Virtual Private Cloud

Best Practices

 Issue
 40

 Date
 2024-04-26





HUAWEI TECHNOLOGIES CO., LTD.

Copyright © Huawei Technologies Co., Ltd. 2024. 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

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

Notice

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

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

Security Declaration

Vulnerability

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

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

For vulnerability information, enterprise customers can visit the following web page: <u>https://securitybulletin.huawei.com/enterprise/en/security-advisory</u>

Contents

1 VPC and Subnet Planning Suggestions	1
2 VPC Connectivity	6
3 Private Network Access	10
4 Public Network Access	14
5 Lower Network Costs	19
6 Using Access Control to Protect Your Resources	21
7 Using IP Address Groups to Reduce the Number of Security Group Rules	25
8 Using a Third-Party Firewall to Scrub Traffic for VPCs Connected by VPC F Connections	Peering 29
9 Using Third-Party Firewalls When Connecting an On-premises Data Center Cloud	er to the 40
10 Building Highly Available Web Server Clusters with Keepalived	44
11 Deploying Containers that Can Communicate with Each Other on Huaw ECSs	ei Cloud 55
12 Configuring Policy-based Routes for an ECS with Multiple NICs	
12.1 Overview	
	59
12.2 Collecting ECS Network Information	59 60
12.2 Collecting ECS Network Information 12.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6)	59 60 64
 12.2 Collecting ECS Network Information 12.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6) 12.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/IPv6) 	59 60 64 74
 12.2 Collecting ECS Network Information 12.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6) 12.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/IPv6) 13 VPC Peering Configurations 	
 12.2 Collecting ECS Network Information 12.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6) 12.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/IPv6) 13 VPC Peering Configurations	
 12.2 Collecting ECS Network Information	
 12.2 Collecting ECS Network Information	
 12.2 Collecting ECS Network Information	

1 VPC and Subnet Planning Suggestions

Before creating your VPCs, determine how many VPCs, the number of subnets, and what IP address ranges or connectivity options you will need.

- How Do I Determine How Many VPCs I Need?
- How Do I Plan Subnets?
- How Do I Plan Routing Policies?
- How Do I Connect to an On-Premises Data Center?
- How Do I Access the Internet?

How Do I Determine How Many VPCs I Need?

VPCs are region-specific. By default, networks in VPCs in different regions or even in the same region are not connected.

One VPC

If your services do not require network isolation, a single VPC should be enough.

• Multiple VPCs

If you have multiple service systems in a region and each service system requires an isolated network, you can create a separate VPC for each service system.

If you require network connectivity between separate VPCs in the same account or in different accounts, you can use VPC peering connections or Cloud Connect.

- If two VPCs are in the same region, use a VPC peering connection.
- If two VPCs are in different regions, use Cloud Connect.

NOTE

By default, you can create a maximum of five VPCs in each region. If this cannot meet your service requirements, request a quota increase. For details, see **How Do I Apply for a Higher Quota?**

The following table lists the private CIDR blocks that you can specify when creating a VPC. Consider the following when selecting a VPC CIDR block:

 Number of IP addresses: Reserve sufficient IP addresses in case of business growth. • IP address range: Avoid IP address conflicts if you need to connect a VPC to an on-premises data center or connect two VPCs.

 Table 1-1 lists the supported VPC CIDR blocks.

Table	1-1	VPC	CIDR	blocks
-------	-----	-----	------	--------

VPC CIDR Block	IP Address Range	Maximum Number of IP Addresses
10.0.0/8-24	10.0.0.0-10.255.255.255	2^24-2=16777214
172.16.0.0/12-24	172.16.0.0-172.31.255.25 5	2^20-2=1048574
192.168.0.0/16-24	192.168.0.0-192.168.255. 255	2^16-2=65534

NOTICE

Table 1-1 describes the primary IPv4 CIDR blocks available for a VPC. After a VPC is created, its primary CIDR block cannot be modified. If the primary CIDR block is insufficient, you can **add a secondary IPv4 CIDR block for the VPC**.

How Do I Plan Subnets?

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.

• After a subnet is created, its CIDR block cannot be modified. Subnets in the same VPC cannot overlap.

When you create a VPC, a default subnet will be created together. If you need more subnets, see **Creating a Subnet for the VPC**.

A subnet mask can be between the netmask of its VPC CIDR block and /28 netmask. If a VPC CIDR block is 10.0.0/16, its subnet mask can be between 16 and 28.

For example, if the CIDR block of VPC-A is 10.0.0/16, you can specify 10.0.0/24 for subnet A01, 10.0.1.0/24 for subnet A02, and 10.0.2.0/24 for subnet A03.

NOTE

By default, you can create a maximum of 100 subnets in each region. If this cannot meet your service requirements, request a quota increase by referring to **How Do I Apply for a Higher Quota?**

When planning subnets, consider the following:

• You create different subnets for different modules in a VPC. For example, in VPC-A, you can create subnet A01 for web services, subnet A02 for management services, and subnet A03 for data services. You can leverage network ACLs to control access to each subnet.

• If your VPC needs to communicate with an on-premises data center through VPN or Direct Connect, ensure that the VPC subnet and the CIDR block used for communication in the data center do not overlap.

How Do I Plan Routing Policies?

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. A route table contains a set of routes that are used to determine where network traffic from your subnets in a VPC is directed. The default route table ensures that subnets in a VPC can communicate with each other.

If you do not want to use the default route table, you can now create a custom route table and associate it with the subnets. The custom route table associated with a subnet affects only the outbound traffic. The default route table controls the inbound traffic.

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

- System routes: Routes that are automatically added by the system and cannot be modified or deleted. System routes allow instances in a VPC to communicate with each other.
- Custom routes: Routes that can be modified and deleted. The destination of a custom route cannot overlap with that of a system route.

You cannot add two routes with the same destination to a VPC route table even if their next hop types are different, because the destination determines the route priority. According to the longest match routing rule, the destination with a higher matching degree is preferentially selected for packet forwarding.

How Do I Connect to an On-Premises Data Center?

If you require interconnection between a VPC and an on-premises data center, ensure that the VPC does not have an overlapping IP address range with the onpremises data center to be connected.

As shown in **Figure 1-1**, you have VPC 1 in region A and VPC 2 and VPC 3 in region B. To connect to an on-premises data center, they can use a VPN, as VPC 1 does in Region A; or a Direct Connect connection, as VPC 2 does in Region B. VPC 2 connects to the data center through a Direct Connect connection, but to connect to another VPC in that region, like VPC 3, a VPC peering connection must be established.



Figure 1-1 Connections to on-premises data centers

When planning CIDR blocks for VPC 1, VPC 2, and VPC 3:

- The CIDR block of VPC 1 cannot overlap with the CIDR block of the onpremises data center in Region A.
- The CIDR block of VPC 2 cannot overlap with the CIDR block of the onpremises data center in Region B.
- The CIDR blocks of VPC 2 and VPC 3 cannot overlap.

How Do I Access the Internet?

Use EIPs to enable a small number of ECSs to access the Internet.

When only a few ECSs need to access the Internet, you can bind the EIPs to the ECSs. This will provide them with Internet access. You can also dynamically unbind the EIPs from the ECSs and bind them to NAT gateways and load balancers instead, which will also provide Internet access. The process is not complicated.

For more information about EIP, see **EIP Overview**.

Use a NAT gateway to enable a large number of ECSs to access the Internet.

When a large number of ECSs need to access the Internet, the public cloud provides NAT gateways for your ECSs. With NAT gateways, you do not need to assign an EIP to each ECS. NAT gateways reduce costs as you do not need so many EIPs. NAT gateways offer 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. SNAT prevents the EIPs of ECSs from being exposed to the Internet. DNAT can implement portlevel 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.

For more information, see **NAT Gateway User Guide**.

Use ELB to access the Internet If there are a large number of concurrent requests.

In high-concurrency scenarios, such as e-commerce, you can use load balancers provided by the ELB service to evenly distribute incoming traffic across multiple ECSs, allowing a large number of users to concurrently access your business system or application. ELB is deployed in the cluster mode. It provides fault tolerance for your applications by automatically balancing traffic across multiple AZs. You can also take advantage of deep integration with Auto Scaling (AS), which enables automatic scaling based on service traffic and ensures service stability and reliability.

For more information, see **Elastic Load Balance User Guide**.

Helpful Links

- Application Scenarios
- Private Network Access
- Public Network Access

2_{VPC Connectivity}

Accessing the Internet

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

Cloud Service	Application Scenario	Description	Reference
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.	Configuring the VPC of ECSs That Access the Internet Using EIPs
		An EIP can be bound to an ECS to enable Internet access, or unbound to disable access.	
		Shared bandwidths and shared data packages can be used to lower costs.	

Tabl	le 2-1	Accessing	the	Internet
140		/ CCCJJIIIQ	unc.	IIIICCI IICC

Cloud Service	Application Scenario	Description	Reference
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 EIPs to access the Internet. In this way, you can reduce management costs and prevent the EIPs of ECSs from being exposed to the Internet. DNAT implements 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. However, DNAT does not balance traffic.	Using SNAT to Access the Internet Using DNAT to Provide Services Accessible from the Internet
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 ELB?

Connecting VPCs

You can connect VPCs using the following cloud services.

Cloud Service	Application Scenario	Description	Reference
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 of charge.	Creating a VPC Peering Connection with Another VPC in Your Account Creating a VPC Peering Connection with a VPC in Another Account
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 will be affected by the Internet quality.	Connecting to a VPC Through a VPN

Table 2-2 Connecting VPCs

Connecting to an On-premises Data Center (IDC)

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

Cloud Service	Applicatio n Scenario	Description	Reference
VPN	Use VPN to connect a VPC to an on- premises data center at 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 will be affected by the Internet quality.	Connecting to a VPC Through a VPN

 Table 2-3 Connecting to an on-premises data center

Cloud Service	Applicatio n Scenario	Description	Reference
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 features low latency and is very secure. Direct Connect is a good choice if you have strict requirements on network transmission quality and security.	Accessing Multiple VPCs Using a Connection

3 Private Network Access

Connecting to an On-premises Data Center

You can connect a VPC to your on-premises data center. Once you have established this secure, reliable connection, you can move at scale to Huawei Cloud, a cloud with massive computing, storage, and network resources. With Huawei Cloud, you will be unaffected by sudden fluctuations in demand for services. Both Direct Connect and VPN support the connections between your data center and your VPCs on the cloud.

Direct Connect

Direct Connect provides high-speed, stable, and secure dedicated network connections that connect your data centers to VPCs. With Direct Connect, you can connect computers in your on-premises data center to cloud servers or hosting servers on Huawei Cloud. It maximizes cloud computing capacities and existing IT facilities to build a flexible, scalable hybrid cloud computing environment.

Figure 3-1 Connecting to an on-premises data center with a Direct Connect connection



VPN

VPN establishes a secure, encrypted communication tunnel between your local data center and your VPC on Huawei Cloud. With VPN, you can connect to a VPC and access the resources deployed there.

Connecting VPCs and Data Centers with Cloud Connect

Cloud Connect allows you to quickly build high-quality networks that can connect VPCs across regions and work with Direct Connect to connect VPCs and onpremises data centers.

With Cloud Connect, you can build a globally connected cloud network with enterprise-class scalability and communications capabilities.

Figure 3-2 Connecting VPCs and data centers with Cloud Connect



Connecting VPCs

If you want to connect VPCs in the same region, you can use VPC peering connections.

If you want to connect VPCs in different regions and construct a service network across regions, you can use Direct Connect, VPN, or Cloud Connect.

• VPC peering

You can use VPC peering connections to connect VPCs in the same region.

Figure 3-3 Connecting VPCs in the same region with a VPC peering connection



• Direct Connect

Direct Connect provides high-speed, stable, and secure dedicated network connections that connect your data centers to VPCs. With Direct Connect, you can connect computers in your on-premises data center to cloud servers or hosting servers on Huawei Cloud. It maximizes cloud computing capacities and existing IT facilities to build a flexible, scalable hybrid cloud computing environment. Direct Connect can also be used to connect VPCs in different regions.

Figure 3-4 Connecting VPCs in different regions with Direct Connect



• VPN

VPN establishes a secure, encrypted communication tunnel between your local data center and your VPC on Huawei Cloud. With VPN, you can connect to a VPC and access the resources deployed there. VPN can connect VPCs in different regions.





Cloud Connect

Cloud Connect allows you to quickly build high-quality networks that can connect VPCs across regions and work with Direct Connect to connect VPCs and on-premises data centers. With Cloud Connect, you can build a globally connected cloud network with enterprise-class scalability and communications capabilities.



Figure 3-6 Connecting VPCs in different regions with Cloud Connect

4 Public Network Access

Products

Cloud services, such as EIP, NAT Gateway, and ELB can be used to connect to the Internet.

• EIP

The EIP service provides independent public IP addresses and bandwidth for Internet access. 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.

• ELB

ELB distributes access traffic among multiple ECSs to balance the application load, improving fault tolerance and expanding service capabilities of applications. You can create a load balancer, configure a listening protocol and port, and add backend servers to a load balancer. You can also check the running state of backend servers to ensure that requests are sent only to healthy servers.

NAT Gateway

NAT Gateway provides both SNAT and DNAT for your servers in a VPC and allows servers in your VPC to access or provide services accessible from the Internet.

Providing Services Accessible from the Internet

- Single ECS provides services accessible from the Internet.
 - If you have only one application and the service traffic is small, you can assign an EIP and bind it to the ECS so that the ECS can provide services accessible from the Internet.





• Multiple ECSs balance workloads.

In high-concurrency scenarios, such as e-commerce, you can use load balancers provided by the ELB service to evenly distribute incoming traffic across multiple ECSs, allowing a large number of users to concurrently access your business system or application. ELB deeply integrates with the Auto Scaling (AS) service, which enables automatic scaling based on service traffic and ensures service stability and reliability.



Accessing the Internet

• Single ECS accesses the Internet.

When an ECS needs to access the Internet, you can bind an EIP to the ECS so that the ECS can access the Internet. Huawei Cloud allows your EIP to be billed on a pay-per-use basis. If you do not need to use the EIP, you can flexibly unbind it.





• Multiple ECSs access the Internet.

If multiple ECSs in your VPC need to access the Internet, you can use a NAT gateway and configure SNAT rules by subnet to allow ECSs in the VPC to access the Internet. If you access to the Internet using an EIP but with no DNAT rules configured, external users cannot directly access the public network address of the NAT gateway through the Internet, ensuring ECS security.



Figure 4-4 NAT gateway

5 Lower Network Costs

You can select a proper product and billing mode based on your service requirements.

Dedicated Bandwidth

If you want to ensure the bandwidth available for a particular EIP, you are advised to purchase dedicated bandwidth. Dedicated bandwidth can only be used for a single, specific EIP. Dedicated bandwidth is not affected by other services.

An EIP can be billed by bandwidth or by traffic:

- Bandwidth: If your services use a large amount of traffic but are stable, an EIP billed by bandwidth is recommended.
- Traffic: If your services only use a relatively small amount of traffic, an EIP billed by traffic combined with a shared data package is recommended for a more favorable price.

If your traffic is stable, the yearly/monthly billing based on the bandwidth is more cost effective.

Shared Bandwidth

When you host a large number of applications on the cloud, if each EIP uses dedicated bandwidth, a lot of bandwidths are required, which incurs high costs. If all EIPs share the same bandwidth, your network operation costs will be lowered and your system O&M as well as resource statistics will be simplified. Multiple EIPs whose billing mode is pay-per-use can be added to a shared bandwidth. You can bind EIPs to products such as ECSs, NAT gateways, and load balancers so that these products can use the shared bandwidth.

Shared Data Package

A shared data package is a prepaid package for public network traffic. The price of the package is lower than that for the postpaid billing by traffic. Shared data packages greatly reduce the cost of traffic on a public network. A shared data package takes effect immediately after being purchased and no additional operations are required. If you have subscribed to pay-per-use EIPs billed by traffic in a region and buy a shared data package in the same region, the EIPs will use the shared data package.

• When to use a shared data package

After a shared data package takes effect for a bandwidth billed by traffic, the traffic used by the bandwidth is deducted from the shared data package first. After the shared data package is used up, the bandwidth is billed by the amount of traffic used. A shared data package saves more if your amount of traffic used is huge.

- Additional notes on shared data packages
 - Only the traffic generated in the region selected when the shared data package is purchased can be deducted.
 - Dynamic and static shared data packages are used to deduct the traffic generated by dynamic BGP and static BGP EIPs, respectively.
 - A shared data package has a validity period of one calendar month or one calendar year from the date of purchase. After this period expires, the unused traffic expires as well and cannot be used. You are advised to evaluate the size of a shared data package required based on the historical usage.
 - If you enable the auto-renew function for a shared data package, the system automatically attempts to renew the subscription within seven days before the shared data package expires. After the renewal is successful, the remaining traffic in the shared data package can be used within the new validity period.
 - After a shared data package is used up, your service will not automatically stop. The system automatically bills you based on traffic, ensuring service system availability.

6 Using Access Control to Protect Your Resources

A VPC is your private network on the cloud. You can configure security groups and network ACL rules to ensure the security of instances, such as ECSs, databases, and containers, running in a VPC.

- A security group protects the instances in it.
- A network ACL protects associated subnets and all the resources in the subnets.

As shown in **Figure 6-1**, security groups A and B protect the network security of ECSs. Network ACLs A and B add an additional layer of defense to ECSs in subnets 1 and 2.



Figure 6-1 Security groups and network ACLs

Here are some common access control configuration examples:

- Security group: Allowing Traffic from Given IP Addresses or a Security Group
- Security group: Allowing Traffic from a Virtual IP Address

- Security group: Allowing Communications Between Instances in Two VPCs Connected by a VPC Peering Connection
- Network ACL: Allowing Traffic from Specific IP Addresses to the Instances in a Subnet

Allowing Traffic from Given IP Addresses or a Security Group

You can add inbound rules to allow traffic from specific IP addresses and other security groups. As you can see in Figure 6-2, there are two subnets (Subnet-A and Subnet-B) in VPC-X. ECSs in Subnet-A are associated with security group Sg-A, and ECSs in Subnet-B are associated with security group Sg-B.

- Add inbound rule A01 to Sg-A to allow traffic from IP addresses in 172.16.0.0/24 to access SSH port 22 on the ECSs in Sg-A for remotely logging in to these ECSs.
- Add inbound rule B01 to **Sg-B** to allow the ECSs in **Sg-A** to access SSH port 22 on the ECSs in **Sg-B** for remotely logging in to the ECSs in **Subnet-B**.

Figure 6-2 Security group examples (allowing traffic from given IP addresses and security groups)



Allowing Traffic from a Virtual IP Address

You can add inbound rules to allow traffic from virtual IP addresses and other security groups. In Figure 6-3, there are two subnets (Subnet-A and Subnet-B) in VPC-X. ECSs in Subnet-A are associated with security group Sg-A, and ECSs in Subnet-B are associated with security group Sg-B.

- Add inbound rule A01 to Sg-A to allow the ECSs in Sg-B to access the ECSs in Sg-A using private IP addresses.
- Add inbound rule B01 to Sg-B to allow traffic from virtual IP address
 192.168.0.21 to the ECSs in Sg-B using any protocol over any port. You can also set the source to the CIDR block of Subnet-A (192.168.0.0/24).

Do not add rules like rule B02. This rule allows the ECSs in **Sg-A** to access the ECSs in **Sg-B** using private IP addresses but not virtual IP address **192.168.0.21**.



Figure 6-3 Security group examples (allowing traffic from a virtual IP address)

Allowing Communications Between Instances in Two VPCs Connected by a VPC Peering Connection

In **Figure 6-4**, **VPC-A** and **VPC-B** are connected by VPC peering connection **peering-AB**. To allow ECSs in **Sg-A** and **Sg-B** to communicate with each other, you can add the following rules:

- Rule A01 with **Source** to **Sg-B** to allow ECSs in **Sg-B** to access ECSs in **Sg-A**.
- Rule B01 with Source to Sg-A to allow ECSs in Sg-A to access ECSs in Sg-B.

Figure 6-4 Security group examples (allowing communications between ECSs in two VPCs connected by a VPC peering connection)



Allowing Traffic from Specific IP Addresses to the Instances in a Subnet

A network ACL controls traffic in and out of a subnet. If both security group and network ACL rules are configured, traffic matches network ACL rules first and then security group rules. You can configure network ACL rules to add an additional layer of protection to your instances.

