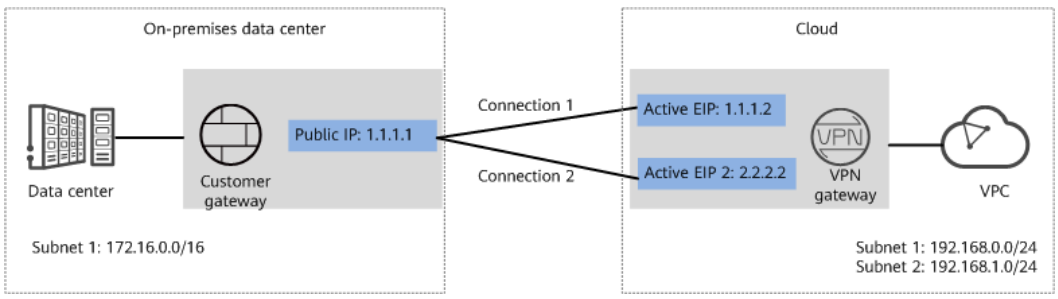

Solution Advantages

- A VPN gateway provides two IP addresses to establish dual independent VPN connections with a customer gateway. If one VPN connection fails, traffic can be quickly switched to the other VPN connection.
- Active-active VPN gateways can be deployed in different AZs to ensure AZ-level high availability.

Limitations and Constraints

- The local and customer subnets of the VPN gateway cannot be the same. That is, the VPC subnet and the data center subnet to be interconnected cannot be the same.
- The IKE policy, IPsec policy, and PSK of the VPN gateway must be the same as those of the customer gateway.
- The local and remote interface address configurations on the VPN gateway and customer gateway are reversed.
- The security groups associated with ECSs in the VPC permit access from and to the on-premises data center.

4.1.1.2 Planning Networks and Resources

Data Plan

Table 4-1 Data plan

Category	Item	Data
VPC	Subnet that needs to access the on-premises data center	• 192.168.0.0/24 • 192.168.1.0/24
VPN gateway	Interconnection subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses. 192.168.2.0/24
	HA mode	Active-active
	EIP	EIPs are automatically generated when you buy them. By default, a VPN gateway uses two EIPs. In this example, the EIPs are as follows: • Active EIP: 1.1.1.2 • Active EIP 2: 2.2.2.2

Category	Item	Data
VPN connection	Tunnel interface addresses under **Connection 1's Configuration**	IP addresses used to establish an IPsec tunnel between a VPN gateway and a customer gateway. At the two ends of the IPsec tunnel, the configured local and remote tunnel interface addresses must be reversed. ● Local tunnel interface address: 169.254.70.1/30 ● Customer tunnel interface address: 169.254.70.2/30
	Tunnel interface addresses under **Connection 2's Configuration**	● Local tunnel interface address: 169.254.71.1/30 ● Customer tunnel interface address: 169.254.71.2/30
On-premises data center	Subnet that needs to access the VPC	172.16.0.0/16
Customer gateway	Public IP address	This public IP address is assigned by a carrier. In this example, the public IP address is as follows: 1.1.1.1
IKE and IPsec policies	PSK	Test@123
	IKE policy	● Version: v2 ● Authentication algorithm: SHA2-256 ● Encryption algorithm: AES-128 ● DH algorithm: Group 15 ● Lifetime (s): 86400 ● Local ID: IP address ● Peer ID: IP address
	IPsec policy	● Authentication algorithm: SHA2-256 ● Encryption algorithm: AES-128 ● PFS: DH Group15 ● Transfer protocol: ESP ● Lifetime (s): 3600

4.1.1.3 Procedure

Prerequisites

- Cloud side
  - A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
  - Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.
  - An enterprise router has been created if you want to use it to connect to a VPN gateway. For details, see the enterprise router documentation.
- Data center side
  - IPsec has been configured on the VPN device in the on-premises data center.

Procedure

VPNs support static routing mode, BGP routing mode, and policy-based mode. The following uses the static routing mode as an example.

Step 1 Log in to the management console.

Step 2 Click **Service List** and choose **Networking** > **Virtual Private Network**.

Step 3 Configure a VPN gateway.

1. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click **Buy S2C VPN Gateway**.
2. Set parameters as prompted.

   Table 4-2 only describes the key parameters for creating a VPN gateway.

Table 4-2 Description of VPN gateway parameters

Parameter	Description	Value
Name	Name of a VPN gateway.	vpngw-001
Network Type	Select **Public network**.	Public network
Associate With	Select **VPC**. If the VPN gateway is associated with an enterprise router, select **Enterprise Router**.	VPC
Enterprise Router	Specify the associated enterprise router only when **Associate With** is set to **Enterprise Router**.	er-001

Parameter	Description	Value
VPC	VPC to which the interconnection subnet belongs. When **Associate With** is set to **Enterprise Router**, the associated enterprise router can be located in the VPC or not.	vpc-001(192.168.0.0/16)
Interconnection Subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	192.168.2.0/24
Local Subnet	This parameter is available only when **Associate With** is set to **VPC**. – Enter CIDR block Enter the subnet that needs to access the on-premises data center. The subnet can belong to the associated VPC or not. – Select subnet Select a subnet that belongs to the associated VPC and needs to access the on-premises data center.	192.168.0.0/24,192.168.1.0/24
BGP ASN	BGP AS number.	64512
HA Mode	Select **Active-active**.	Active-active
Active EIP	EIP 1 used by the VPN gateway to access the on-premises data center.	1.1.1.2
Active EIP 2	EIP 2 used by the VPN gateway to access the on-premises data center.	2.2.2.2

Step 4 Configure the customer gateway.

1. Choose **Virtual Private Network** > **Enterprise – Customer Gateways**, and click **Create Customer Gateway**.
2. Set parameters as prompted.

   Table 4-3 only describes the key parameters for creating a customer gateway.

Table 4-3 Description of customer gateway parameters

Parameter	Description	Value
Name	Name of a customer gateway.	cgw-fw

Parameter	Description	Value
Identifier	IP address used by the customer gateway to communicate with the VPN gateway. Ensure that UDP port 4500 is permitted on the customer gateway device in the on-premises data center.	1.1.1.1

Step 5 Configure VPN connections.

1. Choose **Virtual Private Network** > **Enterprise – VPN Connections**, and click **Create VPN Connection**.
2. Set VPN connection parameters and click **Buy Now**.

   Table 4-4 describes the key parameters for creating VPN connections.

Table 4-4 Description of VPN connection parameters

Parameter	Description	Value
Name	VPN connection name.	vpn-001
VPN Gateway	VPN gateway for which VPN connections are created.	vpngw-001
VPN Gateway IP of Connection 1	Active EIP bound to the VPN gateway.	1.1.1.2
Customer Gateway of Connection 1	Customer gateway of connection 1.	1.1.1.1
VPN Gateway IP of Connection 2	Active EIP 2 bound to the VPN gateway.	2.2.2.2
Customer Gateway of Connection 2	Customer gateway of connection 2.	1.1.1.1
VPN Type	Select **Static routing**.	Static routing

Parameter	Description	Value
Customer Subnet	Subnet in the on-premises data center that needs to access the VPC. – A customer subnet cannot be included in any local subnet or any subnet of the VPC to which the VPN gateway is attached. – Reserved VPC CIDR blocks such as 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 cannot be used as customer subnets. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.	172.16.0.0/16
Connection 1's Configuration	Configure the IP address assignment mode of tunnel interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, and policies for connection 1.	*Set parameters based on the site requirements.*
Interface IP Address Assignment	– Manually specify In this example, select **Manually specify**. – Automatically assign	Manually specify
Local Tunnel Interface Address	Tunnel interface IP address of the VPN gateway.	169.254.70.1/30
Customer Tunnel Interface Address	Tunnel interface IP address of the customer gateway device.	169.254.71.1/30
Link Detection	Whether to enable route reachability detection in multi-link scenarios. When NQA is enabled, ICMP packets are sent for detection and your device needs to respond to these ICMP packets.	**NQA** enabled
PSK, Confirm PSK	The value must be the same as the PSK configured on the customer gateway device.	Test@123

Parameter	Description	Value
Policy Settings	The policy settings must be the same as those on the customer gateway device.	Default
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**. **NOTE** If you disable **Same as that of connection 1**, you are advised to use the same settings as connection 1 for connection 2, except the local and customer tunnel interface addresses.	Disabled
Local Tunnel Interface Address	Tunnel IP address of the VPN gateway.	169.254.70.2/30
Customer Tunnel Interface Address	Tunnel IP address of the customer gateway.	169.254.71.2/30

Step 6 Configure the customer gateway device.

The configuration procedures may vary according to the type of the customer gateway device.

----End

Verification

- About 5 minutes later, check states of the VPN connections.

  Choose **Virtual Private Network** > **Enterprise – VPN Connections**. The states of the two VPN connections are both **Normal**.
- Verify that servers in the on-premises data center and ECSs in the VPC subnet can ping each other.

4.1.2 Connecting an On-premises Data Center to a VPC on the Cloud Through VPN (Active/Standby Mode)

4.1.2.1 Overview

Scenario

VPN can be used to enable communication between an on-premises data center and ECSs in a VPC.

Networking

In this example, two VPN connections working in active/standby mode are set up between an on-premises data center and a VPC to ensure network reliability. If

one VPN connection fails, traffic is automatically switched to the other VPN connection, ensuring service continuity.

Figure 4-2 Networking diagram

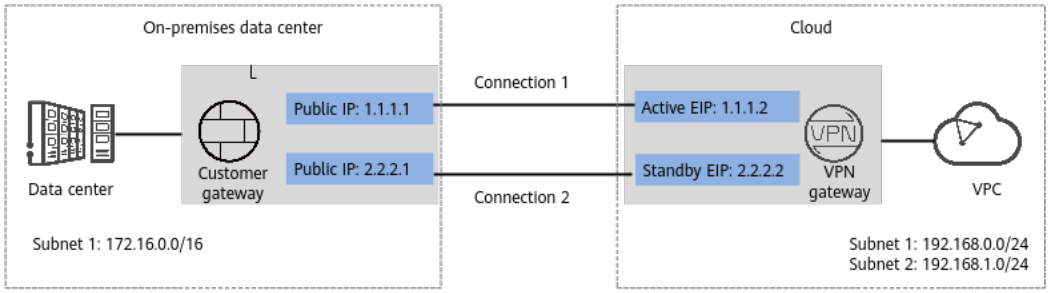

Solution Advantages

- A VPN gateway provides two IP addresses to establish dual independent VPN connections with a customer gateway. If one VPN connection fails, traffic can be quickly switched to the other VPN connection.
- Active/Standby mode: A VPN gateway communicates with a customer gateway through the active connection. If the active connection fails, traffic is automatically switched to the standby VPN connection. After the fault is rectified, traffic is switched back to the original active VPN connection. Traffic leaving the cloud is preferentially transmitted through the active EIP, allowing you to determine the VPN connection through which traffic is transmitted.

Limitations and Constraints

- The local and customer subnets of the VPN gateway cannot be the same. That is, the VPC subnet and the data center subnet to be interconnected cannot be the same.
- The IKE policy, IPsec policy, and PSK of the VPN gateway must be the same as those of the customer gateway.
- The local and remote interface address configurations on the VPN gateway and customer gateway are reversed.
- The security groups associated with ECSs in the VPC permit access from and to the on-premises data center.

