Virtual Private Cloud

Best Practices

Issue	01
Date	2024-10-25

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


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

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process.* For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

 Copyright © Huawei Technologies Co., Ltd.

1 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

 Copyright © Huawei Technologies Co., Ltd.

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.

 Copyright © Huawei Technologies Co., Ltd.

2 VPC Security

2.1 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 2-1 shows how ecs-A01 accesses ecs-C01. You can view the request and response traffic paths.

 Copyright © Huawei Technologies Co., Ltd.

Figure 2-1 Networking planning when a third-party firewall is used for scrubbing traffic across VPCs

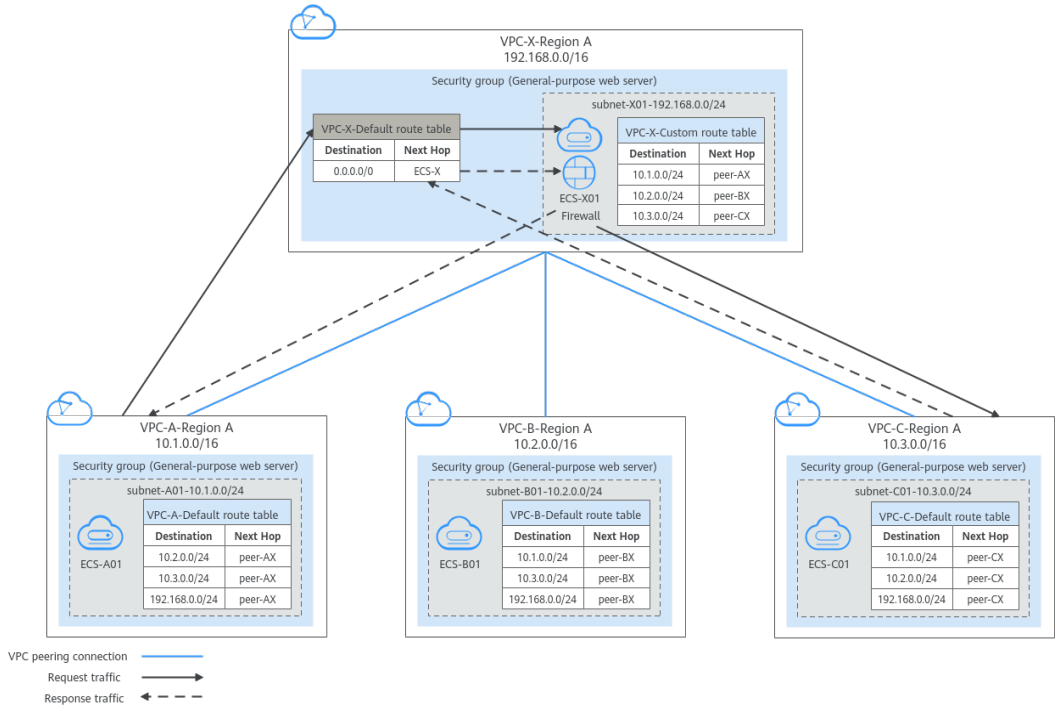

Resource Planning

In this example, you need to create VPCs, ECSs, and VPC peering connections. For details about required resources, see **Table 2-1**.

NOTE

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

 Copyright © Huawei Technologies Co., Ltd.

Table 2-1 Required resources

Resource	Description
VPC	Table 2-2 shows details about the required VPCs. In this example, there are four VPCs, including three VPCs where services are deployed and one VPC where the firewall is deployed. These VPCs are from the same region, and their subnet 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 with each 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 2-3 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 2-4 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.

 Copyright © Huawei Technologies Co., Ltd.

Table 2-2 VPC details

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 2-3 ECS details

ECS Name	VPC Name	Subnet Name	Private IP Address	Image	Security Group	ECS Is Used to Deploy
ecs-A01	VPC-A	subnet-A01	10.1.0.139	Public image: CentOS 8.2 64bit	sg-demo: General-purpose web server	Services
ecs-B01	VPC-B	subnet-B01	10.2.0.93			Services
ecs-C01	VPC-C	subnet-C01	10.3.0.220			Services
ecs-X01	VPC-X	subnet-X01	192.168.0.5			Firewall

Table 2-4 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 2-5**.

