Virtual Private Cloud

Service Overview

Issue	44
Date	2020-09-07

HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2020. All rights reserved.

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

Trademarks and Permissions

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

Notice

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

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

Contents

1 What Is Virtual Private Cloud?

Overview

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

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

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

Product Architecture

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

Figure 1-1 Architecture

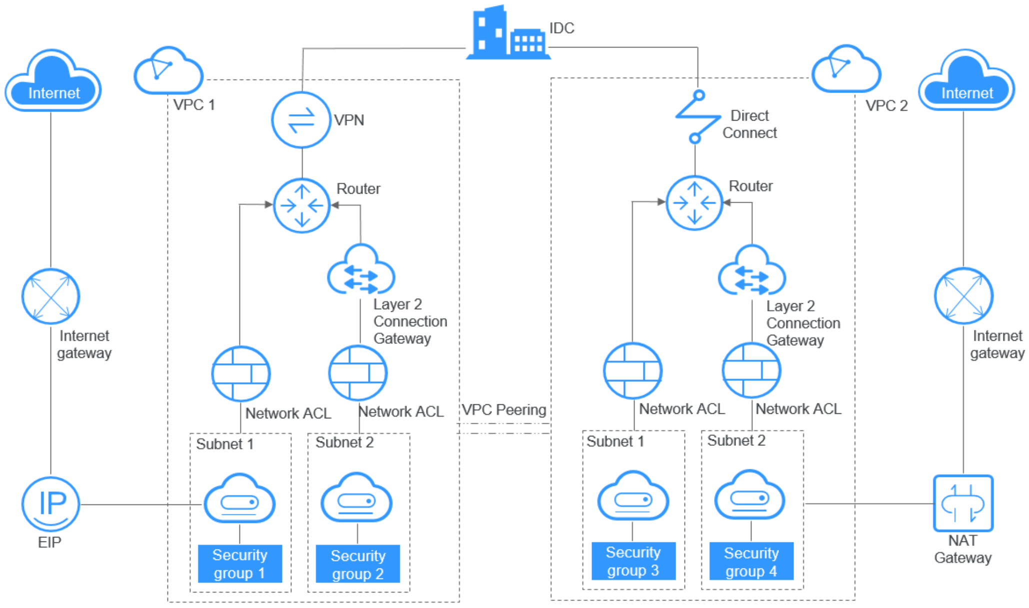

VPC Components

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

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

Security Features

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

VPC Connectivity

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

- VPC Peering allows two VPCs in the same region to communicate with each other using private IP addresses.
- Elastic IP or NAT Gateway allows ECSs in a VPC to communicate with the Internet.
- Virtual Private Network (VPN), Cloud Connect, Direct Connect, or Layer 2 Connection Gateway can connect your data center to VPCs.

Accessing the VPC

You can access the VPC service through the management console or using HTTPS-based APIs.

- Management console

  You can use the console to perform operations on VPC resources directly. To access the VPC service, log in to the **management console** and select **Virtual Private Cloud** from the console homepage.
- API

  If you need to integrate the VPC service provided by the cloud system into a third-party system for secondary development, you can use an API to access the VPC service. For details, see the **Virtual Private Cloud API Reference**.

2 Product Advantages

Flexible Configuration

You can create VPCs, add subnets, specify IP address ranges, and configure DHCP and route tables. You can configure the same VPC for ECSs that are in different availability zones (AZs).

Secure and Reliable

Each VPC is completely logically isolated from other VPCs using the tunneling technology. By default, different VPCs cannot communicate with each other. Network ACLs are provided to protect subnets, and security groups are provided to protect ECSs. The network ACLs and security groups add additional layers of security to your VPCs, making your network very secure.

Figure 2-1 Secure and Reliable

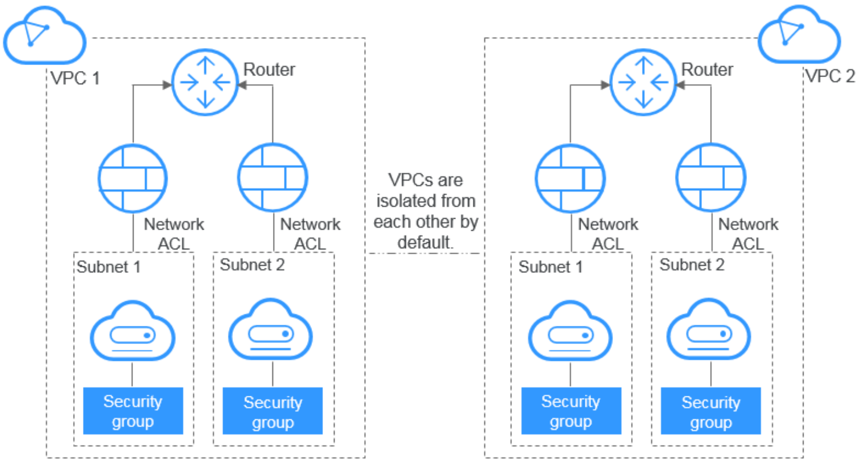

Interconnectivity

By default, instances in a VPC cannot access the Internet. You can leverage EIPs, load balancers, NAT gateways, VPN connections, and Direct Connect connections to enable access to or from the Internet.

By default, instances in two VPCs cannot communicate with each other. You can create a VPC peering connection to enable the instances in the two VPCs in the same region to communicate with each other using private IP addresses.

Layer 2 Connection Gateway can establish network communication between the cloud and on-premises networks and allow you to migrate data center or private cloud services to the cloud without changing subnets.

Multiple connectivity options are provided to meet enterprises' diverse service requirements for the cloud, to allow you to deploy enterprise applications with ease, and to lower enterprise IT operation and maintenance (O&M) costs.

Figure 2-2 Interconnectivity

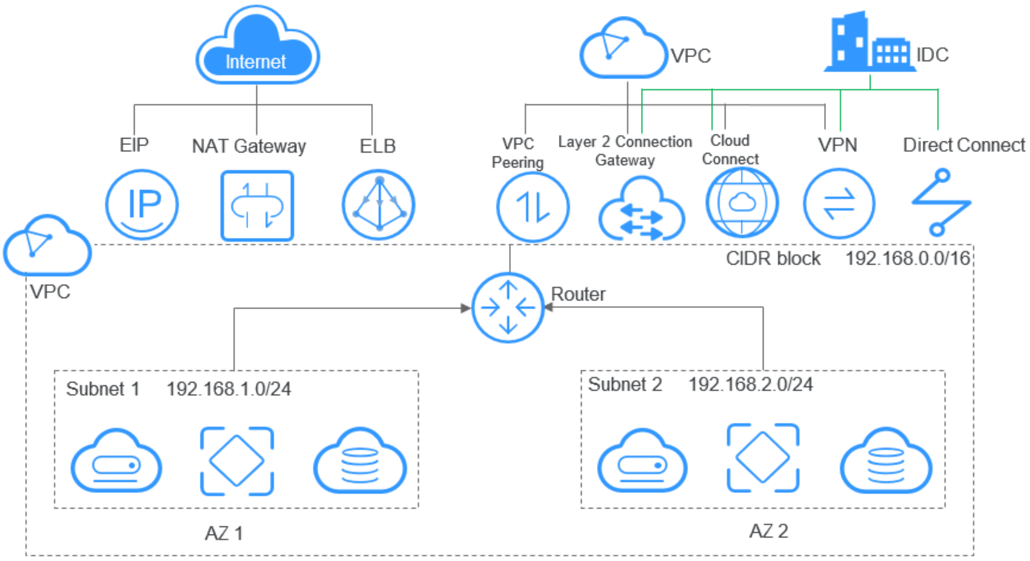

High-Speed Access

Dynamic BGP is used to provide access to various carrier networks. For example, up to 21 dynamic BGP connections are established to multiple carriers. The dynamic BGP connections enable real-time failover based on the preset routing protocols, ensuring high network stability, low network latency, and smooth access to services on the cloud.

Advantage Comparison

Table 2-1 lists the advantages of a VPC over a traditional IDC.

Table 2-1 Comparison between a VPC and a traditional IDC

Item	VPC	Traditional IDC
Deployment cycle	● You do not need to perform complex engineering deployment, such as engineering planning and cabling. ● You can determine your networks, subnets, and routes on HUAWEI CLOUD based on service requirements.	You need to set up networks and perform tests. The entire process takes a long time and requires professional technical support.
Total cost	HUAWEI CLOUD provides flexible billing modes for network services, so you can select the one that can best fit your business needs. In addition, you do not need to pay for upfront costs and network O&M costs, lowering total cost of ownership (TCO).	You need to invest heavily in equipment rooms, power supply, construction, and hardware materials. You also need professional O&M teams to ensure network security. Asset management costs increase with business changes.
Flexibility	A variety of network services are available for you to choose from. If you need more network resources (such as bandwidth), dynamic expansion can be performed conveniently and quickly.	You have to strictly comply with the network plan to complete the service deployment. When there are changes in your service requirements, the network cannot be dynamically adjusted.
Security	VPCs are logically isolated from each other. You can leverage security features such as network ACLs and security groups, and even security services like Advanced Anti-DDoS (AAD) to secure your cloud resources.	The network is difficult to maintain and has poor security. Therefore, you need professional personnel to ensure network security.

