Virtual Private Cloud

Service Overview

Issue   01

Date   2024-11-28

Huawei Cloud Computing Technologies Co., Ltd.

Contents

1 What Is Virtual Private Cloud?

VPC Overview

Virtual Private Cloud (VPC) allows you to provision logically isolated virtual private networks for cloud resources, such as cloud servers, containers, and databases. You can create subnets, security groups, network ACLs, route tables, and more to manage cloud resources flexibly. You can also use EIPs to connect cloud resources in VPCs to the Internet, and use Direct Connect and VPN to connect on-premises data centers to VPCs to build a hybrid cloud network.

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 following describes the basics, security, connectivity, and O&M of VPCs.
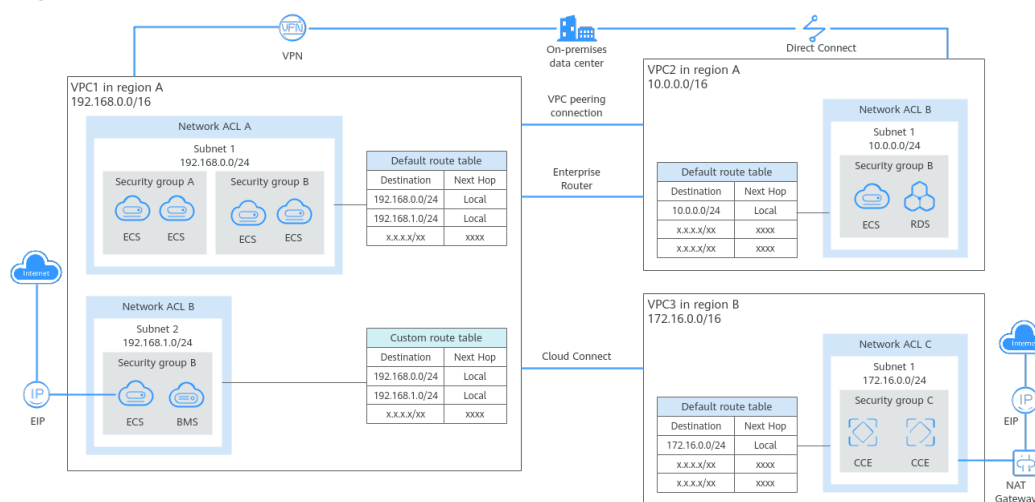
Figure 1-1 VPC architecture

Table 1-1 Architecture description

Item	Brief	Details
VPC basics	A VPC is a logically isolated virtual private network. You can define a CIDR block for each VPC and add one or more subnets. You can also configure route tables to control where the traffic from your subnet is directed. VPCs are logically isolated from each other, but subnets in a VPC can communicate with each other by default.	• IPv4 CIDR block: When creating a VPC, you need to specify an IPv4 CIDR block for it. Supported IPv4 CIDR blocks are 10.0.0.0/8-24, 172.16.0.0/12-24, and 192.168.0.0/16-24. • Subnet: You can divide a VPC into one or more subnets as required to deploy your instances (such as cloud servers, containers, and databases). Private IP addresses are then assigned to your instances from the subnets where they are running. For more information, see **Subnet**. • Route table: Each VPC comes with a default route table that allows communications between subnets in a VPC. You can add routes to the default route table or create a route table to control traffic. For details, see **Route Tables and Routes**.
VPC security	Security groups and network ACLs protect the cloud resources deployed in a VPC.	• Security groups protect instances. You can add inbound and outbound rule to protect all the resources in a security group. For details about security groups, see **Security Groups and Security Group Rules**. • Network ACLs protect associated subnets. You can add inbound and outbound rule to protect all the resources in a subnet. For details, see **Network ACL Overview**. Network ACLs protect subnets, while security groups protect instances in a subnet. If both security group and network ACL rules are configured, traffic matches network ACL rules first and then security group rules. For details, see **What Is Access Control?**

Item	Brief	Details
VPC connectivity	You can combine VPC and other networking services to build networks to meet different requirements. ● Use VPC peering connections or an enterprise router to connect different VPCs in the same region. ● Use Cloud Connect to connect VPCs in different regions. ● Use an EIP or NAT gateway to allow the instances in a VPC to the Internet. ● Use Direct Connect or VPN to connect an on-premises data center to VPCs.	● Connecting VPCs in the same region – VPC peering connections: connect VPCs in a region in the same account or different accounts. For details, see **VPC Peering Connection Overview**. – Enterprise routers: connect multiple VPCs in the same region, as a high-performance centralized router. For details, see **What Is an Enterprise Router?** VPC peering connections are free of charge, while enterprise routers are not free. Compared with VPC peering connections, enterprise routers simplify the network structure and make it easy for scale-out and O&M. ● Connecting VPCs across regions Cloud Connect: connects VPCs in different regions to quickly build cross-region networks. For details, see **What Is Cloud Connect?** ● Connecting a VPC to the Internet – EIPs: enable your cloud resources to communicate with the Internet. For details, see **What Is Elastic IP?** – Public NAT gateways: enables instances (such as ECSs or BMSs) in a VPC to share an EIP to communicate with the Internet. A public NAT gateway supports up to 20 Gbit/s of bandwidth. For details, see **What Is NAT Gateway?** ● Connecting an on-premises data center to a VPC – Direct Connect: allows you to establish a stable, high-speed, low-latency, secure, and dedicated network connection

Item	Brief	Details
		that connects your on-premises data center to the cloud. Direct Connect helps you build a flexible, scalable hybrid cloud computing environment. For details, see **What Is Direct Connect?** – VPN: establishes a secure, encrypted communication tunnel between your on-premises data center and your VPC. For details, see **What Is Virtual Private Network?** Compared with Direct Connect, VPN is cost-effective and can be quickly deployed.
VPC O&M	VPC flow logs and traffic mirroring track traffic in a VPC for network O&M.	• VPC flow logs: records traffic to and from a VPC in real time. VPC flow logs help you monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification. For details, see **VPC Flow Log Overview**. • Traffic mirroring: mirrors traffic that meets a mirror filter from an elastic network interface to a destination, where you can use the mirrored traffic for inspection, audit analysis, and troubleshooting. For details, see **Traffic Mirroring Overview**.

Accessing the VPC Service

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

- Management console

  You can use the console to directly perform operations on VPC resources. 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 a VPC into a third-party system for secondary development, you can use APIs to access the VPC service. For details, see the **Virtual Private Cloud API Reference**.

2 VPC Advantages

Flexible Configuration

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

Secure and Reliable

VPCs are logically isolated through tunneling technologies. By default, different VPCs cannot communicate with each other. You can use network ACLs to protect subnets and security groups to protect ECSs. They provide multiple layers of security for your VPCs.

Figure 2-1 Secure and reliable

Seamless Interconnectivity

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

By default, instances in different 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.

You can use a Layer 2 connection gateway (L2CG) provided by our Enterprise Switch service to establish network communication between the cloud and on-premises networks, and migrate data center or private cloud services to the cloud without changing subnets.

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

Figure 2-2 Interconnectivity
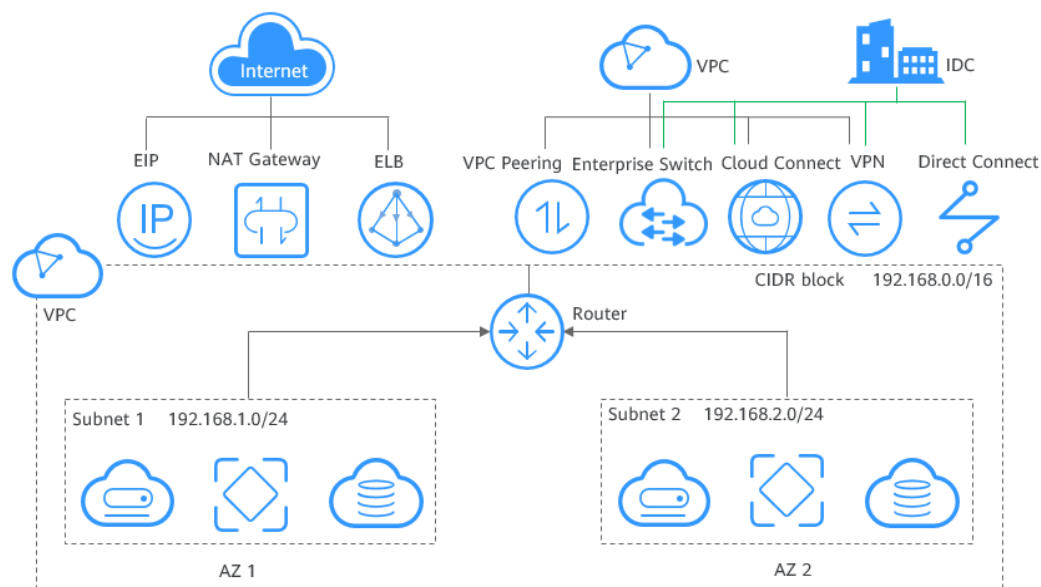
High-Speed Access

Dynamic BGP is used to provide access to various carrier networks. You can establish over 20 dynamic BGP connections to different carriers. Dynamic BGP connections enable real-time failovers based on 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, including 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. You can select whichever one best fits your business needs. There are no upfront costs and network O&M costs, reducing the 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 any change in business requirements.
Flexibility	Huawei Cloud provides a variety of network services for you to choose from. If you need more network resources (for instance, if you need more bandwidth), you can expand resources on the fly.	You have to strictly comply with the network plan to complete the service deployment. If there are changes in your service requirements, it is difficult to dynamically adjust the network.
Security	VPCs are logically isolated from each other. You can use security features such as network ACLs and security groups, and even security services like Advanced Anti-DDoS (AAD) to protect your cloud resources.	The network is insecure and difficult to maintain. You need professional technical personnel to ensure network security.