- Inbound rule A01 of security group **Sg-A** allows traffic from any IP address to the ECSs in **Sg-A** over SSH port 22.
- To enhance security, add inbound rule A01 to network ACL **Fw-A** to allow traffic from 10.0.1.0/24 to the ECSs in **Subnet-A** over SSH port 22.

The default inbound rule of **Fw-A** denies all traffic to the ECSs in **Subnet-A**. So you need to add custom rules to allow access from specific IP addresses. If you associate **Subnet-A** with **Fw-A**, the instances in **Subnet-A** can still communicate with each other.

Figure 6-5 Network ACL examples (allowing traffic from specific IP addresses to the instances in a subnet)



7 Using IP Address Groups to Reduce the Number of Security Group Rules

Scenarios

An IP address group is a collection of one or more IP addresses. You can use IP address groups when configuring security group rules. If you change the IP addresses in an IP address group, the security group rules are changed accordingly. You do not need to modify the security group rules one by one.

Finance and securities enterprises have high security requirements when planning cloud networks. Access to instances is often controlled based on IP addresses. To simplify security group rule configuration and control access based on IP addresses, you can use IP address groups to manage IP address ranges and IP addresses with the same security requirements. For more information about IP address groups, see **IP Address Group Overview**.

Suppose your enterprise has an online office system deployed on the cloud. To provide services for different departments, you associate office servers with different security groups based on security levels. These servers are accessed from a large number of IP addresses that may change from time to time.

- If IP address groups are not used, you need to configure multiple rules to control access from different sources. Once the IP addresses change, you need to adjust the rules in each security group one by one. The management workload increases with the number of security groups and rules.
- If IP address groups are used, you can add the IP addresses with the same security requirements to an IP address group and add rules with source set to this IP address group. When an IP address changes, you only need to change it in the IP address group. Then, the security group rules using the IP address group change accordingly. You do not need to modify the security group rules one by one. This simplifies security group management and improves efficiency.

Solution Architecture

In this practice, the instances are associated with three security groups based on different security requirements. In addition, these instances need to be accessed by specific IP addresses over SSH port 22. To simplify management, you can use IP address groups.

- 1. Create an IP address group and add IP addresses that need to access the instances.
- 2. Add inbound rules to allow traffic from the IP address group to the instances in the three security groups.

Table 7-1	Inbound	rules
-----------	---------	-------

Direction	Action	Туре	Protocol & Port	Source
Inbound	Allow	IPv4	TCP:22	IP address group

3. Change the IP addresses in the IP address group if any IP addresses change. Then, the rules using the IP address group change accordingly.

Constraints

Security group rules using IP address groups do not take effect for the following instances:

- General computing (S1, C1, and C2 ECSs)
- Memory-optimized (M1 ECSs)
- High-performance computing (H1 ECSs)
- Disk-intensive (D1 ECSs)
- GPU-accelerated (G1 and G2 ECSs)
- Large-memory (E1, E2, and ET2 ECSs)

Resource Planning

In this practice, the IP address group and security groups must be in the same region. For details, see **Table 7-2**. The following resource details are only examples. You can modify them as required.

Resource	Quantity	Description	
IP address group	1	Create an IP address group and add IP addresses that need to access the instances.	
		Name: ipGroup-A	
		• Max. IP Addresses: Set it as required. In this practice, 20 is used.	
		• IP Address Version: Set it as required. In this practice, IPv4 is used.	
		IP Addresses:	
		– 11.xx.xx.64/32	
		– 116.xx.xx.252/30	
		– 113.xx.xx.0/25	
		– 183.xx.xx.208/28	
Security group	3	Add inbound rules to allow traffic from ipGroup-A to the instances in the three security groups, as shown in Table 7-3 .	

 Table 7-2 Resource planning

Table 7-3 Inbound rules

Direction	Action	Туре	Protocol & Port	Source
Inbound	Allow	IPv4	TCP:22	ipGroup-A

Procedure

Step 1 Create IP address group **ipGroup-A** and add IP addresses that need to access the instances.

For details, see Creating an IP Address Group.

Step 2 Add inbound rules to allow traffic from **ipGroup-A** to the instances in the three security groups.

For details, see Adding a Security Group Rule.

After the rules are added, traffic from 11.xx.xx.64/32, 116.xx.xx.252/30, 113.xx.xx.0/25, and 183.xx.xx.208/28 are allowed to the Linux ECSs over SSH port 22.

Step 3 Change IP addresses in the IP address group.

After security group rules are added, you can add IP addresses to **ipGroup-A**. For example, you can add 117.xx.xx.0/25 to **ipGroup-A**, and the security groups rule is applied automatically, allowing traffic from 117.xx.xx.0/25 over SSH port 22.

For details, see Managing IP Addresses in an IP Address Group.

----End

8 Using a Third-Party Firewall to Scrub Traffic for VPCs Connected by VPC Peering Connections

Application Scenario

VPC allows you to configure and manage virtual networks. You can use security groups and network ACLs to control network access. You can also use third-party firewalls to ensure the security of cloud services.

This section describes how to use a firewall to scrub traffic across VPCs that are connected using VPC peering connections.

Architecture

In this example, services are deployed in VPC-A, VPC-B, and VPC-C, and the firewall is deployed in VPC-X. These VPCs communicate with each other through VPC peering connections. The traffic across VPC-A, VPC-B, and VPC-C must flow through the firewall in VPC-X. The default route table of VPC-X directs all inbound traffic to the firewall. After being scrubbed by the firewall, the traffic is sent to a service VPC based on the custom route table.

Figure 8-1 shows how ecs-A01 accesses ecs-C01. You can view the request and response traffic paths.



Figure 8-1 Networking planning when a third-party firewall is used for scrubbing traffic across VPCs

Resource Planning

In this example, you need to create VPCs, ECSs, and VPC peering connections. For details about required resources, see **Table 8-1**.

NOTE

The following resource planning details are only examples for your reference. You need to plan resources based on actual service requirements.

Resource	Description			
VPC	Table 8-2 shows details about the required VPCs.In this example, there are four VPCs, including three VPCs whereservices are deployed and one VPC where the firewall isdeployed. These VPCs are from the same region, and theirsubnet CIDR blocks do not overlap.			
	• Services are deployed in VPC-A, VPC-B, and VPC-C, and the firewall is deployed in VPC-X. These VPCs communicate version other through VPC peering connections.			
	• VPC-A, VPC-B, VPC-C and VPC-X each have a subnet.			
	• The subnets of VPC-A, VPC-B, VPC-C are associated with their default route table.			
	• VPC X has a default route table and a custom route table. The subnet of VPC X is associated with the custom route table.			
	The default route table controls the inbound traffic to VPC-X, and the custom route table controls the outbound traffic from VPC-X.			
	NOTICE The subnet CIDR blocks of the VPCs that need to communicate with each other through a VPC peering connection cannot overlap. Otherwise, the VPC peering connection does not take effect. For details, see Unsupported VPC Peering Configurations .			
ECS	Table 8-3 shows details about the required ECSs.			
	The four ECSs are in different VPCs. If the ECSs are in different security groups, add rules to the security groups to allow access to each other.			
VPC peering connection	Table 8-4 shows details about the required VPC peering connections.			
	There are three VPC peering connections.			
	• peer-AX: connects VPC-A and VPC-X			
	peer-BX: connects VPC-B and VPC-X			
	peer-CX: connects VPC-C and VPC-X			
	VPC peering connections are transitive. After routes are configured, VPC-A, VPC-B, and VPC-C can communicate with each other through VPC-X.			

Table 8-1 Required resources

VPC Name	VPC CIDR Block	Subnet Name	Subnet CIDR Block	Route Table	Subnet Is Used to Deploy
VPC-A	10.1.0.0/16	subnet- A01	10.1.0.0/24	Default route table	Services
VPC-B	10.2.0.0/16	subnet- B01	10.2.0.0/24	Default route table	Services
VPC-C	10.3.0.0/16	subnet- C01	10.3.0.0/24	Default route table	Services
VPC-X	192.168.0.0 /16	subnet- X01	192.168.0.0 /24	Custom route table	Firewall

Table 8-2 VPC details

Table 8-3 ECS details

ECS Name	VPC Name	Subnet Name	Private IP Address	lmag e	Securi ty Group	ECS Is Used to Deploy
ecs- A01	VPC-A	subnet- A01	10.1.0.13 9	Public image	sg- demo:	Services
ecs- B01	VPC-B	subnet- B01	10.2.0.93	: Cent OS	Genera l-	Services
ecs- C01	VPC-C	subnet- C01	10.3.0.22 0	8.2 64bit	e web server	Services
ecs- X01	VPC-X	subnet- X01	192.168. 0.5			Firewall

Table 8-4 VPC peering connection details

Connection Name	Local VPC	Peer VPC
peer-AX	VPC-A	VPC-X
peer-BX	VPC-B	VPC-X
peer-CX	VPC-C	VPC-X

Route Configuration

You need to add routes to VPC route tables to allow communication between VPCs and scrub traffic through the firewall. For details, see **Table 8-5**.
NOTE

The following routes are only examples for your reference. You need to plan routes based on actual service requirements.

Table 8-5 Required route tables

Route Table	Description
Route tables of service VPCs	Table 8-6 shows details about route tables of service VPCs.The default route tables of VPC-A, VPC-B, and VPC-C have routeswith destinations set to other VPC subnets and with next hop setto VPC peering connection.
Route tables of firewall VPC	 Table 8-7 shows details about route tables of the firewall VPC-X. 1. In the default route table of VPC-X: If the firewall is deployed on an ECS, add a route with destination set to 0.0.0.0/0 and next hop set to ecs-X01 to direct traffic to the ECS with the firewall deployed
	 If the firewall is deployed on two ECSs and the ECSs communicate with external systems through a virtual IP address, the virtual IP address is dynamically switched to the standby ECS to continue providing services when the active ECS is faulty and cannot provide services. In this scenario, add a route with destination set to 0.0.0.0/0 and next hop set to the virtual IP address to direct traffic to the ECS with the firewall deployed.
	In this example, the firewall is deployed on an ECS. The traffic across VPC-A, VPC-B, and VPC-C needs to pass through VPC-X and be directed to the firewall for scrubbing.
	 In the custom route table of VPC-X, add routes with destination set to subnet CIDR blocks of service VPCs (VPC-A, VPC-B, and VPC-C) and next hop set to VPC peering connection.

VP C N a m e	Route Table	Destinati on	Next Hop Type	Next Hop	Route Type	Route Function
VP C- A	Default route table: rtb-vpc-A	10.2.0.0/2 4	VPC peering connecti on	peer-AX	Custo m	 Destination: subnet-B01 in VPC-B Connects subnet-A01 to subnet-B01
		10.3.0.0/2 4	VPC peering connecti on	peer-AX	Custo m	 Destination: subnet-C01 in VPC-C Connects subnet-A01 to subnet-C01
		192.168.0. 0/24	VPC peering connecti on	peer-AX	Custo m	 Destination: subnet-X01 in VPC-X Connects subnet-A01 to subnet-X01
VP C- B	Default route table: rtb-vpc-B	10.1.0.0/2 4	VPC peering connecti on	peer-BX	Custo m	 Destination: subnet-A01 in VPC-A Connects subnet-A01 to subnet-B01
		10.3.0.0/2 4	VPC peering connecti on	peer-BX	Custo m	 Destination: subnet-C01 in VPC-C Connects subnet-B01 to subnet-C01
		192.168.0. 0/24	VPC peering connecti on	peer-BX	Custo m	 Destination: subnet-X01 in VPC-X Connects subnet-B01 to subnet-X01

Table 8-6 Details about route tables of service VPCs

VP C N a m e	Route Table	Destinati on	Next Hop Type	Next Hop	Route Type	Route Function
VP C- C	Default route table: rtb-vpc-C	10.1.0.0/2 4	VPC peering connecti on	peer-CX	Custo m	 Destination: subnet-A01 in VPC-A Connects subnet-A01 to subnet-C01
		10.2.0.0/2 4	VPC peering connecti on	peer-CX	Custo m	 Destination: subnet-B01 in VPC-B Connects subnet-B01 to subnet-C01
		192.168.0. 0/24	VPC peering connecti on	peer-CX	Custo m	 Destination: subnet-X01 in VPC-X Connects subnet-C01 to subnet-X01

VP C Na me	Route Table	Destinati on	Next Hop Type	Next Hop	Route Type	Route Function
VP C- X	Default route table: rtb-vpc-X	0.0.0/0	Server	ECS-X	Custo m	 Destination: ecs- X with firewall deployed Direct inbound traffic of VPC-X to the firewall. If your firewall is deployed on multiple ECSs and these ECSs communicate with external networks through a virtual IP address, set the next hop of the route to the virtual IP address.
	Custom route table: rtb-vpc- custom- X	10.1.0.0/2 4	VPC peering connecti on	peer-AX	Custo m	 Destination: subnet-A01 in VPC-A Connects subnet-A01 to subnet-X01
		10.2.0.0/2 4	VPC peering connecti on	peer-BX	Custo m	 Destination: subnet-B01 in VPC-B Connects subnet-B01 to subnet-X01
		10.3.0.0/2 4	VPC peering connecti on	peer-CX	Custo m	 Destination: subnet-C01 in VPC-C Connects subnet-C01 to subnet-X01

Table 8-7 Details about route tables of firewall VPC

Notes and Constraints

- A VPC peering connection can only enable communication between VPCs in the same region.
- The subnet CIDR blocks of the VPCs that need to communicate with each other through a VPC peering connection cannot overlap. Otherwise, the VPC

peering connection does not take effect. For details, see **Unsupported VPC Peering Configurations**.

• The subnet where the ECS deployed with a third-party firewall resides needs to be associated with a custom route table. Ensure that the region where your resources are located supports custom route tables.

If **Route Tables** is displayed in the left pane of the network console, custom route tables are supported.

Figure 8-2 Route Tables



Procedure

Step 1 Create four VPCs and their subnets in region A.

For details, see **Creating a VPC**.

For details about VPCs and their subnets, see Table 8-2.

- **Step 2** Create a custom route table in VPC-X and associate subnet-X01 with the custom route table.
 - Create a custom route table in VPC-X.
 For details, see Creating a Custom Route Table.
 - 2. Associate subnet-X01 with the custom route table created in **Step 2.1**.

After subnet-X01 is created, it is automatically associated with the default route table of VPC-X. You need to associate the custom route table created in **Step 2.1** to subnet-X01.

For details, see Changing the Route Table Associated with a Subnet.

Step 3 Create an ECS in each VPC.

For details, see **Purchasing an ECS**.

- **Step 4** Configure the NIC of ecs-X and install the third-party firewall on ecs-X.
 - Disable source/destination check for the NIC of ecs-X.
 For details, see Disabling Source and Destination Check.
 - 2. Install a third-party firewall on ecs-X.
- **Step 5** (Optional) Configure a virtual IP address for ECSs.

You can create two ECSs in VPC-X and bind them to the same virtual IP address so that they can work in the active and standby mode. If the active ECS is faulty and cannot provide services, the virtual IP address will be dynamically switched to the standby ECS to continue providing services. Skip this step if the ECS where the firewall is deployed does not need to work in the active/standby mode.

1. Assign a virtual IP address in the VPC-X subnet.

For details, see Assigning a Virtual IP Address.

2. Bind the virtual IP address to the active and standby ECSs where the firewall is deployed.

For details, see **Binding a Virtual IP Address to an EIP or ECS**.

Step 6 Create three VPC peering connections and configure routes.

- 1. Create three VPC peering connections.
 - If your VPCs are in the same account, see **Creating a VPC Peering Connection with Another VPC in Your Account**.
 - If your VPCs are in different accounts, see Creating a VPC Peering Connection with a VPC in Another Account.

For details about VPC peering connections, see Table 8-4.

2. In the default route tables of the three service VPCs, add routes with destination set to the other three VPCs and with next hop set to the VPC peering connection.

For details, see Adding a Custom Route.

In this example, add the routes planned in **Table 8-6** to the route tables of VPC-A, VPC-B, and VPC-C.

3. Add routes to the default and custom route tables of the firewall VPC.

For details, see Adding a Custom Route.

In this example, add the routes planned in **Table 8-7** to the default and custom route tables of VPC-X.

Step 7 Log in to the ECS and check whether the firewall takes effect.

Multiple methods are available for logging in to an ECS. For details, see **Logging In to an ECS**.

In this example, use VNC provided on the management console to log in to an ECS.

1. Log in to ecs-A01 and verify the network connectivity between VPC-A and VPC-B.

ping Private IP address of ecs-B01

Example command:

ping 10.2.0.93

If information similar to the following is displayed, the two VPCs can communicate with each other.

[root@ecs-A01 ~]# ping 10.2.0.93 PING 10.2.0.93 (10.2.0.93) 56(84) bytes of data. 64 bytes from 10.2.0.93: icmp_seq=1 ttl=64 time=0.849 ms 64 bytes from 10.2.0.93: icmp_seq=2 ttl=64 time=0.455 ms 64 bytes from 10.2.0.93: icmp_seq=3 ttl=64 time=0.385 ms 64 bytes from 10.2.0.93: icmp_seq=4 ttl=64 time=0.372 ms --- 10.2.0.93 ping statistics ---

- 2. Keep the network connectivity between VPC-A and VPC-B in **Step 7.1** and log in to ecs-X01 to verify whether the traffic from VPC-A to VPC-B flows through ecs-X01.
- 3. On ecs-X01, check the traffic change on eth0.

Run the following command at least twice consecutively to check whether the values of RX packets and TX packets change:

ifconfig eth0

If the packets change, the traffic flows through ecs-X01 and is scrubbed by the firewall.

[root@ecs-X01 ~]# ifconfig eth0 eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 192.168.0.5 netmask 255.255.255.0 broadcast 192.168.0.255 inet6 fe80::f816:3eff:feb6:a632 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:b6:a6:32 txqueuelen 1000 (Ethernet) RX packets 726222 bytes 252738526 (241.0 MiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 672597 bytes 305616882 (291.4 MiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0 [root@ecs-X01 ~]# ifconfig eth0 eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 192.168.0.5 netmask 255.255.255.0 broadcast 192.168.0.255 inet6 fe80::f816:3eff:feb6:a632 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:b6:a6:32 txqueuelen 1000 (Ethernet) RX packets 726260 bytes 252748508 (241.0 MiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 672633 bytes 305631756 (291.4 MiB)

- TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
- 4. Repeat **Step 7.1** to **Step 7.3** to check the communication between other VPCs.

----End

9 Using Third-Party Firewalls When Connecting an On-premises Data Center to the Cloud

Scenarios

Your on-premises data center communicates with Huawei Cloud through Direct Connect or VPN. A third-party virtual firewall is deployed on the cloud to filter traffic.

This section describes how to use a third-party virtual firewall when connecting your on-premises data center to multiple VPCs.

Solution Advantages

- You can use third-party firewalls.
- The traffic between the cloud and the on-premises data center will pass through the third-party virtual firewall.
- You can define security rules as required.

Typical Topology

Assume that your services are deployed in VPC 1, VPC 2, VPC 3, and your onpremises data center, and you need to use a third-party virtual firewall on the cloud. You can configure the virtual firewall on ECS 2 in VPC 2 and use VPC peering connections and configure routes to enable communication between the VPCs. In addition, you need to create a Direct Connect connection to enable communication between VPC 3 and the on-premises data center.

The deployment diagram is as follows:



Figure 9-1 Deployment diagram

Prerequisites

The subnet CIDR blocks of VPC 1, VPC 2, and VPC 3 cannot overlap with each other. Otherwise, communication through VPC peering connections will fail.

Procedure

Step 1 Create VPCs.

Create VPC 1, VPC 2, and VPC 3.

For details, see Creating a VPC.

NOTE

The CIDR blocks of VPC 1, VPC 2, and VPC 3 cannot overlap with each other. For example, the CIDR block of VPC 1 is 10.0.1.0/24, VPC 2 is 10.0.2.0/24, and VPC 3 is 172.16.0.0/16.

Step 2 Create ECSs.

1. Create ECS 1 and ECS 2, which belong to the VPC 1 subnet and VPC 2 subnet, respectively.

For details, see Purchasing an ECS.

NOTE

The source/destination check must be disabled for the ECS 2 NIC.

2. Deploy a third-party virtual firewall on ECS 2.

Step 3 Create VPC peering connections.

Create VPC peering connections between VPC 1 and VPC 2, VPC 2 and VPC 3 to enable communications between them.

When creating a VPC peering connection, do not configure routes for the local and peer ends. Configure routes in **Step 7**.