Copyright © Huawei Technologies Co., Ltd.

NOTE

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

Table 2-5 Required route tables

Route Table	Description
Route tables of service VPCs	**Table 2-6** shows details about route tables of service VPCs. The default route tables of VPC-A, VPC-B, and VPC-C have routes with destinations set to other VPC subnets and with next hop set to VPC peering connection.
Route tables of firewall VPC	**Table 2-7** 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. 2. 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.

 Copyright © Huawei Technologies Co., Ltd.

Table 2-6 Details about route tables of service VPCs

VPC Name	Route Table	Destination	Next Hop Type	Next Hop	Route Type	Route Function
VPC-A	Default route table: rtb-vpc-A	10.2.0.0/24	VPC peering connection	peer-AX	Custom	• Destination: subnet-B01 in VPC-B • Connects subnet-A01 to subnet-B01
		10.3.0.0/24	VPC peering connection	peer-AX	Custom	• Destination: subnet-C01 in VPC-C • Connects subnet-A01 to subnet-C01
		192.168.0.0/24	VPC peering connection	peer-AX	Custom	• Destination: subnet-X01 in VPC-X • Connects subnet-A01 to subnet-X01
VPC-B	Default route table: rtb-vpc-B	10.1.0.0/24	VPC peering connection	peer-BX	Custom	• Destination: subnet-A01 in VPC-A • Connects subnet-A01 to subnet-B01
		10.3.0.0/24	VPC peering connection	peer-BX	Custom	• Destination: subnet-C01 in VPC-C • Connects subnet-B01 to subnet-C01
		192.168.0.0/24	VPC peering connection	peer-BX	Custom	• Destination: subnet-X01 in VPC-X • Connects subnet-B01 to subnet-X01

 Copyright © Huawei Technologies Co., Ltd.

VPC Name	Route Table	Destination	Next Hop Type	Next Hop	Route Type	Route Function
VPC-C	Default route table: rtb-vpc-C	10.1.0.0/24	VPC peering connection	peer-CX	Custom	• Destination: subnet-A01 in VPC-A • Connects subnet-A01 to subnet-C01
		10.2.0.0/24	VPC peering connection	peer-CX	Custom	• Destination: subnet-B01 in VPC-B • Connects subnet-B01 to subnet-C01
		192.168.0.0/24	VPC peering connection	peer-CX	Custom	• Destination: subnet-X01 in VPC-X • Connects subnet-C01 to subnet-X01

Copyright © Huawei Technologies Co., Ltd.

Table 2-7 Details about route tables of firewall VPC

VPC Name	Route Table	Destination	Next Hop Type	Next Hop	Route Type	Route Function
VPC-X	Default route table: rtb-vpc-X	0.0.0.0/0	Server	ECS-X	Custom	• 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	10.1.0.0/24	VPC peering connection	peer-AX	Custom	• Destination: subnet-A01 in VPC-A • Connects subnet-A01 to subnet-X01
		10.2.0.0/24	VPC peering connection	peer-BX	Custom	• Destination: subnet-B01 in VPC-B • Connects subnet-B01 to subnet-X01
		10.3.0.0/24	VPC peering connection	peer-CX	Custom	• Destination: subnet-C01 in VPC-C • Connects subnet-C01 to subnet-X01

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

 Copyright © Huawei Technologies Co., Ltd.

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 2-2 Route Tables

  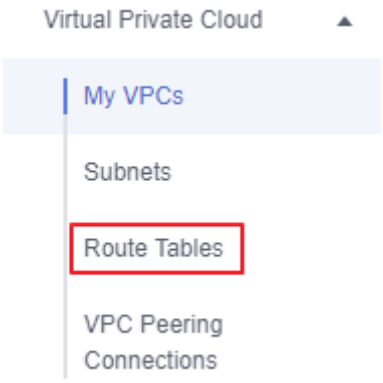

Procedure

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 2-2**.

Step 2 Create a custom route table in VPC-X and associate subnet-X01 with the custom route table.

1. 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 a Custom ECS**.

Step 4 Configure the NIC of ecs-X and install the third-party firewall on ecs-X.

1. Disable source/destination check for the NIC of ecs-X.

   a. In the ECS list, click the name of the target ECS.