3 Application Scenarios

Dedicated Networks on Cloud

Scenario

Each VPC represents a private network and is logically isolated from other VPCs. You can deploy your service system in a VPC to build a private network environment on the cloud. If you have multiple service systems, for example, a production system and a test system, you can deploy them in two different VPCs to isolate them. If you want to establish communication between these two VPCs, you can create a VPC peering connection between them.

Related Services

ECS

Figure 3-1 Dedicated networks on cloud

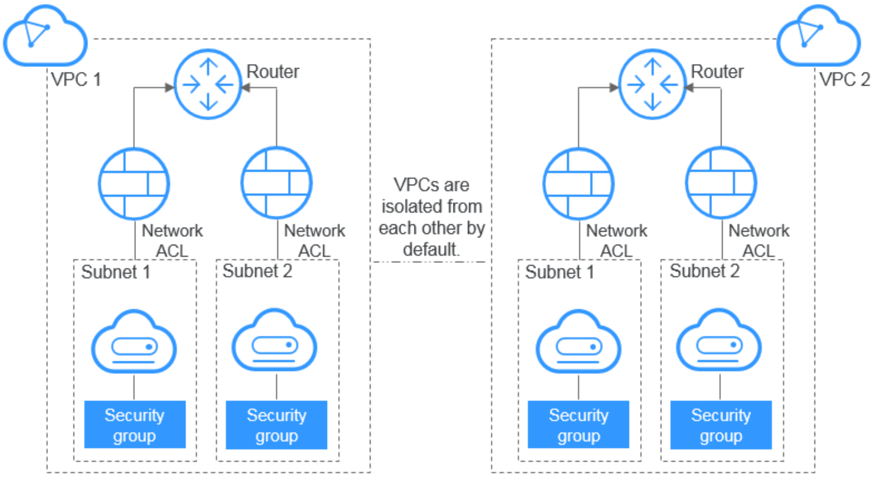

Web Application or Website Hosting

Scenario

You can host web applications and websites in a VPC and use the VPC as a regular network. With EIPs or NAT gateways, you can connect ECSs running your web applications to the Internet. With the load balancers provided by the ELB service, you can evenly distribute traffic across multiple ECSs.

Cloud resources in a VPC can use the following cloud services to connect to the Internet.

Table 3-1 Accessing the Internet

Cloud Service	Application Scenario	Description	Related Operations
EIP	Single ECS accesses the Internet.	You can assign an EIP and bind it to an ECS so that the ECS can access the Internet or provide services accessible from the Internet. An EIP can be bound to an ECS to enable Internet access, or unbound to disable access. Shared bandwidth and shared data packages can be used to lower costs.	**What Are EIPs?**
NAT Gateway	Multiple ECSs share an EIP to access the Internet.	A NAT gateway offers both source network address translation (SNAT) and destination network address translation (DNAT). SNAT allows multiple ECSs in the same VPC to share one or more EIPs to access the Internet. It reduces management costs and prevents the ECS EIPs from being exposed to the Internet. DNAT can implement port-level data forwarding. It maps EIP ports to ECS ports so that the ECSs in a VPC can share the same EIP and bandwidth to provide Internet-accessible services. But DNAT does not balance traffic.	**Using SNAT to Access the Internet** **Using DNAT to Provide Services Accessible from the Internet**

Cloud Service	Application Scenario	Description	Related Operations
ELB	Use load balancers provided by the ELB service to evenly distribute incoming traffic across multiple ECSs in high-concurrency scenarios, such as e-commerce.	Load balancers distribute traffic across multiple backend ECSs, balancing the workload on each ECS (at Layer 4 or Layer 7). You can bind EIPs to ECSs to allow the access from the Internet. ELB expands the service capabilities of your applications and improves availability by eliminating single points of failures.	What Is Elastic Load Balance?

Related Services

ECS, EIP, NAT Gateway, and ELB

Figure 3-2 Web application or website hosting

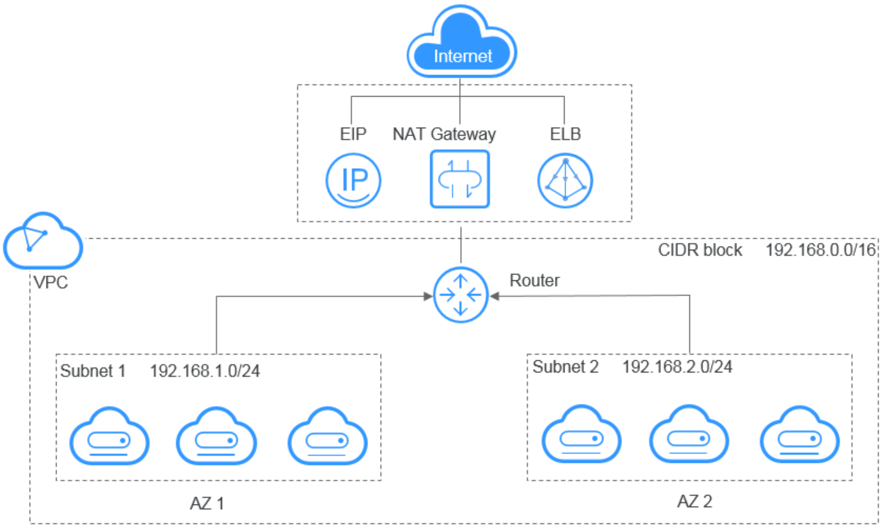

Web Application Access Control

Scenario

You can create a VPC and security groups to host multi-tier web applications in different security zones. You can associate web servers and database servers with different security groups and configure different access control rules for security groups. You can launch web servers in a publicly accessible subnet, but run database servers in subnets that are not publicly accessible. This arrangement ensures high security.

Related Services

ECS

Figure 3-3 Web application access control

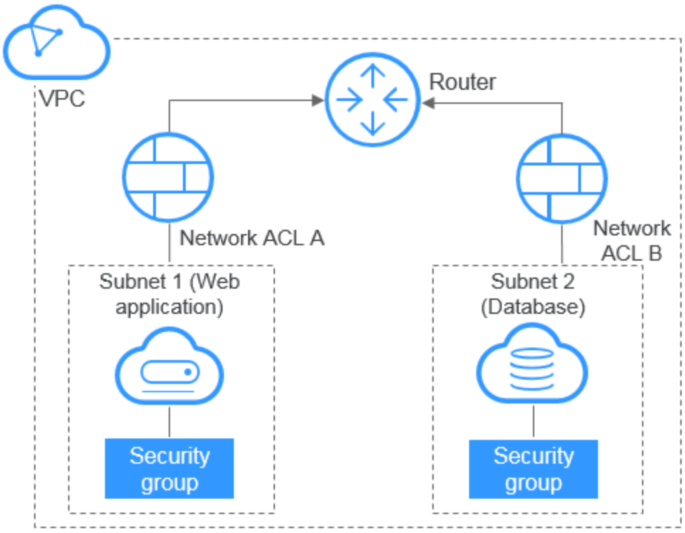

VPC Connectivity Options

Scenario

You can use the following cloud products to allow two VPCs to communicate with each other.

Table 3-2 Connecting VPCs

Cloud Service	Application Scenario	Description	Related Operations
VPC Peering	Connect VPCs in the same region.	You can request a VPC peering connection with another VPC in your account or in another account, but the two VPCs must be in the same region. VPC peering connections are free.	**Creating a VPC Peering Connection with Another VPC in Your Account** **Creating a VPC Peering Connection with a VPC in Another Account**
Cloud Connect	Connect VPCs in different regions.	Cloud Connect allows you to connect two VPCs in the same account or in different accounts even they are in different regions.	**Communication Between VPCs Across Regions**

Cloud Service	Application Scenario	Description	Related Operations
VPN	Use VPN to connect VPCs across regions at a low cost.	VPN uses an encrypted communications tunnel to connect VPCs in different regions and sends traffic over the Internet. It is inexpensive, easy to configure, and easy to use. However, VPN connections may be affected by the Internet quality.	**Connecting to a VPC Through a VPN**

Related Services

ECS, Cloud Connect, and VPN

Figure 3-4 VPC connectivity options

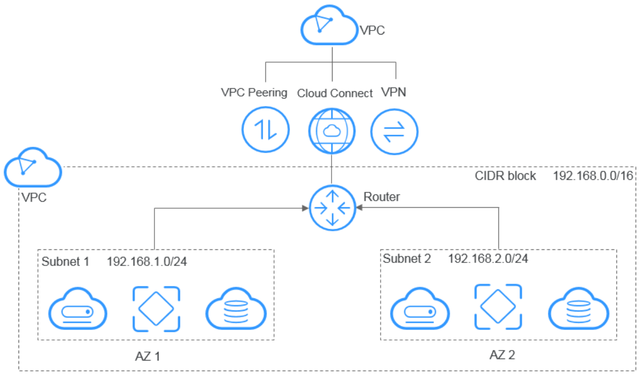

Hybrid Cloud Deployment

Scenario

If you have an on-premises data center and you do not want to migrate all of your business to the cloud, you can build a hybrid cloud, so that you can keep core data in your data center.