For details about creating VPC peering connections, see **Creating a VPC Peering Connection with Another VPC in Your Account**.

Step 4 Create a route table for a subnet.

Create a custom route table and associate it with the VPC 2 subnet to control the outbound traffic.

For details, see **Creating a Custom Route Table**.

Step 5 (Optional) Assign a virtual IP address and bind it to ECSs.

You can create two ECSs in VPC 2 and bind them to the same virtual IP address so that they can work in the active and standby mode. If the active ECS is faulty and cannot provide services, the virtual IP address will be dynamically switched to the standby ECS to continue providing services. Skip this step if the standby ECS is not required.

1. Assign a virtual IP address in the VPC 2 subnet.

For details, see Assigning a Virtual IP Address.

2. Bind the virtual IP address to ECS 2.

For details, see **Binding a Virtual IP Address to an EIP or ECS**.

Step 6 Create a Direct Connect connection.

Use a Direct Connect connection to enable communication between VPC 3 and the on-premises data center. For details, see **Create a Connection**.

Step 7 Configure routes.

You can configure routes to forward traffic to a next hop and finally to a destination.

1. Add the following route to the default route table of VPC 1:

Add a route to forward traffic from VPC 1 to the on-premises data center, set the destination of the route to the CIDR block of the on-premises data center, and the next hop of the route to the VPC peering connection between VPC 1 and VPC 2.

2. Add the following route to the default route table of VPC 2:

Set the destination of the route to 0.0.0/0, and the next hop of the route to ECS 2.

If there are two ECSs that use the same virtual IP address to work in the active and standby mode, the next hop should be the virtual IP address.

- 3. Add the following routes to the route table of VPC 2 subnet:
 - a. Add a route to forward traffic from VPC 2 to VPC 1, set the destination of the route to the CIDR block of VPC 1, and the next hop of the route to the VPC peering connection between VPC 1 and VPC 2.
 - b. Add a route to forward traffic from VPC 2 to the on-premises data center, set the destination of the route to the CIDR block of the on-premises data center, and the next hop of the route to the VPC peering connection between VPC 2 and VPC 3.
- 4. Add the following route to the default route table of VPC 3:

Set the destination of the route to 0.0.0.0/0, and the next hop of the route to the VPC peering connection between VPC 2 and VPC 3.

A Direct Connect connection has been created in **Step 6**. Thus, a route to the Direct Connect connection will be automatically delivered by the system.

----End

Verification

Log in to ECS 1 and then access your on-premises data center from ECS 1. Check whether ECS 2 can receive packets sent from ECS 1 to the data center. Check whether the packets pass through and are filtered by the firewall on ECS 2.

10 Building Highly Available Web Server Clusters with Keepalived

Scenarios

Virtual IP addresses are used for active and standby switchover of ECSs to achieve high availability. This way if one ECS goes down for some reason, the other one can take over and services continue uninterrupted.

This document uses CentOS 7.4 (64-bit) ECSs as an example to describe how to set up highly available web server clusters using Keepalived and Nginx.

Background

A web cluster consists of multiple web servers and a load balancer. Access requests will first be received by the load balancer, which then distributes the requests to backend web servers based on the load balancing policy.

In this document, Nginx is used to implement load balancing.

Network Topology

The data planning is as follows:

No.	Item	Quant ity	Specification
1	VPC	1	192.168.0.0/16
	Subnet	1	192.168.0.0/24
2	ECS	2	1 vCPU, 1 GB, CentOS 7.4 64bit
	IP address	2	ecs-HA1: 192.168.0.10 ecs-HA2: 192.168.0.20
3	EIP	1	122.xx.xx.189

Table 10-1 Data plannir	ng
--------------------------------	----

No.	Item	Quant ity	Specification
	Virtual IP address	1	192.168.0.100

Implementation methods:

- Configure the two ECSs in the same subnet to work in the active/standby mode using Keepalived.
- Bind a single virtual IP address to the two ECSs.
- Bind the virtual IP address to an EIP, then you can access the active and standby ECSs bound with the virtual IP address from the Internet.

Figure 10-1 Network topology



NOTE

- Select a region based on your service requirements.
- All cloud resources must be in the same region.

Procedure

The overall operation process is as follows:

Figure 10-2 Operation process



Step 1 Create a VPC and a subnet.

- 1. Log in to the management console.
- 2. Click Service List. Under Networking, click Virtual Private Cloud.
- 3. Click Create VPC.

Set required parameters as prompted based on Table 10-2.

Table	10-2	Parameter	configurations
-------	------	-----------	----------------

Parameter	Example Value
Name (of the VPC)	vpc-НА
CIDR Block (of the VPC)	192.168.0.0/16
Name (of the subnet)	subnet-HA
CIDR Block (of the subnet)	192.168.0.0/24

4. Click Create Now.

Step 2 Apply for required cloud resources.

- 1. Buy ECSs.
 - a. Log in to the management console.
 - b. Click Service List. Under Compute, click Elastic Cloud Server.
 - c. Click **Buy ECS**.
 - d. On the **Buy ECS** page, set parameters as prompted. For details, see **Table 10-1**.

e. Set the ECS name to ecs-HA1 and ecs-HA2.

D NOTE

In this example, no data disk is purchased. You can buy data disks based on service requirements and ensure their service data consistency.

f. (Optional) Configure security group rules to ensure that the two ECSs can communicate with each other.

In this example, the two ECSs are in the same security group and can communicate with each other through the internal network by default. In this case, you do not need to configure rules.

If two ECSs are in different security groups, you need to add inbound security group rules for the two ECSs. For details, see **Enabling ECSs in Different Security Groups to Communicate with Each Other Through an Internal Network**.

- 2. Buy an EIP.
 - a. Log in to the management console.
 - b. Click Service List. Under Networking, click Elastic IP.
 - c. Click **Buy EIP** and set parameters as prompted. For details, see **Table 10-1**.
- 3. Assign a virtual IP address.
 - a. Log in to the management console.
 - b. Click Service List. Under Networking, click Virtual Private Cloud.
 - c. In the navigation pane on the left, click **Subnets**.
 - d. In the subnet list, locate the target subnet and click its name.
 - e. On the **IP Addresses** tab page, click **Assign Virtual IP Address** and set parameters as prompted.

Step 3 Configure the ECSs.

- 1. Configure the ecs-HA1.
 - a. Bind EIP (122.xx.xx.189) to ecs-HA1.
 - i. Log in to the management console.
 - ii. Click Service List. Under Compute, click Elastic Cloud Server.
 - iii. In the ECS list, click the name of ecs-HA1.
 - iv. Click the EIPs tab and then Bind EIP.
 - v. On the **Bind EIP** page, select a NIC and an EIP, and click **OK**.
 - b. Connect to ecs-HA1 using SSH and run the following command to install the Nginx and Keepalived packages and related dependency packages:

yum install nginx keepalived -y

c. Run the following command to edit the **nginx** configuration file and save it:

vim /etc/nginx/nginx.conf

An example is provided as follows:

user root; worker_processes 1; #error_log logs/error.log;

```
#error_log logs/error.log notice;
#error_log logs/error.log info;
#pid logs/nginx.pid;
events {
    worker_connections 1024;
http {
   include mime.types;
   default_type application/octet-stream;
   #log_format main '$remote_addr - $remote_user [$time_local] "$request" '
   # '$status $body_bytes_sent "$http_referer"
   # '"$http_user_agent" "$http_x_forwarded_for"';
   #access_log logs/access.log main;
   sendfile on;
   #tcp_nopush on;
   #keepalive_timeout 0;
   keepalive_timeout 65;
   #gzip on;
   server {
        listen 80;
        server_name localhost;
        #charset koi8-r;
        #access_log logs/host.access.log main;
        location / {
                root html:
                index index.html index.html;
              }
               #error_page 404 /404.html;
               # redirect server error pages to the static page /50x.html
               error_page 500 502 503 504 /50x.html;
               location = /50x.html {
                              root html;
                              }
      }
 3
```

d. Run the following command to edit the **index.html** file and save the file:

vim /usr/share/nginx/html/index.html

An example is provided as follows: Welcome to ECS-HA1

e. Run the following commands to set the automatic startup of Nginx upon ECS startup:

systemctl enable nginx

systemctl start nginx.service

f. Verify the access to a single Nginx node.

Figure 10-3 ECS-HA1 access verification

Welcome to ECS-HA1

g. Run the following command to edit the **keepalived** configuration file and save it:

vim /etc/keepalived/keepalived.conf

An example is provided as follows:

```
! Configuration File for keepalived
global_defs {
router_id master-node
vrrp_script chk_http_port {
               script "/etc/keepalived/chk_nginx.sh"
               interval 2
               weight -5
               fall 2
               rise 1
              ł
vrrp_instance VI_1 {
           state MASTER
           interface eth0
           mcast_src_ip 192.168.0.10
          virtual_router_id 51
           priority 101
           advert_int 1
           authentication {
                    auth_type PASS
                    auth_pass 1111
                   }
           unicast_src_ip 192.168.0.10
           virtual_ipaddress {
                    192.168.0.100
                    }
track_script {
       chk_http_port
        }
}
```

h. Run the following command to edit the **nginx** monitoring script and save it:

vim /etc/keepalived/chk_nginx.sh

An example is provided as follows:

```
#!/bin/bash
counter=$(ps -C nginx --no-heading|wc -l)
if [ "${counter}" = "0" ]; then
systemctl start nginx.service
sleep 2
counter=$(ps -C nginx --no-heading|wc -l)
if [ "${counter}" = "0" ]; then
systemctl stop keepalived.service
fi
fi
```

chmod +x /etc/keepalived/chk_nginx.sh

i. Run the following commands to set the automatic startup of Keepalived upon ECS startup:

systemctl enable keepalived

systemctl start keepalived.service

- 2. Configure the ecs-HA2.
 - a. Unbind EIP (122.xx.xx.189) from ecs-HA1.
 - i. Log in to the management console.
 - ii. Click Service List. Under Compute, click Elastic Cloud Server.
 - iii. In the ECS list, click the name of ecs-HA1.
 - iv. Click the **EIPs** tab.
 - v. Locate the row that contains the EIP (122.xx.xx.189), and click **Unbind**.

- b. Bind EIP (122.xx.xx.189) to ecs-HA2.
 - i. Log in to the management console.
 - ii. Click Service List. Under Compute, click Elastic Cloud Server.
 - iii. In the ECS list, click the name of ecs-HA2.
 - iv. Click the **EIPs** tab.
 - v. Click Bind EIP.
 - vi. Select a NIC and an EIP and click OK.
- c. Connect to ecs-HA2 using SSH and run the following command to install the Nginx and Keepalived packages and related dependency packages:

```
yum install nginx keepalived -y
```

d. Run the following command to edit the **nginx.conf** configuration file:

vim /etc/nginx/nginx.conf

An example is provided as follows:

```
user root:
worker_processes 1;
#error_log logs/error.log;
#error_log logs/error.log notice;
#error_log logs/error.log info;
#pid logs/nginx.pid;
events {
    worker_connections 1024;
    }
http {
   include mime.types;
   default_type application/octet-stream;
   #log_format main '$remote_addr - $remote_user [$time_local] "$request" '
   # '$status $body_bytes_sent "$http_referer"
   # '"$http_user_agent" "$http_x_forwarded_for";
   #access_log logs/access.log main;
   sendfile on;
   #tcp_nopush on;
   #keepalive timeout 0;
   keepalive_timeout 65;
   #gzip on;
   server {
      listen 80;
      server name localhost;
      #charset koi8-r;
      #access_log logs/host.access.log main;
      location / {
             root html;
             index index.html index.htm;
      #error_page 404 /404.html;
      # redirect server error pages to the static page /50x.html
      error_page 500 502 503 504 /50x.html;
      location = /50x.html {
                    root html;
                    }
      }
```

e. Run the following command to edit the index.html file:

vim /usr/share/nginx/html/index.html

An example is provided as follows:

Welcome to ECS-HA2

f. Run the following commands to set the automatic startup of Nginx upon ECS startup:

systemctl enable nginx systemctl start nginx.service

g. Test the access to a single Nginx node.

Figure 10-4 ECS-HA2 verification result

← → C (i) 122.112.218.189

Welcome to ECS-HA2

h. Run the following command to edit the Keepalived configuration file: vim /etc/keepalived/keepalived.conf

An example is provided as follows:

```
! Configuration File for keepalived
global_defs {
router_id master-node
vrrp_script chk_http_port {
      script "/etc/keepalived/chk_nginx.sh"
       interval 2
       weight -5
       fall 2
      rise 1
     }
vrrp_instance VI_1 {
   state BACKUP
   interface eth0
   mcast_src_ip 192.168.0.20
   virtual_router_id 51
   priority 100
   advert_int 1
   authentication {
          auth_type PASS
          auth_pass 1111
          }
   unicast_src_ip 192.168.0.20
   virtual_ipaddress {
             192.168.0.100
              }
track script {
   chk_http_port
   }
}
```

i. Run the following command to edit the **nginx** monitoring script and add execute permissions:

vim /etc/keepalived/chk_nginx.sh

An example is provided as follows: #!/bin/bash counter=\$(ps -C nginx --no-heading|wc -l) if ["\${counter}" = "0"]; then

```
if [ "${counter}" = "0" ]; then
   systemctl start nginx.service
   sleep 2
   counter=$(ps -C nginx --no-heading|wc -l)
   if [ "${counter}" = "0" ]; then
      systemctl stop keepalived.service
   fi
   fi
```

chmod +x /etc/keepalived/chk_nginx.sh

j. Run the following commands to set the automatic startup of Keepalived upon ECS startup:

systemctl enable keepalived systemctl start keepalived

Step 4 Bind a virtual IP address to an ECS.

- 1. Unbind EIP (122.xx.xx.189) from ecs-HA2.
- 2. Bind the virtual IP address to ecs-HA1.
 - a. Log in to the management console.
 - b. Click Service List. Under Networking, click Virtual Private Cloud.
 - c. In the navigation pane on the left, click **Subnets**.
 - d. In the subnet list, locate the target subnet and click its name.
 - e. Click the **IP Addresses** tab, locate the row that contains the target virtual IP address, and click **Bind to Server** in the **Operation** column.
 - f. On the page that is displayed, select ecs HA1.
- 3. Bind the virtual IP address to ecs-HA2 by referring to **2**.
- 4. Bind the virtual IP address to the EIP 122.xx.xx.189.
 - a. Log in to the management console.
 - b. Click Service List. Under Networking, click Virtual Private Cloud.
 - c. In the navigation pane on the left, click **Subnets**.
 - d. In the subnet list, locate the target subnet and click its name.
 - e. Click the **IP Addresses** tab, locate the row that contains the target virtual IP address, and click **Bind to EIP** in the **Operation** column.
 - f. On the page that is displayed, select the EIP (122.xx.xx.189).
 - g. Click **OK**.

----End

Verification

- 1. Run the **reboot** command to restart ecs-HA1 and ecs-HA2.
- 2. Remotely log in to ecs-HA1 through the management console.
- 3. Run the following command to check whether the virtual IP address is bound to the eth0 NIC of ecs-HA1:

ip addr show

As shown in **Figure 10-5**, the virtual IP address has been bound to the eth0 NIC of ecs-HA1.

Figure 10-5 Virtual IP address of ecs-HA1

[root@ecs-ha1 ~]# ip addr show	
1: lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN qlen 1</loopback,up,lower_up>	
link/loopback 00:00:00:00:00 brd 00:00:00:00:00	
inet 127.0.0.1/8 scope host lo	
valid_lft forever preferred_lft forever	
inet6 ::1/128 scope host	
valid_lft forever preferred_lft forever	
2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc pfifo_fast state UP qlen 1000</broadcast,multicast,up,lower_up>	
link/ether fa:16:3e:a2:c5:72 brd ff:ff:ff:ff:ff	
inet 192.168.0.10/24 brd 192.168.0.255 scope global dynamic eth0	
valid_lft 86066sec preferred_lft 86066sec	
inet 192.168.0.100/32 scope global eth0	
valid_lft forever preferred_lft forever	
inet6 fe80::f816:3eff:fea2:c572/64 scope link	
valid_lft forever preferred_lft forever	
	4

4. Use a browser to access the EIP and check whether the web page on ecs-HA1 can be accessed.

If the information shown in Figure 10-6 is displayed, the access is normal.

Figure 10-6 ecs-HA1 access verification



Welcome to ECS-HA1

- 5. Run the following command to disable Keepalived on ecs-HA1: systemctl stop keepalived.service
- 6. Run the following command to check whether ecs-HA2 has taken over the virtual IP address:

ip addr show

Figure 10-7 Virtual IP address of ecs-HA2

[root@ecs-haZ]]# ip addr show
1: lo: <loopback,up,lower_up> mtu 65536 qdisc noqueue state UNKNOWN glen 1</loopback,up,lower_up>
link/loopback 08:00:00:00:00:00 brd 08:00:00:00:00
inet 127.0.0.1/8 scope host to
valid_lft forever preferred_lft forever
inet6 ::1/128 scope host
valid_lft forever preferred_lft forever
2: eth0: <broadcast,multicast,up,lower_up> mtu 1500 qdisc pfifo_fast state UP qlen 1000</broadcast,multicast,up,lower_up>
link/ether fa:16:3e:79:03:21 brd ff:ff:ff:ff:ff:ff
inet 192.168.0.20/24 brd 192.168.0.255 scope global dynamic eth0
valid_lft 04950sec preferred_lft 04950sec
inet 192.160.0.100/32 scope global eth0
valid_lft forever preferred_lft forever
inet6 fe80::f816:3eff:fe79:321/64 scope link
valid_lft forever preferred_lft forever
[maatBaca_bo2 #]#

7. Use a browser to access the EIP and check whether the web page on ecs-HA2 can be accessed.

If the information shown in Figure 10-8 is displayed, the access is normal.

Figure 10-8 ecs-HA2 access verification

← → C ① 122.1 2 189

Welcome to ECS-HA2

11 Deploying Containers that Can Communicate with Each Other on Huawei Cloud ECSs

Scenarios

You can deploy containers that are not provided by Huawei Cloud container services on Huawei Cloud ECSs and enable the containers on different ECSs but in the same subnet to communicate with each other.

Solution Advantages

- Containers deployed on ECSs can use CIDR blocks that are not from those of the ECS VPCs, but use routes added to VPC route tables for data forwarding.
- You only need to add routes to the route tables to allow communications among containers, which is flexible and convenient.

Typical Topology

The network topology requirements are as follows:

- ECSs are in the same subnet. As shown in the following figure, the VPC subnet is 192.168.0.0/24, and the IP addresses of the ECS 1 and ECS 2 are 192.168.0.2 and 192.168.0.3, respectively.
- Containers are on CIDR blocks that are not from those of the VPC subnets that the ECSs belong to. Containers on the same ECS are on the same CIDR block, but containers on different ECSs are on different CIDR blocks. As shown in the following figure, the CIDR block of containers on ECS 1 is 10.0.2.0/24, and that on ECS 2 is 10.0.3.0/24.
- The next hop of the data packets sent to a container is the ECS where the container is deployed. As shown in the following figure, the next hop of the packets sent to CIDR block 10.0.2.0/24 is 192.168.0.2, and that of the packets sent to CIDR block 10.0.3.0/24 is 192.168.0.3.





Procedure

Step 1 Create VPCs.

For details, see Creating a VPC.

Step 2 Create ECSs.

For details, see **Purchasing an ECS**.

After the ECS is created, disable source/destination check on the ECS NIC, as shown in **Figure 11-2**.

Figure 11-2 Disabling source/destination check



Step 3 Deploy containers on ECSs.

You can use Docker CE to deploy containers. For details, see the documentation of Docker CE.

NOTE

Containers on the same ECS must be on the same CIDR block and the CIDR blocks of containers on different ECSs cannot overlap.

Step 4 Add routes to the VPC route table.

Set the next hop of the packets sent to CIDR block 10.0.2.0/24 to 192.168.0.2, and set the next hop of the packets sent to CIDR block 10.0.3.0/24 to 192.168.0.3.

NOTE

- By default, a VPC supports containers from a maximum of 50 different CIDR blocks. If containers from more different CIDR blocks need to be deployed in a VPC, apply for more route tables for the VPC.
- After a container is migrated to another ECS, you need to add routes to the route table of the ECS VPC.

Step 5 Add security group rules.

To use ping and traceroute commands to check the communications between containers, add the rules shown in **Table 11-1** to the security group of the ECSs to allow ICMP and UDP traffic.