      The ECS details page is displayed.

   b. On the **Network Interfaces** tab, click ⌄ to expand the details area and check whether **Source/Destination Check** is disabled.

 Copyright © Huawei Technologies Co., Ltd.

If the information shown in **Figure 2-3** is displayed, **Source/Destination Check** is disabled.

Figure 2-3 Disabling **Source/Destination Check**

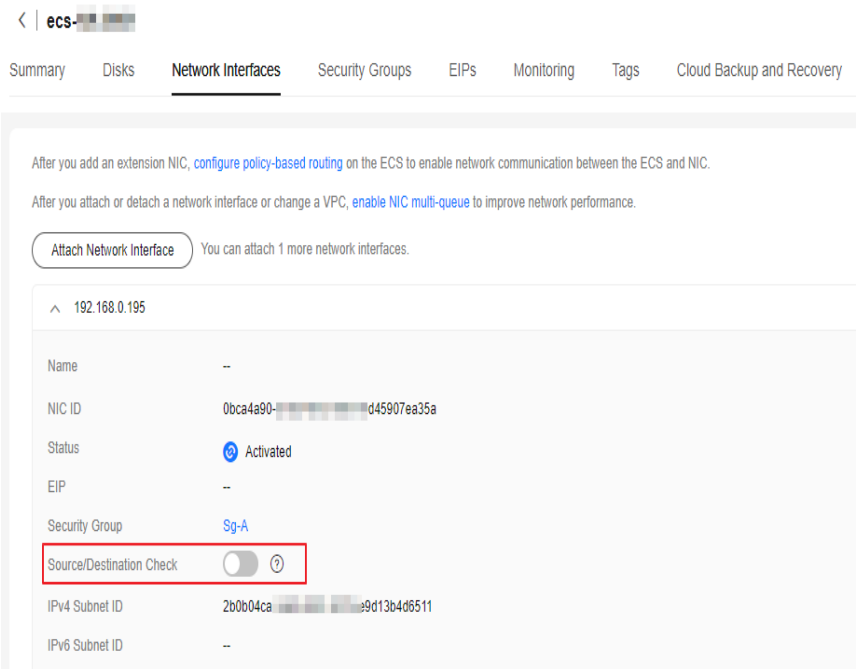

2. Install a third-party firewall on ecs-X.

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 2-4**.

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**.

 Copyright © Huawei Technologies Co., Ltd.

In this example, add the routes planned in **Table 2-6** 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 2-7** to the default and custom route tables of VPC-X.