Table 3-3 Connecting to an on-premises data center

Cloud Service	Application Scenario	Description	Related Operations
VPN	Use VPN to connect a VPC to an on-premises data center with a low cost.	VPN uses an encrypted communications tunnel to connect a VPC on the cloud to an on-premises data center and sends traffic over the Internet. It is inexpensive, easy to configure, and easy to use. However, VPN connections may be affected by the Internet quality.	**Connecting to a VPC Through a VPN** **Layer 2 Connection Gateway**
Direct Connect	Use a physical connection to connect a VPC to an on-premises data center.	Direct Connect provides physical connections between VPCs and data centers. It has the advantages of low latency and is very secure. Direct Connect is a good choice when there are strict requirements on network transmission quality.	**Accessing Multiple VPCs Using a Connection** **Layer 2 Connection Gateway**
Cloud Connect	Connect VPCs in different regions.	Cloud Connect allows the loading of Direct Connect virtual gateways to a Cloud Connect connection, interconnecting an on-premises data center with VPCs across regions.	**Communication Between VPCs Across Regions** **Communication Between Data Centers and VPCs in Different Regions**

Related Services

ECS, Direct Connect, Cloud Connect, and VPN

Figure 3-5 Hybrid cloud deployment

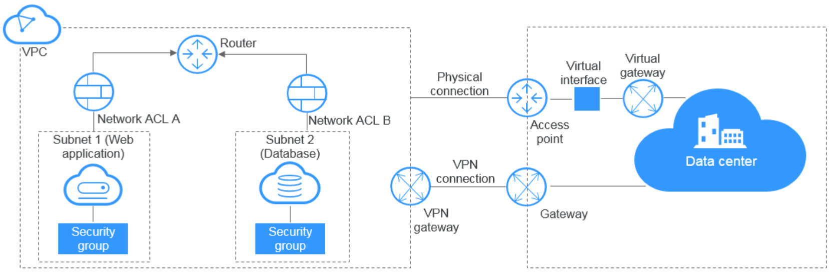

Copyright © Huawei Technologies Co., Ltd.

4 Functions

Table 4-1 lists common VPC functions.

Before using the VPC service, you are advised to learn basic concepts, such as subnets, route tables, security groups, and EIPs, to better understand the functions provided by the VPC service.

Table 4-1 Common VPC functions

Category	Function	Description
VPC and Subnet	VPC	A VPC provides an isolated virtual network for your cloud resources. You can configure and manage the network as required. You can create VPCs, modify basic information about VPCs, delete VPCs, and export the VPC list. For details, see **Creating a VPC**.
	Subnet	A subnet is a network plane for managing cloud resources in a VPC. All cloud resources must be deployed in a subnet. You can create subnets, modify subnet information, and delete subnets. For details, see **Creating a VPC**.

Category	Function	Description
	Route Table	A route table contains routes, which are used to determine where traffic is directed. When you create a VPC, the system automatically creates a default route table. The route table ensures that all subnets in the VPC can communicate with each other. You can also add custom routes to control where traffic is directed. You can add, query, modify, and delete routes. For details, see **Route Table Overview**. In some regions, you can visit the route table module directly from the navigation pane on the left of the network console. You can associate subnets with a route table to facilitate flexible route management. For details, see **Route Table (Route Table Module Can Be Directly Accessed from the Navigation Pane)**.
	Virtual IP Address	A virtual IP address can be shared among multiple ECSs. An ECS can have both private and virtual IP addresses, and you can access the ECS through either IP address. In addition, the virtual IP address has the same network access capability as the private IP address. Virtual IP addresses are used for high availability as they make active/standby ECS switchover possible. You can assign and release virtual IP addresses, bind a virtual IP address to an EIP or ECS, and access a virtual IP address through an EIP, a VPN, Direct Connect, or VPC peering connection. For details, see **Virtual IP Address Overview**.
	IPv4 and IPv6 Dual-Stack Network	IPv4 and IPv6 dual stack allows your resources to use both the IPv4 and IPv6 addresses for private and public network communication. HUAWEI CLOUD supports the IPv4/IPv6 dual stack. You can create an IPv4/IPv6 dual-stack network or add an IPv6 subnet to a VPC to form a dual-stack network. For details, see **IPv4 and IPv6 Dual-Stack Network**.

Category	Function	Description
	VPC Flow Log	A VPC flow log records information about the traffic going to and from a VPC. VPC flow logs help you monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification. You can create, view, enable, disable, and delete VPC flow logs. For details, see **VPC Flow Log Overview**.
Access Control	Security Group	A security group is a collection of access control rules for ECSs that have the same security protection requirements and are mutually trusted within a VPC. After a security group is created, you can create different access rules for the security group to protect the ECSs that it contains. You can create and delete security groups, add multiple security group rules and replicate security group rules, modify, delete, import or export security group rules, view the security group of an ECS, change the security group of an ECS, and add cloud resources to or remove them from a security group. For details, see **Security Group Overview**.
	Network ACL	A network ACL is an optional layer of security for your subnets. You associate one or more subnets with a network ACL. The network ACL can help you control traffic in and out of the subnets. You can create, view, modify, delete, enable, disable network ACLs, associate subnets with or disassociate them from network ACLs, and add, modify, change the sequence of, enable, disable, and delete network ACL rules. For details, see **Network ACL Overview**.
EIP and Bandwidth	EIP	The Elastic IP (EIP) service enables you to use static public IP addresses and scalable bandwidths to connect your cloud resources to the Internet. EIPs can be bound to or unbound from cloud resources. You can assign EIPs, bind EIPs to cloud resources, unbind EIPs from cloud resources, release EIPs, modify EIP bandwidth, and upgrade static BGP to dynamic BGP. For details, see **EIP Overview**.

Category	Function	Description
	Shared Bandwidth	Shared bandwidth allows multiple EIPs to share the same bandwidth. The ECSs, BMSs, and load balancers that are bound with EIPs in the same region can use the same shared bandwidth. You can assign, modify, delete a shared bandwidth, add EIPs to a shared bandwidth, and remove EIPs from a shared bandwidth. For details, see **Shared Bandwidth Overview**.
	Shared Data Package	A shared data package provides a quota for data usage. Such packages are cost-effective and easy to use. Shared data packages take effect immediately after your purchase. If you have subscribed to pay-per-use EIPs using bandwidth billed by traffic in a region and buy a shared data package in the same region, the EIPs will use the shared data package. After the package quota is used up or the package expires, the EIPs will continue to be billed on a pay-per-use basis. For details, see **Shared Data Package Overview**.
	Bandwidth Add-On Package	A bandwidth add-on package is used to temporarily increase the maximum shared or dedicated bandwidth of a yearly/monthly EIP. You can purchase, modify, and unsubscribe from bandwidth add-on packages. For details, see **Bandwidth Add-On Package Overview**.
Resource Interconnection	VPC Peering Connection	A VPC peering connection is a network connection between two VPCs. A VPC peering connection allows two VPCs communicate with each other using private IP addresses as if they were in the same VPC. You can create a VPC peering connection between your own VPCs, or between your VPC and a VPC of another account within the same region. A VPC peering connection between VPCs in different regions will not take effect. You can create a VPC peering connection with another VPC in your account or with a VPC in another account. You can also view, modify, and delete VPC peering connections. For details, see **Creating a VPC Peering Connection**.

Category	Function	Description
	Layer 2 Connection Gateway	A Layer 2 connection gateway (L2CG) is a virtual tunnel gateway that can work with a Direct Connect or VPN connection to establish network communication between the cloud and on-premises networks. The gateway allows you to migrate data center or private cloud services to the cloud without changing subnets and IP addresses. You can buy, query, modify, and delete L2CGs, and create, query, and delete Layer 2 connections. For details, see **Layer 2 Connection Gateway**.
Monitoring	Viewing Metrics	After the VPC service becomes available to you, you can view the bandwidth and EIP usage through Cloud Eye, create and set alarm rules, and customize the monitored objects and notification policies without adding plug-ins. For details, see **Supported Metrics**.
Auditing	Viewing Audit Logs	With CTS, you can record operations performed on the VPC service for further query, audit, and backtracking purposes. You can view and export operation records of the last seven days on the CTS console. For details, see **Supported VPC Operations**.
Tag	Tag Management	Tags help you identify and manage cloud resources. You can manage VPC tags, subnet tags, and tags on HUAWEI CLOUD.
Permissions	Permissions Management	You can use Identity and Access Management (IAM) to implement fine-grained permissions management for your VPCs, allowing enterprises to set different access permissions based on organizations and responsibilities. You can create an IAM user, grant permissions to the user, and create custom VPC policies. For details, see **Permissions Management**.

5 Notes and Constraints

VPC

Table 5-1 lists the quotas for VPC resources per region for your account.

Table 5-1 VPC resource quotas

Resource	Default Quota	How to Increase Quota
VPCs per account	5	Submit a service ticket.
Subnets per account	100	Submit a service ticket.
Security groups per account	100	Submit a service ticket.
Security group rules per account	5000	Submit a service ticket.
Routes per route table	100	This quota cannot be increased.
Default route table per VPC	1	This quota cannot be increased.
VPC peering connections per region	50	This quota cannot be increased.
Network ACLs per account	200	Submit a service ticket.
Layer 2 connection gateways per account	5	Submit a service ticket.