For details, see Adding a Security Group Rule.

Table 11-1	Security	group	rules
------------	----------	-------	-------

Direction	Protocol	Port	Source
Inbound	ICMP	All	0.0.0/0
Inbound	UDP	All	0.0.0/0

----End

Verification

Use the ping command to check whether the containers deployed on two different ECSs can communicate with each other.

Run the following commands to create a network connection **my-net** on ECS 1, set the CIDR block to be used by a container on ECS 1 to 10.0.2.0/24, and create the container that uses **my-net**.

\$ docker network create --subnet 10.0.2.0/24 my-net \$ docker run -d --name nginx --net my-net -p 8080:80 nginx:alpine

Run the following commands to create a network connection and container on ECS 2, and set the CIDR block to be used by the container to 10.0.3.0/24.

```
$ docker network create --subnet 10.0.3.0/24 my-net
$ docker run -d --name nginx --net my-net -p 8080:80 nginx:alpine
```

Run the following command to set the default policy of the FORWARD chain in the filter table of iptables on the ECS to ACCEPT.

NOTE

This operation is required because Docker sets the default policy of the FORWARD chain in the filter table of iptables to DROP for security purposes.

\$ iptables -P FORWARD ACCEPT

Ping and traceroute 10.0.3.2 from 10.0.2.2. The ping and traceroute operations are successful, and the packet is tracerouted in the following sequence: 10.0.2.2 ->

10.0.2.1 -> 192.168.0.3 -> 10.0.3.2, which is consistent with the configured route forwarding rules.

[root@ecs1 ~]# docker exec -it nginx /bin/sh / # traceroute -d 10.0.3.2 traceroute to 10.0.3.2 (10.0.3.2), 30 hops max, 46 byte packets 1 10.0.2.1 (10.0.2.1) 0.007 ms 0.004 ms 0.007 ms 2 192.168.0.3 (192.168.0.3) 0.232 ms 0.165 ms 0.248 ms 3 10.0.3.2 (10.0.3.2) 0.366 ms 0.308 ms 0.158 ms / # ping 10.0.3.2 PING 10.0.3.2 (10.0.3.2): 56 data bytes 64 bytes from 10.0.3.2: seq=0 ttl=62 time=0.570 ms 64 bytes from 10.0.3.2: seq=1 ttl=62 time=0.343 ms 64 bytes from 10.0.3.2: seq=3 ttl=62 time=0.319 ms

12 Configuring Policy-based Routes for an ECS with Multiple NICs

12.1 Overview

Background

If an ECS has multiple NICs, the primary NIC can communicate with external networks by default, but the extension NICs cannot. To enable extension NICs to communicate with external works either, you need to configure policy-based routes for these NICs.

Scenarios

This example describes how to configure policy-based routes for an ECS with two NICs. Figure 12-1 shows the networking. The details are as follows:

- The primary and extension NICs on the source ECS are in different subnets of the same VPC.
- The source and destination ECSs are in different subnets of the same VPC and the two ECSs can communicate with each other through primary NICs without configuring policy-based routes.
- After policy-based routes are configured for the two NICs of the source ECS, both the primary and extension NICs can communicate with the destination ECS.

NOTICE

You can select a destination IP address based on service requirements. Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.





Operation Guide

You can follow the following operations to configure policy-based routes for Linux and Windows ECSs. For details, see **Table 12-1**.

Table 12-1	Operation	instructions
------------	-----------	--------------

OS	IP Address Version	Description
Linux	IPv4	Take an ECS running CentOS 8.0 (64-bit) as an example.
IPv6		Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6)
Windows	IPv4	Take an ECS running Windows Server 2012
	IPv6	Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/ IPv6)

12.2 Collecting ECS Network Information

Scenarios

Before configuring policy-based routes for a multi-NIC ECS, you need to collect network information about the ECS.

• Table 12-2 lists the information to be collected for a Linux ECS using IPv4.

Table 12-2 Linux ECS using IPV4	Table	12-2	Linux	ECS	using	IPv4
---------------------------------	-------	------	-------	-----	-------	------

EC S	Primary NIC	Extension NIC	How to Obtain
So ur ce	 NIC address: 10.0.0.115 Subnet: 10.0.0.0/24 Subnet gateway: 10.0.0.1 	 NIC address: 10.0.1.183 Subnet: 10.0.1.0/24 Subnet gateway: 10.0.1.1 	 Obtaining ECS NIC Addresses Obtaining Subnet CIDR
De sti na tio n	NIC address: 10.0.2.12	N/A	Blocks and Gateway Addresses

• Table 12-3 lists the information to be collected for a Linux ECS using IPv6.

EC S	Primary NIC	Extension NIC	How to Obtain
So urc e	 IPv4 address: 10.0.0.102 IPv6 address: 2407:c080:1200:1dd8: 859c:e5d5:8b3d:a2d9 IPv6 subnet: 2407:c080:1200:1dd8:: /64 IPv6 subnet gateway: 2407:c080:1200:1dd8:: 1 	 IPv4 address: 10.0.1.191 IPv6 address: 2407:c080:1200:1a9c:7 cc0:63b5:8e65:4dd8 IPv6 subnet: 2407:c080:1200:1a9c::/ 64 IPv6 subnet gateway: 2407:c080:1200:1a9c:: 1 	 Obtaining ECS NIC Addresses Obtaining Subnet CIDR Blocks and Gateway Addresses
De sti na tio n	 IPv4 address: 10.0.2.3 IPv6 address: 2407:c080:1200:1dd9: 16a7:fe7a:8f71:7044 	N/A	

 Table 12-3
 Linux ECS using IPv6

• Table 12-4 lists the information to be collected for a Windows ECS using IPv4.

Table 12-4 Windows	ECS	using	IPv4
--------------------	-----	-------	------

EC S	Primary NIC	Extension NIC	How to Obtain
So ur ce	 NIC address: 10.0.0.59 Subnet gateway: 10.0.0.1 	 NIC address: 10.0.1.104 Subnet gateway: 10.0.1.1 	 Obtaining ECS NIC Addresses Obtaining
De sti na tio n	NIC address: 10.0.2.12	N/A	Subnet CIDR Blocks and Gateway Addresses

• Table 12-5 lists the information to be collected for a Windows ECS using IPv6.

Table 12-5	Windows	ECS	using	IPv6
------------	---------	-----	-------	------

EC S	Primary NIC	Extension NIC	How to Obtain
So ur ce	NIC address: 2407:c080:802:aba:6788:f b94:d71f:8deb	NIC address: 2407:c080:802:be6:71c8: 42e0:d44e:eeb4	Obtaining ECS NIC Addresses
De sti na tio n	NIC address: 2407:c080:802:be7:c2e6:d 99c:b685:c6c8	N/A	

Obtaining ECS NIC Addresses

- 1. Log in to the management console.
- 2. Click \bigcirc in the upper left corner and select the desired region and project.
- 3. Click Service List and choose Compute > Elastic Cloud Server.
- 4. In the ECS list, click the target ECS name.

The **Summary** tab page of the ECS is displayed.

5. In the **NICs** area, view the IP addresses of the primary and extension NICs. You can view the IPv4 and IPv6 addresses of NICs.

Figure 12-2 IPv4 and IPv6 addresses of NICs

< ecs	ource	🙂 Feedback Rentode Login Start Stop Restart More 💌 C
Summary Disks	Network Interfaces Security Groups EIPs Monitoring Tags Cloud Bar	ckup and Recovery Host Security
ECS Information ID Name	1721206-6704-664956-336693749554 463	Running Montering Heat Security © Double Enable
Description Region AZ Specifications	– 🖉 AZ3 General computing 1 vCPU 2 GB s5 međuri 2	Dicks System Diok eco instruct-volume-0000 General Purpose SSD 40 OB Expand Capacity
Image VPC Global EIP	CentOS 8.0 6401 (Public image vpcPho — Bind	VICs Primary NC 10.0.0.102 2407-08011200140818564:45691803412499 (IP46)
Billing Mode Obtained Launched Defellion Time	Pay-ser-ure Ner 14, 2023 1852:17 GMT-00:00 Ner 14, 2023 1852:45 GMT-00:00 = Monte	Celevisor NC watered 10 0 1 191 2407-c001 1201 191-7c08 2305 8e95 4d98 (IPVC) Security Groups
Constant (The	- moury	

Obtaining Subnet CIDR Blocks and Gateway Addresses

- 1. Log in to the management console.
- 2. Click 💿 in the upper left corner and select the desired region and project.
- 3. Click Service List and choose Compute > Elastic Cloud Server.
- In the ECS list, click the target ECS name.
 The **Summary** tab page of the ECS is displayed.
- 5. In the **ECS Information** area, click the VPC hyperlink. The **Virtual Private Cloud** page is displayed.
- Locate the target VPC and click the number in the Subnets column. The Subnets page is displayed.
- In the subnet list, view the CIDR blocks of the subnets.
 You can view the IPv4 and IPv6 CIDR blocks of subnets.

Figure 12-3 IPv4 and IPv6 CIDR blocks of subnets

Subnets 🔿										de Overview	Create Subnet
(Export ~										
	Q. Select a property or enter a keyword.										0
	□ Name/ID ⇔	VPC 🖯	IPv4 CIDR \varTheta	IPv6 CIDR Block \varTheta		Status 🖯	AZ ⊖	Network A 😌	Route Table \varTheta	Owner Project ID \varTheta	Operation
	c102f0 93-de990e4198c3	vpc	192.168.0.0/20	Enable IPv6		Available	-	Associate	rtb-vpc-default-s Default	148a37dea0315d99	Change Rc
	subnet 6-a297424a13bd	vpc	172.17.16.0/20	Enable IPv6		Available	AZ1	Associate	rtb-vpc-B Default	48a37dea0315d99	Change Rc
	subnet 0-a73ea84774b0	VPC.	192.168.64.0/19	2407:c080:802:1670::/64		Available	AZ1	Associate	rtb-vpc-cce Default	148a37dea0315d99	Change Rc
	subnet 350%2 12-bbc82d9e1168	vpc	192.168.0.0/24	2407:c080.802:16a6::/64		Available	AZ1	Associate	rtb-vpc- Default	48a37dea0315d99	Change Rc

- 8. In the subnet list, click the subnet name. The **Summary** page is displayed.
- 9. Click the **IP Addresses** tab and view the gateway addresses of the subnet. You can view the IPv4 and IPv6 addresses of a gateway.

Figure 12-4 IPv4 and IPv6 addresses of a gateway

< subnet.												
ummany IPAddresses Sharing Tags												
Assign Virtual IP Address Unbind EIP Expo Select a property or enter a keyword.	rt v Learn how to bind a virtual IP address 🕻											
Virtual IP Address 🔶	Bound EIP 🕀	Bound Global EIP 🔶	Bound Instance \ominus	Operation								
Total Records: 0 10 🗸 (1)		↓ 1 /										
Export ~												
Q. Select a property or enter a keyword.												
□ IP Address ⇔	Resource ID \ominus	Used By 🕀		Operation								
IPv4: 192.168.0.1 IPv6: 2407:c080:802:1686::1	d6777907-fc59-4d4f-b463-c37b17d0a4e3	Gateway		Release								
Pv4: 192.168.0.253 IPv6: 2407:c080.802:16a6:##################################	d6777907-fc59-4d4f-b463-c37b17d0a4e3	System interface		Release								
IPv4: 192.168.0.254	-	DHCP		Release								

12.3 Configuring Policy-based Routes for a Linux ECS with Multiple NICs (IPv4/IPv6)

Scenarios

This section describes how to configure policy-based routes for a dual-NIC ECS running CentOS 8.0 (64-bit).

- IPv4: Procedure (Linux ECS Using IPv4)
- IPv6: Procedure (Linux ECS Using IPv6)

For details about the background knowledge and networking of dual-NIC ECSs, see **Overview**.

Procedure (Linux ECS Using IPv4)

1. Collect the ECS network information required for configuring policy-based routes.

For details, see Collecting ECS Network Information.

2. Log in to an ECS.

Multiple methods are available for logging in to an ECS. For details, see **Logging In to an ECS**.

3. Check whether the source ECS can use its primary NIC to communicate with the destination ECS:

ping -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

In this example, run the following command:

ping -I 10.0.0.115 10.0.2.12

If information similar to the following is displayed, the source ECS can use its primary NIC to communicate with the destination ECS.

```
[root@ecs-resource ~]# ping -I 10.0.0.115 10.0.2.12
PING 10.0.2.12 (10.0.2.12) from 10.0.0.115 : 56(84) bytes of data.
64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=0.775 ms
64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.268 ms
64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.220 ms
64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.167 ms
^C
--- 10.0.2.12 ping statistics ---
```

NOTE

Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

4. Query the NIC names of the ECS:

ifconfig

Search for the NIC name based on the NIC address.

- 10.0.0.115 is the IP address of the primary NIC, and the NIC name is eth0.
- 10.0.1.183 is the IP address of the extension NIC, and the NIC name is eth1.

[root@ecs-resource ~]# ifconfig
eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500
inet 10.0.0.115 netmask 255.255.0 broadcast 10.0.0.255
inet6 fe80::f816:3eff:fe92:6e0e prefixlen 64 scopeid 0x20<link>
ether fa:16:3e:92:6e:0e txqueuelen 1000 (Ethernet)
RX packets 432288 bytes 135762012 (129.4 MiB)
RX errors 0 dropped 0 overruns 0 frame 1655
TX packets 423744 bytes 106716932 (101.7 MiB)
TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500

inet 10.0.1.183 netmask 255.255.255.0 broadcast 10.0.1.255 inet6 fe80::f816:3eff:febf:5818 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:bf:58:18 txqueuelen 1000 (Ethernet) RX packets 9028 bytes 536972 (524.3 KiB) RX errors 0 dropped 0 overruns 0 frame 1915 TX packets 6290 bytes 272473 (266.0 KiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

5. Configure temporary routes for the ECS.

NOTICE

Temporary routes take effect immediately after being configured and will be lost after the ECS is restarted. To prevent network interruptions after the ECS is restarted, perform **6** after this step to configure persistent routes.

- a. Configure policy-based routes for both the primary and extension NICs:
 - Primary NIC **ip route add default via** *Subnet gateway* **dev** *NIC name* **table** *Route table name*

ip route add *Subnet CIDR block* **dev** *NIC name* **table** *Route table name*

ip rule add from NIC address table Route table name

 Extension NIC
 ip route add default via Subnet gateway dev NIC name table Route table name **ip route add** *Subnet CIDR block* **dev** *NIC name* **table** *Route table name*

ip rule add from *NIC address* **table** *Route table name* Configure the parameters as follows:

- NIC name: Enter the name obtained in 4.
- Route table name: Customize a route table name using a number.
- Other network information: Enter the IP addresses collected in 1.

In this example, run the following commands:

Primary NIC

ip route add default via 10.0.0.1 dev eth0 table 10 ip route add 10.0.0.0/24 dev eth0 table 10 ip rule add from 10.0.0.115 table 10

Extension NIC

ip route add default via 10.0.1.1 dev eth1 table 20 ip route add 10.0.1.0/24 dev eth1 table 20 ip rule add from 10.0.1.183 table 20

If the ECS has multiple NICs, configure policy-based routes for all NICs one by one.

b. Check whether the policy-based routes are successfully added.

ip rule

ip route show table *Route table name of the primary NIC* **ip route show table** *Route table name of the extension NIC*

The route table name is customized in **5.a**.

In this example, run the following commands:

- ip rule
- ip route show table 10

ip route show table 20

If information similar to the following is displayed, the policy-based routes have been added.

- [root@ecs-resource ~]# ip rule 0: from all lookup local 32764: from 10.0.1.183 lookup 20 32765: from 10.0.0.115 lookup 10 32766: from all lookup main 32767: from all lookup default [root@ecs-resource ~]# ip route show table 10 default via 10.0.0.1 dev eth0 10.0.0.0/24 dev eth0 scope link [root@ecs-resource ~]# ip route show table 20 default via 10.0.1.1 dev eth1 10.0.1.0/24 dev eth1 scope link
- c. Check whether the source ECS and the destination ECS can communicate with each other.

ping -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

ping -I *IP address of the extension NIC on the source ECS IP address of the destination ECS*

In this example, run the following commands:

ping -I 10.0.0.115 10.0.2.12

ping -I 10.0.1.183 10.0.2.12

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS.

[root@ecs-resource ~]# ping -I 10.0.0.115 10.0.2.12 PING 10.0.2.12 (10.0.2.12) from 10.0.0.115 : 56(84) bytes of data. 64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=0.775 ms 64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.268 ms 64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.220 ms 64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.167 ms ^C --- 10.0.2.12 ping statistics ---4 packets transmitted, 4 received, 0% packet loss, time 102ms rtt min/avg/max/mdev = 0.167/0.357/0.775/0.244 ms [root@ecs-resource ~]# ping -I 10.0.1.183 10.0.2.12 PING 10.0.2.12 (10.0.2.12) from 10.0.1.183 : 56(84) bytes of data. 64 bytes from 10.0.2.12: icmp_seq=1 ttl=64 time=2.84 ms 64 bytes from 10.0.2.12: icmp_seq=2 ttl=64 time=0.258 ms 64 bytes from 10.0.2.12: icmp_seq=3 ttl=64 time=0.234 ms 64 bytes from 10.0.2.12: icmp_seq=4 ttl=64 time=0.153 ms ^C --- 10.0.2.12 ping statistics ---4 packets transmitted, 4 received, 0% packet loss, time 92ms rtt min/avg/max/mdev = 0.153/0.871/2.840/1.137 ms

- 6. Configure persistent routes for the ECS.
 - a. Run the following command to open the /etc/rc.local file:

vi /etc/rc.local

- b. Press **i** to enter the editing mode.
- Add the following content to the end of the file: c. # wait for nics up sleep 5 # Add v4 routes for eth0 ip route flush table 10 ip route add default via 10.0.0.1 dev eth0 table 10 ip route add 10.0.0.0/24 dev eth0 table 10 ip rule add from 10.0.0.115 table 10 # Add v4 routes for eth1 ip route flush table 20 ip route add default via 10.0.1.1 dev eth1 table 20 ip route add 10.0.1.0/24 dev eth1 table 20 ip rule add from 10.0.1.183 table 20 # Add v4 routes for cloud-init ip rule add to 169.254.169.254 table main

Parameters are described as follows:

- wait for nics up: file startup time. Set the value to be the same as that in the preceding configurations.
- Add v4 routes for eth0: policy-based routes of the primary NIC. Set the value to be the same as that configured in 5.a.
- Add v4 routes for eth1: policy-based routes of the extension NIC. Set the value to be the same as that configured in 5.a.

- Add v4 routes for cloud-init: Configure the Cloud-Init address. Set the value to be the same as that in the preceding configurations.
- d. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Run the following command to assign execute permissions to the **/etc/ rc.local** file:

chmod +x /etc/rc.local

D NOTE

If your operating system is Red Hat or EulerOS, run the following command after you perform **6.e**:

chmod +x /etc/rc.d/rc.local

f. Run the following command to restart the ECS:

reboot

NOTICE

Policy-based routes added to the **/etc/rc.local** file take effect only after the ECS is restarted. Ensure that workloads on the ECS will not be affected before restarting the ECS.

g. Repeat **5.b** to **5.c** to check whether the policy-based routes are added and whether the source ECS and the destination ECS can communicate with each other.

Procedure (Linux ECS Using IPv6)

1. Collect the ECS network information required for configuring policy-based routes.

For details, see Collecting ECS Network Information.

2. Log in to an ECS.

Multiple methods are available for logging in to an ECS. For details, see **Logging In to an ECS**.

3. Check whether an ECS has IPv6 enabled and can obtain IPv6 addresses.

NOTICE

Perform this step for both the source and destination ECSs to ensure that the ECSs have obtained IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

ECSs in this example run CentOS 8.0 (64-bit). For details about how to obtain IPv6 addresses for ECSs running other OSs, see **Dynamically Assigning IPv6 Addresses**.

a. Run the following command to check whether the ECS has IPv6 addresses:

ip addr
In the following command output, eth0 and eth1 are the NICs of the ECS. Each NIC has one **inet6** followed by an IP address starting with **fe80**. This indicates that the ECS has IPv6 enabled but has not obtained IPv6 addresses. In this case, perform **3.b** to **3.g** obtain IPv6 addresses. [root@ecs-resource ~]# ip addr

2: **eth0**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

link/ether fa:16:3e:22:22:88 brd ff:ff:ff:ff:ff
inet 10.0.0.102/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0
valid_lft 107943256sec preferred_lft 107943256sec
inet6 fe80::f816:3eff:fe22:2288/64 scope link
valid_lft forever preferred_lft forever

3: **eth1**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000 link/ether fa:16:3e:22:23:e1 brd ff:ff:ff:ff:ff:ff

inet 10.0.1.191/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1 valid_lft 107943256sec preferred_lft 107943256sec inet6 fe80::f816:3eff:fe22:23e1/64 scope link

valid_lft forever preferred_lft forever

b. Query the NIC names of the ECS:

ifconfig

Search for the NIC name based on the NIC address.