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

For details, see **How Do I Log In to My 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
```

 Copyright © Huawei Technologies Co., Ltd.

4. Repeat **Step 7.1** to **Step 7.3** to check the communication between other VPCs.

----End

2.2 Using Third-Party Firewalls to Filter Traffic 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 on-premises 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 2-4 Deployment diagram

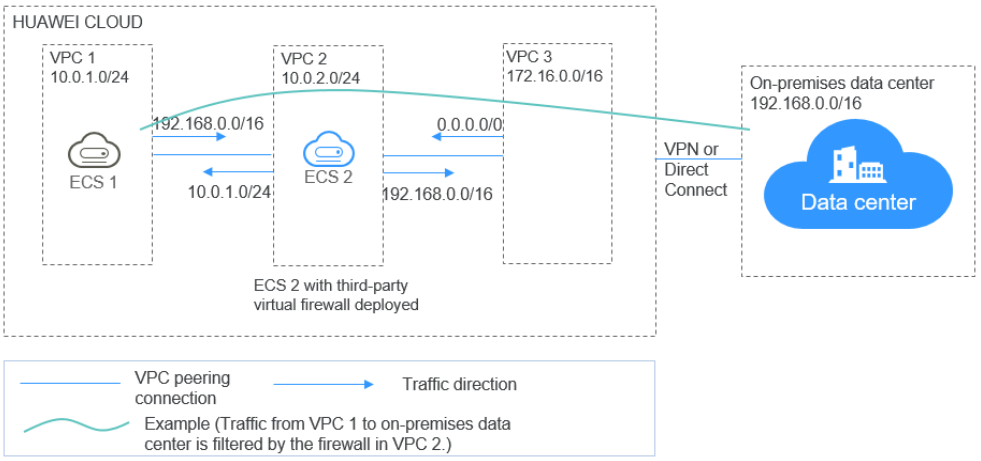

 Copyright © Huawei Technologies Co., Ltd.

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 a Custom 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**.

 Copyright © Huawei Technologies Co., Ltd.

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 0.0.0.0/0, 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.

 Copyright © Huawei Technologies Co., Ltd.

3 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.

 Copyright © Huawei Technologies Co., Ltd.

Figure 3-1 Network topology

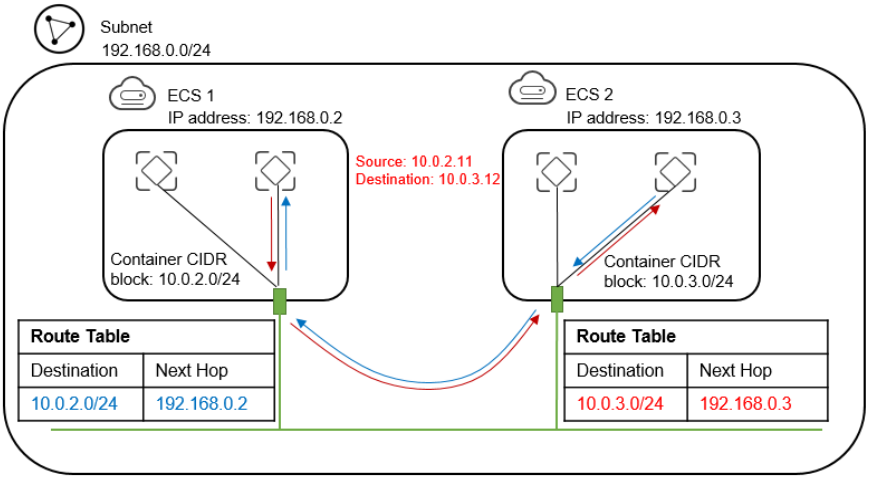

Procedure

Step 1 Create VPCs.

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

Step 2 Create ECSs.

For details, see **Purchasing a Custom ECS**.

After the ECS is created, disable source/destination check on the ECS NIC, as shown in **Figure 3-2**.

Figure 3-2 Disabling source/destination check

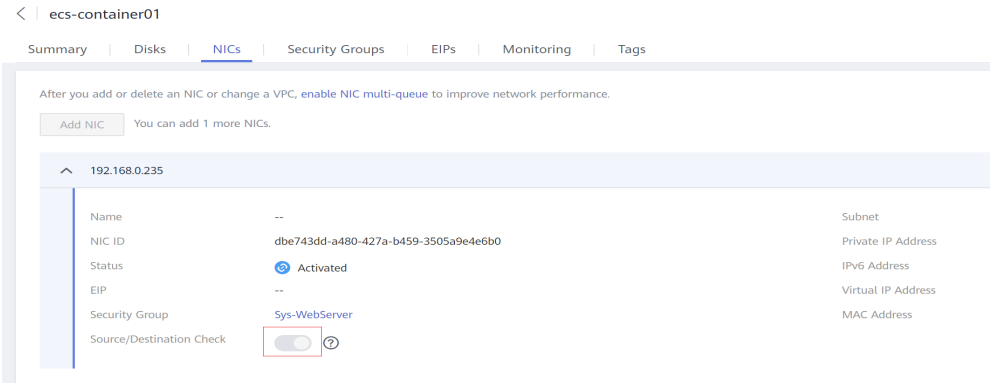

Step 3 Deploy containers on ECSs.

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.

 Copyright © Huawei Technologies Co., Ltd.

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 3-1** to the security group of the ECSs to allow ICMP and UDP traffic.

For details, see **Adding a Security Group Rule**.

Table 3-1 Security group rules

Direction	Protocol	Port	Source
Inbound	ICMP	All	0.0.0.0/0
Inbound	UDP	All	0.0.0.0/0

----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 ->

 Copyright © Huawei Technologies Co., Ltd.

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=2 ttl=62 time=0.304 ms
64 bytes from 10.0.3.2: seq=3 ttl=62 time=0.319 ms
```

 Copyright © Huawei Technologies Co., Ltd.