4.1.2.2 Planning Networks and Resources

Data Plan

Table 4-5 Data plan

Category	Item	Data
VPC	Subnet that needs to access the on-premises data center	• 192.168.0.0/24 • 192.168.1.0/24

Category	Item	Data
VPN gateway	Interconnection subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses. 192.168.2.0/24
	HA mode	Active/Standby
	EIP	EIPs are automatically generated when you buy them. By default, a VPN gateway uses two EIPs. In this example, the EIPs are as follows: ● Active EIP: 1.1.1.2 ● Standby EIP: 2.2.2.2
VPN connection	Tunnel interface addresses under **Connection 1's Configuration**	IP addresses used to establish an IPsec tunnel between a VPN gateway and a customer gateway. At the two ends of the IPsec tunnel, the configured local and remote tunnel interface addresses must be reversed. ● Local tunnel interface address: 169.254.70.1/30 ● Customer tunnel interface address: 169.254.70.2/30
	Tunnel interface addresses under **Connection 2's Configuration**	● Local tunnel interface address: 169.254.71.1/30 ● Customer tunnel interface address: 169.254.71.2/30
On-premises data center	Subnet that needs to access the VPC	172.16.0.0/16
Customer gateway	Public IP address	This public IP address is assigned by a carrier. In this example, the public IP addresses are as follows: ● 1.1.1.1 ● 2.2.2.1
IKE and IPsec policies	PSK	Test@123

Category	Item	Data
	IKE policy	• Version: v2 • Authentication algorithm: SHA2-256 • Encryption algorithm: AES-128 • DH algorithm: Group 15 • Lifetime (s): 86400 • Local ID: IP address • Peer ID: IP address
	IPsec policy	• Authentication algorithm: SHA2-256 • Encryption algorithm: AES-128 • PFS: DH Group15 • Transfer protocol: ESP • Lifetime (s): 3600

4.1.2.3 Procedure

Prerequisites

- Cloud side
  - A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
  - Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.
  - An enterprise router has been created if you want to use it to connect to a VPN gateway. For details, see the enterprise router documentation.
- Data center side
  - IPsec has been configured on the VPN device in the on-premises data center.

Procedure

VPNs support static routing mode, BGP routing mode, and policy-based mode. The following uses the static routing mode as an example.

Step 1 Log in to the management console.

Step 2 Click **Service List** and choose **Networking** > **Virtual Private Network**.

Step 3 Configure a VPN gateway.

1. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click **Buy S2C VPN Gateway**.
2. Set parameters as prompted.

   Table 4-6 only describes the key parameters for creating a VPN gateway.

Table 4-6 Description of VPN gateway parameters

Parameter	Description	Value
Name	Name of a VPN gateway.	vpngw-001
Network Type	Select **Public network**.	Public network
Associate With	Select **VPC**. If the VPN gateway is associated with an enterprise router, select **Enterprise Router**.	VPC
Enterprise Router	Specify the associated enterprise router only when **Associate With** is set to **Enterprise Router**.	er-001
VPC	VPC to which the interconnection subnet belongs. When **Associate With** is set to **Enterprise Router**, the associated enterprise router can be located in the VPC or not.	vpc-001(192.168.0.0/16)
Interconnection Subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	192.168.2.0/24
Local Subnet	This parameter is available only when **Associate With** is set to **VPC**. – Enter CIDR block Enter the subnet that needs to access the on-premises data center. The subnet can belong to the associated VPC or not. – Select subnet Select a subnet that belongs to the associated VPC and needs to access the on-premises data center.	192.168.0.0/24,192.168.1.0/24
BGP ASN	BGP AS number.	64512
HA Mode	Select **Active/Standby**.	Active/Standby
Active EIP	Active EIP used by the VPN gateway to access the on-premises data center.	1.1.1.2
Standby EIP	Standby EIP used by the VPN gateway to access the on-premises data center.	2.2.2.2

Step 4 Configure the customer gateway.

1. Choose **Virtual Private Network** > **Enterprise – Customer Gateways**, and click **Create Customer Gateway**.
2. Set parameters as prompted.

   Table 4-7 only describes the key parameters for creating a customer gateway.

Table 4-7 Description of customer gateway parameters

Parameter	Description	Value
Name	Name of a customer gateway.	cgw-fw
Identifier	IP address used by the customer gateway to communicate with the VPN gateway. Ensure that UDP port 4500 is permitted on the customer gateway device in the on-premises data center.	1.1.1.1

Step 5 Configure VPN connections.

1. Choose **Virtual Private Network** > **Enterprise – VPN Connections**, and click **Create VPN Connection**.
2. Set VPN connection parameters and click **Buy Now**.

   Table 4-8 only describes the key parameters for creating VPN connections.

Table 4-8 Description of VPN connection parameters

Parameter	Description	Value
Name	VPN connection name.	vpn-001
VPN Gateway	VPN gateway for which VPN connections are created.	vpngw-001
VPN Gateway IP of Connection 1	Active EIP of the VPN gateway.	1.1.1.2
Customer Gateway of Connection 1	Customer gateway of connection 1.	1.1.1.1
VPN Gateway IP of Connection 2	Standby EIP of the VPN gateway.	2.2.2.2
Customer Gateway of Connection 2	Customer gateway of connection 2.	2.2.2.1
VPN Type	Select **Static routing**.	Static routing

Parameter	Description	Value
Customer Subnet	Subnet in the on-premises data center that needs to access the VPC. – A customer subnet cannot be included in any local subnet or any subnet of the VPC to which the VPN gateway is attached. – Reserved VPC CIDR blocks such as 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 cannot be used as customer subnets. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.	172.16.0.0/16
Connection 1's Configuration	Configure the IP address assignment mode of tunnel interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, and policies for connection 1.	*Set parameters based on the site requirements.*
Interface IP Address Assignment	– Manually specify In this example, select **Manually specify**. – Automatically assign	Manually specify
Local Tunnel Interface Address	Tunnel interface IP address of the VPN gateway.	169.254.70.1/30
Customer Tunnel Interface Address	Tunnel interface IP address of the customer gateway device.	169.254.70.2/30
Link Detection	Whether to enable route reachability detection in multi-link scenarios. When NQA is enabled, ICMP packets are sent for detection and your device needs to respond to these ICMP packets.	**NQA** enabled
PSK, Confirm PSK	The value must be the same as the PSK configured on the customer gateway device.	Test@123

Parameter	Description	Value
Policy Settings	The policy settings must be the same as those on the customer gateway device.	Default
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**. **NOTE** If you disable **Same as that of connection 1**, you are advised to use the same settings as connection 1 for connection 2, except the local and customer tunnel interface addresses.	Disabled
Local Tunnel Interface Address	Tunnel IP address of the VPN gateway.	169.254.71.1/30
Customer Tunnel Interface Address	Tunnel IP address of the customer gateway.	169.254.71.2/30

Step 6 Configure the customer gateway device.

The configuration procedures may vary according to the type of the customer gateway device.

----End

Verification

- About 5 minutes later, check states of the VPN connections.

  Choose **Virtual Private Network** > **Enterprise – VPN Connections**. The states of the two VPN connections are both **Available**.
- Verify that servers in the on-premises data center and ECSs in the VPC subnet can ping each other.

4.1.3 Connecting Multiple On-premises Branch Networks Through a VPN Hub

4.1.3.1 Overview

Scenario

To meet service requirements, enterprise A needs to implement communication between its two on-premises data centers.

Networking

Figure 4-3 shows the networking where the VPN service is used to connect the two on-premises data centers.

Figure 4-3 Networking diagram

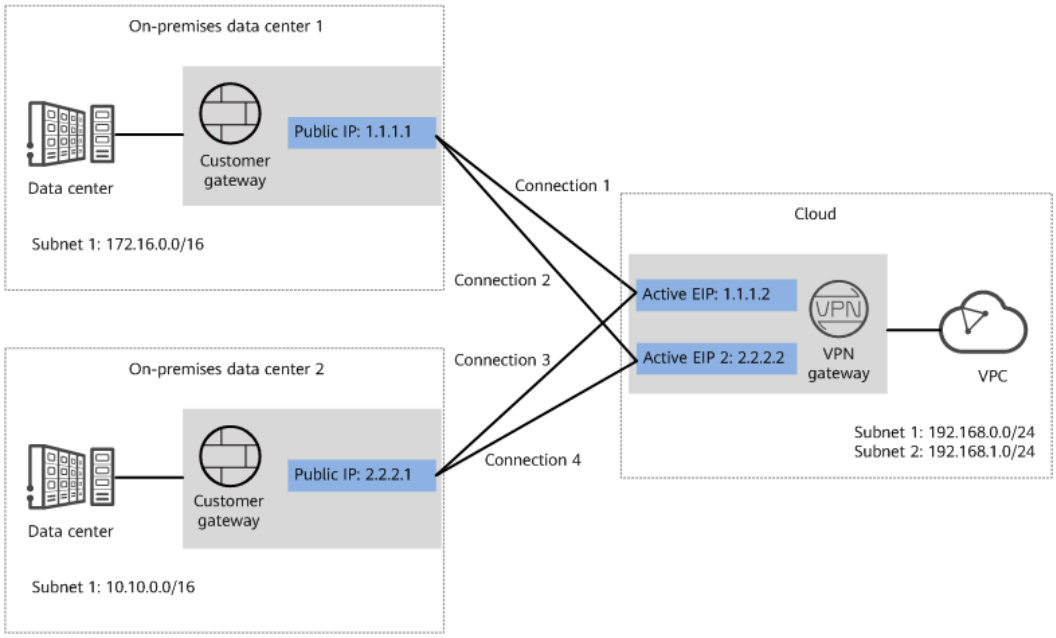

Solution Advantages

- A VPN gateway on the cloud can function as a VPN hub to enable communication between on-premises branch sites. This eliminates the need to configure VPN connections between every two sites.
- A VPN gateway provides two IP addresses to establish dual independent VPN connections with each customer gateway. If one VPN connection fails, traffic can be quickly switched to the other VPN connection, ensuring reliability.

Limitations and Constraints

- The local and customer subnets of the VPN gateway cannot be the same. That is, the VPC subnet and the data center subnet to be interconnected cannot be the same.
- The IKE policy, IPsec policy, and PSK of the VPN gateway must be the same as those of the customer gateway.
- The local and remote interface address configurations on the VPN gateway and customer gateway are reversed.
- The security groups associated with ECSs in the VPC permit access from and to the on-premises data center.