- 10.0.0.102 is the IP address of the primary NIC, and the NIC name is eth0.
- 10.0.1.191 is the IP address of the extension NIC, and the NIC name is eth1.

[root@ecs-resource ~]# ifconfig

eth0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.0.0.102 netmask 255.255.255.0 broadcast 10.0.0.255 inet6 fe80::f816:3eff:fe22:2288 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:22:22:88 txqueuelen 1000 (Ethernet) RX packets 135116 bytes 132321802 (126.1 MiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 60963 bytes 23201005 (22.1 MiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0

- eth1: flags=4163<UP,BROADCAST,RUNNING,MULTICAST> mtu 1500 inet 10.0.1.191 netmask 255.255.255.0 broadcast 10.0.1.255 inet6 fe80::f816:3eff:fe22:23e1 prefixlen 64 scopeid 0x20<link> ether fa:16:3e:22:23:e1 txqueuelen 1000 (Ethernet) RX packets 885 bytes 97676 (95.3 KiB) RX errors 0 dropped 0 overruns 0 frame 0 TX packets 47 bytes 4478 (4.3 KiB) TX errors 0 dropped 0 overruns 0 carrier 0 collisions 0
- c. Configure the **ifcfg** file of the primary NIC.
 - i. Run the following command to open the **ifcfg** file of the primary NIC:

vi /etc/sysconfig/network-scripts/ifcfg-Primary NIC name

The name of the primary NIC is obtained in **3.b**.

In this example, run the following command:

vi /etc/sysconfig/network-scripts/ifcfg-eth0

- ii. Press i to enter the editing mode.
- Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"

- iv. Press **ESC** to exit and enter **:wq!** to save the configuration.
- d. Configure the **ifcfg** file of the extension NIC.
 - i. Run the following command to open the **ifcfg** file of the extension NIC:

vi /etc/sysconfig/network-scripts/ifcfg-Extension NIC name

The name of the extension NIC is obtained in **3.b**.

In this example, run the following command:

vi /etc/sysconfig/network-scripts/ifcfg-eth1

- ii. Press i to enter the editing mode.
- iii. Add the following content to the end of the file: IPV6INIT="yes" DHCPV6C="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- e. Edit the /etc/sysconfig/network file.
 - i. Run the following command to open the **/etc/sysconfig/network** file:

vi /etc/sysconfig/network

- ii. Press i to enter the editing mode.
- iii. Add the following content to the end of the file: NETWORKING_IPV6="yes"
- iv. Press **ESC** to exit and enter :wq! to save the configuration.
- f. Run the following command to restart the network service for the configuration to take effect:

systemctl restart NetworkManager

g. Run the following command to check whether the ECS has IPv6 addresses:

ip addr

In the following command output, each NIC has one more **inet6** followed by an IP address starting with **2407** in addition to the one followed by an IP address starting with **fe80**. In this case, the ECS has obtained IPv6 addresses.

[root@ecs-resource ~]# ip addr

2: **eth0**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

link/ether fa:16:3e:22:22:88 brd ff:ff:ff:ff:ff:ff

inet 10.0.0.102/24 brd 10.0.0.255 scope global dynamic noprefixroute eth0

valid_lft 107999994sec preferred_lft 107999994sec inet6 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9/128 scope global dynamic noprefixroute

valid_lft 7195sec preferred_lft 7195sec

inet6 fe80::f816:3eff:fe22:2288/64 scope link noprefixroute

valid_lft forever preferred_lft forever

3: **eth1**: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000

link/ether fa:16:3e:22:23:e1 brd ff:ff:ff:ff:ff:ff

inet 10.0.1.191/24 brd 10.0.1.255 scope global dynamic noprefixroute eth1

valid_lft 107999994sec preferred_lft 107999994sec

inet6 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8/128 scope global dynamic noprefixroute valid_lft 7198sec preferred_lft 7198sec

inet6 fe80::f816:3eff:fe22:23e1/64 scope link noprefixroute

valid_lft forever preferred_lft forever

4. Check whether the source ECS can use its primary NIC to communicate with the destination ECS:

ping6 -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

In this example, run the following command:

ping6 -I 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

If information similar to the following is displayed, the source ECS can use its primary NIC to communicate with the destination ECS. [root@ecs-resource ~]# ping6 -I 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 : 56 data bytes 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.635 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.320 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.287 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=4 ttl=64 time=0.193 ms ~C --- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics ---4 packets transmitted, 4 received, 0% packet loss, time 3074ms

rtt min/avg/max/mdev = 0.193/0.358/0.635/0.167 ms

NOTE

Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

5. Configure temporary routes for the ECS.

NOTICE

Temporary routes take effect immediately after being configured and will be lost after the ECS is restarted. To prevent network interruptions after the ECS is restarted, perform **6** after this step to configure persistent routes.

- a. Configure policy-based routes for both the primary and extension NICs:
 - Primary NIC

ip -6 route add default via *Subnet gateway* **dev** *NIC name* **table** *Route table name*

ip -6 route add *Subnet CIDR block* **dev** *NIC name* **table** *Route table name*

ip -6 rule add from NIC address table Route table name

Extension NIC

ip -6 route add default via *Subnet gateway* **dev** *NIC name* **table** *Route table name*

ip -6 route add *Subnet CIDR block* **dev** *NIC name* **table** *Route table name*

ip -6 rule add from *NIC address* **table** *Route table name*

Configure the parameters as follows:

NIC name: Enter the name obtained in 3.b.

- Route table name: Customize a route table name using a number.
- Other network information: Enter the IP addresses collected in 1.

In this example, run the following commands:

Primary NIC

ip -6 route add default via 2407:c080:1200:1dd8::1 dev eth0 table 10

ip -6 route add 2407:c080:1200:1dd8::/64 dev eth0 table 10 ip -6 rule add from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 table 10

Extension NIC

ip -6 route add default via 2407:c080:1200:1a9c::1 dev eth1 table 20

ip -6 route add 2407:c080:1200:1a9c::/64 dev eth1 table 20 ip -6 rule add from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 table 20

D NOTE

If the ECS has multiple NICs, configure policy-based routes for all NICs one by one.

b. Check whether the policy-based routes are successfully added.

ip -6 rule

ip -6 route show table Route table name of the primary NIC

ip -6 route show table *Route table name of the extension NIC*

The route table name is customized in **5.a**.

In this example, run the following commands:

- ip -6 rule
- ip -6 route show table 10
- ip -6 route show table 20

If information similar to the following is displayed, the policy-based routes have been added.

[root@ecs-resource ~]# ip -6 rule
0: from all lookup local
32764: from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 lookup 20
32765: from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 lookup 10
32766: from all lookup main
[root@ecs-resource ~]# ip -6 route show table 10
2407:c080:1200:1dd8::/64 dev eth0 metric 1024 pref medium
default via 2407:c080:1200:1dd8::1 dev eth0 metric 1024 pref medium
[root@ecs-resource ~]# ip -6 route show table 20
2407:c080:1200:1a9c::/64 dev eth1 metric 1024 pref medium
default via 2407:c080:1200:1a9c::/1 dev eth1 metric 1024 pref medium

c. Check whether the source ECS and the destination ECS can communicate with each other.

ping -6 -I *IP address of the primary NIC on the source ECS IP address of the destination ECS*

ping -6 -I *IP address of the extension NIC on the source ECS IP address of the destination ECS*

In this example, run the following commands:

ping -6 -l 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

ping -6 -l 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS. [root@ecs-resource ~]# ping -6 -I 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 : 56 data bytes 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.770 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.295 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.245 ms ^C --- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2080ms rtt min/avg/max/mdev = 0.245/0.436/0.770/0.237 ms [root@ecs-resource ~]# ping -6 -I 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 PING 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044(2407:c080:1200:1dd9:16a7:fe7a:8f71:7044) from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 : 56 data bytes 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=1 ttl=64 time=0.922 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=2 ttl=64 time=0.307 ms 64 bytes from 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044: icmp_seq=3 ttl=64 time=0.174 ms ^C

--- 2407:c080:1200:1dd9:16a7:fe7a:8f71:7044 ping statistics ---3 packets transmitted, 3 received, 0% packet loss, time 2059ms rtt min/avg/max/mdev = 0.174/0.467/0.922/0.326 ms

- 6. Configure persistent routes for the ECS.
 - a. Run the following command to open the /etc/rc.local file:

vi /etc/rc.local

- b. Press **i** to enter the editing mode.
- c. Add the following content to the end of the file:
 - # wait for nics up
 - sleep 5 # Add v6 routes for eth0
 - ip -6 route flush table 10
 - ip -6 route add default via 2407:c080:1200:1dd8::1 dev eth0 table 10
 - ip -6 route add 2407:c080:1200:1dd8::/64 dev eth0 table 10
 - ip -6 rule add from 2407:c080:1200:1dd8:859c:e5d5:8b3d:a2d9 table 10
 - # Add v6 routes for eth1
 - ip -6 route flush table 20
 - ip -6 route add default via 2407:c080:1200:1a9c::1 dev eth1 table 20
 - ip -6 route add 2407:c080:1200:1a9c::/64 dev eth1 table 20
 - ip -6 rule add from 2407:c080:1200:1a9c:7cc0:63b5:8e65:4dd8 table 20

Parameters are described as follows:

- wait for nics up: file startup time. Set the value to be the same as that in the preceding configurations.
- Add v6 routes for eth0: policy-based routes of the primary NIC. Set the value to be the same as that configured in 5.a.
- Add v6 routes for eth1: policy-based routes of the extension NIC. Set the value to be the same as that configured in 5.a.
- d. Press ESC to exit and enter :wq! to save the configuration.
- e. Run the following command to assign execute permissions to the **/etc/ rc.local** file:

chmod +x /etc/rc.local

D NOTE

If your operating system is Red Hat or EulerOS, run the following command after you perform **6.e**:

chmod +x /etc/rc.d/rc.local

f. Run the following command to restart the ECS:

reboot

NOTICE

Policy-based routes added to the **/etc/rc.local** file take effect only after the ECS is restarted. Ensure that workloads on the ECS will not be affected before restarting the ECS.

g. Repeat **5.b** to **5.c** to check whether the policy-based routes are added and whether the source ECS and the destination ECS can communicate with each other.

12.4 Configuring Policy-based Routes for a Windows ECS with Multiple NICs (IPv4/IPv6)

Scenarios

This section describes how to configure policy-based routes for a dual-NIC ECS running Windows Server 2012 (64-bit).

- IPv4: Procedure (Windows ECS Using IPv4)
- IPv6: Procedure (Windows ECS Using IPv6)

For details about the background knowledge and networking of dual-NIC ECSs, see **Overview**.

Procedure (Windows ECS Using IPv4)

1. Collect the ECS network information required for configuring policy-based routes.

For details, see Collecting ECS Network Information.

- Log in to an ECS.
 Multiple methods are available for logging in to an ECS. For details, see Logging In to an ECS.
- 3. Check whether the source ECS can use its primary NIC to communicate with the destination ECS:

ping -S *IP* address of the primary NIC on the source ECS IP address of the destination ECS

In this example, run the following command:

ping -S 10.0.0.59 10.0.2.12

If information similar to the following is displayed, the source ECS can use its primary NIC to communicate with the destination ECS.

C:\Us	ers\A	dministi	ator>pi	ng -S	10.0.0.	59 10.0.	2.12
Pingi Reply Reply Reply	ng 10 from from from	0.2.12 10.0.2 10.0.2 10.0.2	from 10 .12: byt .12: byt .12: byt .12: byt	es=32 es=32 es=32 es=32	9 with time<1n time<1n time<1n	32 bytes s TIL=64 s TIL=64 s TIL=64 s TIL=64	of data:
Reply	from	10.0.2	.12: byt	es=32	time<1m	s TTL=64	

NOTE

Before configuring policy-based routes, ensure that the source ECS can use its primary NIC to communicate with the destination ECS.

4. Configure a policy-based route for the extension NIC.

route add -p 0.0.0.0 mask 0.0.0.0 Subnet gateway of the extension NIC metric Route priority

Configure the parameters as follows:

- **0.0.0.0/0**: Default route. Do not change it.
- Subnet gateway of the extension NIC: Enter the IP address collected in 1.
- Route priority: Set its value to 261. The priority of the extension NIC must be lower than that of the primary NIC. A larger value indicates a lower priority.

In this example, run the following command:

route add -p 0.0.0.0 mask 0.0.0.0 10.0.1.1 metric 261

NOTE

- The primary NIC already has policy-based routes and you do not need to configure again.
- If the ECS has multiple extension NICs, configure policy-based routes for all extension NICs one by one.
- 5. Check whether the policy-based route is successfully added.

route print

If information similar to the following is displayed, the policy-based route has been added. The route is persistent and will not be lost after the ECS is restarted.

C:\Users\Administrator>route	print			
Interface List	Rod H	t Hintlû Ethennet	Adapton #2	
14fa 16 3e 5d 3e b6	Red H	at UirtIO Ethernet	: Adapter #3	
1	.Softw	are Loonback Inter	face 1	
1600 00 00 00 00 00 00 00 00	Micros	soft ISATAP Adapte	er #2	
	=======			
IPv4 Route Table				
Active Routes:				
Network Destination Network	etmask	Gateway	Interface	Metric
0.0.0.0 0.	0.0.0	10.0.1.1	10.0.1.104	266
0.0.0.0 0.	0.0.0	10.0.0.1	10.0.0.59	5
10.0.0.0 255.255	255.0	On-link	10.0.0.59	261
10.0.0.59 255.255.25	5.255	On-link	10.0.0.59	261
10.0.0.255 255.255.25	5.255	On-link	10.0.0.59	261
10.0.1.0 255.255	.255.0	On-link	10.0.1.104	261
10.0.1.104 255.255.25	55.255	On-link	10.0.1.104	261
10.0.1.255 255.255.25	5.255	On-link	10.0.1.104	261
127.0.0.0 255.	0.0.0	On-link	127.0.0.1	306
127.0.0.1 255.255.25	55.255	On-link	127.0.0.1	306
127.255.255.255 255.255.25	55.255	On-link	127.0.0.1	306
169.254.169.254 255.255.25	55.255	10.0.0.254	10.0.0.59	6
224.0.0.0 240.	0.0.0	On-link	127.0.0.1	306
224.0.0.0 240.	.0.0.0	On-link	10.0.0.59	261
224.0.0.0240.	0.0.0	On-link	10.0.1.104	261
255.255.255.255 255.255.25	55.255	On-link	127.0.0.1	306
255.255.255.255 255.255.25	5.255	On-link	10.0.0.59	261
255.255.255.255 255.255.25	5.255	0n-link	10.0.1.104	261
Persistent Routes:				
Network Address Ne	etmask	Gateway Address	Metric	
0.0.0.0 0.	0.0.0	10.0.1.1	261	
IPv6 Route Table				
Active Boutes:				
If Metric Network Destination	n	Gateway		
1 306 ::1/128		0n-link		
14 261 fe80::/64		0n-link		
19 261 fe80::/64		0n-link		
19 261 fe80::197b:3504:e6	15:5a4d	/128		
		On-link		
14 261 fe80::e115:8e6a:5c	lcc:671	5/128		
		On-link		
1 306 ff00::/8		On-link		
14 261 ff00::/8		On-link		
19 261 ff00::/8		On-link		
Persistent Routes:				
None				

6. Check whether the source ECS and the destination ECS can communicate with each other.

ping -S *IP* address of the primary NIC on the source ECS IP address of the destination ECS

ping -S *IP* address of the extension NIC on the source ECS IP address of the destination ECS

In this example, run the following commands:

ping -S 10.0.0.59 10.0.2.12

ping -S 10.0.1.104 10.0.2.12

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS.

C:\Users\Administrator>ping -S 10.0.0.59 10.0.2.12 Pinging 10.0.2.12 from 10.0.0.59 with 32 bytes of data: Reply from 10.0.2.12: bytes=32 time<1ms TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=64 Ping statistics for 10.0.2.12: Packets: Sent = 4, Received = 4, Lost = 0 <0% loss>, Approximate round trip times in milli-seconds: Minimum = 0ms, Maximum = 0ms, Average = 0ms C:\Users\Administrator>ping -S 10.0.1.104 10.0.2.12 Pinging 10.0.2.12: bytes=32 time<1ms TIL=64 Reply from 10.0.2.12: bytes=32 time=4ms TIL=64 Reply from 10.0.2.12: bytes=32 time=4ms TIL=64 Reply from 10.0.2.12: bytes=32 time<1ms TIL=64 Ping statistics for 10.0.2.12: Packets: Sent = 4, Received = 4, Lost = 0 (0% loss), Approximate round trip times in milli-seconds: Minimum = 0ms, Maximum = 4ms, Average = 1ms

Procedure (Windows ECS Using IPv6)

1. Collect the ECS network information required for configuring policy-based routes.

For details, see Collecting ECS Network Information.

- Log in to an ECS.
 Multiple methods are available for logging in to an ECS. For details, see Logging In to an ECS.
- 3. Run the following command to check whether the ECS has IPv6 enabled and can obtain IPv6 addresses.

ipconfig

If information similar to the following is displayed, each NIC has an IPv6 address starting with 2407, which indicates that the ECS can obtain IPv6 addresses.

:\Users\Administrator>ipconfig
indows IP Configuration
thernet adapter Ethernet 4:
Connection-specific DNS_Suffix _: openstacklocal
IPv6 Address 2407:c080:802:be6:ec23:ec4:c886:cc1
Link-local 1Pvb Hadress : fe80::8836:ab93:1003:a170%19
IPv4 Address
Subnet Mask
Default Gateway fe80::f816:3eff:fe3e:1e1e%19
thernet adapter Ethernet 2:
Connection-specific DNS Suffix . : openstacklocal
IPv6 Address
Link-local lPvb Hddress : te80::180d:t3b5:27ac:2acb/14
IPv4 Address
Subnet Mask
Default Gateway fe80::f816:3eff:fede:c837%14
192.168.0.1
unnel adapter isatap.openstacklocal:
Media State : Media disconnected Connection-specific DNS Suffix . : openstacklocal
·· Nleewe \ Administration \

NOTICE

Perform this step for both the source and destination ECSs to ensure that the ECSs have obtained IPv6 addresses. Otherwise, the ECSs cannot communicate with each other using IPv6 addresses.

ECSs in this example run Windows Server 2012 (64-bit). No additional configuration is required for such ECSs because they can automatically obtain IPv6 addresses. If your ECS cannot automatically obtain IPv6 addresses, see **Dynamically Assigning IPv6 Addresses**.

4. Check whether the source and destination ECSs can communicate with each other.

ping -6 -S *IP address of the primary NIC on the source ECS IP address of the destination ECS*

ping -6 -S *IP* address of the extension NIC on the source ECS IP address of the destination ECS

In this example, run the following commands:

ping -6 -S 2407:c080:802:aba:8999:5e61:e19:cf7e 2407:c080:802:be7:c2e6:d99c:b685:c6c8

ping -6 -S 2407:c080:802:be6:ec23:ec4:c886:cc1 2407:c080:802:be7:c2e6:d99c:b685:c6c8

If information similar to the following is displayed, both the NICs of the source ECS can communicate with the destination ECS.