NOTE

- The preceding quota applies to a single account.
- It is recommended that a network ACL contain no more than 20 rules in one direction. Otherwise, its performance may deteriorate.
- By default, each account can create only one Layer 2 connection gateway during the open beta test.

EIP

Note the following:

- Each EIP can be used by only one cloud resource at a time. The EIP and the cloud resource must be in the same region.
- Each cloud account can have a maximum of 20 EIPs. If you need more EIPs, submit a service ticket to request quota increase.
- An EIP that has already been bound to a cloud resource cannot be bound to another resource.
- If an EIP is billed by bandwidth, the maximum bandwidth is 2000 Mbit/s. If you need more bandwidth, submit a service ticket or contact your account manager.
- If an EIP is billed by traffic, the maximum bandwidth is 300 Mbit/s.
- You cannot bind or unbind EIPs that are frozen. (For example, EIPs may be frozen if unexpected operations are performed or your account balance is insufficient.)
- To request quota increase, your account must have valid orders and must be continuously using cloud resources. If you have released resources immediately after subscribing to them multiple times, your request for quota increase will be declined.
- You can only release unbound EIPs.
- You cannot buy an EIP that you have released if it is currently being used by another user.
- The price of a pay-per-use EIP includes the retention fee and the bandwidth price. If you unbind an EIP but do not release it, the EIP will continue to be billed and the price includes the retention fee and the bandwidth price. The moment you bind an EIP to an instance, the retention fee is no longer included in the EIP price.
- If you have increased the EIP quota but you have not used the quota for a long time, HUAWEI CLOUD will reduce the quota to the default value. If you want to increase the quota again, submit a service ticket.
- If you use EIP resources in violation of applicable laws and regulations, HUAWEI CLOUD has the right to reclaim the EIP resources and stop providing services to you.
- EIPs cannot be transferred across accounts.

Bandwidth

- You can only add pay-per-use EIPs to a shared bandwidth.
- Shared bandwidths are billed by bandwidth or 95th percentile bandwidth (enhanced). If you select the bandwidth option, the minimum shared bandwidth that you can purchase is 5 Mbit/s. If you select the enhanced 95th percentile bandwidth option, the minimum shared bandwidth that you can purchase is 300 Mbit/s.
- A maximum of 20 EIPs that are billed on a pay-per-use basis can be added to a shared bandwidth. If you want to add more EIPs to a shared bandwidth, submit a service ticket to request for a quota increase.
- Each account can have a maximum of 5 shared bandwidths. If you need more shared bandwidths, submit a service ticket to request quota increase.

- For a yearly/monthly bandwidth, you can only increase the bandwidth size within the bandwidth validity period. You can specify a smaller bandwidth size only when you renew the bandwidth.
- You will be billed only for the bandwidth in the outbound direction.

📖 NOTE

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

Shared Data Package

- A shared data package takes effect only for EIPs using bandwidth billed by traffic. Two types of shared data packages are available: static BGP (for static BGP bandwidth) and dynamic BGP (for dynamic BGP bandwidth).
- A shared data package cannot be unsubscribed.

Bandwidth Add-On Package

- A bandwidth add-on package takes effect only for bandwidth that is billed on a yearly/monthly basis.
- The minimum duration of a bandwidth add-on package is one day. The start time, end time, and bandwidth size of an obtained bandwidth add-on package cannot be changed.

📖 NOTE

For details about how to submit a service ticket, see **Submitting a Service Ticket**.

6 VPC and Other Services

Figure 6-1 shows the relationship between VPC and other services.

Figure 6-1 VPC and other services

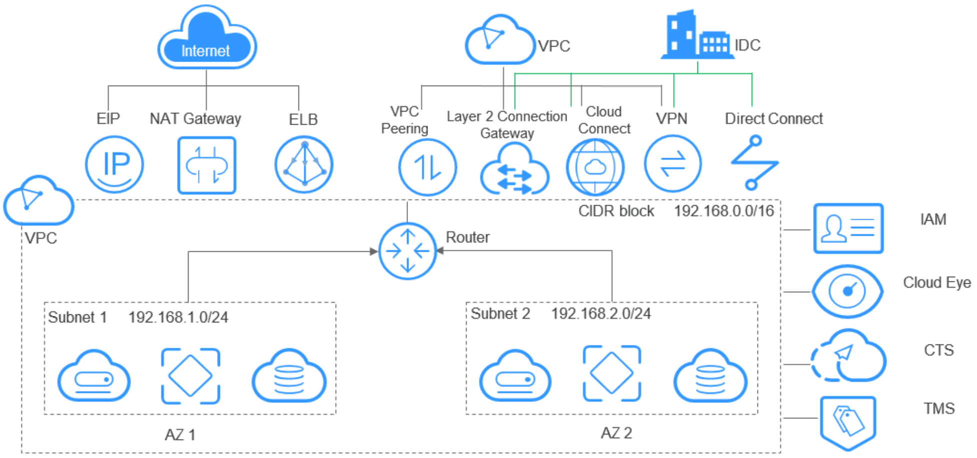

Table 6-1 Related services

Interactive Function	Service	Reference
Secure networks for ECSs.	Elastic Cloud Server (ECS)	**Adding a Security Group Rule**
Connect ECSs in a VPC to the Internet.	Elastic IP (EIP)	**Configuring the VPC of ECSs That Access the Internet Using EIPs**
	NAT Gateway	**Using SNAT to Access the Internet**
Connect a VPC to a local data center.	Virtual Private Network (VPN)	**Virtual Private Network**

Interactive Function	Service	Reference
	Direct Connect	**Direct Connect**
Connect VPCs across regions.	Cloud Connect	**Communication Between VPCs Across Regions**
Distribute incoming traffic to multiple ECSs in a VPC.	Elastic Load Balance (ELB)	**Elastic Load Balance**
If you need to assign different permissions to employees in your enterprise to access your VPC resources, IAM is a good choice for fine-grained permissions management.	Identity and Access Management (IAM)	**Identity and Access Management**
Check the bandwidth and traffic usage.	Cloud Eye	**Viewing Metrics**
Record VPC-related operations for later query, audit, and backtracking.	Cloud Trace Service (CTS)	**Viewing Audit Logs**
Tags identify VPC resources for purposes of easy categorization and quick search.	Tag Management Service (TMS)	**Managing VPC Tags**

7 Billing

Billing Item

The VPC service is free of charge.

Table 7-1 Billing items

Billing Item	Description
EIP	EIPs are required if your resources need to access the Internet.
L2CG	L2CGs can be billed based on a pay-per-use basis. L2CGs are free for use during OBT.

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

The EIP price varies according to the billing mode.

Table 7-2 EIP billing details

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

NOTE

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

Billing Options

When you use an instance, such as an Elastic Cloud Server (ECS), both outbound bandwidth and inbound bandwidth are used. **However, HUAWEI CLOUD only charges fees for outbound bandwidth usage.**

The public network bandwidth can be billed by fixed bandwidth or by traffic usage.

- By fixed bandwidth: You will be billed based on the bandwidth you specify. During the usage, the actual outbound bandwidth will not exceed the specified bandwidth.
- By traffic usage: You will be billed based on a pay-per-use basis for actual traffic usage. To prevent excess charges due to sudden traffic spikes, you can set a peak value for the outbound bandwidth.
- You can select the billing mode based on bandwidth usage. If your required bandwidth is less than or equal to 5 Mbit/s, billing by fixed bandwidth monthly is more favorable than billing by traffic usage. If your required bandwidth is higher than 5 Mbit/s and the bandwidth usage is greater than 20%, billing by fixed bandwidth is more favorable.

Figure 7-1 EIP billing modes

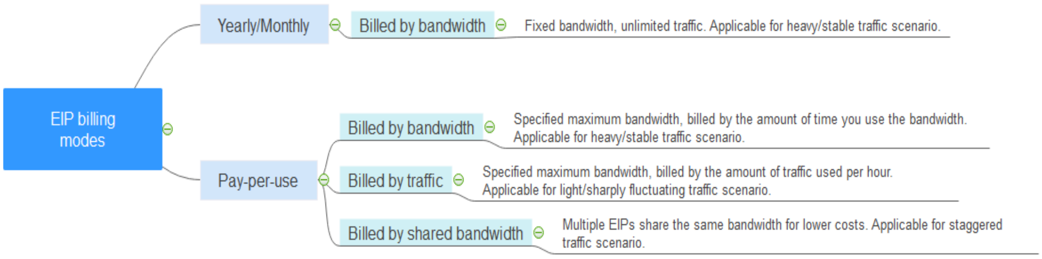
EIP billing modes
Yearly/Monthly
Billed by bandwidth
Fixed bandwidth, unlimited traffic. Applicable for heavy/stable traffic scenario.
Pay-per-use
Billed by bandwidth
Specified maximum bandwidth, billed by the amount of time you use the bandwidth. Applicable for heavy/stable traffic scenario.
Billed by traffic
Specified maximum bandwidth, billed by the amount of traffic used per hour. Applicable for light/sharply fluctuating traffic scenario.
Billed by shared bandwidth
Multiple EIPs share the same bandwidth for lower costs. Applicable for staggered traffic scenario.