3 Application Scenarios

VPC allows you to provision logically isolated virtual private networks for cloud resources, such as cloud servers, containers, and databases. Additionally, you can use VPC and other networking services to set up networks to meet different requirements.

Building a Highly Secure Cloud Network

You can deploy applications on instances in a VPC and configure security groups and network ACLs to protect these instances.

- A security group protects the instances in it.
- A network ACL protects the entire subnet. After a subnet is associated with a network ACL, all instances in the subnet are protected by the network ACL.

In the figure below, your application is deployed on the ECSs in a subnet (Subnet-A01), and the database servers are deployed in another subnet (Subnet-A02) that is isolated from the Internet. To protect these servers, you can configure security groups and network ACLs to control inbound and outbound traffic.

Figure 3-1 Building a secure and private cloud network
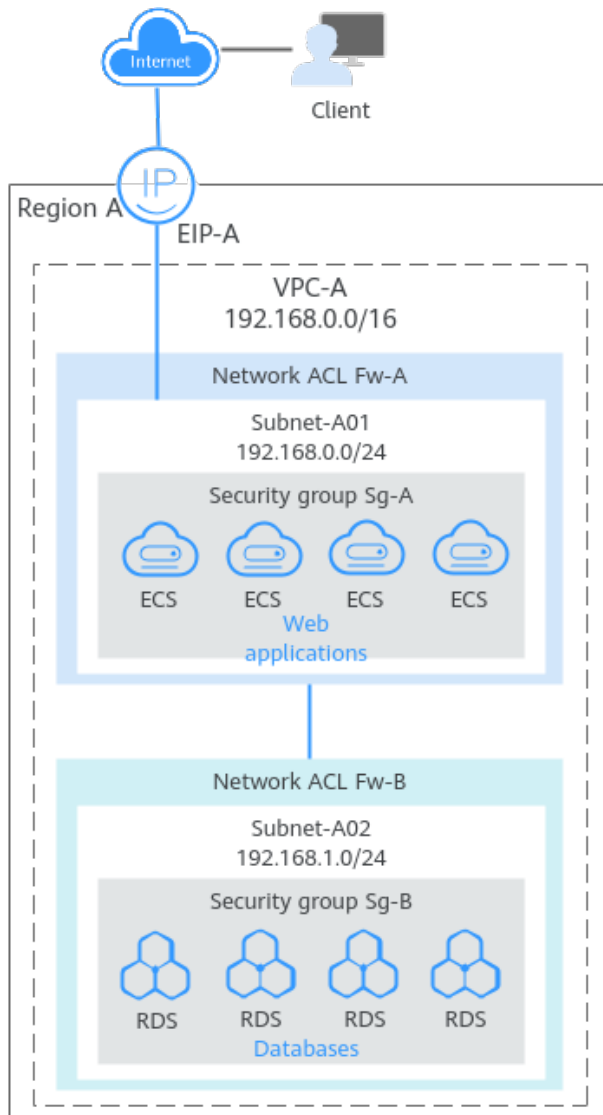
☐ **NOTE**

For more information, see **VPC Access Control Overview**.

Building a Cloud Network for Isolating Services

If you want to isolate services, you can deploy them in different VPCs, as resources in separated VPCs cannot directly communicate with each other.

Figure 3-2 Building a cloud network for isolating services
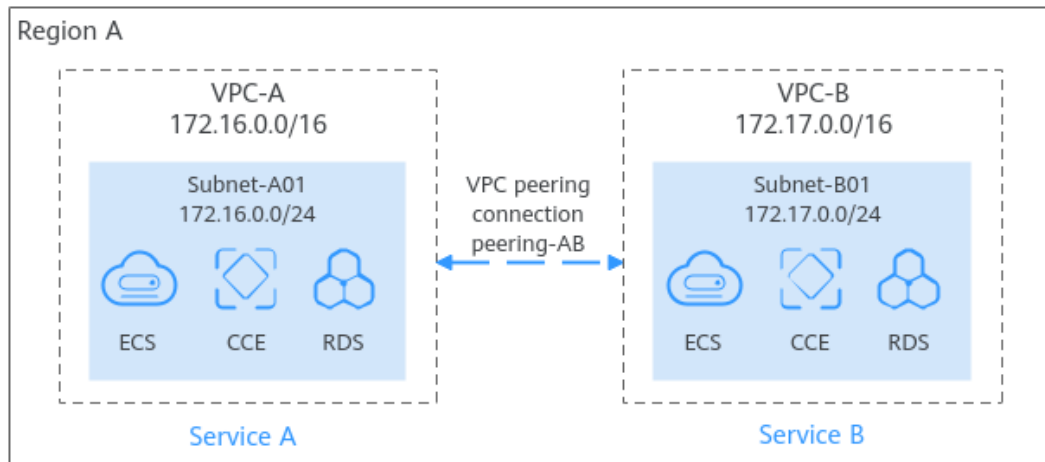
📖 **NOTE**

If two VPCs need to communicate with each other, you can use a VPC peering connection or enterprise router to connect the two VPCs. For details, see **Connecting VPCs**.

Building a Cross-Region DR Network

If you deploy services in multiple cities, you can use VPC, Enterprise Router, and Cloud Connect Central Network to build a cross-region cloud network. This allows you to deploy services in multiple regions to minimize network latency, ensure nearby access, enhance disaster recovery capabilities, and ensure service availability and continuity.

For example, you can attach VPCs in each region to an enterprise router and use a central network to connect enterprise routers in different regions. In this way, instances in different regions can communicate with each other.

Figure 3-3 Building a cross-region DR network
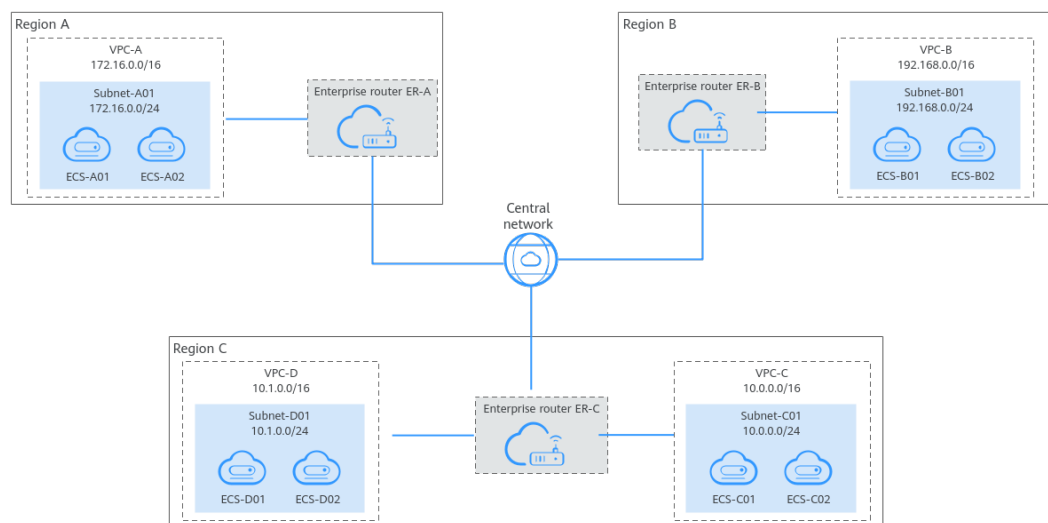
📖 **NOTE**

For more information, see **Connecting a VPC to Other VPCs**.

Building a High-Availability Load Balancing Network

To handle a large number of concurrent requests from the Internet, you can deploy multiple ECSs in a VPC and use ELB to distribute requests across these servers to improve service stability and availability.

Figure 3-4 Building a high-availability load balancing network
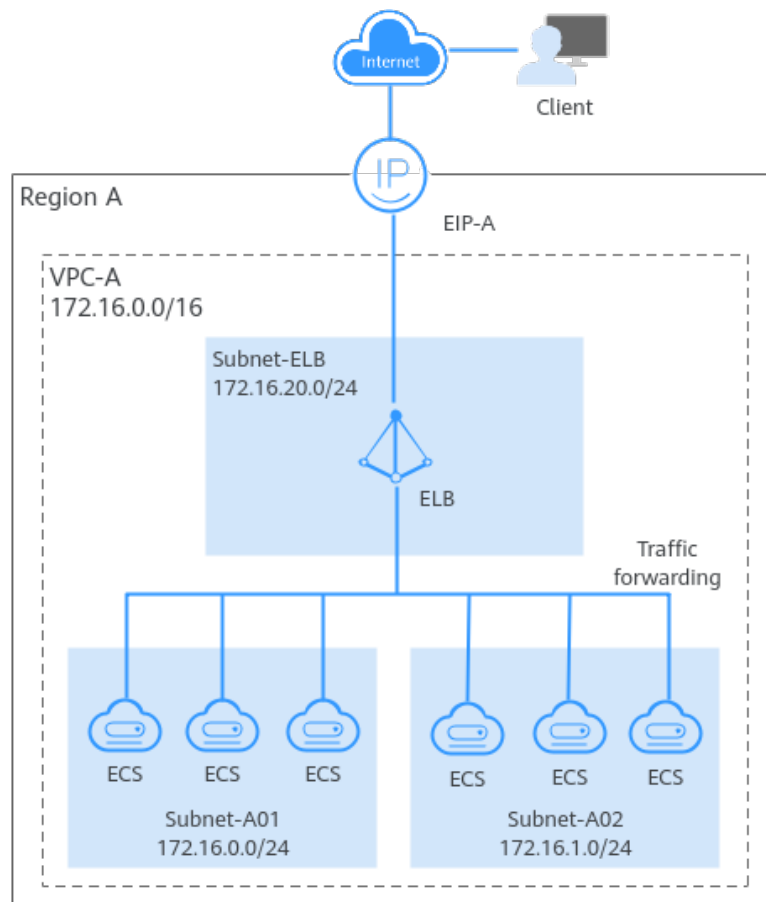
📖 **NOTE**

For more information, see **Using ELB to Distributing Traffic over the Internet**.