「「「「「「」」」」、「「」」」、「「」」、「」」、「」」、「」」、「」、「」」、「」、「
C:\Users\Administrator>ping -6 -\$ 2407:c080:802:aba:8999:5e61:e19:cf7e 2407:c080 :802:be7:c2e6:d99c:b685:c6c8
Pinging 2407:c080:802:be7:c2e6:d99c:b685:c6c8 from 2407:c080:802:aba:8999:5e61:e 19:cf7e with 32 butes of data:
Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms
Ping statistics for 2407:c080:802:be7:c2e6:d99c:b685:c6c8:
Packets: Sent = 4. Received = 4. Lost = 0 (0% loss).
Approximate round trip times in milli-seconds:
Mininum = Oms, Maximum = Oms, Average = Oms
C:\Users\Administrator>ping -6 -8 2407:c080:802:be6:ec23:ec4:c886:cc1 2407:c080: 802:be7:c2e6:d99c:b685:c6c8
Pinging 2407:c080:802:be7:c2e6:d99c:b685:c6c8 from 2407:c080:802:be6:ec23:ec4:c8 86:cc1 with 32 bytes of data:
Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time=3ms
Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms
Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ms
Reply from 2407:c080:802:be7:c2e6:d99c:b685:c6c8: time<1ns
Ping statistics for 2407:c080:802:be7:c2e6:d99c:b685:c6c8:
Packets: Sent = 4. Received = 4. Lost = 0 (0% loss).
Approximate round trip times in milli-seconds:
Mininum = Øms, Maxinum = 3ms, Average = Øms

NOTICE

ECSs in this example run Windows Server 2012 (64-bit). You do not need to configure policy-based routes for these ECSs because both the NICs of such an ECS can communicate with others using IPv6.

13 VPC Peering Configurations

13.1 Overview

VPCs are isolated from each other. To connect two VPCs in the same region, you can use a VPC peering connection to route traffic between them using private IP addresses. To connect VPCs in different regions, you can use **Cloud Connect**.

Here are some scenarios to help you determine which configuration is best suited to your networking requirements.

Scenario	Configuration Description
 VPC CIDR blocks do not overlap. 	Create VPC peering connections to connect entire CIDR blocks of VPCs.
 Subnet CIDR blocks do not overlap. 	For details, see Connecting Entire CIDR Blocks of VPCs.
 VPC CIDR blocks overlap. Some subnet CIDR blocks overlap. 	 Create VPC peering connections to connect specific subnets or ECSs from different VPCs. To connect specific subnets from two VPCs, the subnet CIDR blocks cannot overlap. For details, see Connecting Specific Subnets from Different VPCs. To connect specific ECSs from two VPCs, each ECS must have a unique private IP address. For details, see Connecting Specific ECSs from Different VPCs.
 VPC CIDR blocks overlap. All subnet CIDR blocks overlap. 	VPC peering connections are not usable. For details, see Unsupported VPC Peering Configurations.

Table 13-1 VPC peering scenarios

13.2 Connecting Entire CIDR Blocks of VPCs

Scenarios

You can configure VPC peering connections and set the destination of the routes added to VPC route tables to the peer VPC CIDR block. In this way, a VPC can have access to the entire CIDR block of the peer VPC. **Table 13-2** shows example scenarios.

Scenario	Scenario Description	IP Add ress Ver sion	Networking Configuration	Related Reference
Two VPCs peered together	You have two VPCs that require full access to each others' resources.		Two VPCs Peered Together (IPv4)	 Creating a VPC Peering
	For example, your company has VPC-A for the human resource department, and VPC-B for the finance department. The two departments require full access to each others' resources.	IPv6	Two VPCs Peered Together (IPv6)	Connecti on with Another VPC in Your Account • Creating a VPC Peering
Multiple VPCs peered	You have multiple VPCs that require access to each others' resources. For example, your company has VPC-A for the human resource department, VPC-B for the finance department, and VPC-C for the marketing department. These departments	IPv4	Multiple VPCs Peered Together (IPv4)	on with a VPC in Another
togetner		IPv4	Multiple VPCs Peered Together Through Transitive Peering Connections (IPv4)	Account
	each others' resources.	IPv6	Multiple VPCs Peered Together (IPv6)	

Table 13-2	Scenario	description
------------	----------	-------------

Scenario	Scenario Description	IP Add ress Ver sion	Networking Configuration	Related Reference
One central VPC peered with two VPCs	You have a central VPC that requires access to two peer VPCs, and similarly, the peer VPCs	IPv4	One Central VPC Peered with Two VPCs (IPv4)	
	the two peer VPCs need to be isolated from each other.	IPv6	One Central VPC Peered with Two VPCs (IPv6)	
	For example, public services (such as databases) are deployed on VPC-A. Both VPC-B and VPC-C need to access the databases, but they do not need to access each other.			
One central VPC with primary and secondary CIDR blocks peered with two VPCs	You have a central VPC that has both primary and secondary CIDR blocks. The central VPC needs to communicate with two peer VPCs, but the peer VPCs need to be isolated from each other.	IPv4	One Central VPC with Primary and Secondary CIDR Blocks Peered with Two VPCs (IPv4)	

Scenario	Scenario Description	IP Add ress Ver sion	Networking Configuration	Related Reference
One central VPC peered with multiple	You have a central VPC that requires access to the multiple peer VPCs, and similarly, the peer	IPv4	One Central VPC Peered with Multiple VPCs (IPv4)	
VPCs	VPCs require access to the central VPC. However, the peer VPCs need to be isolated from each other.	IPv6	One Central VPC Peered with Multiple VPCs (IPv6)	
	For example, public services (such as databases) are deployed on your central VPC-A. VPC-B, VPC-C, VPC-D, VPC-E, VPC-F, and VPC-G need to access the databases, but these VPCs do not need to access each other.			

Notes and Constraints

 If you create a VPC peering connection that connects entire CIDR blocks of two VPCs, the VPC CIDR blocks cannot overlap. Otherwise, the VPC peering connection does not take effect. For details, see Invalid VPC Peering for Overlapping VPC CIDR Blocks.

Even if you intend to use the VPC peering connection for IPv6 communication only, you cannot create a VPC peering connection if the VPCs have matching or overlapping IPv4 CIDR blocks. In all examples in this section, the IPv4 CIDR blocks of any VPCs connected by a VPC peering connection do not overlap.

 If the ECSs in VPCs connected by a VPC peering connections are in different security groups, you need to add rules to the security groups to allow access to each other. For details, Enabling ECSs in Different Security Groups to Communicate with Each Other Through an Internal Network.

In all examples in this section, the ECSs in local and peer VPCs are in the same security group. No additional security group rule is required.

- All route tables in a VPC can contain a maximum of 1,000 routes. If you want to establish VPC peering connections between multiple VPCs, consider this restriction when planning networking.
- In a VPC route table, the route priority is as follows:
 - Local route: A route that is automatically added by the system for communication within a VPC. It has a higher priority than a custom route.

 Custom route: A route added by a user or routes that are delivered during instance creation. It uses the longest prefix match rule to find a destination for packet forwarding.

Figure 13-1 VPC route table

<	tb-vpc-A						
Summa	Associated Subnets						
Nam	n murpa 🆉		Type Default				
ID	1439787-1404-4110-87be-388ae311017d		VPC vpc-A				
Crea	Ted Feb 08, 2023 16:37:41 GMT+08:00		Description - Z				
Rou	tes						
	Interie Add Route Replicate Route Expant ~	Learn how to configure routes 🕑					
	Select a property or enter a keyword.						Q®
0	Destination	IP Addresses 😣 🛛 Next Hop Type 🖯	Next Hop	Type Θ	Description (+)	Operation	
	Local	5 Local	Local	System	Default route that enables insta	Modify Delete	
0	172.16.0.0/16	1 VPC peering connection	and the second sec	Custom	-	Modify Delete	
0	192.168.0.0/16	1 VPC peering connection	10070-01	Custom	-	Modity Delete	
0	10.0.0.025	1 Cloud container		Custom	-	Modify Delete	

Two VPCs Peered Together (IPv4)

Create Peering-AB between VPC-A and VPC-B. The CIDR blocks of VPC-A and VPC-B do not overlap.

- For details about resource planning, see Table 13-3.
- For details about VPC peering relationships, see Table 13-4.

Figure 13-2 Networking diagram (IPv4)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb-VPC- B	ECS- B02		10.0.1.167

 Table 13-3 Resource planning details (IPv4)

Table 13-4 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B

After the VPC peering connection is created, add the following routes to the route tables of the local and peer VPCs:

Table 13-5 VPC route tables (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	10.0.1.0/24	Local	Syste m	

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

D NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

Two VPCs Peered Together (IPv6)

Create Peering-AB between VPC-A and VPC-B. The subnets of VPC-A and VPC-B have both IPv4 and IPv6 CIDR blocks and their IPv4 CIDR blocks do not overlap.

- For details about resource planning, see Table 13-6.
- For details about VPC peering relationships, see Table 13-7.

Figure 13-3 Networking diagram (IPv6)



VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -A	VPC 172.1 -A 6.0.0/ 16	Subne t-A01	 IPv4: 172.1 6.0.0 /24 IPv6: 2407: c080: 802:c 34::/ 64 	rtb- VPC-A	ECS- A01	general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	 IPv4: 172.1 6.1.0 /24 IPv6: 2407: c080: 802:c 37::/ 64 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0.0 .0/16	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407: c080: 802:c 35::/ 64 	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
		Subne t-B02	 IPv4: 10.0. 1.0/2 4 IPv6: 2407: c080: 802:c 38::/ 64 	rtb- VPC-B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700:c1cf

Table 13-6 Resource	planning	details	(IPv6)
---------------------	----------	---------	--------

Table 13-7 Peering relationships (IPv6)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B

After the VPC peering connection is created, add the following routes to the route tables of the local and peer VPCs:

 Table 13-8 VPC route tables (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description			
rtb- VPC-	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.			
A	2407:c080:802:c 34::/64	Local	Syst em				
	172.16.1.0/24	Local	Syst em				
	2407:c080:802:c 37::/64	Local	Syst em				
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.			
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the			
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-AB	Cust om	next hop for IPv6 communication.			
rtb- VPC-	10.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.			
В	2407:c080:802:c 35::/64	Local	Syst em				
	10.0.1.0/24	Local	Syst em				

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	2407:c080:802:c 38::/64	Local	Syst em	
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	communication.

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

Multiple VPCs Peered Together (IPv4)

If multiple VPCs need to communicate with each other, their CIDR blocks cannot overlap and you need to create a VPC peering connection between every two VPCs.

- For details about resource planning, see Table 13-9.
- For details about VPC peering relationships, see Table 13-10.



Figure 13-4 Networking diagram (IPv4)

Table 13-9 Resource planning details (IPv4)

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	192.1 68.0.	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194
0/16	Subne t-C02	192.168. 1.0/24	rtb- VPC-C	ECS- C02		192.168.1.200	

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-B is peered with VPC-C.	Peering-BC	VPC-B	VPC-C

 Table 13-10
 Peering relationships (IPv4)

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC -A	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
	172.16.1.0/24	Local	Syste m		
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.	
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.	
rtb- VPC	10.0.0.0/24	Local	Syste m	Local routes are automatically addec for communications within a VPC.	
-В	10.0.1.0/24	Local	Syste m		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.	
	192.168.0.0/16 (VPC-C)	Peerin g-BC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-BC as the next hop.	

Table 13-11 VPC route tables (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
-C	192.168.1.0/24	Local	Syste m		
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.	
	10.0.0.0/16 (VPC-B)	Peerin g-BC	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop.	

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

Multiple VPCs Peered Together Through Transitive Peering Connections (IPv4)

VPC peering connections are transitive. As shown in **Figure 13-5**, there is a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. To enable communication between VPC-B and VP-C, you can use either of the following methods:

- Create a VPC peering connection between VPC-B and VPC-C. For details, see Multiple VPCs Peered Together (IPv4).
- Add routes to direct traffic between VPC-B and VPC-C based on VPC-A. For details, see Table 13-14.



Figure 13-5 Transitive VPC peering connections

Table 13-12 Resource planning details (IPv4)

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02	purpose web server	172.16.1.91
VPC -B	/PC 10.0. B 0.0/1	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	VPC 192.1 -C 68.0. 0/16	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194
		Subne t-C02	192.168. 1.0/24	rtb- VPC-C	ECS- C02		192.168.1.200

Table 13-13 Peering	relationships	(IPv4)
---------------------	---------------	--------

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-A	172.16.1.0/24	Local	Syste m			
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.		
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.		
rtb- VPC	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-В	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.		
	192.168.0.0/16 (VPC-C)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AB as the next hop.		
rtb- VPC -C	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		

 Table 13-14
 VPC route tables (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	192.168.1.0/24	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.
	10.0.0.0/16 (VPC-В)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AC as the next hop.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

Multiple VPCs Peered Together (IPv6)

If multiple VPCs need to communicate with each other, you need to create a VPC peering connection between every two VPCs. In this example, subnets in VPC-A, VPC-B, and VPC-C have IPv6 CIDR blocks and the IPv4 CIDR blocks of VPC-A, VPC-B, and VPC-C cannot overlap.

- For details about resource planning, see Table 13-15.
- For details about VPC peering relationships, see Table 13-16.

Figure 13-6 Networking diagram (IPv6)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC 172.1 -A 16.0.0/ 16	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c080 :802: c34::/ 64 	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c080 :802: c37::/ 64 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	VPC 10.0. -B 0.0/1 6	10.0. Subne IPv4: rtb- 0.0/1 t-B01 10.0. VPC-B 6 IPv4: 10.0. VPC-B 6 IPv6: 2407 10.0. 10.0 IPv6: 10.0. 10.0. 10.0 IPv6: 10.0.	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162 	
		Subne t-B02	 IPv4: 10.0. 1.0/2 4 IPv6: 2407 :c080 :802: c38::/ 64 	rtb- VPC-B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700:c1cf

 Table 13-15 Resource planning details (IPv6)

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -C	192.1 68.0. 0/16	Subne t-C01	 IPv4: 192. 168. 0.0/2 4 IPv6: 2407 :c080 :802: c3c::/ 64 	rtb- VPC-C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af
		Subne t-C02	 IPv4: 192. 168. 1.0/2 4 IPv6: 2407 :c080 :802: c3d:: /64 	rtb- VPC-C	ECS- C02		 IPv4: 192.168.1.200 IPv6: 2407:c080:80 2:c3d:e9ca:16 9a:390c:74d1

Table 13-16 Peering relationships (IPv6)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-B is peered with VPC-C.	Peering-BC	VPC-B	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
A	2407:c080:802:c 34::/64	Local	Syste m			
	172.16.1.0/24	Local	Syste m			
	2407:c080:802:c 37::/64	Local	Syste m			
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.		
	2407:c080:802:c 35::/64 (Subnet- B01)	2:c Peerin Custo Add routes with the IP g-AB m of Subnet-B01 and Sul destinations and Peeri		Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the		
	2407:c080:802:c 38::/64 (Subnet- B02)		Custo m	next hop for IPv6 communication.		
-	192.168.0.0/16 (VPC-C)	Peerin g-AC	Custo m	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.		
	2407:c080:802:c 3c::/64 (Subnet- C01)	Peerin g-AC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the		
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-AC	Custo m	next hop for IPv6 communication.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
В	2407:c080:802:c 35::/64	Local	Syste m			
	10.0.1.0/24	Local	Syste m			
	2407:c080:802:c 38::/64	Local	Syste m			

Table 13-17 VPC	route tables (IPv6)
-----------------	---------------------

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Custo m	the next hop for IPV6 communication.
	192.168.0.0/16 (VPC-C)	Peerin g-BC	Custo m	Add a route with the CIDR block of VPC-C as the destination and Peering-BC as the next hop for IPv4 communication.
	2407:c080:802:c 3c::/64 (Subnet- C01)	Peerin g-BC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-BC as the
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-BC	Custo m	next hop for IPv6 communication.
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
C	2407:c080:802:c 3c::/64	Local	Syste m	
	192.168.1.0/24	Local	Syste m	
	2407:c080:802:c 3d::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Custo m	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Custo m	the next hop for IPv6 communication.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	10.0.0.0/16 (VPC-В)	Peerin g-BC	Custo m	Add a route with the CIDR block of VPC-B as the destination and Peering-BC as the next hop for IPv4 communication.
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-BC	Custo m	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-BC as the
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-BC	Custo m	next hop for IPv6 communication.

NOTE

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

One Central VPC Peered with Two VPCs (IPv4)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. The three VPCs do not have overlapping CIDR blocks.

- For details about resource planning, see Table 13-18.
- For details about VPC peering relationships, see Table 13-19.



Figure 13-7 Networking diagram (IPv4)

Table 13-18 Resource planning details (IPv4)

VP C Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	16	Subne t-A02	172.16.1 .0/24	rtb- VPC-A	ECS- A02		172.16.1.91
VPC -B	10.0.0 .0/16	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
		Subne t-B02	10.0.1.0 /24	rtb- VPC-B	ECS- B02		10.0.1.167
VPC -C	192.1 68.0.0	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194
	/16	Subne t-C02	192.168. 1.0/24	rtb- VPC-C	ECS- C02		192.168.1.200

Table 13-19	Peering	relationships	(IPv4)
-------------	---------	---------------	--------

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
А	172.16.1.0/24	Local	Syste m			
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.		
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
В	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.		
rtb- VPC- C	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
	192.168.1.0/24	Local	Syste m			

Table 13-20 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

One Central VPC Peered with Two VPCs (IPv6)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. Each VPC has IPv6 subnets. The IPv4 CIDR blocks of the three VPCs do not overlap with each other.

- For details about resource planning, see **Table 13-21**.
- For details about VPC peering relationships, see Table 13-22.

Figure 13-8 Networking diagram (IPv6)



VPC Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c08 0:80 2:c3 4::/6 4 	rtb-VPC- A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c08 0:80 2:c3 7::/6 4 	rtb-VPC- A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/ 24 IPv6: 2407 :c08 0:80 2:c3 5::/6 4 	rtb-VPC- B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
		Subne t-B02	 IPv4: 10.0. 1.0/ 24 IPv6: 2407 :c08 0:80 2:c3 8::/6 4 	rtb-VPC- B	ECS- B02		 IPv4: 10.0.1.167 IPv6: 2407:c080:80 2:c38:b9a9:aa 03:2700:c1cf

 Table 13-21 Resource planning details (IPv6)

VPC Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -C	192.1 68.0. 0/16	Subne t-C01 • IPv4: 192. 168. 0.0/ 24 • IPv6: 2407 :c08 0:80 2:c3c ::/64	rtb-VPC- C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af 	
		Subne t-C02	 IPv4: 192. 168. 1.0/ 24 IPv6: 2407 :c08 0:80 2:c3 d::/6 4 	rtb-VPC- C	ECS- C02		 IPv4: 192.168.1.200 IPv6: 2407:c080:80 2:c3d:e9ca:16 9a:390c:74d1

 Table 13-22
 Peering relationships (IPv6)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:
Rou te Tabl e	Destination	Next Hop	Rout e Type	Description			
rtb- VPC	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
-A	2407:c080:802:c 34::/64	Local	Syste m				
	172.16.1.0/24	Local	Syste m				
	2407:c080:802:c 37::/64	Local	Syste m				
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering- AB as the next hop for IPv4 communication.			
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-B01 and Subnet-B02 as the destinations and Peering-AB as the			
	2407:c080:802:c 38::/64 (Subnet- B02)	Peerin g-AB	Cust om	next hop for IPv6 communication.			
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering- AC as the next hop for IPv4 communication.			
	2407:c080:802:c 3c::/64 (Subnet- C01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-C01 and Subnet-C02 as the destinations and Peering-AC as the			
	2407:c080:802:c 3d::/64 (Subnet- C02)	Peerin g-AC	Cust om	next hop for IPV6 communication.			
rtb- VPC	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
-В	2407:c080:802:c 35::/64	Local	Syste m				
	10.0.1.0/24	Local	Syste m				
	2407:c080:802:c 38::/64	Local	Syste m				

Table 13-23 VPC route table details (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.		
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	next hop for IPv6 communication.		
rtb- VPC -C	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
	2407:c080:802:c 3c::/64	Local	Syste m			
	192.168.1.0/24	Local	Syste m			
	2407:c080:802:c 3d::/64	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.		
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Cust om	next hop for IPv6 communication.		

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

One Central VPC with Primary and Secondary CIDR Blocks Peered with Two VPCs (IPv4)

Create Peering-AB between VPC-A and VPC-B, and Peering-AC between VPC-A and VPC-C. VPC-A has both primary and secondary CIDR blocks. The three VPCs do not have overlapping CIDR blocks.

- For details about resource planning, see Table 13-24.
- For details about VPC peering relationships, see Table 13-25.