4.1.3.2 Planning Networks and Resources

Data Plan

Table 4-9 Data plan

Category	Item	Data
VPC	Subnet that needs to access the on-premises data centers	• 192.168.0.0/24 • 192.168.1.0/24
VPN gateway	Interconnection subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses. 192.168.2.0/24
	HA Mode	Active-active
	EIP	EIPs are automatically generated when you buy them. By default, a VPN gateway uses two EIPs. In this example, the EIPs are as follows: • Active EIP: 1.1.1.2 • Active EIP 2: 2.2.2.2
On-premises data center 1	Subnet that needs to access the VPC	172.16.0.0/16
Customer gateway in on-premises data center 1	Public IP address	This public IP address is assigned by a carrier. In this example, the public IP address is as follows: 1.1.1.1
VPN connections of on-premises data center 1	Tunnel interface addresses under **Connection 1's Configuration**	IP addresses used to establish an IPsec tunnel between a VPN gateway and a customer gateway. At the two ends of the IPsec tunnel, the configured local and remote tunnel interface addresses must be reversed. • Local tunnel interface address: 169.254.70.1/30 • Customer tunnel interface address: 169.254.70.2/30

Category	Item	Data
	Tunnel interface addresses under **Connection 2's Configuration**	• Local tunnel interface address: 169.254.71.1/30 • Customer tunnel interface address: 169.254.71.2/30
On-premises data center 2	Subnet that needs to access the VPC	10.10.0.0/16
Customer gateway in on-premises data center 2	Public IP address	This public IP address is assigned by a carrier. In this example, the public IP address is as follows: 2.2.2.1
VPN connections of on-premises data center 2	Tunnel interface addresses under **Connection 1's Configuration**	IP addresses used to establish an IPsec tunnel between a VPN gateway and a customer gateway. At the two ends of the IPsec tunnel, the configured local and remote tunnel interface addresses must be reversed. • Local tunnel interface address: 169.254.72.1/30 • Customer tunnel interface address: 169.254.72.2/30
	Tunnel interface addresses under **Connection 2's Configuration**	• Local tunnel interface address: 169.254.73.1/30 • Customer tunnel interface address: 169.254.73.2/30
IKE and IPsec policies	PSK	Test@123
	IKE policy	• Authentication algorithm: SHA2-256 • Encryption algorithm: AES-128 • DH algorithm: Group 15 • Version: v2 • Lifetime (s): 86400 • Local ID: IP address • Peer ID: IP address

Category	Item	Data
	IPsec policy	• Authentication algorithm: SHA2-256 • Encryption algorithm: AES-128 • PFS: DH Group15 • Transfer protocol: ESP • Lifetime (s): 3600

4.1.3.3 Procedure

Prerequisites

- Cloud side
  - A VPC has been created. For details about how to create a VPC, see the *Virtual Private Cloud User Guide*.
  - Security group rules have been configured for the VPC, and ECSs can communicate with other devices on the cloud. For details about how to configure security group rules, see the *Virtual Private Cloud User Guide*.
  - An enterprise router has been created if you want to use it to connect to a VPN gateway. For details, see the enterprise router documentation.
- Data center side
  - IPsec has been configured on the VPN devices in the two on-premises data centers.
  - The remote subnets of the VPN device in on-premises data center 1 must contain the local subnet of the VPC and the subnet to be interconnected in on-premises data center 2. The remote subnets of the VPN device in on-premises data center 2 must contain the local subnet of the VPC and the subnet to be interconnected in on-premises data center 1.

Procedure

VPNs support static routing mode, BGP routing mode, and policy-based mode. The following uses the static routing mode as an example.

Step 1 Configure a VPN gateway.

1. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click **Buy S2C VPN Gateway**.
2. Set parameters as prompted.

   Table 4-10 only describes the key parameters for creating a VPN gateway.

   Table 4-10 Description of VPN gateway parameters

Parameter	Description	Value
Name	Name of a VPN gateway.	vpngw-001

Parameter	Description	Value
Network Type	Select **Public network**.	Public network
Associate With	Select **VPC**.	VPC
VPC	VPC that the on-premises data centers need to access.	vpc-001(192.168.0.0/16)
Local Subnet	VPC subnets that the on-premises data centers need to access.	192.168.0.0/24,192.168.1.0/24
Interconnection Subnet	This subnet is used for communication between the VPN gateway and VPC. Ensure that the selected interconnection subnet has four or more assignable IP addresses.	192.168.2.0/24
BGP ASN	BGP AS number.	64512
HA Mode	Select **Active-active**.	Active-active
Active EIP	EIP 1 used by the VPN gateway to access the on-premises data center.	1.1.1.2
Active EIP 2	EIP 2 used by the VPN gateway to access the on-premises data center.	2.2.2.2

Step 2 Configure customer gateways.

1. Choose **Virtual Private Network** > **Enterprise – Customer Gateways**, and click **Create Customer Gateway**.
2. Set parameters as prompted.

   Table 4-11 only describes the key parameters for creating a customer gateway.

   Table 4-11 Description of customer gateway parameters

Parameter	Description	Value
Name	Name of a customer gateway.	cgw-fw1
Identifier	IP address used by the customer gateway in on-premises data center 1 to communicate with the VPN gateway. Ensure that UDP port 4500 is permitted on the customer gateway device in the on-premises data center.	1.1.1.1

3. Repeat the preceding operations to configure the customer gateway (2.2.2.1) in on-premises data center 2.

Step 3 Configure VPN connections between the cloud side and on-premises data center 1.

1. Choose **Virtual Private Network** > **Enterprise – VPN Connections**, and click **Create VPN Connection**.
2. Set VPN connection parameters and click **Buy Now**.

   Table 4-12 only describes the key parameters for creating VPN connections.

Table 4-12 Description of VPN connection parameters

Parameter	Description	Value
Name	VPN connection name.	vpn-001
VPN Gateway	VPN gateway for which VPN connections are created.	vpngw-001
VPN Gateway IP of Connection 1	Active EIP of the VPN gateway.	1.1.1.2
Customer Gateway of Connection 1	Customer gateway of connection 1.	1.1.1.1
VPN Gateway IP of Connection 2	Active EIP 2 of the VPN gateway.	2.2.2.2
Customer Gateway of Connection 2	Customer gateway of connection 2.	1.1.1.1
VPN Type	Select **Static routing**.	Static routing
Customer Subnet	Subnet in on-premises data center 1 that needs to access the VPC. – A customer subnet cannot be included in any local subnet or any subnet of the VPC to which the VPN gateway is attached. – Reserved VPC CIDR blocks such as 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 cannot be used as customer subnets. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.	172.16.0.0/16

Parameter	Description	Value
Connection 1's Configuration	Configure the IP address assignment mode of tunnel interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, and policies for connection 1.	*Set parameters based on the site requirements.*
Interface IP Address Assignment	– Manually specify In this example, select **Manually specify**. – Automatically assign	Manually specify
Local Tunnel Interface Address	Tunnel interface IP address of the VPN gateway.	169.254.70.1/30
Customer Tunnel Interface Address	Tunnel interface IP address of the customer gateway device.	169.254.70.2/30
Link Detection	Whether to enable route reachability detection in multi-link scenarios. When NQA is enabled, ICMP packets are sent for detection and your device needs to respond to these ICMP packets.	**NQA** enabled
PSK, Confirm PSK	The value must be the same as the PSK configured on the customer gateway device.	Test@123
Policy Settings	The policy settings must be the same as those on the customer gateway device.	Default
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**. **NOTE** If you disable **Same as that of connection 1**, you are advised to use the same settings as connection 1 for connection 2, except the local and customer tunnel interface addresses.	Disabled
Local Tunnel Interface Address	Tunnel IP address of the VPN gateway.	169.254.71.1/30

Parameter	Description	Value
Customer Tunnel Interface Address	Tunnel IP address of the customer gateway.	169.254.71.2/30

Step 4 Configure VPN connections between the cloud side and on-premises data center 2.

1. Choose **Virtual Private Network** > **Enterprise – VPN Connections**, and click **Create VPN Connection**.
2. Set VPN connection parameters and click **Buy Now**.

   Table 4-13 only describes the key parameters for creating VPN connections.

Table 4-13 Description of VPN connection parameters

Parameter	Description	Value
Name	VPN connection name.	vpn-002
VPN Gateway	VPN gateway for which VPN connections are created.	vpngw-001
VPN Gateway IP of Connection 1	Active EIP of the VPN gateway.	1.1.1.2
Customer Gateway of Connection 1	Customer gateway of connection 1.	2.2.2.1
VPN Gateway IP of Connection 2	Active EIP 2 of the VPN gateway.	2.2.2.2
Customer Gateway of Connection 2	Customer gateway of connection 2.	2.2.2.1
VPN Type	Select **Static routing**.	Static routing

Parameter	Description	Value
Customer Subnet	Subnet in on-premises data center 2 that needs to access the VPC. – A customer subnet cannot be included in any local subnet or any subnet of the VPC to which the VPN gateway is attached. – Reserved VPC CIDR blocks such as 100.64.0.0/10, 100.64.0.0/12, and 214.0.0.0/8 cannot be used as customer subnets. The reserved CIDR blocks vary according to regions and are subject to those displayed on the console. If you need to use 100.64.0.0/10 or 100.64.0.0/12, submit a service ticket.	10.10.0.0/16
Connection 1's Configuration	Configure the IP address assignment mode of tunnel interfaces, local tunnel interface address, customer tunnel interface address, link detection, PSK, confirm PSK, and policies for connection 1.	*Set parameters based on the site requirements.*
Interface IP Address Assignment	– Manually specify In this example, select **Manually specify**. – Automatically assign	Manually specify
Local Tunnel Interface Address	Tunnel interface IP address of the VPN gateway.	169.254.72.1/30
Customer Tunnel Interface Address	Tunnel interface IP address of the customer gateway device.	169.254.72.2/30
Link Detection	Whether to enable route reachability detection in multi-link scenarios. When NQA is enabled, ICMP packets are sent for detection and your device needs to respond to these ICMP packets.	**NQA** enabled
PSK, Confirm PSK	The value must be the same as the PSK configured on the customer gateway device in on-premises data center 2.	Test@123

Parameter	Description	Value
Policy Settings	The policy settings must be the same as those configured on the customer gateway device in on-premises data center 2.	Default
Connection 2's Configuration	Determine whether to enable **Same as that of connection 1**. **NOTE** If you disable **Same as that of connection 1**, you are advised to use the same settings as connection 1 for connection 2, except the local and customer tunnel interface addresses.	Disabled
Local Tunnel Interface Address	Tunnel IP address of the VPN gateway.	169.254.73.1/30
Customer Tunnel Interface Address	Tunnel IP address of the customer gateway.	169.254.73.2/30

Step 5 Configure customer gateway devices in on-premises data centers 1 and 2.

The configuration procedures may vary according to the type of the customer gateway device.

----End

Verification

- About 5 minutes later, check states of the VPN connections.

  Choose **Virtual Private Network** > **Enterprise – VPN Connections**. The states of the four VPN connections are all **Normal**.
- Verify that servers in on-premises data center 1 and servers in on-premises data center 2 can ping each other.

4.2 S2C Classic VPN

4.2.1 Connecting an On-Premises Data Center to a VPC Through a VPN

Scenario

By default, ECSs in a VPC cannot communicate with devices in your on-premises data center or private network. To enable communication between them, you can configure VPN. After that, you need to configure security group rules and check

subnet connectivity to ensure that the VPN is available. VPNs can be classified into the following two types:

- A site-to-site VPN functions as a communication tunnel between a VPC and a single on-premises data center.
- By contrast, a hub-and-spoke VPN is between a VPC and multiple on-premises data centers.

Pay attention to the following when you configure a VPN:

- The local and remote subnets cannot conflict.
- The IKE policies, IPsec policies, and PSKs configured on the cloud and in the on-premises data center must be the same.
- The parameters configured for the local and remote subnets and gateways must be symmetric.
- Security group rules permit access to and from the ECSs in the VPC.
- The status of a VPN changes to **Normal** only after ECSs and on-premises servers access each other.

Prerequisites

You have created the VPC and subnets that the on-premises data center wants to access.

Procedure

1. On the management console, select the appropriate IKE and IPsec policies to create a VPN.
2. Check the IP address pools for the local and remote subnets.

   In **Figure 4-4**, the VPC has subnets 192.168.1.0/24 and 192.168.2.0/24. Your on-premises data center has subnets 192.168.3.0/24 and 192.168.4.0/24. You can set up a VPN to connect these subnets.