Configuration Changes

- Modifying the bandwidth size has the following impacts on the price:

Table 7-3 Impacts after change

Current Billing Mode	Change Scenario	Impact After Successful Change
Pay-per-use	Increase or decrease bandwidth size	The change will take effect immediately.

Current Billing Mode	Change Scenario	Impact After Successful Change
Yearly/ Monthly	Increase bandwidth	The increased bandwidth size will take effect immediately and the price difference will be billed accordingly. The following prices are for reference only. The actual prices are subject to **Pricing Details**. On November 1, 2018, a customer purchases a 1 Mbit/s bandwidth for one month. The price is 18.4 CNY per month. The customer pays for the bandwidth with the account balance of 18.4 CNY. On November 24, 2018, the customer increases the bandwidth to 5 Mbit/s, at a price of 92 CNY per month. In this case, the number of remaining days is 6 (30 - 24), and the price for the upgrade is 92/30 x 6 - 18.4/30 x 6 = 14.72 CNY. For more information, see **Pricing of a Changed Specification**.
	Decrease bandwidth after renewal	The decreased bandwidth size will take effect in the first billing cycle after a successful renewal. ● A renewal cannot be canceled after the payment has been made. ● After a successful renewal in case of a decreased bandwidth size, the bandwidth cannot be modified in the first billing cycle.
	Buy a bandwidth add-on package to temporarily increase the available bandwidth	A bandwidth add-on package is used to temporarily increase the yearly/monthly bandwidth. The validity period of the add-on package can be any number of days (full days only) that fall within the bandwidth subscription period. When the add-on package expires, the bandwidth will automatically default to its original specifications.

- Modifying the EIP billing modes has the following impacts on the price.

Figure 7-2 Billing change

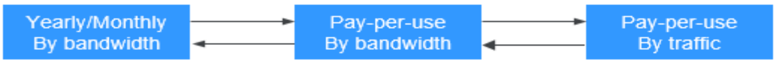

Table 7-4 Flexible billing

Current Billing Mode	Change	Impact After Successful Change
Pay-per-use	Changing the billing option between by traffic and by bandwidth	The change will take effect immediately.
	Changing the billing mode to yearly/monthly	You can change the billing mode from pay-per-use to yearly/monthly on the EIP console or in the billing center. After the change is successful, the new billing mode will take effect immediately. An EIP that is billed by traffic on a pay-per-use basis cannot be changed to an EIP that is billed by traffic on a yearly/monthly basis. You must first change the EIP to be billed by bandwidth and then change its billing mode from pay-per-use to yearly/monthly.
Yearly/Monthly	Changing the billing mode to pay-per-use	You can change the billing mode from yearly/monthly to pay-per-use in the billing center. The new billing mode takes effect only after the yearly/monthly mode expires. An EIP billed by bandwidth on a yearly/monthly basis can only be changed to an EIP billed by bandwidth on a pay-per-use basis. After you change the billing mode, you can configure the EIP to be billed by traffic.

Renewal

You can renew your resources on the **Renewals** page of the management console. For more details, see **Renewal Management**.

Expiration and Overdue Payment

If your account is in arrears, you can view the arrears details in the Billing Center. To prevent related resources from being stopped or released, you need to top up your account within the specified period. For details, see **Repaying Arrears**.

8 Permissions Management

If you need to assign different permissions to employees in your enterprise to access your VPC resources, 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 HUAWEI CLOUD resources.

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

If your HUAWEI CLOUD account does not need individual IAM users for permissions management, you may skip over this section.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see the **IAM Service Overview**.

VPC Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups to which they are added and can perform specified operations on cloud services based on the permissions.

VPC is a project-level service deployed and accessed in specific physical regions. To assign VPC permissions to a user group, specify the scope as region-specific projects and select projects for the permissions to take effect. If **All projects** is selected, the permissions will take effect for the user group in all region-specific projects. When accessing VPC, the users need to switch to a region where they have been authorized to use VPC.

You can grant users permissions by using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that defines permissions related to user responsibilities. This mechanism provides only a limited number of service-level roles for authorization. When using roles to grant permissions, you need to also assign other roles on which the permissions depend to take effect. 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, you can grant VPC users only the permissions for managing a certain type of resources. Most policies define permissions based on APIs. For the API actions supported by VPC, see **Permissions Policies and Supported Actions**.

Table 8-1 lists all the system-defined roles and policies supported by VPC.

Table 8-1 System-defined roles and policies supported by VPC

Policy Name	Description	Policy Type	Dependencies
VPC FullAccess	All operations on VPC.	System-defined policy	None
VPC ReadOnlyAccess	Read-only permissions on VPC.	System-defined policy	None
VPC Administrator	All operations on VPC. To be granted this permission, users must also have the **Tenant Guest** permission.	System-defined role	Dependent on the **Tenant Guest** policy.

Table 8-2 lists the common operations supported by each system-defined policy or role of VPC. Select the policies or roles as required.

Table 8-2 Common operations supported by each system-defined policy or role of VPC

Operation	VPC ReadOnlyAccess	VPC Administrator	VPC FullAccess
Creating a VPC	x	√	√
Modifying a VPC	x	√	√
Deleting a VPC	x	√	√
Viewing VPC information	√	√	√
Creating a subnet	x	√	√

Operation	VPC ReadOnlyAccess	VPC Administrator	VPC FullAccess
Viewing subnet information	√	√	√
Modifying a subnet	x	√	√
Deleting a subnet	x	√	√
Creating a security group	x	x	√
Viewing security group information	√	x	√
Modifying a security group	x	x	√
Deleting a security group	x	x	√
Adding a security group rule	x	x	√
Viewing a security group rule	√	x	√
Modifying a security group rule	x	x	√
Deleting a security group rule	x	x	√
Creating a network ACL	x	√	√
Viewing a network ACL	√	√	√
Modifying a network ACL	x	√	√
Deleting a network ACL	x	√	√

Operation	VPC ReadOnlyAccess	VPC Administrator	VPC FullAccess
Adding a network ACL rule	x	√	√
Modifying a network ACL rule	x	√	√
Deleting a network ACL rule	x	√	√
Creating a VPC peering connection	x	√	√
Modifying a VPC peering connection	x	√	√
Deleting a VPC peering connection	x	√	√
Creating a route table	x	√	√
Deleting a route table	x	√	√
Adding a route	x	√	√
Modifying a route	x	√	√
Deleting a route	x	√	√

Helpful Links

- **What Is IAM?**
- **Creating a User and Granting VPC Permissions**
- **Permissions Policies and Supported Actions**

9 Basic Concepts

9.1 Subnet

A subnet is a range of IP addresses in your VPC and provides IP address management and DNS resolution functions for ECSs in it. The IP addresses of all ECSs in a subnet belong to the subnet.

Figure 9-1 Subnet

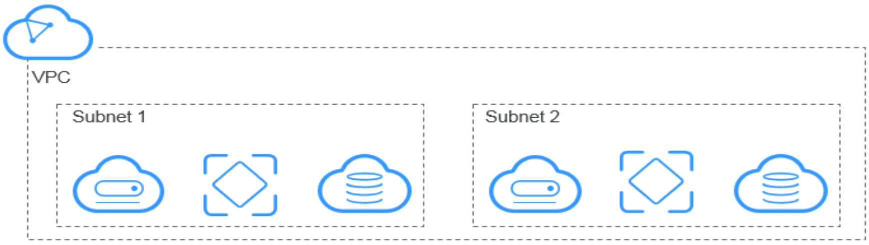

By default, ECSs in all subnets of the same VPC can communicate with one another, but ECSs in different VPCs cannot.

To enable ECSs in different VPCs but in the same region to communicate with one another, you can create a VPC peering connection. For details, see **VPC Peering Connection**.

9.2 Elastic IP

The Elastic IP (EIP) service enables you to use static public IP addresses and scalable bandwidths to connect your cloud resources to the Internet. EIPs can be bound to or unbound from ECSs, BMSs, virtual IP addresses, NAT gateways, or load balancers. Various billing modes are provided to meet diverse service requirements.

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

Figure 9-2 Accessing the Internet using an EIP

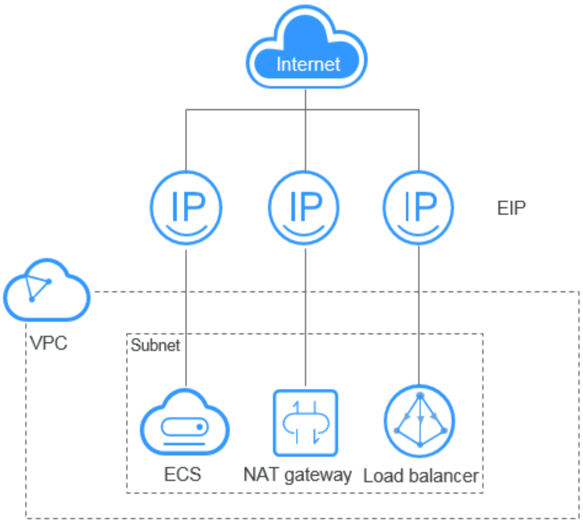

9.3 Route Table

In some regions, you can visit the route table module directly from the navigation pane on the left of the network console. You can associate subnets with a route table to facilitate flexible route management.

For details about the regions that you can visit the route table module directly from the navigation pane, see **Route Table (Route Table Module Can Be Directly Accessed from the Navigation Pane)**, **Default Route Table and Custom Route Table**, and **Route**.