Figure 13-9 Networking diagram (IPv4)

VP C Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	Prima ry CIDR block: 172.1 6.0.0/ 16	Subne t-A01	172.16.0 .0/24	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	Secon dary CIDR block: 192.1 67.0.0 /16	Subne t-A- Exten d01	192.167. 0.0/24	rtb- VPC-A	ECS- A- Exte nd0 1		192.167.0.100
VPC -B	10.0.0 .0/16	Subne t-B01	10.0.0.0 /24	rtb- VPC-B	ECS- B01		10.0.0.139
VPC -C	192.1 68.0.0 /16	Subne t-C01	192.168. 0.0/24	rtb- VPC-C	ECS- C01		192.168.0.194

 Table 13-24 Resource planning details

Table 13-25 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
A	192.167.0.0/24	Local	Syste m			
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.		
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
B	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peerin g-AB	Cust om	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AB as		
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peerin g-AB	Cust om	the next hop.		
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
C	172.16.0.0/16 (Primary CIDR block of VPC-A)	Peerin g-AC	Cust om	Add routes with the primary and secondary CIDR blocks of VPC-A as the destinations and Peering-AC as		
	192.167.0.0/16 (Secondary CIDR block of VPC-A)	Peerin g-AC	Cust om	the next hop.		

Table 13-26 VPC route table details (IPv4)

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

One Central VPC Peered with Multiple VPCs (IPv4)

Create a VPC peering connection between VPC-A and VPC-B, between VPC-A and VPC-C, between VPC-A and VPC-D, between VPC-A and VPC-E, between VPC-A

and VPC-F, and between VPC-A and VPC-G. The CIDR blocks of these VPCs do not overlap.

- For details about resource planning, see Table 13-27.
- For details about VPC peering relationships, see Table 13-28.

Figure 13-10 Networking diagram (IPv4)



Table 13-27 Resource planning details (IPv4)

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02	purpose web server	172.16.1.91

VPC Na me	VPC CIDR Bloc k	Subn et Name	Subnet CIDR Block	Subnet Route Table	ECS Na me	Security Group	Private IP Address
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	192.1 68.0. 0/16	Subne t-C01	192.168 .0.0/24	rtb-VPC- C	ECS- C01		192.168.0.194
VPC -D	10.2. 0.0/1 6	Subne t-D01	10.2.0.0 /24	rtb-VPC- D	ECS- D01		10.2.0.237
VPC -E	10.3. 0.0/1 6	Subne t-E01	10.3.0.0 /24	rtb-VPC- E	ECS- E01		10.3.0.87
VPC -F	172.1 7.0.0/ 16	Subne t-F01	172.17. 0.0/24	rtb-VPC- F	ECS- F01		172.17.0.103
VPC -G	10.4. 0.0/1 6	Subne t-G01	10.4.0.0 /24	rtb-VPC- G	ECS- G01		10.4.0.10

Table 13-28 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-A is peered with VPC-D.	Peering-AD	VPC-A	VPC-D
VPC-A is peered with VPC-E.	Peering-AE	VPC-A	VPC-E
VPC-A is peered with VPC-F.	Peering-AF	VPC-A	VPC-F
VPC-A is peered with VPC-G.	Peering-AG	VPC-A	VPC-G

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
A	172.16.1.0/24	Local	Syste m		
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.	
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.	
	10.2.0.0/16 (VPC-D)	Peerin Cust Ac g-AD om VF Pe		Add a route with the CIDR block of VPC-D as the destination and Peering-AD as the next hop.	
	10.3.0.0/16 (VPC-E)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-E as the destination and Peering-AE as the next hop.	
	172.17.0.0/16 (VPC-F)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-F as the destination and Peering-AF as the next hop.	
	10.4.0.0/16 (VPC-G)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-G as the destination and Peering-AG as the next hop.	
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
В	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.	
rtb- VPC-	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
C	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.	
rtb- VPC- D	10.2.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	

Table 13-29 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	172.16.0.0/16 (VPC-A)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop.		
rtb- VPC-	10.3.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
E	172.16.0.0/16 (VPC-A)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.		
rtb- VPC-	172.17.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
F	172.16.0.0/16 (VPC-A)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.		
rtb- VPC-	10.4.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
G	172.16.0.0/16 (VPC-A)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.		

D NOTE

If the destination of a route in a route table of a VPC is set to the CIDR block of a peer VPC and the CIDR blocks of the two VPCs do not overlap, the VPCs can have full access to each others' resources.

One Central VPC Peered with Multiple VPCs (IPv6)

Create a VPC peering connection between VPC-A and VPC-B, between VPC-A and VPC-C, between VPC-A and VPC-D, between VPC-A and VPC-E, between VPC-A and VPC-F, and between VPC-A and VPC-G. Each VPC has IPv6 subnets. The IPv4 CIDR blocks of these VPCs do not overlap.

- For details about resource planning, see **Table 13-30**.
- For details about VPC peering relationships, see Table 13-31.



Figure 13-11 Networking diagram (IPv6)

Table 13-30 Resource planning details (IPv6)

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c08 0:80 2:c3 4::/6 4 	rtb- VPC-A	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
		Subne t-A02	 IPv4: 172. 16.1. 0/24 IPv6: 2407 :c08 0:80 2:c3 7::/6 4 	rtb- VPC-A	ECS- A02		 IPv4: 172.16.1.91 IPv6: 2407:c080:80 2:c37:594b:4c 0f:2fcd:8b72
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c3 5::/6 4 	rtb- VPC-B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
VPC -C	192.1 68.0. 0/16	Subne t-C01	 IPv4: 192. 168. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c3c ::/64 	rtb- VPC-C	ECS- C01		 IPv4: 192.168.0.194 IPv6: 2407:c080:80 2:c3c:d2f3:d89 1:24f5:f4af
VPC -D	10.2. 0.0/1 6	Subne t-D01	 IPv4: 10.2. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 5::/6 4 	rtb- VPC-D	ECS- D01		 IPv4: 10.2.0.237 IPv6: 2407:c080:80 2:c45:6bb7:f1 61:3596:6e4c

VP C Na me	VPC CIDR Block	Subne t Name	Subnet CIDR Block	Subnet Route Table	ECS Nam e	Security Group	Private IP Address
VPC -E	10.3. 0.0/1 6	Subne t-E01	 IPv4: 10.3. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 6::/6 4 	rtb- VPC-E	ECS- E01		 IPv4: 10.3.0.87 IPv6: 2407:c080:80 2:c46:2a2f:55 8a:85da:4c70
VPC -F	172.1 7.0.0/ 16	Subne t-F01	 IPv4: 172. 17.0. 0/24 IPv6: 2407 :c08 0:80 2:c4 7::/6 4 	rtb- VPC-F	ECS- F01		 IPv4: 172.17.0.103 IPv6: 2407:c080:80 2:c47:b5e2:e6f 0:c42b:44fd
VPC -G	10.4. 0.0/1 6	Subne t-G01	 IPv4: 10.4. 0.0/2 4 IPv6: 2407 :c08 0:80 2:c4 8::/6 4 	rtb- VPC-G	ECS- G01		 IPv4: 10.4.0.10 IPv6: 2407:c080:80 2:c48:3020:f4 8c:4e54:aa17

Table 13-31 Peering relationships (IPv6)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-B.	Peering-AB	VPC-A	VPC-B

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C
VPC-A is peered with VPC-D.	Peering-AD	VPC-A	VPC-D
VPC-A is peered with VPC-E.	Peering-AE	VPC-A	VPC-E
VPC-A is peered with VPC-F.	Peering-AF	VPC-A	VPC-F
VPC-A is peered with VPC-G.	Peering-AG	VPC-A	VPC-G

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC	172.16.0.0/24	16.0.0/24 Local Syste Local ro m for com		Local routes are automatically added for communications within a VPC.
-A	2407:c080:802:c 34::/64	Local	Syste m	
	172.16.1.0/24	Local	Syste m	
	2407:c080:802:c 37::/64	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add a route with the IPv6 CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop for IPv6 communication.

Table 13-32 VPC route table details (IPv6)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	192.168.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering- AC as the next hop for IPv4 communication.
	2407:c080:802:c 3c::/64 (Subnet- C01)	Peerin g-AC	Cust om	Add a route with the IPv6 CIDR block of Subnet-C01 as the destination and Peering-AC as the next hop for IPv6 communication.
	10.2.0.0/16 (VPC-D)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-D as the destination and Peering-AD as the next hop for IPv4 communication.
	2407:c080:802:c 45::/64 (Subnet- D01)	Peerin g-AD	Cust om	Add a route with the IPv6 CIDR block of Subnet-D01 as the destination and Peering-AD as the next hop for IPv6 communication.
	10.3.0.0/16 (VPC-E)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-E as the destination and Peering- AE as the next hop for IPv4 communication.
	2407:c080:802:c 46::/64 (Subnet- E01)	Peerin g-AE	Cust om	Add a route with the IPv6 CIDR block of Subnet-E01 as the destination and Peering-AE as the next hop for IPv6 communication.
	172.17.0.0/16 (VPC-F)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-F as the destination and Peering-AF as the next hop for IPv4 communication.
	2407:c080:802:c 47::/64 (Subnet- F01)	Peerin g-AF	Cust om	Add a route with the IPv6 CIDR block of Subnet-F01 as the destination and Peering-AF as the next hop for IPv6 communication.
	10.4.0.0/16 (VPC-G)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-G as the destination and Peering-AG as the next hop for IPv4 communication.
	2407:c080:802:c 48::/64 (Subnet- G01)	Peerin g-AG	Cust om	Add a route with the IPv6 CIDR block of Subnet-G01 as the destination and Peering-AG as the next hop for IPv6 communication.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
rtb- VPC	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-В	2407:c080:802:c 35::/64	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop for IPv4 communication.		
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AB as the		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AB	Cust om	next hop for IPv6 communication.		
rtb- VPC	192.168.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-C -	2407:c080:802:c 3c::/64	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop for IPv4 communication.		
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AC	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AC as the		
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AC	Cust om	next hop for IPv6 communication.		
rtb- VPC	10.2.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
-D	2407:c080:802:c 45::/64	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AD	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AD as the next hop for IPv4 communication.		

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AD	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AD as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AD	Cust om	next hop for IPv6 communication.
rtb- VPC	10.3.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-E	2407:c080:802:c 46::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AE	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AE as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AE	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AE as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AE	Cust om	next hop for IPv6 communication.
rtb- VPC	172.17.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-F	2407:c080:802:c 47::/64	Local	Syste m	
	172.16.0.0/16 (VPC-A)	Peerin g-AF	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AF as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AF	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AF as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AF	Cust om	next nop for IPV6 communication.
rtb- VPC	10.4.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
-G	2407:c080:802:c 48::/64	Local	Syste m	

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/16 (VPC-A)	Peerin g-AG	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AG as the next hop.
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AG	Cust om	Add routes with the IPv6 CIDR blocks of Subnet-A01 and Subnet-A02 as the destinations and Peering-AG as the
	2407:c080:802:c 37::/64 (Subnet- A02)	Peerin g-AG	Cust om	next hop for IPv6 communication.

D NOTE

You can view IPv6 addresses of VPC subnets on the management console. To enable communication between entire CIDR blocks of two VPCs, you need to add routes with IPv6 CIDR blocks of all subnets in the VPCs as the destinations one by one.

13.3 Connecting Specific Subnets from Different VPCs

Scenarios

You can configure VPC peering connections and set the destination of the routes added to VPC route tables to a subnet CIDR block of a peer VPC. In this way, specific subnets from different VPCs can communicate with each other. Table 13-33 shows example scenarios.

Scenario	Scenario Description	IP Add ress Ver sion	Networking Configuration	Related Reference
Two VPCs peered to two subnets in a central VPC	 You have a central VPC that requires access to the multiple other VPCs. The other VPCs need to be isolated from each other. The central VPC has separate sets of resources in different subnets. The other VPCs require access to some of the resources, but not all of them. 	IPv4 IPv6 / IPv4	Two VPCs Peered to Two Subnets in a Central VPC (IPv4) Two VPCs Peered to Two Subnets in a Central VPC (IPv6/IPv4)	 Creating a VPC Peering Connecti on with Another VPC in Your Account Creating a VPC Peering Connecti on with a VPC in
One central VPC peered to specific subnets in two VPCs	 You have a central VPC that requires access to two other VPCs. The other VPCs need to be isolated from each other. The central VPC has public resources deployed and the other VPCs require access to all resources in the central VPC. Other VPCs have multiple subnets and only one in each VPC is used for accessing resources in the central VPC. 	IPv4	One Central VPC Peered to Specific Subnets in Two VPCs (IPv4)	VPC in Another Account
One central VPC peered to overlappin g subnets from two VPCs	This scenario is similar to the preceding one. If two VPCs with overlapping subnets need to peer with the central VPC, traffic may fail to be forwarded to the required destination. To prevent this, plan the network according to this example.	IPv4	One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)	

Table 13-33 Scenario description

Notes and Constraints

 If the ECSs in VPCs connected by a VPC peering connections are in different security groups, you need to add rules to the security groups to allow access to each other. For details, Enabling ECSs in Different Security Groups to Communicate with Each Other Through an Internal Network.

In all examples in this section, the ECSs in local and peer VPCs are in the same security group. No additional security group rule is required.

- All route tables in a VPC can contain a maximum of 1,000 routes. If you want to establish VPC peering connections between multiple VPCs, consider this restriction when planning networking.
- In a VPC route table, the route priority is as follows:
 - Local route: A route that is automatically added by the system for communication within a VPC. It has a higher priority than a custom route.
 - Custom route: A route added by a user or routes that are delivered during instance creation. It uses the longest prefix match rule to find a destination for packet forwarding.

Figure 13-12 VPC route table

< rtb-	vpc-A								
Summary	Associated Subnets								
Name	rtb-wpc-A 🖉			Type	Default				
ID	14397ff7-1404-4f10-87be-388ae31fd17d			VPC	vpc-A				
Created	Feb 08, 2023 16:37:41 GMT+08:00			Description	- a				
Route	5								
Dek	Add Route Replicate Route Expart - Learn	how to configure routes 🕑							
Qs	elect a property or enter a keyword.								Q®
	Destination	IP Addresses $ \Theta $	Next Hop Type 😑	Next Hop		Type Θ	Description Θ	Operation	
	Local	5	Local	Local		System	Default route that enables insta	Modify Delete	
	172.16.0.0/16	1	VPC peering connection			Custom	-	Modify Delete	
	192.168.0.0/16	1	VPC peering connection	100000-00-		Custom	-	Modify Delete	
	10 0 0 0 25	1	Cloud container			Custom	-	Modily Delete	

Two VPCs Peered to Two Subnets in a Central VPC (IPv4)

The central VPC-A has two subnets, Subnet-A01 and Subnet-A02. The subnets are associated with different route tables. You need to create Peering-AB between Subnet-A01 and VPC-B, and Peering-AC between Subnet-A02 and VPC-C. VPC-B and VPC-C have the same CIDR block. However, there will be no route conflicts because the two subnets in VPC-A are associated with different route tables.

- For details about resource planning, see Table 13-34.
- For details about VPC peering relationships, see Table 13-35.



Figure 13-13 Networking diagram (IPv4)

Table 13-34 Resource planning details (IPv4)

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A01	ECS- A01	sg-web: general- purpose web server	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A02	ECS- A02		172.16.1.91
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71

NOTE

VPC-A has two route tables. Route table rtb-VPC-A01 is associated with Subnet-A01, and route table rtb-VPC-A02 is associated with Subnet-A02. The two subnets can communicate with each other.

Table 13-35 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
Subnet-A01 of VPC-A is peered to VPC-B.	Peering-AB	VPC-A	VPC-B
Subnet-A02 of VPC-A is peered to VPC-C.	Peering-AC	VPC-A	VPC-C

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
AUT	172.16.1.0/24	Local	Syste m	
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A02	A02 172.16.1.0/24 Local Syst		Syste m	
	10.0.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.
rtb- VPC- B	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.

Table 13-36 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
	172.16.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24 (Subnet-A02)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AC as the next hop.

Two VPCs Peered to Two Subnets in a Central VPC (IPv6/IPv4)

The central VPC-A has two subnets, Subnet-A01 and Subnet-A02. The subnets are associated with different route tables. You need to create Peering-AB between Subnet-A01 and VPC-B for IPv6 communication, and Peering-AC between Subnet-A02 and VPC-C for IPv4 communication. VPC-B and VPC-C have the same CIDR block. However, there will be no route conflicts because the two subnets in VPC-A are associated with different route tables.

- For details about resource planning, see Table 13-37.
- For details about VPC peering relationships, see Table 13-38.



Figure 13-14 Networking diagram (IPv6/IPv4)

Table 13-37	Resource	planning	details	(IPv6/IPv4	.)

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	 IPv4: 172. 16.0. 0/24 IPv6: 2407 :c080 :802: c34:: /64 	rtb-VPC- A01	ECS- A01	sg-web: general- purpose web server	 IPv4: 172.16.0.111 IPv6: 2407:c080:80 2:c34:a925:f1 2e:cfa0:8edb
		Subne t-A02	172.16. 1.0/24	rtb-VPC- A02	ECS- A02		172.16.1.91

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -B	10.0. 0.0/1 6	Subne t-B01	 IPv4: 10.0. 0.0/2 4 IPv6: 2407 :c080 :802: c35:: /64 	rtb-VPC- B	ECS- B01		 IPv4: 10.0.0.139 IPv6: 2407:c080:80 2:c35:493:33f 4:4531:5162
VPC -C	10.0. 0.0/1 6	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71

NOTE

VPC-A has two route tables. Route table rtb-VPC-A01 is associated with Subnet-A01, and route table rtb-VPC-A02 is associated with Subnet-A02. The two subnets can communicate with each other.

Table	13-38	Peering	relationships	(IPv6/IPv4)
-------	-------	---------	---------------	-------------

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
Subnet-A01 of VPC-A is peered to VPC-B. (IPv6)	Peering-AB	VPC-A	VPC-B
Subnet-A02 of VPC-A is peered to VPC-C. (IPv4)	Peering-AC	VPC-A	VPC-C

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
A01	2407:c080:802:c 34::/64	Local	Syste m		
	172.16.1.0/24	Local	Syste m		
	10.0.0.0/16 (VPC-B)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop for IPv4 communication.	
	2407:c080:802:c 35::/64 (Subnet- B01)	Peerin g-AB	Cust om	Add a route with the IPv6 CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop for IPv6 communication.	
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
A02	2407:c080:802:c 34::/64	Local	Syste m		
	172.16.1.0/24	Local	Syste m		
	10.0.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop for IPv4 communication.	
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
В	2407:c080:802:c 35::/64	Local	Syste m		
	172.16.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop for IPv4 communication.	
	2407:c080:802:c 34::/64 (Subnet- A01)	Peerin g-AB	Cust om	Add a route with the IPv6 CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop for IPv6 communication.	

Table 13-39 VPC route table details (IPv6/IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.1.0/24 (Subnet-A02)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-A02 as the destination and Peering-AC as the next hop for IPv4 communication.

One Central VPC Peered to Specific Subnets in Two VPCs (IPv4)

You need to create Peering-AB between central VPC-A and Subnet-B01 in VPC-B, and Peering-AC between central VPC-A and Subnet-C02 in VPC-C. VPC-B and VPC-C have the same CIDR block, but the CIDR blocks of Subnet-B01 and Subnet-C02 do not overlap. Therefore, there will be no route conflicts.

- For details about resource planning, see Table 13-40.
- For details about VPC peering relationships, see Table 13-41.

Figure 13-15 Networking diagram (IPv4)



VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general- purpose	172.16.0.111
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01	web server	10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb-VPC- B	ECS- B02		10.0.1.167
VPC -C	10.0. 0.0/1	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71
	6	Subne t-C02	10.0.1.0 /24	rtb-VPC- C	ECS- C02		10.0.1.116

 Table 13-40 Resource planning details (IPv4)

 Table 13-41
 Peering relationships (IPv4)

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered to Subnet-B01 of VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered to Subnet-C02 of VPC-C.	Peering-AC	VPC-A	VPC-C

Table 13-42 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC- A	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description		
	10.0.0.0/24 (Subnet-B01)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-B01 as the destination and Peering-AB as the next hop.		
	10.0.1.0/24 (Subnet-C02)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
В	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.		
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.		
C	10.0.1.0/24	Local	Syste m			
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.		

One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)

If you want to create VPC peering connections between a VPC and multiple overlapping subnets from different VPCs, ensure that the destinations of the routes added for the peering connections do not conflict and traffic can be correctly forwarded.

In this example, you need to create Peering-AB between central VPC-A and Subnet-B02 in VPC-B, and Peering-AC between central VPC-A and Subnet-C02 in VPC-C. Subnet-B02 and Subnet-C02 have the same CIDR block, and ECS-B02 and ECS-C02 have the same private IP address (10.0.1.167/32).