   Figure 4-4 IPsec VPN

   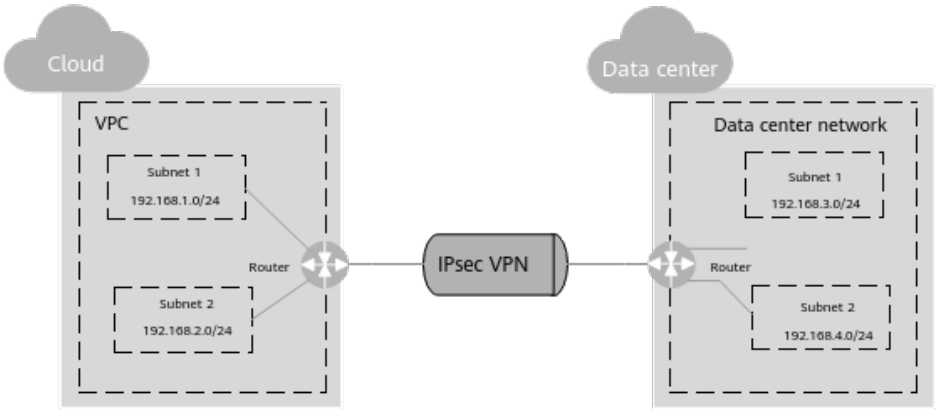

   The IP address pools for the local subnets cannot overlap with those for the remote subnets. Like in this example, the IP address pool for the remote subnets cannot contain the two subnets of the VPC.
3. Configure security group rules for the ECSs to allow packets from and to the on-premises data center over the VPN.

4. Ping the ECSs from the on-premises data center to verify that the security group allows packets from and to the on-premises data center over the VPN.
5. Check the on-premises network configuration.

   A route must be configured for the on-premises network to enable traffic to be forwarded to network devices on the network over the VPN. If the data transmitted through the VPN cannot be forwarded to the network devices, check whether the remote LAN has rules configured to refuse the traffic.

5 Troubleshooting

5.1 The State of a VPN Connection Is Not connected

Symptom

On the **Enterprise – VPN Connections** page of the VPN console, the state of a VPN connection is displayed as **Not connected**.

Possible Causes

- The configurations at the two ends of the VPN connection are incorrect.
- The security group configuration on the management console or the ACL configuration on the customer gateway device is incorrect.
- The IPsec VPN connection negotiation fails or the connection is disconnected.

Procedure

1. Check the configurations at the two ends of the VPN connections.
   a. Check whether the gateway IP addresses configured at the two ends of the VPN connection are reversed.
      i. To check the active and standby EIPs of the VPN gateway, choose **Virtual Private Network** > **Enterprise – VPN Gateways** and view the IP addresses in the **Gateway IP Address** column.
      ii. To check the IP address of the customer gateway, choose **Virtual Private Network** > **Enterprise – Customer Gateways** and view the IP address in the **Identifier** column.
         - If the gateway IP addresses at both ends are not reversed, modify the corresponding settings.
         - If the gateway IP addresses at both ends are reversed, go to **b**.
   b. Check whether the IKE and IPsec policies at both ends are consistent.

      Choose **Virtual Private Network** > **Enterprise – VPN Connections**, locate the target VPN connection, and choose **More** > **Modify Policy Settings**. View the IKE and IPsec policy settings.

If the negotiation parameters in the IKE or IPsec policy at both ends are inconsistent, modify the corresponding policy.

If the negotiation parameters in the IKE and IPsec policies at both ends are consistent, go to **c**.

c. Check whether the PSKs at both ends are the same.

The PSK cannot be checked on the VPN console. If you are not sure whether the PSK configured on the VPN console is correct, you are advised to change it to be the same as that configured on the customer gateway device.

To change the PSK on the VPN console, choose **Virtual Private Network** > **Enterprise – VPN Connections**, locate the target VPN connection, and choose **More** > **Reset PSK**.

d. If the policy-based mode is used, check whether the source and destination CIDR blocks in the policy rules at the two ends of the VPN connection are reversed.

To check policy rules on the VPN console, choose **Virtual Private Network** > **Enterprise – VPN Connections**, locate the target VPN connection, and click **Modify VPN Connection**.

e. If the static routing mode is used and the NQA function is enabled on the VPN console, check whether tunnel interface IP addresses are correctly configured on the customer gateway device.

- To check whether NQA is enabled on the VPN console, choose **Virtual Private Network** > **Enterprise – VPN Connections**, click the name of the target VPN connection, and view the value of **Link Detection** on the **Summary** tab page.
- To check the tunnel interface IP addresses configured on the VPN console, choose **Virtual Private Network** > **Enterprise – VPN Connections**, click **Modify VPN Connection**, and view the values of **Local Interface IP Address** and **Customer Interface IP Address**.

  The local and remote interface IP addresses configured on the customer gateway device must be the same as the values of **Customer Interface IP Address** and **Local Interface IP Address** configured on the VPN console, respectively.

f. If the BGP routing mode is used, check whether the BGP ASNs at the two ends of the VPN connection are reversed.

- To check the BGP ASN of the VPN gateway, choose **Virtual Private Network** > **Enterprise – VPN Gateways**, click the VPN gateway name, and view the BGP ASN in the **Basic Information** area.
- To check the BGP ASN of the customer gateway, choose **Virtual Private Network** > **Enterprise – Customer Gateways** and view the value in the **BGP ASN** column.

2. Check the security group configuration on the management console and the ACL configuration on the customer gateway device.

a. Check whether the default security group on the management console permits the ports corresponding to the public IP addresses of the customer gateway.

b. To check the default security group on the management console, perform the following steps:

i. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click the name of the VPC associated with the VPN gateway.

ii. On the **Virtual Private Cloud** page, click the number in the **Route Tables** column.

iii. On the **Route Tables** page, click the name of the route table.

iv. Locate and click the next hop of the active or standby EIP of the VPN gateway.

v. On the **Associated Security Groups** tab page, check whether the security group permits traffic of the ports.

c. Verify that an ACL on the customer gateway device permits the ports corresponding to the active and standby EIPs of the VPN gateway.

3. Check IPsec connection logs.

a. Log in to the management console.

b. Click in the upper left corner and select the desired region and project.

c. Click in the upper left corner, and choose > **Virtual Private Network**.

d. In the navigation pane on the left, choose **Virtual Private Network** > **Enterprise – VPN Connections**.

e. On the **VPN Connection** page, locate the target VPN connection, and click **View Logs** to view connection logs.

Check IPsec connection logs, and locate the fault based on the log keywords and error codes listed in **Table 5-1**.

Table 5-1 Common causes of VPN disconnection

Category	Error Code	Description
IPsec VPN negotiation failure	phase1 proposal mismatch	IKE proposal parameters on both ends do not match.
	phase2 proposal or pfs mismatch	IPsec proposal parameters, PFS algorithms, or security ACLs on both ends do not match.
	responder dh mismatch	The DH algorithm of the responder does not match that of the initiator.
	initiator dh mismatch	The DH algorithm of the initiator does not match that of the responder.
	encapsulation mode mismatch	Encapsulation modes on both ends do not match.

Category	Error Code	Description
	flow mismatch	Security ACLs on both sides do not match.
	version mismatch	IKE versions on both ends do not match.
	peer address mismatch	IKE peer addresses on both ends do not match.
	config ID mismatch	No IKE peer with the specified ID is found.
	exchange mode mismatch	Negotiation modes on both ends do not match.
	authentication fail	The identity authentication fails.
	construct local ID fail	A local ID fails to be constructed.
	rekey no find old sa	The old SA fails to be found during renegotiation.
	rekey fail	The old SA is going offline during renegotiation.
	first packet limited	First packets are rate limited.
	unsupported version	The IKE version is not supported.
	malformed message	There is a malformed message.
	malformed payload	There is a malformed payload.
	critical drop	The critical payload is not recognized.
	cookie mismatch	The cookies do not match.
	invalid cookie	The cookie is invalid.
	invalid length	The packet length is invalid.
	unknown exchange type	The negotiation mode is unknown.
	short packet	Packets are ultra-short.
	unknown exchange type	The negotiation mode is unknown.
	uncritical drop	The non-critical payload is not identified.
	route limit	The number of imported routes reaches the upper limit.

Category	Error Code	Description
	ip assigned fail	IP address assignment fails.
	eap authentication timeout	EAP authentication times out.
	eap authentication fail	EAP authentication fails.
	xauth authentication fail	XAUTH authentication fails.
	xauth authentication timeout	XAUTH authentication times out.
	license or specification limited	There is license control.
	local address mismatch	The local IP address and interface IP address in IKE negotiation do not match.
	dynamic peers number reaches limitation	The number of IKE peers reaches the upper limit.
	ipsec tunnel number reaches limitation	The number of IPsec tunnels reaches the upper limit.
	netmask mismatch	The mask does not match the configured one after the IPsec mask filtering function is enabled.
	flow conflict	A data flow conflict exists.
	proposal mismatch or use sm in ikev2	IPsec proposals on both ends do not match or IKEv2 uses an SM algorithm.
	ikev2 not support sm in ipsec proposal ikev2	IKEv2 does not support the SM algorithm used in the IPsec proposal.
	no policy applied on interface	No policy is applied to an interface.
	nat detection fail	NAT detection fails.
	fragment packet limit	The number of fragments exceeds the upper limit.
	fragment packet reassemble timeout	Fragment reassembly times out.
	peer cert is expired	The peer certificate has expired.
	peer cert is revoked by CRL	The peer certificate is revoked.

Category	Error Code	Description
	sa with same user exists	The SA is the same as that of another user.
	max transmit reached	The number of retransmitted packets reaches the maximum.
IPsec VPN connection disconnection	dpd timeout	DPD detection times out.
	peer request	The remote end sends a message, asking the local end to tear down a tunnel.
	config modify or manual offline	The SA is automatically deleted due to configuration modification, or is manually deleted.
	phase1 hard expiry	In phase 1, hard timeout (no new SA negotiation succeeds) occurs.
	phase2 hard expiry	A hard timeout occurs in phase 2.
	heartbeat timeout	Heartbeat detection times out.
	re-auth timeout	The SA is deleted because the re-authentication times out.
	aaa cut user	The SA is deleted because the AAA module logs out the user.
	ip address syn failed	IP addresses fail to be synchronized.
	hard expiry triggered by port mismatch	Hard timeout occurs due to a NAT port number mismatch.
	kick old sa with same flow	The old SA is deleted when the same flow is transmitted.
	cpu table updated	When an SPU is removed and inserted, the SAs of CPUs other than the one on the SPU are deleted.
	flow overlap	The IP address in the encrypted data flow conflicts with the remote IP address.
	spi conflict	An SPI conflict occurs.
	admin down	The VPN tunnel status is admin down.
	peer address switch	The peer address is changed.
	forward down	The fwd group goes down.

Category	Error Code	Description
	sa with same user exists	The SA is the same as that of another user.
	reset sa by ike user	A user resets the SA.
	phase1 sa replace	A new IKE SA replaces the old one.
	phase2 sa replace	A new IPsec SA replaces the old one.
	manual offline	The connection is manually torn down.
	nhrp notify	The NHRP module notifies the device of SA deletion.
	receive backup delete info	The standby device receives an SA backup deletion message from the active device.
	eap delete old sa	When the peer device performs EAP authentication repeatedly, the local device deletes the old SA.
	receive invalid spi notify	The device receives an invalid SPI notification.
	dns resolution status change	The DNS resolution status is changed.
	ikev1 phase1-phase2 sa dependent offline	The device deletes the associated IPsec SA when deleting an IKEv1 SA.
	exchange timeout	Packet exchange times out.
	hash gene adjusted	The hash factor is adjusted, causing the IPsec tunnel to be deleted.