Building a Hybrid Cloud Network

If your cloud and on-premises services want to communicate with each other, you can use VPN or Direct Connect to build a hybrid cloud network.

- Hybrid cloud networking using VPN and VPC

  In **Figure 3-5**, some workloads have been migrated to a VPC (VPC-A), and some workloads are still running on on-premises servers. With a VPN connection, on-premises servers can quickly access the cloud resources in the

VPC. Compared with Direct Connect, VPN is easier to configure and cost-effective.

Figure 3-5 Connecting a VPC to an on-premises data center using VPN
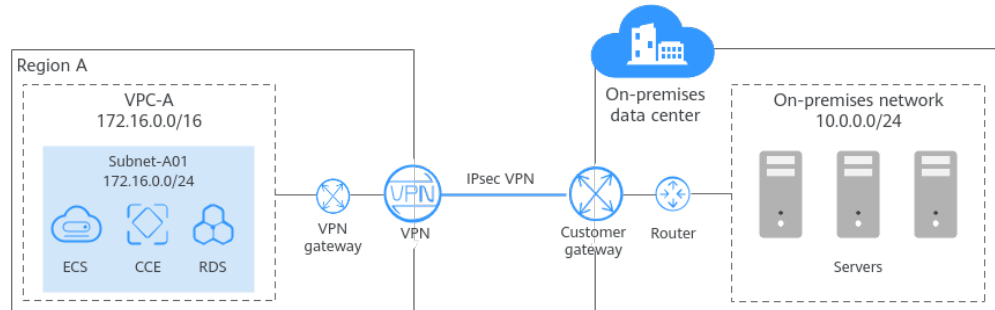
- Hybrid cloud networking using Direct Connect and VPC

  In **Figure 3-6**, some workloads are running in a VPC (VPC-A) on the cloud, and some are running in the on-premises data center. A Direct Connect connection connects the on-premises data center to the cloud. Direct Connect connections are faster and more stable than VPN connections.

Figure 3-6 Connecting a VPC to an on-premises data center using Direct Connect
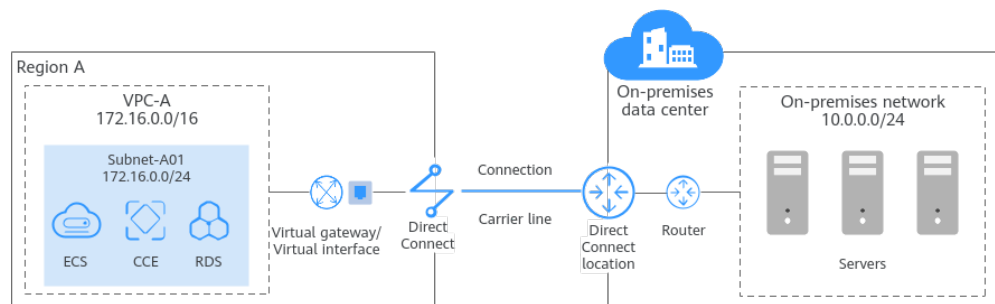
📖 **NOTE**

For more information, see **Connecting VPCs to On-Premises Data Centers**.

4 VPC Functions

VPC provides various functions for you to flexibly configure services and build diversified networks. For details, see **Table 4-1**.

- Basic VPC functions: VPC, subnet, route table, route, virtual IP address, network interface, and supplementary network interface
- Network security: security group, network ACL, and IP address group
- Network connectivity: VPC peering connection
- Network O&M: VPC flow log and traffic mirroring

Table 4-1 VPC functions

Functio n	Description	Reference
VPC	VPC allows you to provision logically isolated virtual private networks for cloud resources, such as cloud servers, containers, and databases. You can create subnets, security groups, network ACLs, route tables, and more to manage cloud resources flexibly. You can also use EIPs to connect cloud resources in VPCs to the Internet, and use Direct Connect and VPN to connect on-premises data centers to VPCs to build a hybrid cloud network.	**Creating a VPC and Subnet**
Subnet	A subnet is a unique CIDR block with a range of IP addresses in a VPC. All resources in a VPC must be deployed on subnets. Subnets in a VPC cannot overlap with each other.	**Creating a Subnet for an Existing VPC**

Function	Description	Reference
Route table and route	A route table contains a set of routes that are used to control the traffic in and out of your subnets in a VPC. Each subnet must be associated with a route table. A subnet can only be associated with one route table, but a route table can be associated with multiple subnets.	**Route Table and Route Overview**
Virtual IP address	A virtual IP address is a private IP address that can be independently assigned from and released to a VPC subnet. You can: ● Bind one or more virtual IP addresses to a cloud server so that you can use either the virtual IP address or private IP address to access the server. If you have multiple services running on a cloud server, you can use different virtual IP addresses to access them. ● Bind a virtual IP address to multiple cloud servers. You can use a virtual IP address and an HA software (such as Keepalived) to set up a high-availability active/standby cluster.	**Virtual IP Address Overview**
Elastic network interface	An elastic network interface is a virtual network card. You can create network interfaces and attach them to your cloud servers to obtain flexible and highly available network configurations.	**Elastic Network Interface Overview**
Supplementary network interface	Supplementary network interfaces are a supplement to elastic network interfaces. If the number of elastic network interfaces that can be attached to your cloud server cannot meet your requirements, you can use supplementary network interfaces, which can be attached to VLAN subinterfaces of elastic network interfaces.	**Supplementary Network Interface Overview**
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. You can create a security group and define different access rules to protect the ECSs that it contains.	**Security Group and Security Group Rule Overview**

Function	Description	Reference
Network ACL	A network ACL is an optional layer of security for your subnets. After you add inbound and outbound rules to a network ACL and associate subnets with it, you can control traffic in and out of the subnets.	**Network ACL Overview**
IP address group	An IP address group is a collection of IP addresses. It can be associated with security groups and network ACLs to simplify IP address configuration and management in networking.	**IP Address Group Overview**
VPC peering connection	A VPC peering connection enables two VPCs in the same region to communicate using private IP addresses. The VPCs to be connected can be from the same account or different accounts.	**VPC Peering Connection Overview**
VPC sharing	VPC sharing allows multiple accounts to create and manage cloud resources, such as ECSs, load balancers, and RDS instances, in one VPC. With Resource Access Manager (RAM), you can share subnets in a VPC with one or more accounts so you can centrally manage resources in multiple accounts, which improves resource management efficiency and reduces O&M costs.	**VPC Sharing Overview**
Edge gateway	An edge gateway can connect subnets in the same VPC but from both edge and central AZs or from different edge AZs.	**Edge Gateway Overview**
IPv4/IPv6 dual stack network	IPv4/IPv6 dual stack allows your resources to use both IPv4 and IPv6 addresses for private and public network communications.	**IPv4/IPv6 Dual-Stack Network**
VPC flow log	A VPC flow log records information about the traffic going to and from a VPC. You can use flow logs to monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification.	**VPC Flow Log Overview**

Functio n	Description	Reference
Traffic mirrorin g	Traffic Mirroring can be used to mirror traffic that meets a mirror filter from an elastic network interface. You can configure inbound and outbound rules for a mirror filter to determine which traffic from an elastic network interface will be mirrored to a network interface or load balancer. You can then send the traffic for inspection, audit analysis, and troubleshooting.	**Traffic Mirroring Overview**

5 Security

5.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 5-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud**: Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.

- **Tenant**: Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 5-1 Huawei Cloud shared security responsibility model
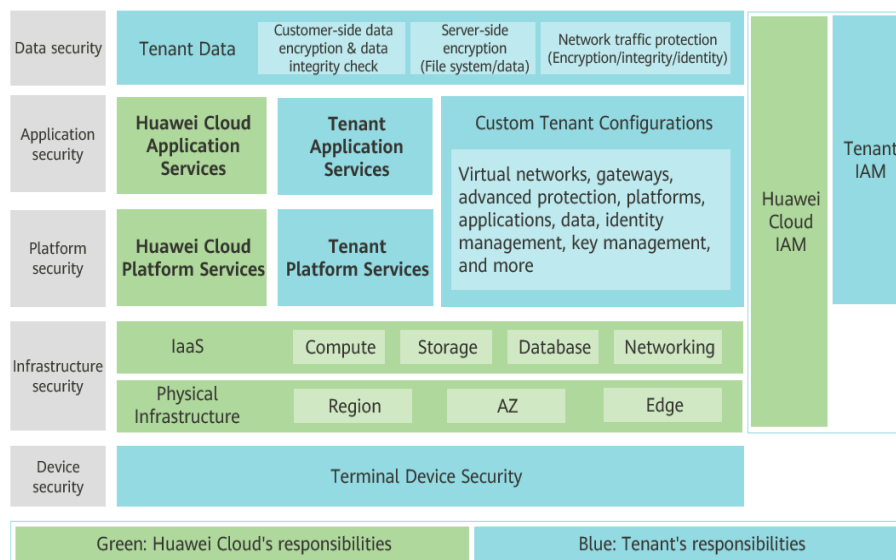
5.2 Identity Authentication and Access Control

Identity Authentication

Identity and Access Management (IAM) enables you to easily manage users and control their access to Huawei Cloud services and resources.