For details about the regions that you have to visit the route table module through the VPC details page, see **Route Table (Route Table Module Can Be Accessed Through the VPC Details Page)**.

Route Table (Route Table Module Can Be Directly Accessed from the Navigation Pane)

A route table contains a set of rules, called routes, that are used to control where inbound and outbound subnet traffic is forwarded within a VPC. Each subnet in a VPC must be associated with a route table. A subnet can only be associated with one route table at a time, but you can associate multiple subnets with the same route table.

Figure 9-3 Route Table

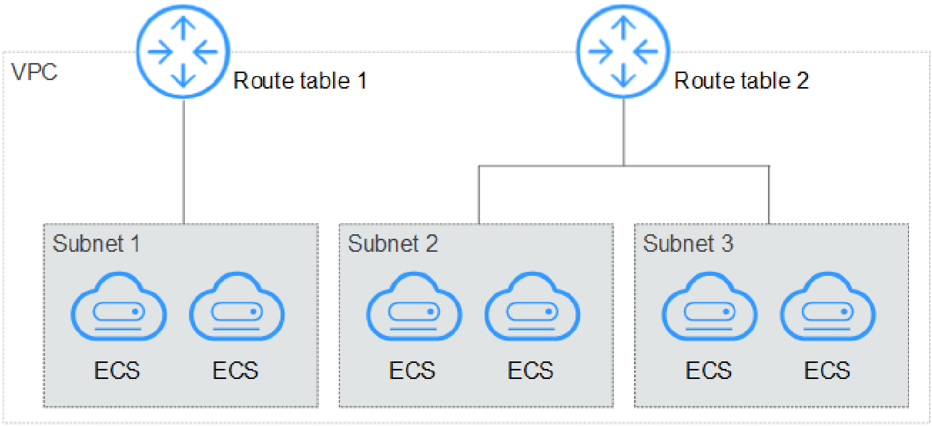

Default Route Table and Custom Route Table

When you create a VPC, the system automatically generates a default route table for the VPC. If you create a subnet in the VPC, the subnet automatically associates with the default route table. You can add, delete, and modify routes in the default route table, but cannot delete the table. When you create a VPN, Direct Connect, or Cloud Connect connection, the default route table automatically delivers a route that cannot be deleted or modified. If you want to modify or delete the route, you can associate your subnet with a custom route table and replicate the route to the custom route table to modify or delete it.

You can also create a custom route table and associate subnets that have the same routing requirements with this table. Custom route tables can be deleted if they are no longer required.

NOTE

To use a custom route table, you need to submit a service ticket. You need to click **Increase quota** on the **Create Route Table** page or choose **More** > **Service Tickets** > **Create Service Ticket** in the upper right corner of the page. For more information, see **Submitting a Service Ticket**.

Route

A route is configured with the destination, next hop type, and next hop to determine where the network traffic is directed. Routes are classified into system routes and custom routes.

- System route: Routes that are automatically added by the system and cannot be modified or deleted.

  After a route table is created, the system automatically adds the following system routes to the route table, so that instances in a VPC can communicate with each other.

  – Routes whose destination is 100.64.0.0/10 or 198.19.128.0/20.
  – Routes whose destination is a subnet CIDR block.

  NOTE

  In addition to the preceding system routes, the system automatically adds a route whose destination is 127.0.0.0/8. This is the local loopback address.

- Custom route: A route that can be modified and deleted. The destination of a custom route cannot overlap with that of a system route.

  You can add a custom route and configure the destination, next hop type, and next hop in the route to determine where the network traffic is directed. **Table 9-1** lists the supported types of next hops.

Table 9-1 Next hop type

Next Hop Type	Description
ECS	Traffic intended for the destination is forwarded to an ECS in the VPC.
Extension NIC	Traffic intended for the destination is forwarded to the extension NIC of an ECS in the VPC.
VPN gateway	Traffic intended for the destination is forwarded to a VPN gateway.
Cloud connection	Traffic intended for the destination is forwarded to a cloud connection.
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.
NAT gateway	Traffic intended for the destination is forwarded to a NAT gateway.
VPC peering connection	Traffic intended for the destination is forwarded to a VPC peering connection.
Virtual IP address	Traffic intended for the destination is forwarded to a virtual IP address and then sent to active and standby ECSs to which the virtual IP address is bound.
VPC endpoint	Traffic intended for the destination is forwarded to a VPC endpoint.
Cloud container	Traffic intended for the destination is forwarded to a cloud container.

NOTE

- When you add a custom route to a default route table, the next hop type cannot be set to VPN gateway.
- If you have to specify the destination when creating a service, a system route is delivered. If you do not need to specify a destination when creating a service, a custom route that can be modified or deleted is delivered automatically.

  For example, you do not need to specify a destination when creating a NAT gateway, the system automatically delivers a custom route that you can modify or delete. However, when you create a VPN gateway, you need to specify the remote subnet, that is, the destination of a route. In this case, the system delivers a system route. If the route destination can be modified on the **Route Tables** page, the destination will be inconsistent with that configured remote subnet. To modify the destination, you can go to the specific service page to modify the remote subnet, then the route destination will be changed accordingly.

Route Table (Route Table Module Can Be Accessed Through the VPC Details Page)

A route table contains a set of rules that determine where network traffic is directed. You can add routes to a route table to enable other ECSs in a VPC to access the Internet through the ECS that has a bound EIP.

You can use a route table configured in standalone or active/standby mode.

- **Figure 9-4** shows the route table configured in standalone mode.

Figure 9-4 Route table configured in standalone mode

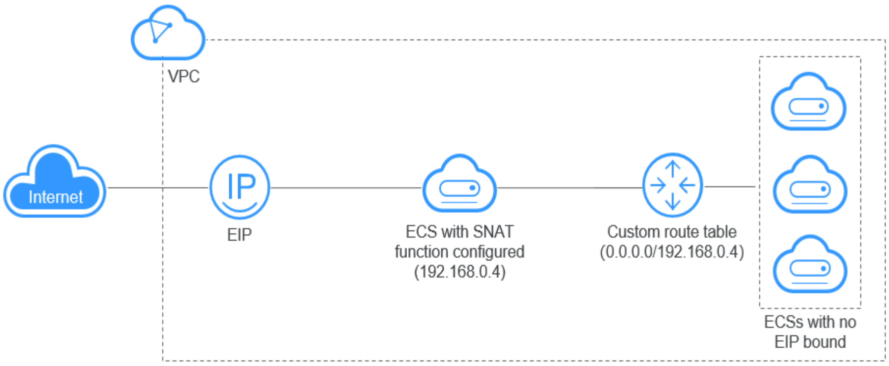
VPC
Internet
EIP
ECS with SNAT function configured (192.168.0.4)
Custom route table (0.0.0.0/192.168.0.4)
ECSs with no EIP bound

In standalone mode, ECSs in a VPC that do not have EIPs bound access the Internet through an ECS that has an EIP bound and has SNAT function configured.

You can create a route table for the VPC used by ECSs that do not have EIPs bound to enable these ECSs to access the Internet. The next hop in the route table is the private IP address of the ECS that has an EIP bound (the private IP address of the SNAT server).

- **Figure 9-5** shows the route table configured in active/standby mode.

Figure 9-5 Route table configured in active/standby mode

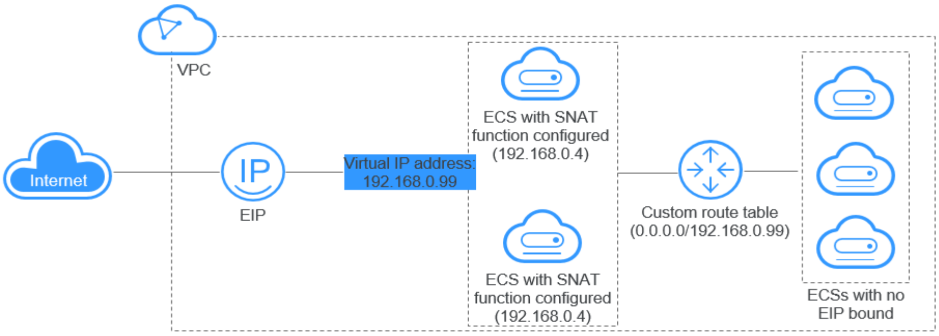
VPC
Internet
IP
EIP
Virtual IP address: 192.168.0.99
ECS with SNAT function configured (192.168.0.4)
ECS with SNAT function configured (192.168.0.4)
Custom route table (0.0.0.0/192.168.0.99)
ECSs with no EIP bound

In active/standby mode, ECSs in a VPC that do not have EIPs bound access the Internet through two ECSs that have EIPs bound and have the SNAT function configured.

In active/standby mode, you can add a route table for the VPC used by ECSs that do not have EIPs bound to enable these ECSs to access the Internet. The next hop in the route table is the virtual IP address of the two ECSs that have EIPs bound.

9.4 Security Group

A security group is a collection of access control rules for ECSs that have the same security protection requirements and are mutually trusted. After a security group is created, you can create different access rules for the security group, these rules will apply to any ECS that the security group contains.