5.2 Ping Tests Between Cloud and On-premises Networks Fail

Symptom

- Servers in an on-premises data center cannot ping ECSs in a VPC.
- ECSs in a VPC cannot ping the servers in an on-premises data center.

Possible Causes

- The security group configuration on the management console is incorrect.

- The ACL rule associated with the interconnection subnet is incorrectly configured.
- The ACL configuration on the customer gateway device is incorrect.
- The route configuration on the customer gateway device is incorrect.

Procedure

1. Check the security group configuration on the management console.
   a. Verify that the default security group on the management console permits data flows destined for the customer subnet.
   b. To check the default security group on the management console, perform the following steps:
      i. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click the name of the VPC associated with the VPN gateway.
      ii. Click the number of route tables corresponding to the VPC.
      iii. On the **Route Tables** page, click the name of the route table.
      iv. Locate and click the next hop of the active or standby EIP of the VPN gateway.
      v. On the **Associated Security Groups** tab page, check the ports permitted by the security group.
   c. Verify that the default security group on the management console permits data flows originated from the customer subnet.
   d. Verify that the default security group on the management console permits data flows destined for the local subnet.
   e. Verify that the default security group on the management console permits data flows originated from the local subnet.
   f. Verify that a security group permits data flows from the ECSs to the customer subnet.
   g. To check whether such a security group has been configured, choose **Compute** > **Elastic Cloud Server**, click an ECS name, click the **Security Groups** tab, and click **Manage Rule**.
   h. Verify that a security group permits data flows from the customer subnet to the ECSs.
2. The ACL rule associated with the interconnection subnet is incorrectly configured.
   a. Check whether the ACL rule associated with the interconnection subnet permits the ports between all local and customer subnets.
      i. Choose **Virtual Private Network** > **Enterprise – VPN Gateways**, and click the name of target VPN gateway.
      ii. In the **Basic Information** area, check and record the interconnection subnet.
      iii. In the **Basic Information** area, click the name of the associated VPC.
      iv. On the **Summary** tab page of the VPC, click the number of subnets in the **Networking Components** area.
      v. Find the interconnection subnet in the subnet list, and click the ACL name in the **Network ACL** column.

vi. Permit the ports between all local and customer subnets.

3. Check the ACL configuration on the customer gateway device.
   a. Verify that an ACL rule on the customer gateway device permits data flows destined for the local subnet of the VPN gateway.
   b. Verify that an ACL rule on the customer gateway device permits data flows originated from the local subnet of the VPN gateway.
4. To check the local subnet of the VPN gateway, choose **Virtual Private Network** > **Enterprise – VPN Gateways**, click the VPN gateway name, and view the value of **Local Subnet** in the **Basic Information** area.
5. Check the route configuration on the customer gateway device.
   a. Verify that the public network route is correctly configured. That is, the destination address is an EIP of the VPN gateway, and the next hop is the egress interface address of the customer gateway device.
   b. Verify that the private network route is correctly configured. That is, the destination address is the local subnet of the VPN gateway, and the next hop is the egress interface address of the customer gateway device.
   c. To check the local subnet of the VPN gateway, choose **Virtual Private Network** > **Enterprise – VPN Gateways**, click the VPN gateway name, and view the value of **Local Subnet** in the **Basic Information** area.

5.3 Packet Loss Occurs

Symptom

- Packet loss occurs when a server in an on-premises data center pings an ECS in a VPC.
- Packet loss occurs when an ECS in a VPC pings a server in an on-premises data center.

Procedure

- Check the customer-side networking and bandwidth.
  - Check whether the customer network has multiple egresses working in load balancing mode and whether traffic destined for the cloud is distributed to a non-VPN egress. Ensure that the traffic destined for the cloud is transmitted through the same egress.
  - Ping the IP address of the VPN gateway at the cloud side and other public IP addresses (for example, 114.114.114.114) from the customer gateway to check the delay and packet loss rate on the public network.

    If the quality of the public network is poor, you are advised to seek help from the corresponding carrier.
  - Check whether traffic on the customer gateway device exceeds the bandwidth limit.
- Check the cloud-side networking and bandwidth.
  - Check whether traffic exceeds the bandwidth of the VPN gateway.

    If traffic exceeds the bandwidth of the VPN gateway, increase the bandwidth.

6 FAQs

6.1 FAQs - S2C Enterprise Edition VPN

6.1.1 Popular Questions

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

Table 6-1 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

Protocol	Parameter	Value
	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 16 • Group 19 • Group 20 • Group 21
	Version	• v1 (not recommended due to security risks) • v2 (default 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.

Protocol	Parameter	Value
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
	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.

6.1.1.2 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.

6.1.1.3 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.

6.1.1.4 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.

6.1.1.5 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.

6.1.1.6 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.

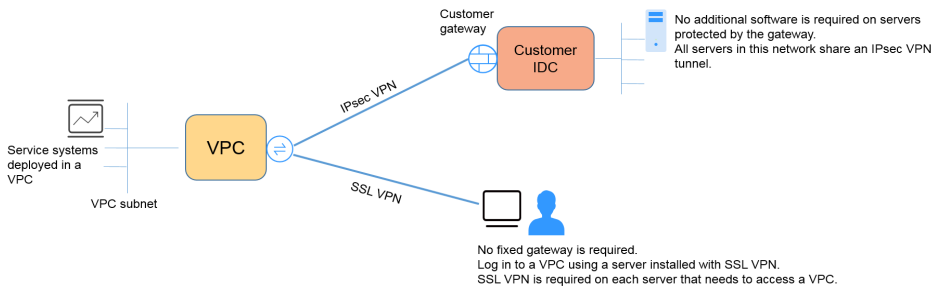

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.1.1.7 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

6.1.1.8 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.

6.1.1.9 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.

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.

6.1.1.10 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.

6.1.1.11 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.

NOTE

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 6-1 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 client.

Figure 6-1 Test result for 200 Mbit/s bandwidth (iPerf3 client)

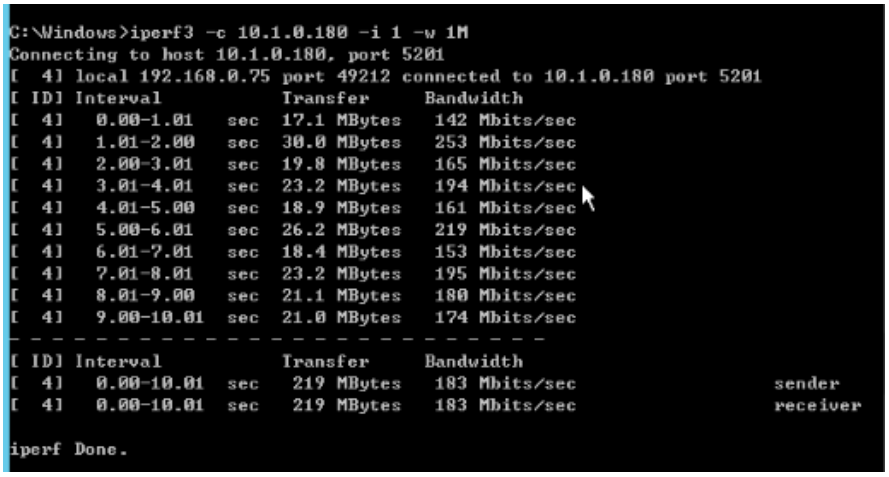

Figure 6-2 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 server.

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

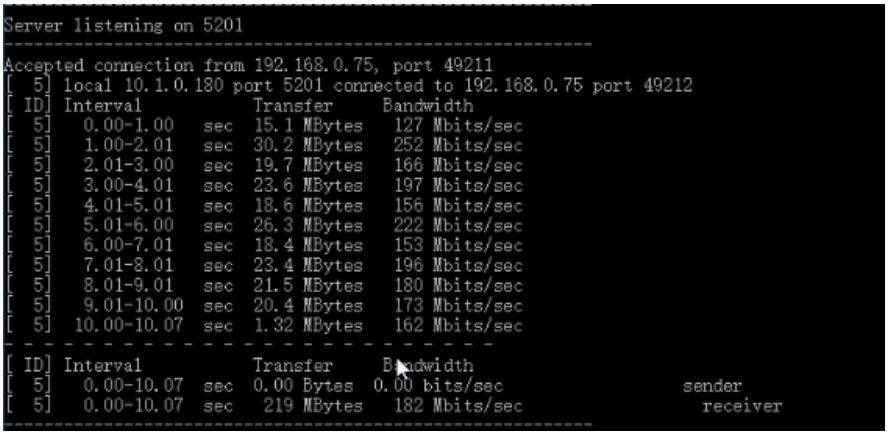

2. When the ECSs at the two ends of the VPN connection run CentOS 7, iPerf3 is used to test the network speed. The test result is 180 Mbit/s, meeting requirements.
3. When the ECS functioning as a server runs CentOS 7 and the ECS functioning as a client runs Windows, iPerf3 and FileZilla are used to 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 6-3 shows the result of using iPerf3 to test the network speed between two ECSs running different operating systems.

Figure 6-3 Test result on iPerf3

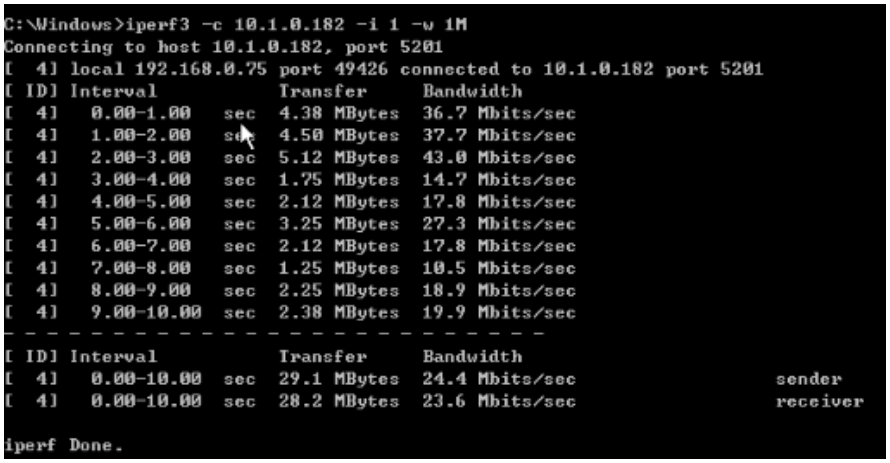

When the bandwidth of a purchased VPN gateway is 1000 Mbit/s:

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.

6.1.1.12 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.

6.1.1.13 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 on-premises 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.

6.1.1.14 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.

6.1.1.15 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 on-premises 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.

6.1.1.16 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.

6.1.1.17 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).

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

Yes.

6.1.1.19 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 on-premises 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.

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

1. 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.

6.1.1.21 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.

6.1.1.22 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.

6.1.1.23 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**.

6.1.2 General Consulting

6.1.2.1 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 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.
- 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.
  - A VPN cannot be used between the cloud and your home network that uses PPPoE dial-up.
  - A VPN cannot be used between the cloud and 4G/5G routers.
  - A VPN cannot be used between the cloud and your personal terminals.