You can use IAM to control access to your VPC resources. IAM permissions define which actions on your cloud resources are allowed or denied.

After creating an IAM user, the administrator needs to add it to a user group and grant the permissions required by VPC to the user group. Then, all users in this group automatically inherit the granted permissions.

- **IAM Functions**
- **Permissions**

Access Control

- **Security Groups**

  A security group is a collection of access control rules for cloud resources, such as cloud servers, containers, and databases, that have the same security protection requirements and that are mutually trusted within a VPC. After a security group is created, you can create various access rules for the security group, these rules will apply to all cloud resources added to this security group.

  You can create and delete security groups, fast-add, replicate, modify, delete, import or export security group rules, view or change the security group of an ECS, and add ECSs to or remove them from a security group.

  You can define access rules for a security group. Then these rules will apply to all cloud resources added to this security group.

  For details, see **Security Group Overview**.

- **Network ACLs**

  A network ACL is an optional layer of security for your subnets. You can associate one or more subnets with a network ACL to 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, add, modify, change the sequence of, enable, disable, and delete network ACL rules.

  You can define network ACL rules to control traffic in and out of the subnets.

  For details, see **Network ACL Overview**.

5.3 Auditing and Logging

Auditing

Cloud Trace Service (CTS) records operations on cloud resources in your account. You can use the logs to perform security analysis, track resource changes, audit compliance, and locate faults.

After CTS is enabled, it can record VPC operations.

- If you want to enable and configure CTS, refer to **CTS Getting Started**.
- If you want to know supported operations on VPCs, refer to **Supported VPC Operations**.
- If you want to view traces, refer to **Viewing Traces**.

Logging

A VPC flow log records information about traffic going to and from your VPC. VPC flow logs help you monitor network traffic, analyze network attacks, and determine whether security group and network ACL rules require modification.

- **VPC Flow Log Overview**
- **Creating a VPC Flow Log**

5.4 Risk Monitoring

Cloud Eye is a multi-dimensional resource monitoring platform. You can use Cloud Eye to monitor the utilization of service resources, track the running status of cloud services, configure alarm rules and notifications, and quickly respond to resource changes.

With Cloud Eye, you can view the bandwidth and EIP usage. You can also create alarm rules and configure monitoring thresholds and alarm notifications. This will ensure you learn about VPC resource status in a timely manner.

- **What Is Cloud Eye?**
- **Supported Metrics**
- **Viewing Metrics**

6 Notes and Constraints

Constraints on VPC Resources

Note the restrictions on the following VPC resources before using them.

- **VPC and subnet CIDR blocks**
- **Secondary IPv4 CIDR blocks**
- **Route table and routes**
- **Virtual IP addresses**
- **Network interfaces**
- **Supplementary network interfaces**
- **Security groups**
- **Network ACLs**
- **IP address groups**
- **VPC peering connections**
- **Shared VPCs**
- **IPv4/IPv6 dual-stack networks**
- **VPC flow logs**
- **Traffic mirrors**

VPC Resource Quotas

A quota defines the maximum number of resources of a certain type that can be created in a region.

Suppose the VPC quota in CN-Hong Kong is 5. If two VPCs have been created in this region, the remaining quota is 3.

To help you save quotas, Huawei Cloud sets limit on the maximum number of cloud resources that you can create in each region.

You can **log in to the console** to view the default quotas for each resource. To increase the resource quota, you can refer to **Applying for a Higher Quota**.

Table 6-1 lists the quotas about VPC resources. Resource quotas are displayed by region. By default, the quotas are the same across regions.

Table 6-1 VPC resource quotas

Item	Default Quota	Adjustable
Maximum number of VPCs per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of subnets per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of route tables that can be associated with a VPC in a region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of routes per route table in a region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	No
Maximum number of security groups per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of security group rules per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of network ACLs per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of IP address groups per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	Yes **Submit a service ticket**
Maximum number of VPC peering connections per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	No
Maximum number of VPC flow logs per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	No
Maximum number of mirror sessions per region	Quotas vary by your account type and service level. You can go to the **console** to check your quotas.	No

7 VPC and Related Services

Figure 7-1 shows the relationship between VPC and other services.

Figure 7-1 VPC and related services
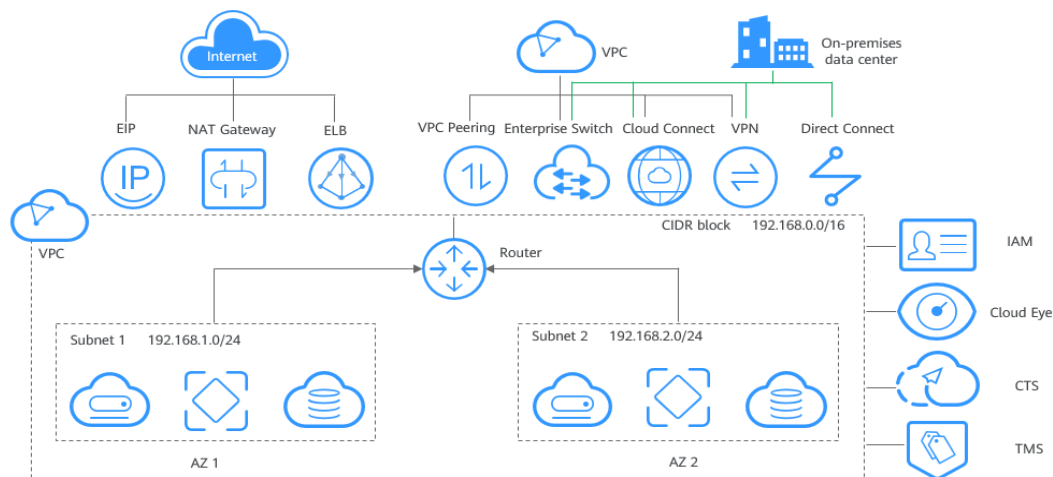
Table 7-1 Related services

Service	Interaction Function
Elastic Cloud Server (ECS)	Using Security Groups to Protect ECS Network Security
Elastic IP (EIP)	Using an EIP to Connect a VPC to the Internet
NAT Gateway	Using a Public NAT Gateway to Connect a VPC to the Internet
Virtual Private Network (VPN)	Using VPN to Connect a VPC to an On-Premises Data Center
Direct Connect	Using Direct Connect to Connect a VPC to an On-Premises Data Center

Service	Interaction Function
Enterprise Router	**Using an Enterprise Router to Connect VPCs in the Same Region**
Cloud Connect	**Using Cloud Connect to Connect VPCs in Different Regions**
Elastic Load Balance (ELB)	**Using ELB to Distribute Traffic Across Multiple Backend Servers in a VPC**
Identity and Access Management (IAM)	**Creating an IAM User and Granting VPC Permissions**
Cloud Eye	**Cloud Eye Monitoring**
Cloud Trace Service (CTS)	**CTS Auditing**
Tag Management Service (TMS)	**Using TMS to Identify VPC Resources**

8 Billing

VPC provides a wide range of cloud resources. Some are free, while some are not. **Table 8-1** describes how these resources are billed.

Table 8-1 VPC resource billing

Resource	Billing Description
VPC	Free
Subnet	Free
Route table	Free
VPC peering connection	Free
Elastic network interface	Free
Supplementary network interface	Free
IP address group	Free
Security group	Free
Network ACL	Free
Edge gateway	Free
VPC flow log	Free
Traffic Mirroring	Free

Resource	Billing Description

EIP and bandwidth	If you use EIPs and bandwidths, you need to pay for their prices.
	
	● EIP reservation price
	  If your pay-per-use EIP has no instance bound, you will be billed for the EIP reservation price.
	
	● Fixed bandwidth
	  – EIP bandwidth prices: bandwidth prices of yearly/monthly EIPs and pay-per-use EIPs (by bandwidth); traffic price of pay-per-use EIPs (by traffic)
	  – Shared bandwidth price
	  – Shared data package price
	
	For details, see **EIP Billing**.

VPC endpoint	If you use VPC endpoints, you need to pay for them. For details, see **VPC Endpoint Billing**.

□ NOTE

Currently, free resources are not billed. You will be notified in advance if the billing starts.

9 Permissions

If you need to assign different permissions to personnel in your enterprise to access your VPCs, IAM is a good choice. IAM provides identity authentication, permissions management, and access control, helping you to securely access your Huawei Cloud resources.

With IAM, you can create IAM users, and assign permissions to control their access to specific resources. For example, if you want some software developers in your enterprise to use VPCs but do not want them to delete VPCs or perform any other high-risk operations, you can grant permissions to use VPCs but not permissions to delete them.

If your HUAWEI ID does not require IAM for permissions management, you can skip this section.

IAM is a free service. You only pay for the resources in your account. For more information, see **IAM Service Overview**.

VPC 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.