- For details about resource planning, see Table 13-43.
- For details about VPC peering relationships, see Table 13-44.



Figure 13-16 Networking diagram (IPv4)

Table 13-43 Resource planning details (IPv4)

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general- purpose web server	172.16.0.111
VPC -B	10.0. 0.0/1	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
	6	Subne t-B02	10.0.1.0 /24	rtb-VPC- B	ECS- B02		10.0.1.167
VPC -C	10.0. 0.0/1	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71
	6	Subne t-C02	10.0.1.0 /24	rtb-VPC- C	ECS- C02		10.0.1.167

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered to Subnet-B02 of VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered to Subnet-C02 of VPC-C.	Peering-AC	VPC-A	VPC-C

 Table 13-44
 Peering relationships (IPv4)

If you add routes to the route tables of the local and peer VPCs according to **Table 13-45**, the response traffic cannot be correctly forwarded. The details are as follows:

- 1. ECS-B02 in Subnet-B02 of VPC-B sends request traffic to VPC-A through the route with Peering-AB as the next hop in the rtb-VPC-B route table.
- 2. VPC-A receives the request traffic from ECS-B02 and expects to send the response traffic to ECS-B02. The rtb-VPC-A route table has the route with 10.0.1.167/32 as the destination, but its next hop is Peering-AC. The response traffic is incorrectly sent to VPC-C.
- 3. ECS-C02 in Subnet-C02 of VPC-C has the same private IP address (10.0.1.167/32) as ECS-B02. The response traffic is incorrectly sent to ECS-C02, and ECS-B02 cannot receive the response traffic.

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description	
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
A	10.0.1.0/24 (Subnet-C02)	Peerin g-AC	Cust om	Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.	
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.	
В	10.0.1.0/24	Local	Syste m		
	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.	

Table 13-45 VPC route table details (IPv4)

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description			
rtb- VPC- C	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.			
	10.0.1.0/24	Local	Syste m				
	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.			

If there are overlapping subnets, configure routes as follows to prevent traffic from being incorrectly forwarded:

• Suggestion 1: In the rtb-VPC-A route table, add a route with Peering-AB as the next hop and the private IP address of ECS-B02 (10.0.1.167/32) as the destination. The route with 10.0.1.167/32 as the destination is preferentially matched based on the longest prefix match rule to ensure that VPC-A sends the response traffic to ECS-B02. For more configurations, see **Connecting Specific ECSs from Different VPCs**.

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description	
rtb- VPC -A	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.	
	10.0.1.167/32 (ECS-B02)	Peerin g-AB	Cust om	Add a route with the private IP address of ECS-B02 as the destination and Peering-AB as the next hop.	
	10.0.1.0/24 (Subnet-C02)	0.0.1.0/24 Peerin Cust Subnet-C02) g-AC om		Add a route with the CIDR block of Subnet-C02 as the destination and Peering-AC as the next hop.	

 Table 13-46
 VPC route table details

• Suggestion 2: In the rtb-VPC-A route table, change the destination of the route with Peering-AC as the next hop from Subnet-C02 to Subnet-C01. Add a route with Peering-AB as the next hop and Subnet-B02 as the destination to ensure that VPC-A can send the response traffic to Subnet-B02 in VPC-B.

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description
rtb- VPC -A	172.16.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.
	10.0.1.0/24 (Subnet-B02)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.
	10.0.0.0/24 Peer (Subnet-C01) g-AC		Cust om	Add a route with the CIDR block of Subnet-C01 as the destination and Peering-AC as the next hop.

 Table 13-47
 VPC route table details

13.4 Connecting Specific ECSs from Different VPCs

Scenarios

If your VPCs with the same CIDR block also include subnets that overlap, you can configure VPC peering connections that connect specific ECSs from these VPCs.

To enable traffic forwarding among these ECSs, you need to add routes with private IP addresses of these ECSs as the destinations and a VPC peering connection as the next hop to VPC route tables. **Table 13-48** shows example scenarios.

	sion		
ou want a central VPC communicate with the her two VPCs. However, ou do not want the her two VPCs to ommunicate with each her. ne other two VPCs have e same CIDR block and so include subnets that verlap. To prevent route onflicts in the central PC, you can configure PC peering connections connect to specific CSs in the other two PCs.	IPv4	ECS in a Central VPC Peered to ECSs in Two Other VPCs (IPv4)	 Creating a VPC Peering Connecti on with Another VPC in Your Account Creating a VPC Peering Connecti on with a VPC in Another Account
his scenario is similar to e preceding one. In Idition to peering ecific ECSs, you can eate the following VPC eering connections ased on the longest efix match rule: Create a VPC peering connection between the central VPC and an ECS in VPC-B Create a VPC peering connection between the central VPC and a subnet in VPC-C his configuration	IPv4	A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)	
	a want a central VPC communicate with the her two VPCs. However, a do not want the her two VPCs to mmunicate with each her. e other two VPCs have e same CIDR block and o include subnets that erlap. To prevent route offlicts in the central C, you can configure C peering connections connect to specific Ss in the other two Cs. s scenario is similar to e preceding one. In dition to peering ecific ECSs, you can ate the following VPC ering connections sed on the longest efix match rule: Create a VPC peering connection between the central VPC and an ECS in VPC-B Create a VPC peering connection between the central VPC and a subnet in VPC-C is configuration bands the mmunication scope.	JumpJumpa want a central VPC communicate with the her two VPCs. However, a do not want the her two VPCs to nmunicate with each her.IPv4a o not want the her two VPCs to nmunicate with each her.IPv4a other two VPCs have e same CIDR block and o include subnets that erlap. To prevent route officts in the central C, you can configure C peering connections connect to specific Ss in the other two Cs.IPv4s scenario is similar to e preceding one. In dition to peering ecific ECSs, you can ate the following VPC ering connections sed on the longest fix match rule:IPv4Create a VPC peering connection between the central VPC and an ECS in VPC-BIPv2Create a VPC peering connection between the central VPC and a subnet in VPC-CIPv2	J want a central VPC communicate with the her two VPCs. However, a do not want the her two VPCs to nmunicate with each her.IPv4ECS in a Central VPC Peered to ECSs in Two Other VPCs (IPv4)e other two VPCs have e same CIDR block and o include subnets that erlap. To prevent route filicts in the central C, you can configure C peering connections connect to specific Ss in the other two Cs.IPv4A Central VPC Peered with Two Other VPCs (IPv4)s scenario is similar to e preceding one. In dition to peering ceific ECSs, you can ate the following VPC ering connections sed on the longest efix match rule:IPv4A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)Create a VPC peering connection between the central VPC and an ECS in VPC-BIPv4A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)Create a VPC peering connection between the central VPC and a subnet in VPC-C is configuration pands the nmunication scope.IPv4

Table 13-48 Scenario description

Notes and Constraints

• If the ECSs in VPCs connected by a VPC peering connections are in different security groups, you need to add rules to the security groups to allow access to each other. For details, Enabling ECSs in Different Security Groups to Communicate with Each Other Through an Internal Network.

In all examples in this section, the ECSs in local and peer VPCs are in the same security group. No additional security group rule is required.

- All route tables in a VPC can contain a maximum of 1,000 routes. If you want to establish VPC peering connections between multiple VPCs, consider this restriction when planning networking.
- In a VPC route table, the route priority is as follows:
 - Local route: A route that is automatically added by the system for communication within a VPC. It has a higher priority than a custom route.
 - Custom route: A route added by a user or routes that are delivered during instance creation. It uses the longest prefix match rule to find a destination for packet forwarding.

Figure	13-17	VPC	route	table
--------	-------	-----	-------	-------

<	rth-spic A									
Sumn	nary	Associated Subnets								
N	arne	m-upc-A d			Type	Default				
ID		1439717-1404-4110-87be-388ae311d17d			VPC	vpc-A				
Cr	reated	Feb 08, 2023 16:37:41 GMT+08:00			Description	- a				
R	outes									
	Delete	Add Route Replicate Route Expart - Les	n how to configure routes 🕑							
	C. Select	a property or enter a keyword.								0
	Des	tination	IP Addresses $ \hat{\Theta} $	Next Hop Type (Next Hop		Type \ominus	Description Θ	Operation	
	Loc	ы	5	Local	Local		System	Default route that enables insta	Modify Delete	
	172	16.0.0/16	1	VPC peering connection			Custom	-	Modify Delete	
	192	168.0.0/16	1	VPC peering connection			Custom	-	Modify Delete	
	10.0	0.025	1	Cloud container			Custom	-	Modify Delete	

ECS in a Central VPC Peered to ECSs in Two Other VPCs (IPv4)

You want to create a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. VPC-B and VPC-C have matching CIDR blocks. You can set the destinations of routes to private IP addresses of specific ECSs to limit traffic to these ECSs. If the destination of a route is not properly planned, traffic cannot be correctly forwarded. For details, see **One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)**.

In this example, you need to create Peering-AB between ECS-A01-1 in VPC-A and ECS-B01 in VPC-B, and Peering-AC between ECS-A01-2 in VPC-A and ECS-C01 in VPC-C. Subnet-B01 and Subnet-C01 have matching CIDR blocks. The private IP addresses of ECS-B01 and ECS-C01 must be different. Otherwise, there will be route conflicts because the route table of VPC-A will have routes with the same destination.

- For details about resource planning, see **Table 13-49**.
- For details about VPC peering relationships, see Table 13-50.



Figure 13-18 Networking diagram (IPv4)

Table 13-49 Resour	rce planning de	etails (IPv4)
--------------------	-----------------	---------------

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/ 16	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01- 1	sg-web: general- purpose web server	172.16.0.111
					ECS- A01- 2		172.16.0.218
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01	-	10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
ECS-A01-1 in VPC-A is peered with ECS-B01 in VPC-B.	Peering-AB	VPC-A	VPC-B
ECS-A01-2 in VPC-A is peered with ECS-C01 in VPC-C.	Peering-AC	VPC-A	VPC-C

 Table 13-50 Peering relationships (IPv4)

Table 13-31 VPC Toule lable details (IPV4)	Table	13-51	VPC	route	table	details	(IPv4)
--	-------	-------	-----	-------	-------	---------	--------

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC- A	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	10.0.0.139/32 (ECS-B01)	Peerin g-AB	Cust om	Add a route with the private IP address of ECS-B01 as the destination and Peering-AB as the next hop.
	10.0.0.71/32 (ECS-C01)	Peerin g-AC	Cust om	Add a route with the private IP address of ECS-C01 as the destination and Peering-AC as the next hop.
rtb- VPC- B	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
	172.16.0.111/32 (ECS-A01-1)	Peerin g-AB	Cust om	Add a route with the private IP address of ECS-A01-1 as the destination and Peering-AB as the next hop.
rtb- VPC- C	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
------------------------	--------------------------------	----------------	-------------------	--
	172.16.0.218/32 (ECS-A01-2)	Peerin g-AC	Cust om	Add a route with the private IP address of ECS-A01-2 as the destination and Peering-AC as the next hop.

A Central VPC Peered with Two Other VPCs Using Longest Prefix Match (IPv4)

You want to create a VPC peering connection between VPC-A and VPC-B, and between VPC-A and VPC-C. VPC-B and VPC-C have matching CIDR blocks. You can set the destinations of routes to private IP addresses of specific ECSs to limit traffic to these ECSs. If the destination of a route is not properly planned, traffic cannot be correctly forwarded. For details, see **One Central VPC Peered to Overlapping Subnets from Two VPCs (IPv4)**.

In this example, you need to create Peering-AB between central VPC-A and ECS-B01 in VPC-B, and Peering-AC between central VPC-A and VPC-C. Subnet-B01 and Subnet-C01 have matching CIDR blocks. You can use the longest prefix match rule to control traffic forwarding.

- For details about resource planning, see Table 13-52.
- For details about VPC peering relationships, see Table 13-53.



Figure 13-19 Networking diagram (IPv4)

Table 13-52 Resource	planning	details	(IPv4)
----------------------	----------	---------	--------

VPC Na me	VPC CIDR Block	Subn et Name	Subnet CIDR Block	VPC Route Table	ECS Na me	Security Group	Private IP Address
VPC -A	172.1 6.0.0/	Subne t-A01	172.16. 0.0/24	rtb-VPC- A	ECS- A01	sg-web: general-	172.16.0.111
	16	Subne t-A02	172.16. 1.0/24	rtb-VPC- A	ECS- A02	purpose web server	172.16.1.91
VPC -B	10.0. 0.0/1 6	Subne t-B01	10.0.0.0 /24	rtb-VPC- B	ECS- B01		10.0.0.139
VPC -C	10.0. 0.0/1 6	Subne t-C01	10.0.0.0 /24	rtb-VPC- C	ECS- C01		10.0.0.71

Table	13-53	Peering	relationships	(IPv4)
-------	-------	---------	---------------	--------

Peering Relationship	Peering Connection Name	Local VPC	Peer VPC
VPC-A is peered with ECS-B01 in VPC-B.	Peering-AB	VPC-A	VPC-B
VPC-A is peered with VPC-C.	Peering-AC	VPC-A	VPC-C

After the VPC peering connections are created, add the following routes to the route tables of the local and peer VPCs:

Rou te Tabl e	Destination	Next Hop	Rout e Type	Description
rtb- VPC-	172.16.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
A	172.16.1.0/24	Local	Syste m	
	10.0.0.139/32 (ECS-B01)	Peerin g-AB	Cust om	Add a route with the private IP address of ECS-B01 as the destination and Peering-AB as the next hop.
	10.0.0.0/16 (VPC-C)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-C as the destination and Peering-AC as the next hop.
rtb- VPC-	10.0.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
В	172.16.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.
rtb- VPC-	10.0.0/24	Local	Syste m	Local routes are automatically added for communications within a VPC.
L	172.16.0.0/16 (VPC-A)	Peerin g-AC	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AC as the next hop.

Table 13-54 VPC route table details (IPv4)

13.5 Unsupported VPC Peering Configurations

Scenarios

The VPC peering connection configurations are not supported in Table 13-55.

Scenario	Example
 If VPCs with the same CIDR block also include subnets that overlap, VPC peering connections are not usable. If two VPCs have overlapping CIDR blocks but some of their subnets do not overlap, you cannot create a VPC peering connection to connect specific subnets that do not overlap. 	 Invalid VPC Peering for Overlapping VPC CIDR Blocks VPCs with the same CIDR block also include subnets that overlap. Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.
VPC peering connections cannot enable ECSs in their VPCs to share an EIP to access the Internet. If VPC-A and VPC-B are peered and ECS- A01 in VPC-A has an EIP, ECS-B01 in VPC- B cannot access the Internet using the EIP bound to ECS-A01.	Invalid VPC Peering for Sharing an EIP

Notes and Constraints

 If the ECSs in VPCs connected by a VPC peering connections are in different security groups, you need to add rules to the security groups to allow access to each other. For details, Enabling ECSs in Different Security Groups to Communicate with Each Other Through an Internal Network.

In all examples in this section, the ECSs in local and peer VPCs are in the same security group. No additional security group rule is required.

- All route tables in a VPC can contain a maximum of 1,000 routes. If you want to establish VPC peering connections between multiple VPCs, consider this restriction when planning networking.
- In a VPC route table, the route priority is as follows:
 - Local route: A route that is automatically added by the system for communication within a VPC. It has a higher priority than a custom route.
 - Custom route: A route added by a user or routes that are delivered during instance creation. It uses the longest prefix match rule to find a destination for packet forwarding.

Figure 13-20 VPC route table

< rtb.	vpc-A						
Summary	Associated Subnets						
Name	rtb-vpc-A		Type Default				
ID	1439787-1404-4f10-87be-388ae311417d		VPC vpc-A				
Created	Feb 08, 2023 16:37:41 GMT+08:00		Description - R				
Route							
Dek	Add Route Replicate Route Export - Learn ho	er to configure routes 🔝					
Q S	alect a property or enter a keyword.						00
	Destination	IP Addresses 😑 👘 Next Hop Type 😔	Next Hop	Type 🖯	Description (+)	Operation	
	Local	5 Local	Local	System	Default route that enables insta	Modily Delete	
	172.16.0.0/18	1 VPC peering connection		Custom	-	Modify Delete	
0	192.168.0.0/16	1 VPC peering connection		Custom	-	Modify Delete	
	10.0.0.025	1 Cloud container		Custom	-	Modily Delete	

Invalid VPC Peering for Overlapping VPC CIDR Blocks

If two VPCs have overlapping CIDR blocks, the VPC peering connection may not take effect due to route conflicts. The following describes the reasons and configuration suggestions.

• VPCs with the same CIDR block also include subnets that overlap.

VPC peering connections are not usable. As shown in **Table 13-56**, VPC-A and VPC-B, and their subnets have the same CIDR block. If you create a VPC peering connection between VPC-A and VPC-B, their route tables are shown in **Table 13-56**.

In the rtb-VPC-A route table, the custom route for routing traffic from VPC-A to VPC-B and the local route have overlapping destinations. The local route has a higher priority and traffic will be forwarded within VPC-A and cannot reach VPC-B.



Figure 13-21 Networking diagram (IPv4)

Rou te Tabl e	Destination	Next Hop	Rou te Type	Description
rtb- VPC	10.0.0/24	Local	Syst em	Local routes are automatically added for communications within
-A	10.0.1.0/24	Local	Syst em	a VPC.
	10.0.0.0/16 (VPC-В)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-B as the destination and Peering-AB as the next hop.
rtb- VPC	10.0.0.0/24	Local	Syst em	Local routes are automatically added for communications within
-В	10.0.1.0/24	Local	Syst em	a VPC.
	10.0.0.0/16 (VPC-A)	Peerin g-AB	Cust om	Add a route with the CIDR block of VPC-A as the destination and Peering-AB as the next hop.

Table 13-56 VPC route table details

If two VPCs want to use their IPv6 CIDR blocks for communication by a VPC peering connection but the IPv4 CIDR blocks of the VPCs or subnets overlap, the connection is not usable.





• Two VPCs have overlapping CIDR blocks but some of their subnets do not overlap.

VPC peering connections will not take effect in the following scenarios:

- Connecting overlapping CIDR blocks of VPCs

As shown in **Figure 13-23**, if you create a VPC peering connection between VPC-A and VPC-B, the VPC peering connection will not take effect because the two VPCs have the same CIDR block.

- Connecting overlapping subnets from different VPCs

If you create a VPC peering connection between Subnet-A01 and Subnet-B02, the route tables are shown in **Table 13-57**. In the rtb-VPC-B route table, the custom route for routing traffic from Subnet-B02 to Subnet-A01 and the local route have overlapping destinations. The local route has a higher priority and traffic will be forwarded within Subnet-B02 and cannot reach Subnet-A01.



Figure 13-23 Networking diagram (IPv4)

Table 13	3-57 VPC	route	table	details

Ro ute Tab le	Destination	Next Hop	Rou te Typ e	Description
rtb- VPC	10.0.0.0/24	Local	Syst em	Local routes are automatically added for communications
-A	10.0.1.0/24	Local	Syst em	within a VPC.
	10.0.2.0/24 (Subnet-B02)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-B02 as the destination and Peering-AB as the next hop.

Ro ute Tab le	Destination	Next Hop	Rou te Typ e	Description
rtb- VPC -B	10.0.0.0/24	Local	Syst em	Local routes are automatically added for communications within a VPC.
	10.0.2.0/24	Local	Syst em	
	10.0.0.0/24 (Subnet-A01)	Peerin g-AB	Cust om	Add a route with the CIDR block of Subnet-A01 as the destination and Peering-AB as the next hop.

If the subnets connected by a VPC peering connection do not overlap, the connection will take effect. As shown in **Figure 13-24**, you can create a VPC peering connection between Subnet-A02 and Subnet-B02. In this case, the routes do not conflict and the VPC peering connection takes effect.



Figure 13-24 Networking diagram (IPv4)

If two VPCs want to use their IPv6 CIDR blocks for communication by a VPC peering connection but the IPv4 CIDR blocks of the VPCs or subnets overlap, the connection is not usable.



Figure 13-25 Networking diagram (IPv6)

Invalid VPC Peering for Sharing an EIP

As shown in **Figure 13-26**, although VPC-A and VPC-B are peered and ECS-A01 in VPC-A has an EIP, ECS-B01 in VPC-B cannot access the Internet using the EIP bound to ECS-A01. If you want multiple resources to share an EIP, refer to **Using NAT Gateway and VPC Peering to Enable Communication Between VPCs and Internet**.



Figure 13-26 Networking diagram