6.1.2.2 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 on-premises 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.

6.1.2.3 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 on-premises 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.

6.1.2.4 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 on-premises data centers. Each VPN connection can connect a VPC to only one on-premises data center.

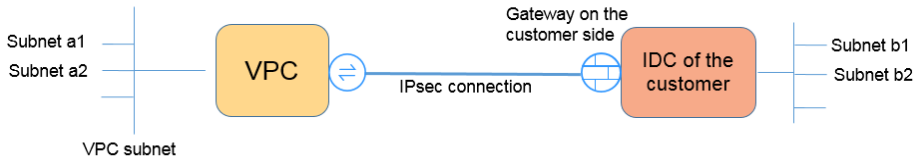

NOTE

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.

6.1.2.5 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.

6.1.2.6 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.

6.1.2.7 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

6.1.2.8 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.

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

Table 6-2 VPN negotiation parameters

Protocol	Parameter	Value
IKE	Version	• v1 (v1 has low security. If the device supports v2, v2 is recommended.) • v2 (default value)

Protocol	Parameter	Value
	Negotiation Mode	• Main (default value) • 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
	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 19 • Group 20 • Group 21
	Lifetime (s)	**86400** (default value) Unit: second Value range: **60** to **604800**

Protocol	Parameter	Value
	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 (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.

6.1.2.10 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.

6.1.2.11 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

NOTE

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

6.1.2.12 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.

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.

6.1.2.13 Can EIPs Be Used as VPN Gateway IP Addresses?

Yes.

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

6.1.2.14 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.

6.1.2.15 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.

6.1.2.16 Does VPN Support IPv6?

No.

6.1.2.17 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.

6.1.2.18 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 brute-force 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.

    NOTE

    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.

6.1.2.19 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 6-3 lists the number of bits corresponding to the DH groups used by VPN.

Table 6-3 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

DH Group	Modulus
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.

6.1.2.20 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.

6.1.2.21 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 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.

6.1.2.22 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.

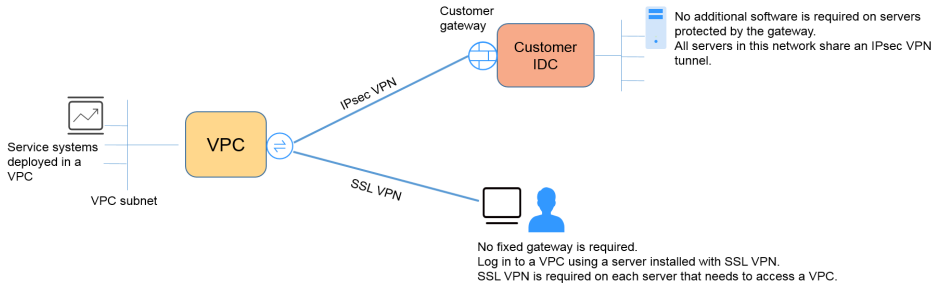

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.1.2.23 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.

6.1.2.24 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.

6.1.2.25 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.

6.1.2.26 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.

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

1. 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.

6.1.2.28 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.

6.1.2.29 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 Professional 1 and Professional 2.

6.1.3 Networking and Application Scenarios

6.1.3.1 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.

6.1.3.2 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 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.

6.1.3.3 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 on-premises 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.

6.1.3.4 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.

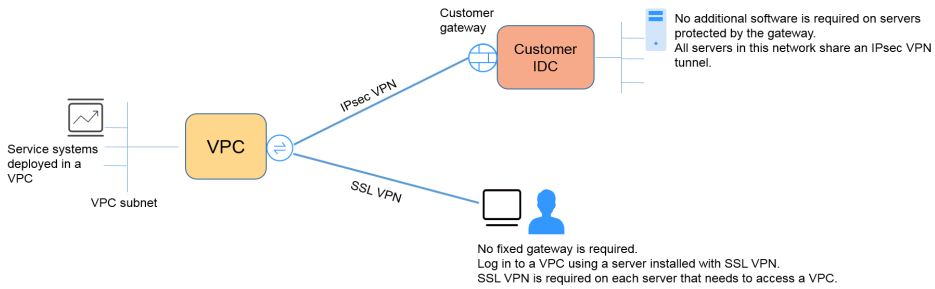

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.1.3.5 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.

6.1.3.6 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).

6.1.3.7 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.
  - 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.

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

Yes.

6.1.3.9 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.

6.1.3.10 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.

6.1.3.11 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.
   - 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.

6.1.3.12 How Do I Connect Four Subnets?

Figure 6-4 shows the network topology.

Figure 6-4 Network topology

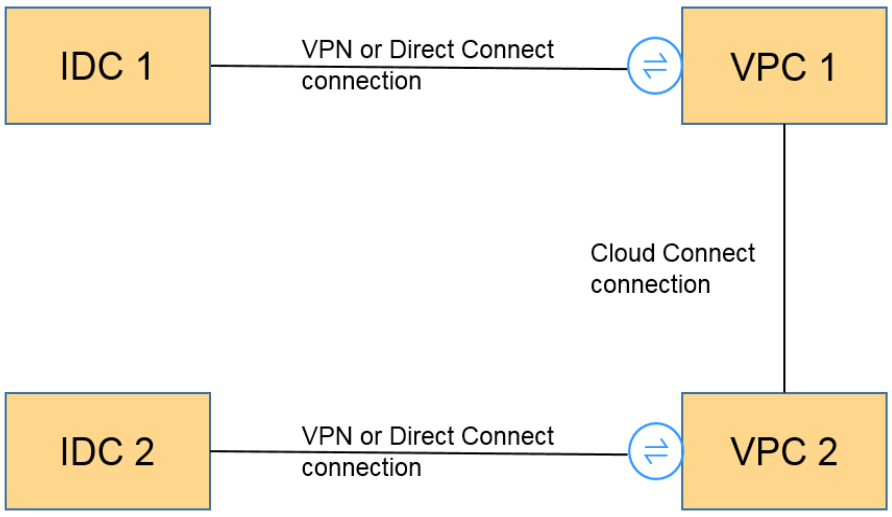

1. Use a VPN connection or Direct Connect connection to connect IDC 1 to VPC 1.
2. Use a Cloud Connect connection to connect VPC 1 to VPC 2. (You can also use a VPC peering connection to connect VPC 1 to VPC 2 if they are in the same region.)
3. Use a VPN connection or Direct Connect connection to connect 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.

6.1.3.13 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.

6.1.3.14 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.

6.1.3.15 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.

6.1.3.16 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.

6.1.3.17 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.

6.1.4 Operations on the Console

6.1.4.1 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 on-premises 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.

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

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

NOTE

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

6.1.4.3 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.

6.1.4.4 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.

6.1.4.5 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**.

6.1.4.6 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.

6.1.4.7 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.

6.1.4.8 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.

6.1.4.9 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.

6.1.4.10 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.

6.1.4.11 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.)

6.1.4.12 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.

6.1.4.13 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.

6.1.4.14 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 brute-force 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.

    NOTE

    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.

6.1.4.15 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.

6.1.4.16 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.

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.

6.1.4.17 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.

6.1.5 VPN Negotiation and Interconnection

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

Table 6-4 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**.
	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**.

Protocol	Parameter	Value
	Version	• v1 (For security reasons, IKEv1 is not recommended. If your devices support IKEv2, select IKEv2.) • 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.
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**.

Protocol	Parameter	Value
	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**

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.

6.1.5.2 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

6.1.5.3 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.

6.1.5.4 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.

6.1.5.5 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

NOTE

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

6.1.5.6 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.

6.1.5.7 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.

6.1.5.8 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.

6.1.5.9 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.

6.1.5.10 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.

6.1.5.11 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.

6.1.5.12 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 6-5 lists the number of bits corresponding to the DH groups used by VPN.

Table 6-5 Number of bits corresponding to each DH group

DH Group	Modulus
1	768 bits

DH Group	Modulus
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.

6.1.6 Connection or Ping Failure

6.1.6.1 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.

6.1.6.2 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 on-premises 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.

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

1. 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.

6.1.6.4 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.

6.1.6.5 Is an IPsec VPN Connection Automatically Established?

Yes. An IPsec VPN connection is automatically established.

6.1.6.6 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.

6.1.6.7 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.

6.1.6.8 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.

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.

6.1.6.9 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.

6.1.6.10 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.

6.1.6.11 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.

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

1. 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.

6.1.6.13 What Should I Do If I Cannot Access 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.

6.1.6.14 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.

6.1.6.15 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.

6.1.7 Public Addresses

6.1.7.1 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.

6.1.7.2 Can EIPs Be Used as VPN Gateway IP Addresses?

Yes.

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

6.1.7.3 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.

6.1.7.4 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.

6.1.8 Route Configurations

6.1.8.1 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.

6.1.8.2 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.

6.1.8.3 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 on-premises network, add a route to the on-premises network with the gateway address of the non-primary NIC as the next hop.

6.1.9 Subnet Configurations

6.1.9.1 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.)

6.1.9.2 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.

6.1.9.3 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.

6.1.9.4 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.

6.1.9.5 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.

6.1.9.6 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.

6.1.10 VPN Interesting Traffic

6.1.10.1 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**.

6.1.10.2 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.

6.1.11 Keeping VPN Connections Alive

6.1.11.1 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 on-premises 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.

6.1.12 Monitoring

6.1.12.1 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.

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.

6.1.12.2 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.

6.1.12.3 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.

6.1.13 Bandwidth and Network Speed

6.1.13.1 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.

   NOTE

   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 6-5 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 client.

   Figure 6-5 Test result for 200 Mbit/s bandwidth (iPerf3 client)

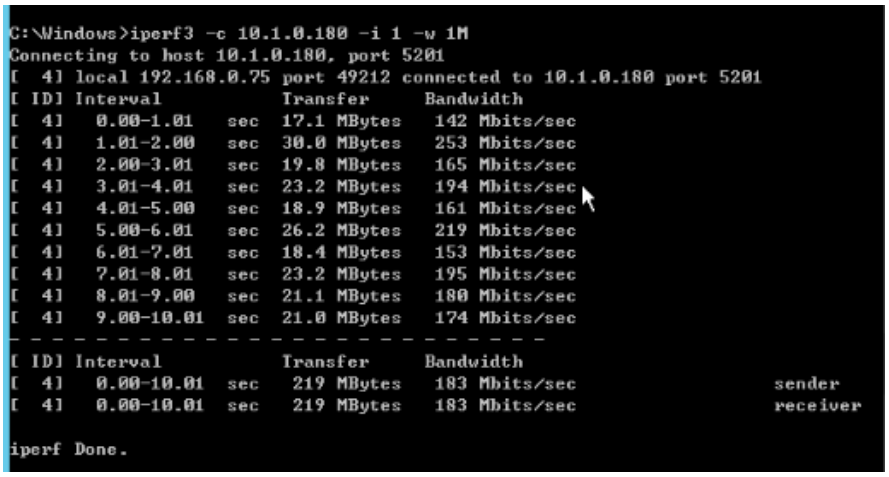

   Figure 6-6 shows the result of testing the 200 Mbit/s bandwidth on the iPerf3 server.