Your account automatically comes with a default security group (**Sys-default**). The default security group allows all outbound traffic, denies all inbound traffic, and allows all traffic between ECSs in the group. Your ECSs in this security group can communicate with each other already. You do not need to add additional rules.

Figure 9-6 illustrates how the default security group works.

Figure 9-6 Default security group

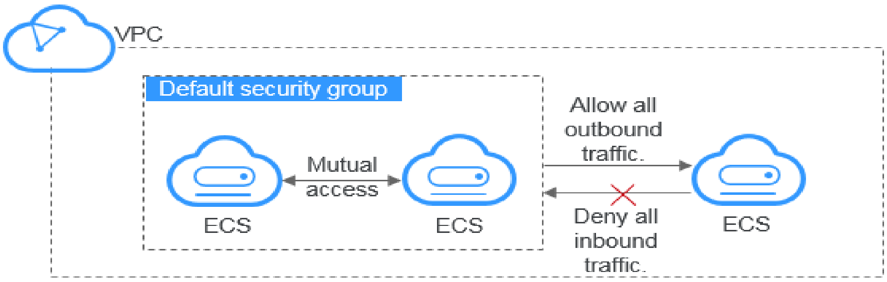

Table 9-2 describes the default rules for the default security group.

Table 9-2 Rules in the default security group (**Sys-default**)

Direction	Protocol	Port/ Range	Source/ Destination	Description
Outbound	All	All	Destination: 0.0.0.0/0	Allows all outbound traffic.
Inbound	All	All	Source: the current security group (for example, sg-*xxxxx*)	Allows communication among ECSs within the security group and denies all inbound traffic (incoming data packets).
Inbound	TCP	22	Source: 0.0.0.0/0	Allows all IP addresses to access Linux ECSs over SSH.
Inbound	TCP	3389	Source: 0.0.0.0/0	Allows all IP addresses to access Windows ECSs over RDP.

You can also create custom security groups and rules as required.

NOTE

If two ECSs are in the same security group but in different VPCs, the ECSs cannot communicate with each other. You can use VPC peering connections to enable communication between ECSs in different VPCs so that security groups can control traffic between the ECSs. For details about VPC connectivity, see **Application Scenarios**.

9.5 VPC Peering Connection

A VPC peering connection is a network connection between two VPCs in one region that enables you to route traffic between them using private IP addresses. ECSs in either VPC can communicate with each other just as if they were in the same region. You can create a VPC peering connection between your own VPCs, or between your VPC and another account's VPC within the same region. A VPC peering connection between VPCs in different regions will not take effect.

- Creating a VPC peering connection between VPCs in your account

Figure 9-7 Creating a VPC peering connection between VPCs in your account

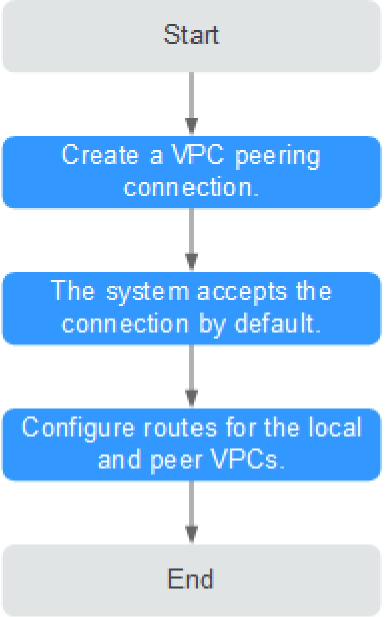
Start
Create a VPC peering connection.
The system accepts the connection by default.
Configure routes for the local and peer VPCs.
End

If you create a VPC peering connection between two VPCs in your account, the system automatically accepts the connection by default. You need to add routes for the local and peer VPCs to enable communication between the two VPCs.

- Creating a VPC peering connection with a VPC in another account

Figure 9-8 Creating a VPC peering connection with a VPC in another account

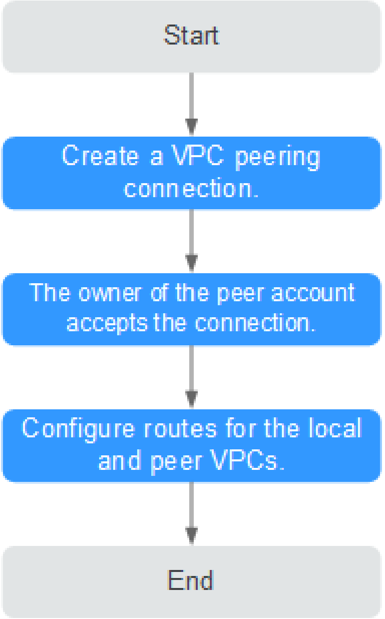
Start
Create a VPC peering connection.
The owner of the peer account accepts the connection.
Configure routes for the local and peer VPCs.
End

If you create a VPC peering connection between your VPC and a VPC that is in another account, the VPC peering connection will be in the **Awaiting acceptance** state. After the owner of the peer account accepts the connection, the connection status changes to **Accepted**. The owners of both the local and peer accounts must configure the routes required by the VPC peering connection to enable communication between the two VPCs.

If the local and peer VPCs have overlapping CIDR blocks, the routes added for the VPC peering connection may be invalid. Before creating a VPC peering connection between two VPCs that have overlapping CIDR blocks, ensure that none of the subnets in the two VPCs overlap. In this case, the created VPC peering connection enables communication between two subnets in the two VPCs.

You can run the **ping** command to check whether the two VPCs can communicate with each other.

9.6 Network ACL

A network ACL is an optional layer of security for your subnets. After you associate one or more subnets with a network ACL, the network ACL can help you control traffic in and out of the subnets.

Figure 9-9 Security groups and network ACLs

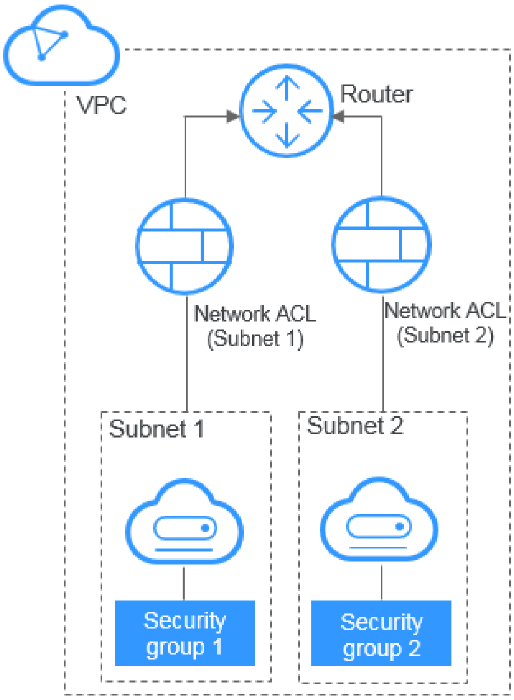

Similar to security groups, network ACLs control access to subnets and add an additional layer of defense to your subnets. Security groups only have the "allow" rules, but network ACLs have both "allow" and "deny" rules. You can use network ACLs together with security groups to implement access control that is both comprehensive and fine-grained.

Network ACL Basics

- Your VPC does not come with a default network ACL, but you can create one and associate it with a subnet if required. By default, each network ACL denies all inbound traffic to and outbound traffic from the associated subnet until you add rules.

- You can associate a network ACL with multiple subnets. However, a subnet can only be associated with one network ACL at a time.
- Each newly created network ACL is in the **Inactive** state until you associate subnets with it.
- Network ACLs are stateful. If you send a request from your instance and the outbound traffic is allowed, the response traffic for that request is allowed to flow in regardless of inbound network ACL rules. Similarly, if inbound traffic is allowed, responses to allowed inbound traffic are allowed to flow out, regardless of outbound rules.

  The timeout period of connection tracking varies according to the protocol. The timeout period of a TCP connection in the established state is 600s, and the timeout period of an ICMP connection is 30s. For other protocols, if packets are received in both directions, the connection tracking timeout period is 180s. If one or more packets are received in one direction but no packet is received in the other direction, the connection tracking timeout period is 30s. For protocols other than TCP, UDP, and ICMP, only the IP address and protocol number are tracked.

Default Network ACL Rules

By default, each network ACL has preset rules that allow the following packets:

- Packets whose source and destination are in the same subnet
- Broadcast packets with the destination 255.255.255.255/32, which is used to configure host startup information.
- Multicast packets with the destination 224.0.0.0/24, which is used by routing protocols.
- Metadata packets with the destination 169.254.169.254/32 and TCP port number 80, which is used to obtain metadata.
- Packets from CIDR blocks that are reserved for public services (for example, packets with the destination 100.125.0.0/16)
- A network ACL denies all traffic in and out of a subnet excepting the preceding ones. Table 9-3 shows the default network ACL rules. The default rules cannot be modified or deleted.

Table 9-3 Default network ACL rules

Direction	Priority	Action	Protocol	Source	Destination	Description
Inbound	*	Deny	All	0.0.0.0/0	0.0.0.0/0	Denies all inbound traffic.
Outbound	*	Deny	All	0.0.0.0/0	0.0.0.0/0	Denies all outbound traffic.