VPC is a project-level service deployed for specific regions. When you set **Scope** to **Region-specific projects** and select the specified projects (for example, **ap-southeast-1**) in the specified regions (for example, **CN-Hong Kong**), the users only have permissions for VPCs in the selected projects. If you set **Scope** to **All resources**, users have permissions for VPCs in all region-specific projects. When accessing VPCs, the users need to switch to the authorized region.

You can grant permissions by using roles and policies.

- Roles: A coarse-grained authorization strategy provided by IAM to assign permissions based on users' job responsibilities. Only a limited number of service-level roles are available for authorization. When you grant permissions using roles, you also need to attach dependent roles. Roles are not ideal for fine-grained authorization and least privilege access.

- Policies: A fine-grained authorization strategy that defines permissions required to perform operations on specific cloud resources under certain

conditions. This type of authorization is more flexible and is ideal for least privilege access. For example, you can grant VPC users only the permissions for managing a certain type of resources. A majority of fine-grained policies contain permissions for specific APIs, and permissions are defined using API actions. For the API actions supported by VPC, see **Permissions Policies and Supported Actions**.

Table 9-1 lists all the system-defined permissions for VPC.

Table 9-1 System-defined permissions for VPC

Policy Name	Description	Policy Type	Dependencies
VPC FullAccess	Full permissions for VPC.	System-defined policy	To use the VPC flow log function, users must also have the **LTS ReadOnlyAccess** permission.
VPC ReadOnlyAccess	Read-only permissions on VPC.	System-defined policy	None
VPC Administrator	Most permissions on VPC, excluding creating, modifying, deleting, and viewing security groups and security group rules. To be granted this permission, users must also have the **Tenant Guest** permission.	System-defined role	**Tenant Guest** policy, which must be attached in the same project as **VPC Administrator**.

Table 9-2 lists the common operations supported by system-defined permissions for VPC.

Table 9-2 Common operations supported by system-defined permissions

Operation	VPCReadOnlyAccess	VPC Administrator	VPC FullAccess
Creating a VPC	Not supported	Supported	Supported
Modifying a VPC	Not supported	Supported	Supported
Deleting a VPC	Not supported	Supported	Supported

Operation	VPCReadOnlyAccess	VPC Administrator	VPC FullAccess
Viewing VPC information	Supported	Supported	Supported
Creating a subnet	Not supported	Supported	Supported
Viewing subnet information	Supported	Supported	Supported
Modifying a subnet	Not supported	Supported	Supported
Deleting a subnet	Not supported	Supported	Supported
Creating a security group	Not supported	Not supported	Supported
Viewing security group information	Supported	Not supported	Supported
Modifying a security group	Not supported	Not supported	Supported
Deleting a security group	Not supported	Not supported	Supported
Adding a security group rule	Not supported	Not supported	Supported
Viewing a security group rule	Supported	Not supported	Supported
Modifying a security group rule	Not supported	Not supported	Supported
Deleting a security group rule	Not supported	Not supported	Supported
Creating a network ACL	Not supported	Supported	Supported
Viewing a network ACL	Supported	Supported	Supported

Operation	VPCReadOnlyAccess	VPC Administrator	VPC FullAccess
Modifying a network ACL	Not supported	Supported	Supported
Deleting a network ACL	Not supported	Supported	Supported
Adding a network ACL rule	Not supported	Supported	Supported
Modifying a network ACL rule	Not supported	Supported	Supported
Deleting a network ACL rule	Not supported	Supported	Supported
Creating a VPC peering connection	Not supported	Supported	Supported
Modifying a VPC peering connection	Not supported	Supported	Supported
Deleting a VPC peering connection	Not supported	Supported	Supported
Querying a VPC peering connection	Supported	Supported	Supported
Accepting a VPC peering connection request	Not supported	Supported	Supported
Rejecting a VPC peering connection request	Not supported	Supported	Supported
Creating a route table	Not supported	Supported	Supported
Deleting a route table	Not supported	Supported	Supported
Modifying a route table	Not supported	Supported	Supported

Operation	VPCReadOnlyAccess	VPC Administrator	VPC FullAccess
Associating a route table with a subnet	Not supported	Supported	Supported
Adding a route	Not supported	Supported	Supported
Modifying a route	Not supported	Supported	Supported
Deleting a route	Not supported	Supported	Supported
Creating a VPC flow log	Not supported	Supported	Supported
Viewing a VPC flow log	Supported	Supported	Supported
Enabling or disabling a VPC flow log	Not supported	Supported	Supported
Deleting a VPC flow log	Not supported	Supported	Supported
Creating an IP address group	Not supported	Supported	Supported
Associating an IP address group with resources	Not supported	Supported	Supported
Disassociating an IP address group from resources	Not supported	Supported	Supported
Adding IP addresses to an IP address group	Not supported	Supported	Supported
Deleting IP addresses from an IP address group	Not supported	Supported	Supported

Operation	VPCReadOnlyAccess	VPC Administrator	VPC FullAccess
Modifying an IP address group	Not supported	Supported	Supported
Deleting an IP address group	Not supported	Supported	Supported

Helpful Links

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

10 Basic Concepts

10.1 Subnet

A subnet is a unique CIDR block with a range of IP addresses in a VPC. All resources in a VPC must be deployed on subnets. Subnets in a VPC cannot overlap with each other.

- All instances in different subnets of the same VPC can communicate with each other by default, and the subnets can be located in different AZs. If you have a VPC with two subnets in it and they are located in different AZs, they can communicate with each other by default.

- A cloud resource can be in a different AZ from its subnet. For example, a cloud server in AZ 1 can be in a subnet in AZ 3. If AZ 3 becomes faulty, cloud servers in AZ 1 can still use the subnet in AZ 3, and your services are not interrupted.

Figure 10-1 Subnet

Managing Subnets

- **Creating a VPC and Subnet**
- **Creating a Subnet for an Existing VPC**

10.2 Route Table

What Is a Route Table?

A route table contains a set of routes that are used to control the traffic in and out of your subnets in a VPC. Each subnet must be associated with a route table. A subnet can only be associated with one route table, but a route table can be associated with multiple subnets.

Both IPv4 and IPv6 routes are supported.

Figure 10-2 Route tables
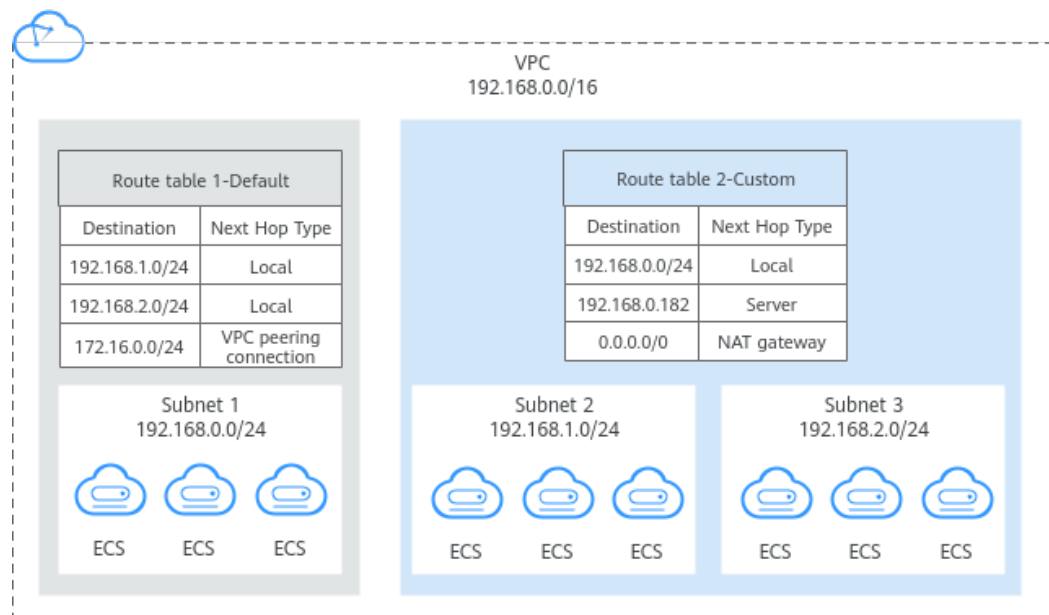
- Default route table: Each VPC comes with a default route table. If you create a subnet in a VPC, the subnet associates with the default route table. The default route table ensures that subnets in a VPC can communicate with each other.

    - You can add routes to, delete routes from, and modify routes in the default route table, but cannot delete the table.

    - When you create a VPN, Cloud Connect, or Direct Connect connection, the default route table automatically delivers a route that cannot be deleted or modified.

- Custom route table: If you do not want to use the default route table, you can create a custom route table and associate it with the subnet. Custom route tables can be deleted if they are no longer required.

    The custom route table associated with a subnet only controls the outbound traffic. The default route table of a subnet controls the inbound traffic.

☐ **NOTE**

By default, the quota for custom route tables is 0. To create custom route tables, **apply for a quota increase first**.

Route

You can add routes to both default and custom route tables and configure the destination, next hop type, and next hop for the routes to determine where network traffic is directed. Routes are classified into system routes and custom routes.

- System route: A system route is automatically added by the VPC service or other services (such as VPN and Direct Connect) and cannot be deleted or modified.

  Each route table comes with routes whose next hops are Local. Generally, a route table contains the following local routes:

  - Routes whose destination is 100.64.0.0/10, which is used to deploy public services, for example, the DNS servers. The route directs instances in a subnet to access these services.