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

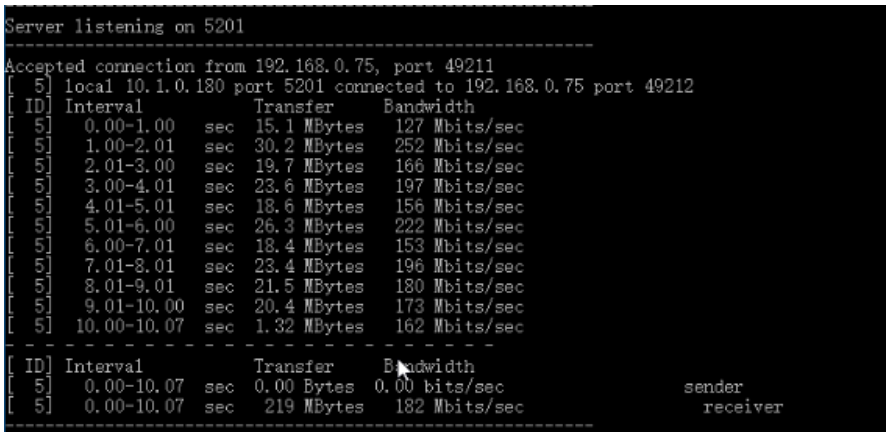

2. When the ECSs at the two ends of the VPN connection run CentOS 7, iPerf3 is used to test the network speed. The test result is 180 Mbit/s, meeting requirements.
3. When the ECS functioning as a server runs CentOS 7 and the ECS functioning as a client runs Windows, iPerf3 and FileZilla are used to 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 6-7 shows the result of using iPerf3 to test the network speed between two ECSs running different operating systems.

Figure 6-7 Test result on iPerf3

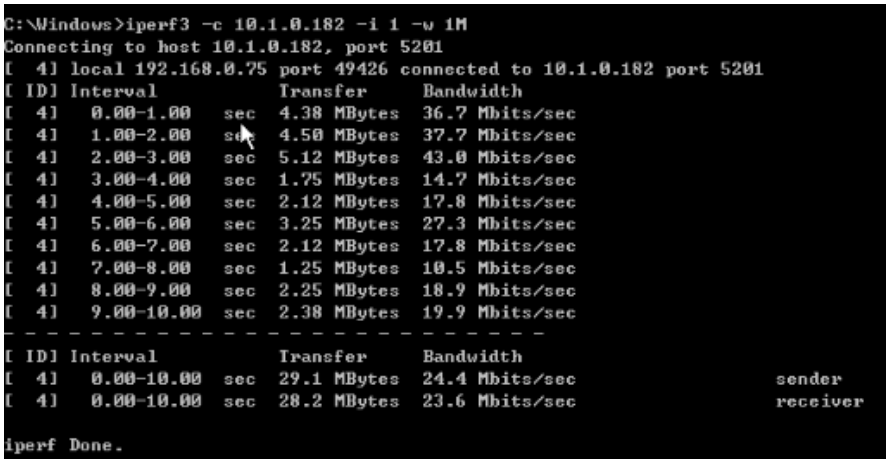

When the bandwidth of a purchased VPN gateway is 1000 Mbit/s:

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.

6.1.13.2 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.

6.1.13.3 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.

6.1.13.4 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.

6.1.13.5 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.

6.1.13.6 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.

6.2 FAQs - S2C Classic VPN

6.2.1 General Questions

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

Table 6-6 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)

Policy	Parameter	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)
	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.)

Policy	Parameter	Value
	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) • AH • AH-ESP
	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.

6.2.1.2 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 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.

6.2.1.3 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.

6.2.1.4 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.

6.2.1.5 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.

6.2.1.6 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.

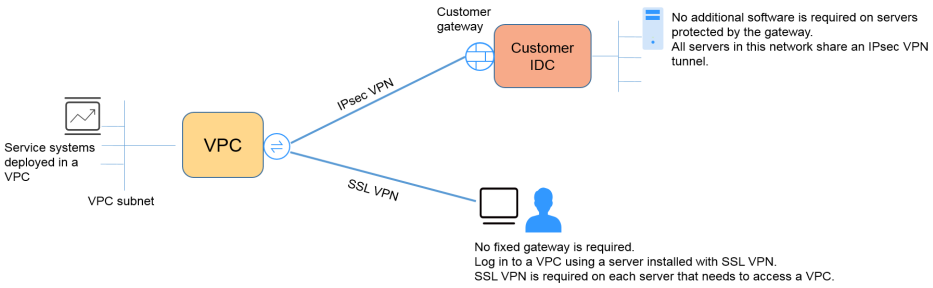

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.2.1.7 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.

6.2.1.8 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.

6.2.1.9 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.

6.2.1.10 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:

1. 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 6-8 shows the test result.

Figure 6-8 Test result for 200 Mbit/s bandwidth (iPerf3 client)

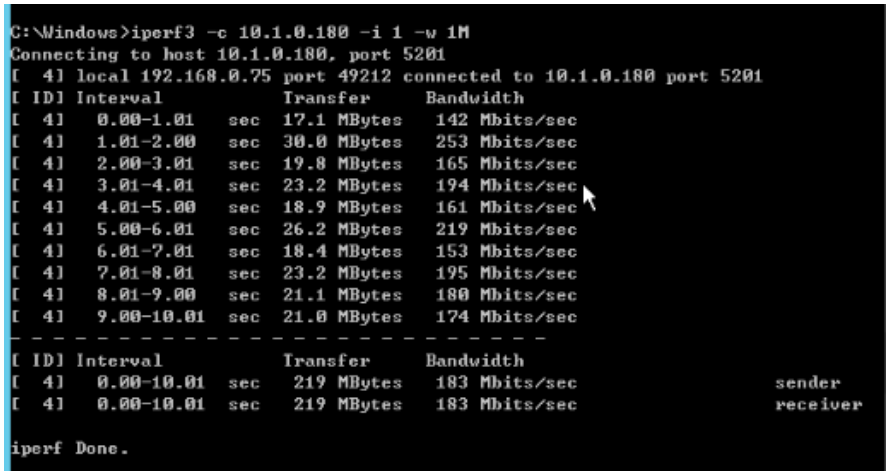

Figure 6-9 shows the test result.

Figure 6-9 Test result for 200 Mbit/s bandwidth (iPerf3 server)

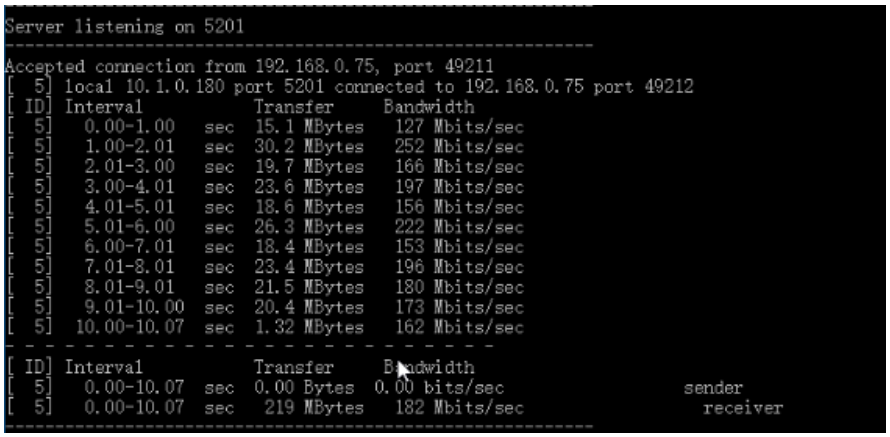

2. If the ECSs at the two ends of the VPN run CentOS 7, use iPerf3 to test the network speed. The network speed can reach 180 Mbit/s.
3. If the ECS functioning as the server runs CentOS 7, and the ECS functioning as the client runs Windows, use iPerf3 and FileZilla to test the network speed.

   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 6-10 shows the test result.

Figure 6-10 Test result when ECSs at the two ends run different OSs (iPerf3)

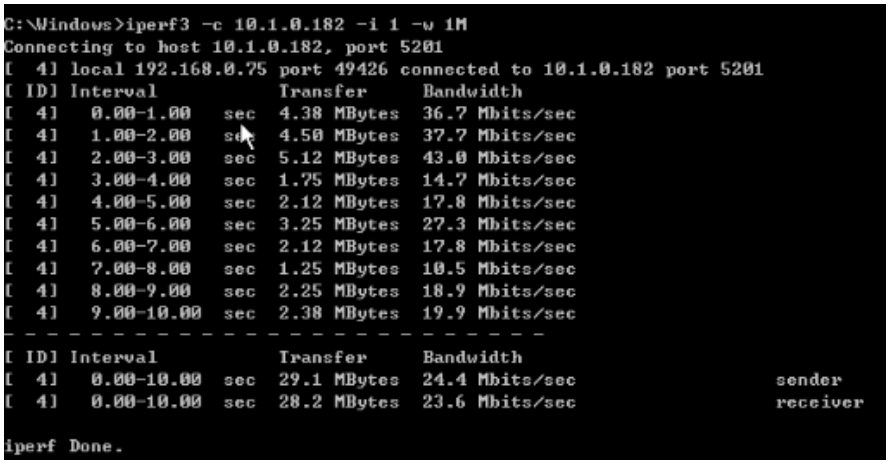

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 1,000 Mbit/s:

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.

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

No.

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

6.2.1.12 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 on-premises 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 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.

6.2.1.13 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.

6.2.1.14 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.

6.2.1.15 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.

6.2.1.16 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.

6.2.1.17 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.

6.2.1.18 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.

6.2.1.19 What Can I Do If 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.

6.2.1.20 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.

6.2.1.21 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.

6.2.1.22 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

6.2.2 Product Consultation

6.2.2.1 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 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.
- 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.
  - A VPN cannot be used between the cloud and your home network that uses PPPoE dial-up.
  - A VPN cannot be used between the cloud and 4G/5G routers.
  - A VPN cannot be used between the cloud and your personal terminals.

6.2.2.2 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 on-premises 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.

6.2.2.3 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 on-premises 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 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.

6.2.2.4 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.

6.2.2.5 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.

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

Table 6-7 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) • AH • AH-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.

6.2.2.7 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.

6.2.2.8 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.

6.2.2.9 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.

6.2.2.10 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.

6.2.2.11 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.

6.2.2.12 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.

6.2.2.13 Does VPN Support IPv6?

No.

Currently, only IPv4 is supported.

6.2.2.14 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.

6.2.2.15 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 re-negotiation 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.

6.2.2.16 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 6-8 lists the bits corresponding to the DH groups used by VPN.

Table 6-8 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.

6.2.2.17 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.

6.2.2.18 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.

6.2.2.19 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.

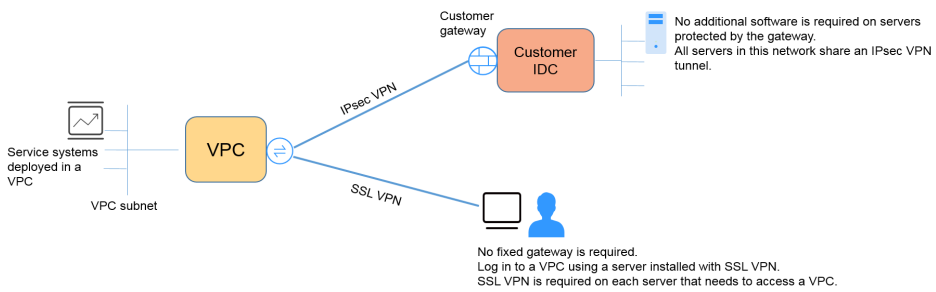

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.2.2.20 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.

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

No.

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

6.2.2.22 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.

6.2.2.23 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.

6.2.2.24 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.

6.2.2.25 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.