Rule Priorities

- Each network ACL rule has a priority value where a smaller value corresponds to a higher priority. Any time two rules conflict, the one with the higher priority is the one that get applied. The rule whose priority value is an asterisk (*) has the lowest priority.
- If multiple network ACL rules conflict, the rule with the highest priority takes effect. If you need a rule to take effect before or after a specific rule, you can insert that rule before or after the specific rule.

Application Scenarios

- If the application layer needs to provide services for users, traffic must be allowed to reach the application layer from all IP addresses. However, you also need to prevent illegal access from malicious users.

  Solution: You can add network ACL rules to deny access from suspect IP addresses.
- How can I isolate ports with identified vulnerabilities? For example, how do I isolate port 445 that can be exploited by WannaCry worm?

  Solution: You can add network ACL rules to deny access traffic from specific port and protocol, for example, TCP port 445.
- No defense is required for the communication within a subnet, but access control is required for communication between subnets.

  Solution: You can add network ACL rules to control traffic between subnets.
- For frequently accessed applications, a security rule sequence may need to be adjusted to improve performance.

  Solution: A network ACL allows you to adjust the rule sequence so that frequently used rules are applied before other rules.

9.7 Virtual IP Address

A virtual IP address can be shared among multiple ECSs. An ECS can have both private and virtual IP addresses, and you can access the ECS through either IP address. A virtual IP address has the same network access capabilities as a private IP address, including layer 2 and layer 3 communication in VPCs, access between VPCs using VPC peering connections, as well as access through EIPs, VPN connections, and Direct Connect connections.

A virtual IP address can be bound to multiple ECSs deployed in active/standby mode. You can bind an EIP to the virtual IP address. When the EIP is accessed from the Internet, the virtual IP address has made it possible to either the active or standby ECS, making ECSs highly fault tolerant.

Networking

Virtual IP addresses are used for high availability as they make active/standby ECS switchover possible. This way if one ECS goes down for some reason, the other one can take over and services continue uninterrupted. ECSs can be configured for HA or as load balancing clusters.

- **Networking mode 1**: HA

If you want to improve service availability and avoid single points of failure, you can deploy ECSs in the active/standby mode or deploy one active ECS and multiple standby ECSs. In this arrangement, the ECSs all use the same virtual IP address. If the active ECS becomes faulty, a standby ECS takes over services from the active ECS and services continue uninterrupted.

Figure 9-10 Networking diagram of the HA mode

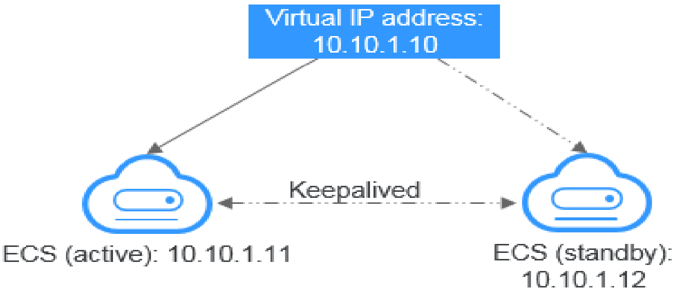

- In this configuration, a single virtual IP address is bound to two ECSs in the same subnet.
- Keepalived is then used to configure the two ECSs to work in the active/standby mode. Follow industry standards for configuring Keepalived. The details are not included here.

- **Networking mode 2**: HA load balancing cluster

  If you want to build a high-availability load balancing cluster, use Keepalived and configure LVS nodes as direct routers.

Figure 9-11 HA load balancing cluster

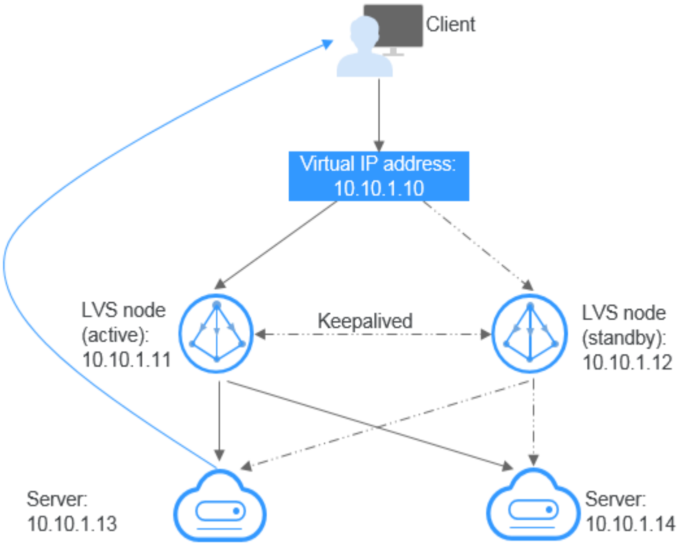
Client
Virtual IP address: 10.10.1.10
LVS node (active): 10.10.1.11
Keepalived
LVS node (standby): 10.10.1.12
Server: 10.10.1.13
Server: 10.10.1.14

- Bind a single virtual IP address to two ECSs.
- Configure the two ECSs as LVS nodes working as direct routers and use Keepalived to configure the nodes in the active/standby mode. The two ECSs will evenly forward requests to different backend servers.

- Configure two more ECSs as backend servers.
- Disable the source/destination check for the two backend servers.

Follow industry standards for configuring Keepalived. The details are not included here.

Application Scenarios

- Accessing the virtual IP address through an EIP

  If your application has high availability requirements and needs to provide services through the Internet, it is recommended that you bind an EIP to a virtual IP address.

- Using a VPN, Direct Connect, or VPC peering connection to access a virtual IP address

  To ensure high availability and access to the Internet, use a VPN for security and Direct Connect for a stable connection. The VPC peering connection is needed so that the VPCs in the same region can communicate with each other.

9.8 Layer 2 Connection Gateway

A Layer 2 connection gateway (L2CG) is a virtual tunnel gateway that can work with a Direct Connect or VPN connection to establish network communication between the cloud and on-premises networks. The gateway allows you to migrate data center or private cloud services to the cloud without changing subnets and IP addresses.

A Direct Connect or VPN connection establishes a Layer 3 network tunnel between the cloud and on-premises networks, but the subnets on the cloud and on-premises networks must not overlap. If the cloud and on-premises networks are on the same subnet but need to communicate with each other, you can use a L2CG to enable the communication at a Layer 2 network.

Figure 9-12 shows the networking diagram of a L2CG.

Figure 9-12 L2CG networking

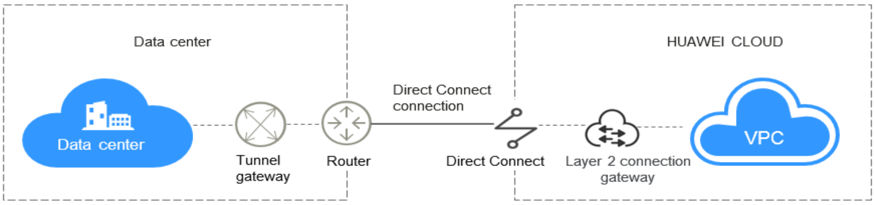

A L2CG is a tunnel gateway of a VPC and corresponds to a tunnel gateway of your data center. A L2CG can work together with a Direct Connect or VPN connection to establish a Layer 2 network between a VPC and your data center.

A Layer 2 connection connects a VPC subnet to a L2CG and specifies the L2CG to connect to the tunnel gateway in an enterprise data center so that the VPC subnet can communicate with the subnet in the enterprise data center at the Layer 2 network.

9.9 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters. AZs within a region are interconnected using high-speed optical fibers to support cross-AZ high-availability systems.

Figure 9-13 shows the relationship between regions and AZs.

Figure 9-13 Regions and AZs

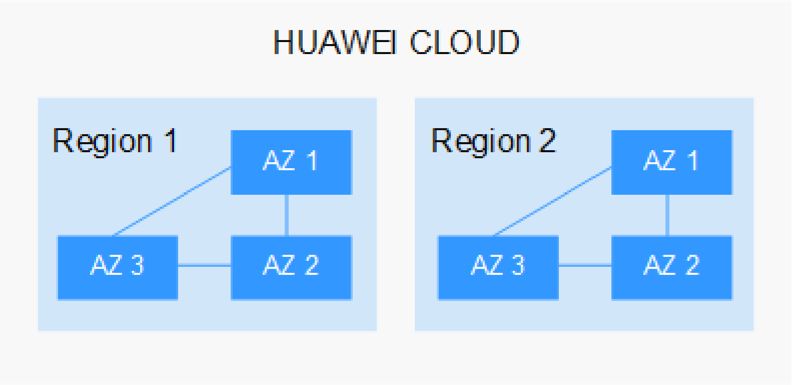

HUAWEI CLOUD provides services in many regions around the world. Select a region and AZ based on requirements.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for low network latency and quick access. Regions within the Chinese mainland provide the same infrastructure, BGP network quality, as well as resource operations and configurations. Therefore, if your target users are on the Chinese mainland, you do not need to consider the network latency differences when selecting a region.

Copyright © Huawei Technologies Co., Ltd.

  - If your target users are in Asia Pacific (excluding the Chinese mainland), select the **AP-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If your target users are in Africa, select the **AF-Johannesburg** region.
  - If your target users are in Europe, select the **EU-Paris** region.
- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For low network latency, deploy resources in the same AZ.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.