  - Routes whose destination is 198.19.128.0/20 (IP address range used by internal services, such as VPC Endpoint).

  - Routes whose destination is 127.0.0.0/8 (local loopback addresses)

  - Routes whose destination is a subnet CIDR block that enables instances in a VPC to communicate with each other.

    If you enable IPv6 when creating a subnet, the system automatically assigns an IPv6 CIDR block to the subnet. Then, you can view IPv6 routes in its route table. Example destinations of subnet CIDR blocks are as follows:

    - IPv4: 192.168.2.0/24

    - IPv6: 2407:c080:802:be7::/64

- Custom route: After a route table is created, you can add custom routes and configure information such as the destination and next hop in the route to determine where network traffic is directed. In addition to manually added custom routes, there are custom routes added by other cloud services, such as Cloud Container Engine (CCE) or NAT Gateway.

  Route tables include default route tables and custom route tables. They support the next hop types described in **Table 10-1** and **Table 10-2**. The default route table supports fewer next hop types than a custom route table. This is because services like VPN, Direct Connect, and Cloud Connect automatically add routes to the default table.

Table 10-1 Next hop types supported by the default route table

Next Hop Type	Description
Server	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.
Supplementary network interface	Traffic intended for the destination is forwarded to the supplementary network interface of an ECS in the VPC.

Next Hop Type	Description
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.
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.
Global internet gateway	Traffic intended for the destination is forwarded to a global internet gateway.

Table 10-2 Next hop types supported by a custom route table

Next Hop Type	Description
Server	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.
BMS user-defined network	Traffic intended for the destination is forwarded to a BMS user-defined network.
VPN gateway	Traffic intended for the destination is forwarded to a VPN gateway.
Direct Connect gateway	Traffic intended for the destination is forwarded to a Direct Connect gateway.
Cloud connection	Traffic intended for the destination is forwarded to a cloud connection.
Supplementary network interface	Traffic intended for the destination is forwarded to the supplementary network interface of an ECS in the VPC.

Next Hop Type	Description
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.
Enterprise router	Traffic intended for the destination is forwarded to an enterprise router.
Cloud firewall	Traffic intended for the destination is forwarded to a cloud firewall.
Global internet gateway	Traffic intended for the destination is forwarded to a global internet gateway.

 NOTE

If you specify the destination when creating a resource, a system route is delivered. If you do not specify a destination when creating a resource, a custom route that can be modified or deleted is delivered.

For example, when you create a NAT gateway, the system automatically delivers a custom route without a specific destination (0.0.0.0/0 is used by default). In this case, you can change the destination. However, when you create a VPN gateway, you need to specify the remote subnet as the destination of a route. In this case, this route will be delivered as a system route. Do not modify the route destination on the **Route Tables** page. If you do, the destination will be inconsistent with the configured remote subnet. To modify the route destination, go to the specific resource page and modify the remote subnet, then the route destination will be changed accordingly.

You cannot add a route whose next hop type is **VPC endpoint** or **Cloud container** to a route table. These routes are automatically added by the VPC Endpoint or CCE service.

Managing Route Tables

- **Creating a Custom Route Table**
- **Adding Routes to a Route Table**
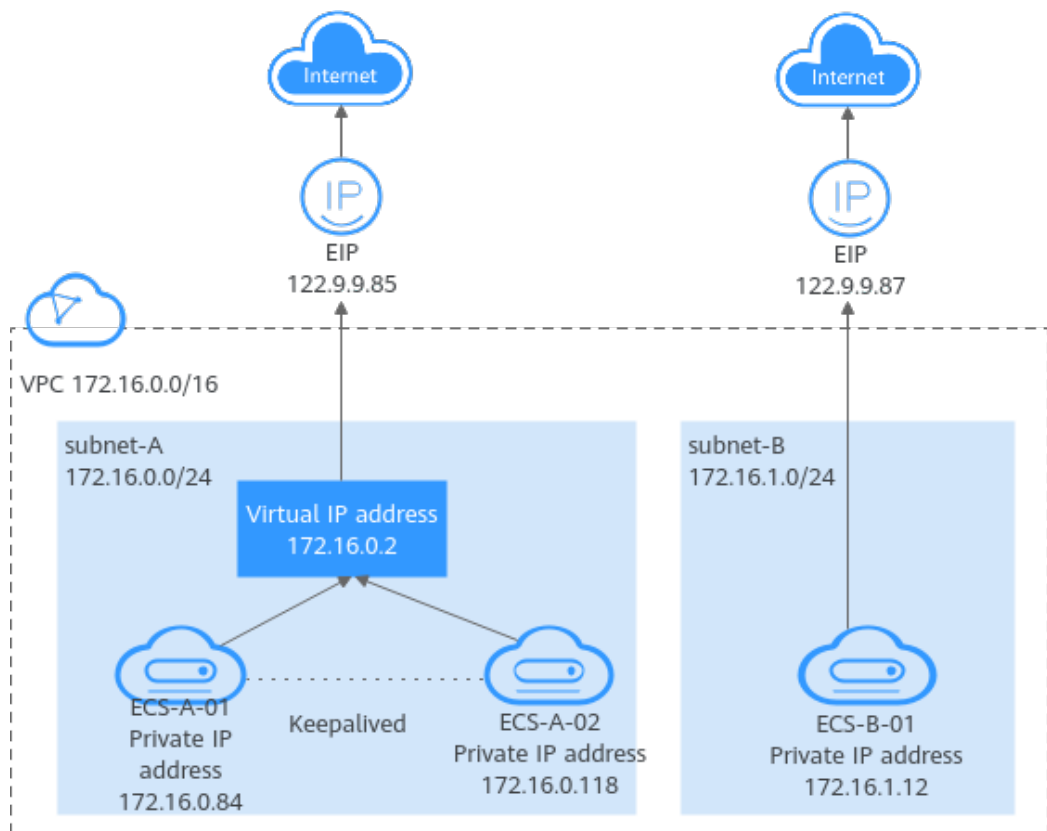
10.3 Virtual IP Address

A virtual IP address is a private IP address that can be independently assigned from and released to a VPC subnet. You can:

- Bind one or more virtual IP addresses to a cloud server so that you can use either the virtual IP address or private IP address to access the server. If you have multiple services running on a cloud server, you can use different virtual IP addresses to access them.

- Bind a virtual IP address to multiple cloud servers. You can use a virtual IP address and an HA software (such as Keepalived) to set up a high-availability active/standby cluster. If you want to improve service availability and eliminate single points of failure, you can deploy cloud servers in the active/ standby pair or deploy one cloud server and multiple standby cloud servers. In this case, the cloud servers can use the same virtual IP address. If the active cloud server goes down, the standby cloud server becomes the active server and continues to provide services.

Generally, cloud servers use private IP addresses for internal network communication. A virtual IP address has the same network access capabilities as a private IP address. You can use either of them to enable layer 2 and layer 3 communications in a VPC, access a different VPC using a peering connection, enable Internet access through EIPs, and connect the cloud and the on-premises servers using VPN connections and Direct Connect connections. **Figure 10-3** describes how private IP addresses, the virtual IP address, and EIPs work together.

- Private IP addresses are used for internal network communication.

- The virtual IP address works with Keepalived to build an HA cluster. ECSs in this cluster can be accessed through one virtual IP address.

- EIPs are used for Internet communication.

Figure 10-3 Different types of IP addresses used by ECSs

Managing Virtual IP Addresses

- **Assigning a Virtual IP Address**
- **Binding a Virtual IP Address to an Instance or EIP**

10.4 Elastic Network Interface

An elastic network interface (referred to as a network interface in this documentation) is a virtual network card. You can create and configure network interfaces and attach them to your instances (ECSs and BMSs) to obtain flexible and highly available network configurations.

Network Interface Types

- A primary network interface is created together with an instance by default, and cannot be detached from the instance.
- An extended network interface is created on the **Network Interfaces** console, and can be attached to or detached from an instance.

Application Scenarios

- Flexible migration

  You can detach a network interface from an instance and then attach it to another instance. The network interface retains its private IP address, EIP, and security group rules. In this way, service traffic on the faulty instance can be quickly migrated to the standby instance, implementing quick service recovery.

- Traffic management

  You can attach multiple network interfaces that belong to different subnets in a VPC to the same instance, and configure the network interfaces to carry the private network traffic, public network traffic, and management network traffic of the instance. You can configure access control policies and routing policies for each subnet, and configure security group rules for each network interface to isolate networks and service traffic.

Managing Network Interfaces

- **Creating a Network Interface**
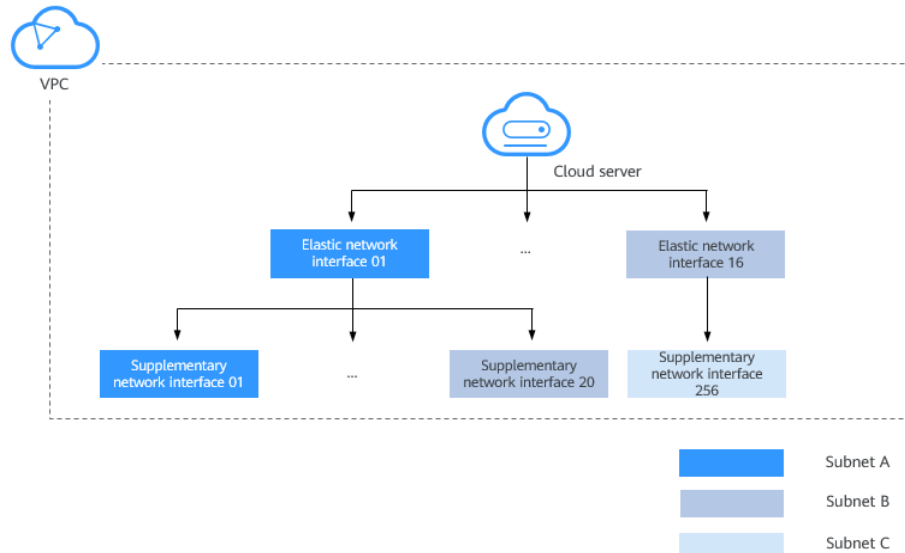- **Attaching a Network Interface to a Cloud Server**

10.5 Supplementary Network Interface

Supplementary network interfaces are a supplement to elastic network interfaces. If the number of elastic network interfaces that can be attached to your cloud server cannot meet your requirements, you can use supplementary network interfaces, which can be attached to VLAN subinterfaces of elastic network interfaces.