6.2.2.26 What Can I Do If 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.

6.2.2.27 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.

6.2.3 Networking and Application Scenarios

6.2.3.1 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.

6.2.3.2 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 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.

6.2.3.3 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.

6.2.3.4 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.

6.2.3.5 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.

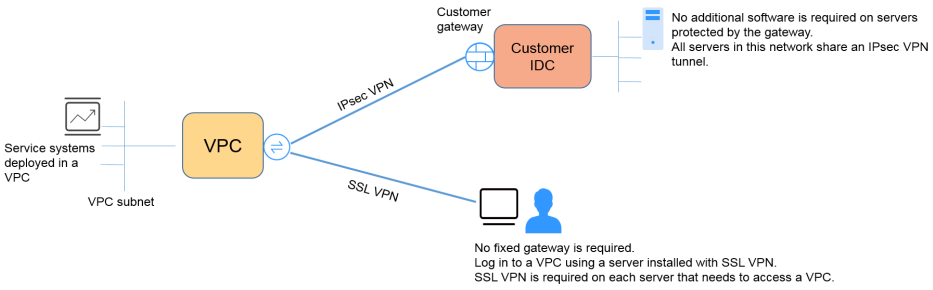

NOTE

Currently, only IPsec VPN is supported; SSL VPN is not supported.

6.2.3.6 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.

6.2.3.7 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.

6.2.3.8 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.

6.2.3.9 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.

6.2.3.10 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.

6.2.3.11 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.

6.2.3.12 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.

6.2.3.13 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 on-premises 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.

6.2.3.14 How Do I Connect Four Subnets?

Figure 6-11 shows the network topology.

Figure 6-11 Network topology

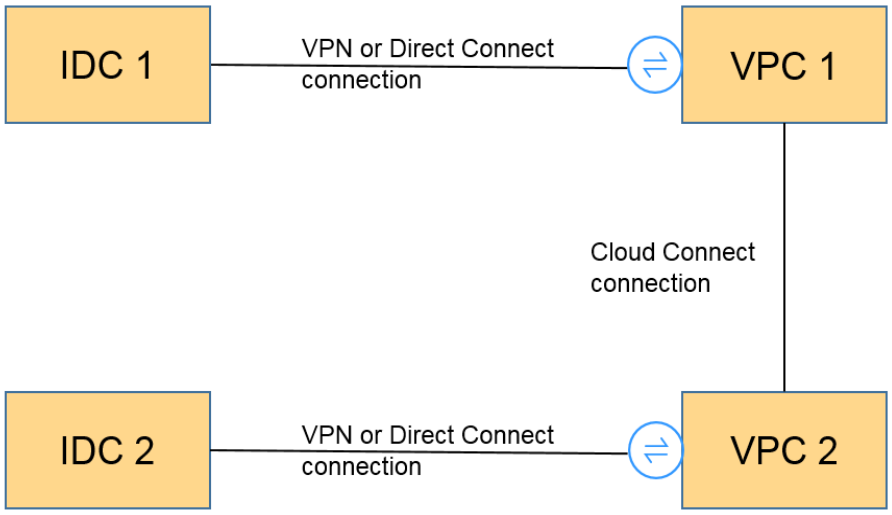

1. Use a VPN or Direct Connect connection to connect IDC 1 to VPC 1.
2. Use a Cloud Connect connection to connect VPC 1 to VPC 2. (You can also use a VPC peering connection to enable communications between VPC 1 and VPC 2 if they are in the same region.)
3. Use a VPN or Direct Connect connection to connect 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.

6.2.3.15 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.

6.2.3.16 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.

6.2.3.17 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.

6.2.3.18 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.

6.2.3.19 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.

6.2.4 Related Operations on the Console

6.2.4.1 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 on-premises 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 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.

6.2.4.2 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.

NOTE

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

6.2.4.3 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.

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.

6.2.4.4 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.

6.2.4.5 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

6.2.4.6 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.

6.2.4.7 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.

6.2.4.8 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.

6.2.4.9 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.

6.2.4.10 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.

6.2.4.11 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.)

6.2.4.12 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**.

6.2.4.13 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.

6.2.4.14 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.

6.2.4.15 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 re-negotiation 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.

6.2.4.16 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.

6.2.4.17 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.

6.2.4.18 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.

6.2.5 VPN Negotiation and Interconnection

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

Table 6-9 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**

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.

6.2.5.2 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.

6.2.5.3 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.

6.2.5.4 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.

6.2.5.5 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.

6.2.5.6 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

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.

6.2.5.7 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.

6.2.5.8 Will a VPN Connection Be Reestablished After Its Configuration Is Modified?

A VPN connection consists of local subnets, remote subnets, remote gateway, pre-shared 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.

6.2.5.9 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.

6.2.5.10 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 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.

6.2.5.11 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.

6.2.5.12 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 6-10 lists the bits corresponding to the DH groups used by VPN.

Table 6-10 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.

6.2.6 Connection or Ping Failure

6.2.6.1 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.

6.2.6.2 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.
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, to customer service personnel.

6.2.6.3 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.

6.2.6.4 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.

6.2.6.5 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.

6.2.6.6 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.

6.2.6.7 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 on the on-premises gateway device to keep the connection alive.

6.2.6.8 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.

6.2.6.9 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.

6.2.6.10 What Can I Do If 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.

6.2.6.11 What Should I Do If I Cannot Access 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 on-premises subnets.

6.2.6.12 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.

6.2.6.13 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

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.

6.2.7 EIPs

6.2.7.1 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.

6.2.7.2 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.

6.2.7.3 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.

6.2.7.4 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.

NOTE

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.

6.2.7.5 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 on-premises 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.

6.2.8 Route Configurations

6.2.8.1 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.

6.2.8.2 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.

6.2.8.3 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.

6.2.9 Subnet Setting

6.2.9.1 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.)

6.2.9.2 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.

6.2.9.3 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.

6.2.9.4 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.

6.2.9.5 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.

6.2.9.6 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.

6.2.9.7 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 .
- 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.

6.2.10 VPN Interesting Traffic

6.2.10.1 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

6.2.10.2 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.

6.2.11 Keeping VPN Connection Alive

6.2.11.1 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 &
```

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. 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.

6.2.12 Monitoring

6.2.12.1 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.

6.2.12.2 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.

6.2.12.3 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.

6.2.13 Bandwidth and Network Speed

6.2.13.1 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:

1. 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 6-12 shows the test result.

   Figure 6-12 Test result for 200 Mbit/s bandwidth (iPerf3 client)

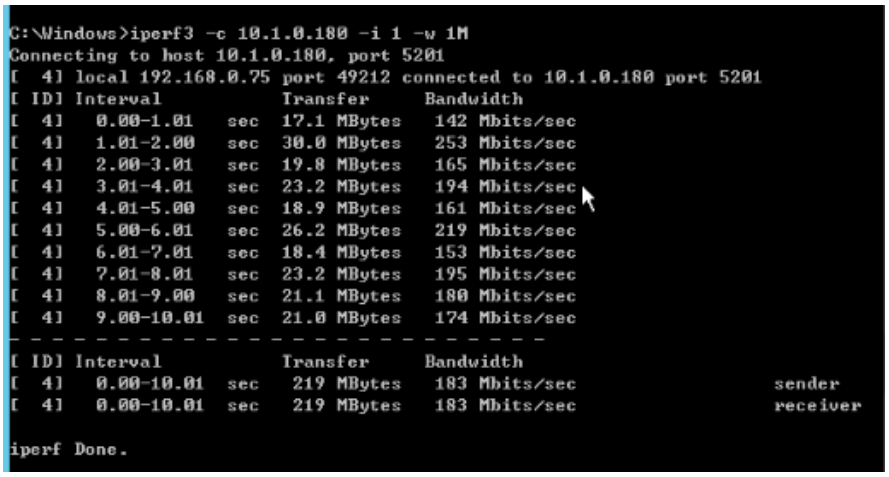

   Figure 6-13 shows the test result.

Figure 6-13 Test result for 200 Mbit/s bandwidth (iPerf3 server)

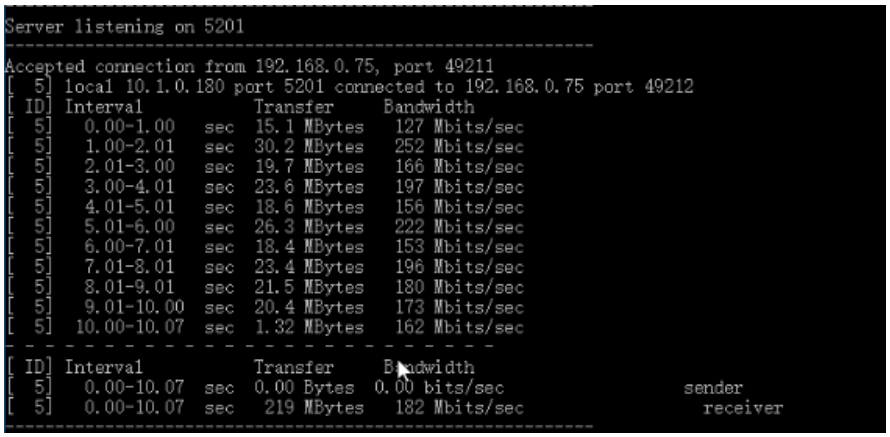

2. If the ECSs at the two ends of the VPN run CentOS 7, use iPerf3 to test the network speed. The network speed can reach 180 Mbit/s.
3. If the ECS functioning as the server runs CentOS 7, and the ECS functioning as the client runs Windows, use iPerf3 and FileZilla to test the network speed.

   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 6-14 shows the test result.

Figure 6-14 Test result when ECSs at the two ends run different OSs (iPerf3)

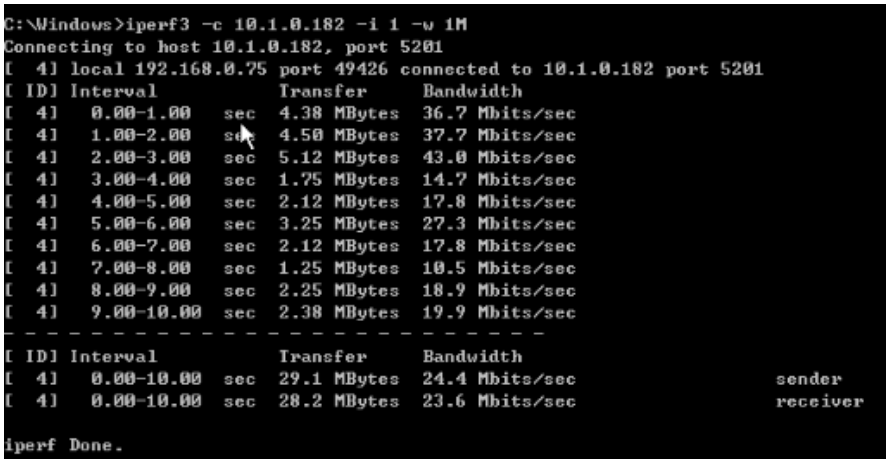

Perform the following steps to test the network speed of your VPN gateway if the bandwidth of your VPN gateway is 1,000 Mbit/s:

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.

6.2.13.2 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.

6.2.13.3 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.

NOTE

The maximum bandwidth of a VPN connection is 300 Mbit/s.

6.2.13.4 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.

NOTE

Changing the VPN bandwidth will not interrupt workload running and networks.

6.2.13.5 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.

6.2.13.6 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 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.

6.2.13.7 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.