Application Scenarios

Supplementary network interfaces are attached to VLAN subinterfaces of elastic network interfaces. **Figure 10-4** shows the networking diagram.

Figure 10-4 Supplementary network interface networking diagram

The number of elastic network interfaces that can be attached to each ECS is limited. If this limit cannot meet your requirements, you can attach supplementary network interfaces to elastic network interfaces.

- You can attach supplementary network interfaces that belong to different subnets in the same VPC to an ECS. Each supplementary network interface has its private IP address and EIP for private or Internet communication.

- You can security group rules for supplementary network interfaces for network isolation.

Managing Supplementary Network Interfaces

- **Creating a Supplementary Network Interface**
- **Binding or Unbinding an EIP to or from a Supplementary Network Interface**
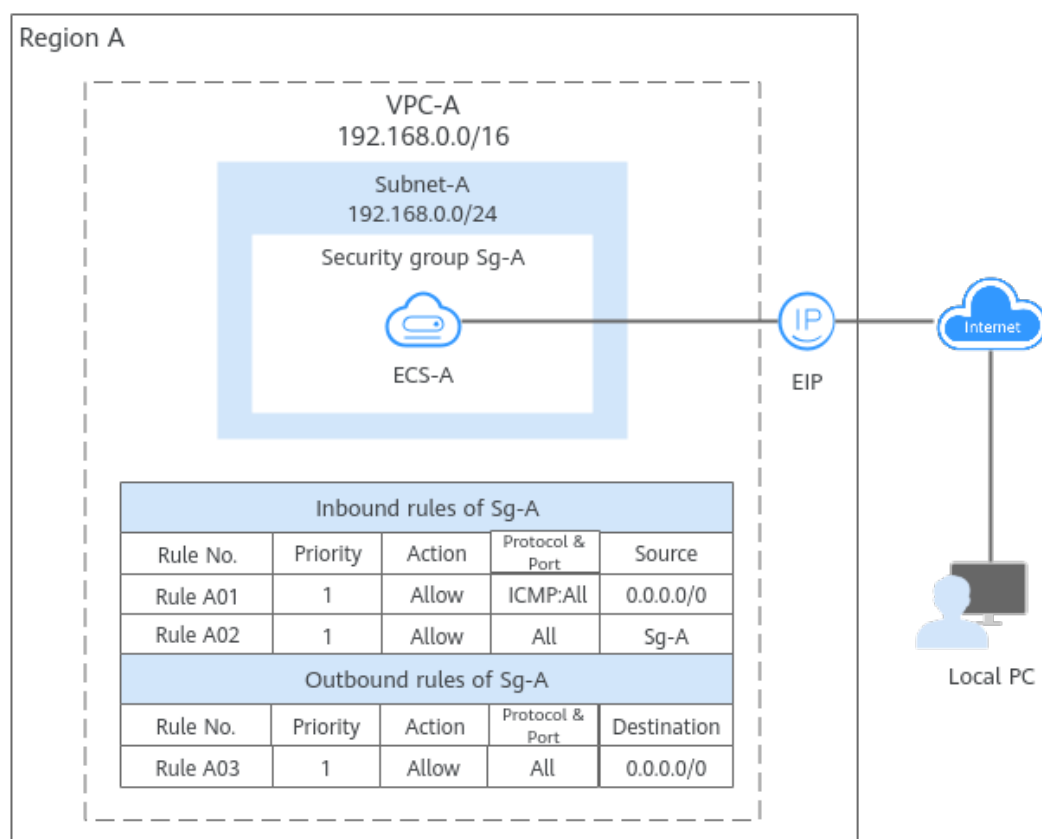
10.6 Security Group

A security group is a collection of access control rules for cloud resources, such as cloud servers, containers, and databases, that have the same security protection requirements and that are mutually trusted. After a security group is created, you can configure access rules that will apply to all cloud resources added to this security group.

Each security group can have both inbound and outbound rules. You need to specify the source, port, and protocol for each inbound rule and specify the destination, port, and protocol for each outbound rule to control the inbound and

outbound traffic to and from the instances in the security group. As shown in **Figure 10-5**, you have a VPC (**VPC-A**) with a subnet (**Subnet-A**) in region A. An ECS (**ECS-A**) is running in **Subnet-A** and associated with security group **Sg-A**.

- Security group **Sg-A** has a custom inbound rule to allow ICMP traffic to **ECS-A** from your PC over all ports. However, the security group does not have rules that allow SSH traffic to **ECS-A** so you cannot remotely log in to **ECS-A** from your PC.
- If **ECS-A** needs to access the Internet through an EIP, the outbound rule of **Sg-A** must allow all traffic from **ECS-A** to the Internet.

Figure 10-5 A security group architecture

Managing Security Groups

- **Creating a Security Group**
- **Adding a Security Group Rule**

10.7 Network ACL

A network ACL is an optional layer of security for your subnets. After you add inbound and outbound rules to a network ACL and associate subnets with it, you can control traffic in and out of the subnets.

A network ACL is different from a security group. A security group protects the instances in it, such as ECSs, databases, and containers, while a network ACL protects the entire subnet. Security groups are a mandatory layer of protection but

network ACLs are optional. Network ACLs and security groups can be used together for fine-grained access control.

You need to specify the protocol, source port and address, and destination port and address for each inbound and outbound rule of the network ACL. Suppose you have two subnets in region A, as shown in **Figure 10-6**. **Subnet-X01** is associated with network ACL **Fw-A**, and ECSs deployed in this subnet provide web services accessible from the Internet. **Subnet-X02** is associated with network ACL **Fw-B**. **Subnet-X02** and **Subnet-Y01** are connected through a VPC peering connection. Now, you need to configure inbound and outbound rules to allow **ECS-C01** in **Subnet-Y01** to remotely log in to ECSs in **Subnet-X02**.

● Inbound and outbound rules on **Fw-A**:

Custom inbound rule **A01** allows any IP address to access the ECSs in **Subnet-X01** through port 80 over TCP or HTTP. If the traffic does not match custom rule **A01**, the default rule is applied and the traffic is denied to flow into the subnet.

Stateful network ACLs allow responses to inbound requests to leave the subnet without being controlled by rules. The responses from ECSs in **Subnet-X01** can go out of the subnet. Other outbound traffic is not allowed to leave **Subnet-X01**, because the default rule is applied.

● Inbound and outbound rules on **Fw-B**:

Custom inbound rule **B01** allows **ECS-C01** in **Subnet-Y01** to use access the ECSs in **Subnet-X02** through port 22 over TCP or SSH.

Custom outbound rule **B02** allows all ICMP traffic over any port. The ping traffic from ECSs in **Subnet-X02** to **ECS-C01** in **Subnet-Y01** can be routed successfully to test the network connectivity.

Figure 10-6 Network ACL rules
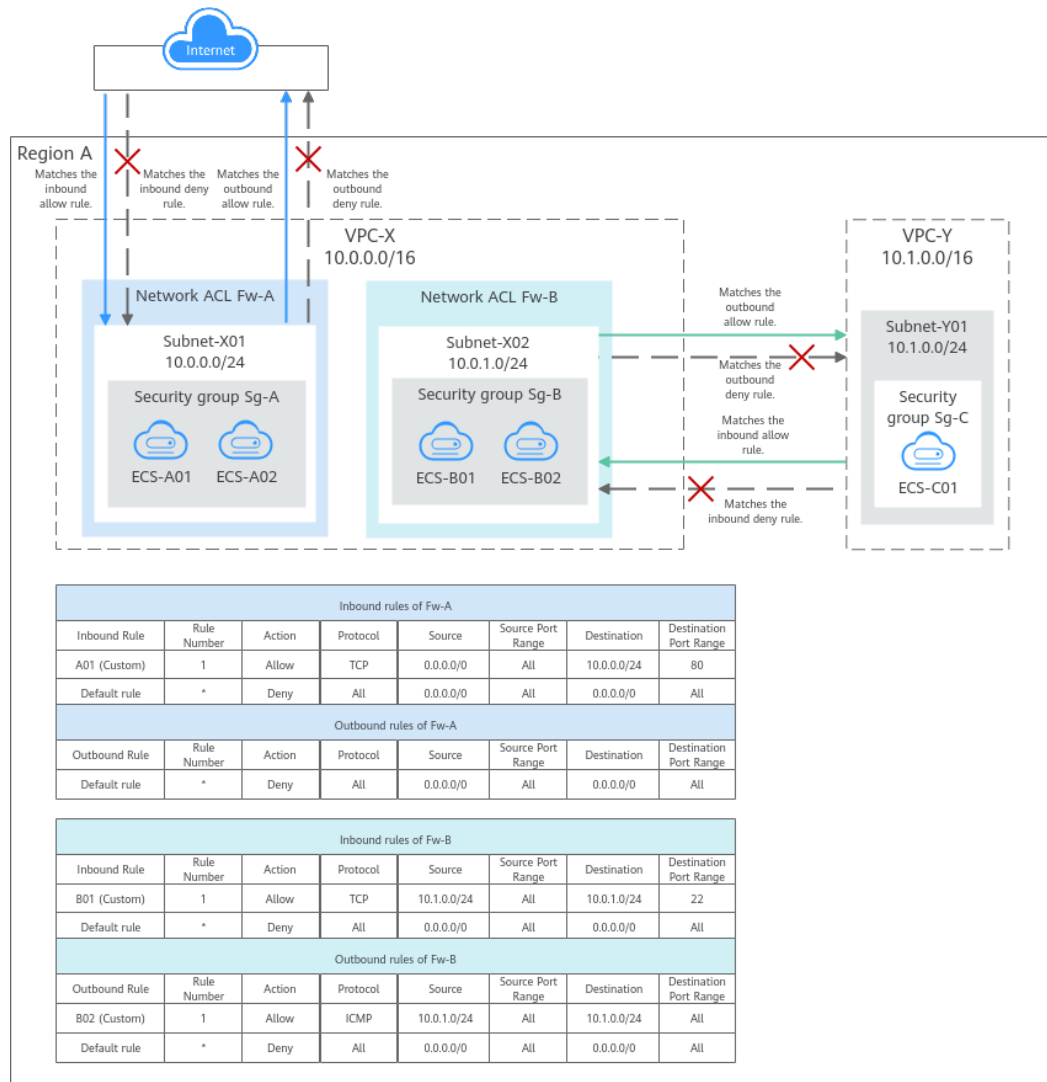
Managing Network ACLs

- **Creating a Network ACL**
- **Adding a Network ACL Rule**

10.8 IP Address Group

An IP address group is a collection of IP addresses. It can be associated with security groups and network ACLs to simplify IP address configuration and management.

You can add IP address ranges and IP addresses that need to be managed in a unified manner to an IP address group. An IP address group can work together with different cloud resources. **Table 10-3** lists the resources that can be associated with an IP address group.

Table 10-3 Resources that can be associated with an IP address group

Resource	Description	Example
Security group	The **Source** or **Destination** of a security group rule can be set to **IP address group**.	As shown in **Figure 10-7**, the inbound rule of security group **sg-A** uses IP address group **ipGroup-A** as the source.
Network ACL	The **Source** or **Destination** of a network ACL is set to **IP address group**.	As shown in **Figure 10-7**, the inbound rule of network ACL **fw-A** uses IP address group **ipGroup-A** as the source.

Figure 10-7 Using IP address group
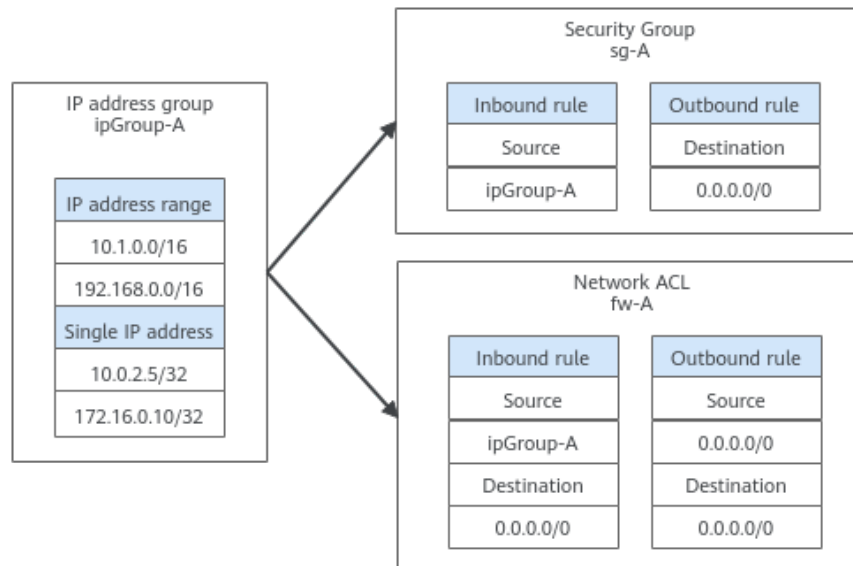
Managing IP Address Groups

- **Creating an IP Address Group**
- **Adding IP Addresses to an IP Address Group**

10.9 VPC Peering Connection

A VPC peering connection enables two VPCs in the same region to communicate using private IP addresses. The VPCs to be connected can be from the same account or different accounts.

- If you want to connect VPCs in different regions, use **Cloud Connect**.
- You can use VPC peering connections to build different networks. For details, see **VPC Peering Connection Usage Examples**.

Figure 10-8 shows an application scenario of VPC peering connections.

- There are two VPCs (VPC-A and VPC-B) in region A that are not connected.

- Service servers (ECS-A01 and ECS-A02) are in VPC-A, and database servers (RDS-B01 and RDS-B02) are in VPC-B. The service servers and database servers cannot communicate with each other.
- You need to create a VPC peering connection (peering-AB) between VPC-A and VPC-B so the service servers and database servers can communicate with each other.

Figure 10-8 Two VPCs connected by a VPC peering connection
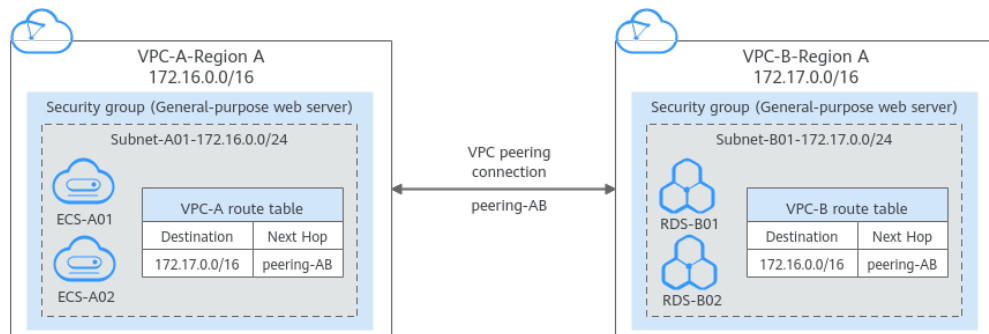
Managing VPC Peering Connections

- **Using a VPC Peering Connection to Connect Two VPCs in the Same Account**
- **Using a VPC Peering Connection to Connect Two VPCs in Different Accounts**

10.10 Elastic IP

The Elastic IP (EIP) service enables your cloud resources to communicate with the Internet using static public IP addresses and scalable bandwidths.

EIPs can be bound to or unbound from ECSs, BMSs, NAT gateways, and virtual IP addresses.

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

Figure 10-9 Accessing the Internet using an EIP
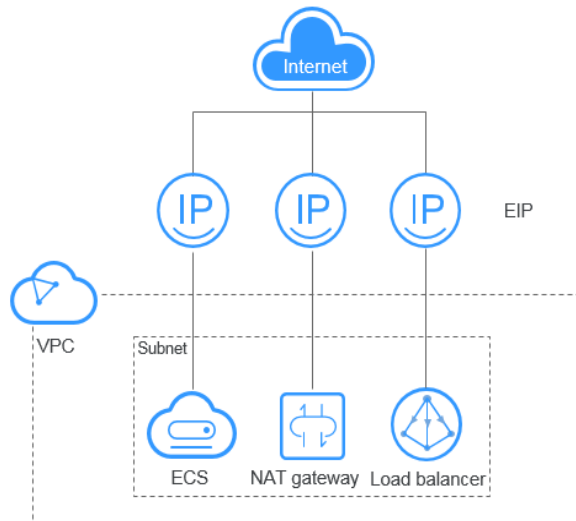
Managing EIPs

- **Assigning an EIP**
- **Binding an EIP to an Instance**

10.11 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.

Figure 10-10 shows the relationship between regions and AZs.

Figure 10-10 Regions and AZs
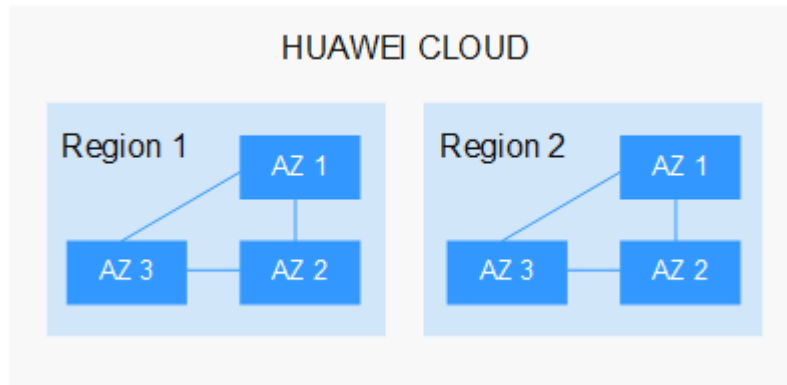
Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for lower network latency and quick access.

  - If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-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 Latin America, select the **LA-Santiago** region.

    📖 NOTE

    The **LA-Santiago** region is located in Chile.

- 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 lower 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**.