

# Virtual Private Cloud

## Best Practices

Issue 47  
Date 2020-12-17



**Copyright © Huawei Technologies Co., Ltd. 2021. All rights reserved.**

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

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

## **Notice**

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

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

---

# Contents

---

|   |           |
|---|-----------|
| <b>1 Network Planning.....</b>  | <b>1</b>  |
| <b>2 VPC Connectivity.....</b>  | <b>5</b>  |
| <b>3 Private Network Access.....</b>  | <b>9</b>  |
| <b>4 Public Network Access.....</b>   | <b>13</b> |
| <b>5 Lower Network Costs.....</b>   | <b>18</b> |
| <b>6 Access Control.....</b>  | <b>20</b> |
| <b>7 Using Third-Party Firewalls When Connecting VPCs.....</b>  | <b>24</b> |
| <b>8 Using Third-Party Firewalls When Connecting an On-premises Data Center to the Cloud.....</b>       | <b>30</b> |
| <b>9 Deploying Containers That Can Communicate With Each Other on ECSs.....</b>                         | <b>35</b> |
| <b>10 Creating a L2CG for a Direct Connect Connection to Migrate Services at a Layer 2 Network.....</b> | <b>39</b> |
| <b>11 Building Highly Available Web Server Clusters with Keepalived.....</b>                            | <b>48</b> |
| <b>12 Using IP Address Groups to Reduce the Number of Security Group Rules.....</b>                     | <b>58</b> |

# 1 Network Planning

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?

VPCs are region-specific. By default, networks in VPCs in different regions or even in the same region are not connected. The communications on these different networks are completely isolated from each other, this is not the case for different AZs in the same VPC. Two networks in the same VPC should be able to communicate with each other even if they are in different AZs.


### 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, you can use a VPC peering connection as shown in [Figure 1-1](#).

**Figure 1-1** VPC peering connection



### Default VPC Quota

By default, you can create a maximum of five VPCs in your account. If this cannot meet your service requirements, request a quota increase. For details, see [What Is a Quota?](#)

## How Do I Plan Subnets?

A subnet is a range of IP addresses in your VPC. All of the resources in a VPC must be deployed on subnets and the subnets on a VPC cannot overlap. Once a subnet has been created, its CIDR block cannot be modified.

The subnets used to deploy your resources must reside within your VPC, and the subnet masks used to define them can be between the netmask of its VPC CIDR block and /29 netmask.

- 10.0.0.0 – 10.255.255.255
- 172.16.0.0 – 172.31.255.255
- 192.168.0.0 – 192.168.255.255

### Subnet Planning

- We recommend that you create different subnets for different service modules in a VPC. For example, you can create different subnets for web, application, and database servers. A web server is in a publicly accessible subnet, and application and database servers are in non-publicly accessible subnets. You can leverage network ACLs to help control access to the servers in each subnet.
- If you only need to plan subnets for VPCs, and communication between VPCs and on-premises data centers are not required, you can create subnets within any of the CIDR blocks listed above.
- If your VPC needs to communicate with an on-premises data center through VPN or Direct Connect, the VPC CIDR block cannot overlap with the CIDR block of the on-premises data center. Therefore, when creating a VPC or subnet, ensure that its CIDR block does not overlap with any CIDR block on the data center.
- When determining the size of a VPC or subnet CIDR block, ensure that the number of available IP addresses on the CIDR block meet your service requirements.

### Default Subnet Quota

By default, you can create up to 100 subnets in your account. If you need more, submit a service ticket to request a quota increase. For details, see [What Is a Quota?](#)

## How Do I Plan Routing Policies?

A route table contains a set of routes that are used to control where inbound and outbound subnet traffic is forwarded within a VPC. When you create a VPC, it automatically has a default route table, which enables internal communication within that VPC.

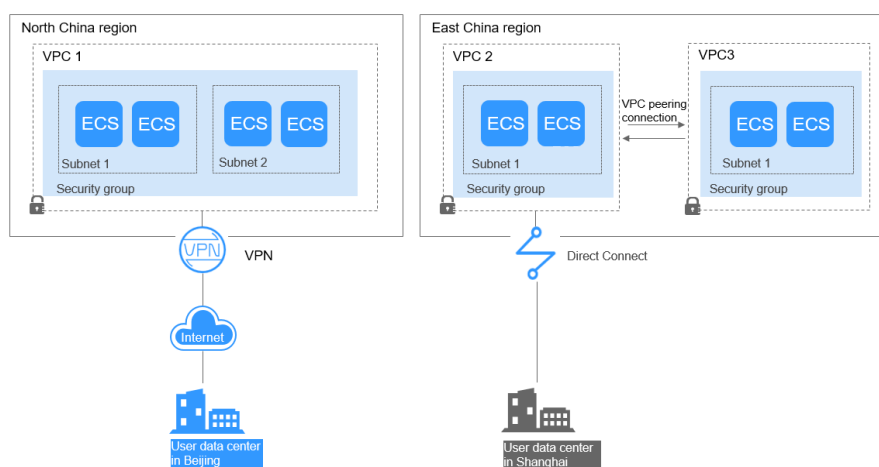
- If you do not need to explicitly control how each subnet routes inbound and outbound traffic, you can use the default route table.
- If you need to explicitly control how each subnet routes inbound and outbound traffic in a VPC, you can add custom routes to the route table.

## 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 on-premises data center to be connected.

In [Figure 1-2](#), VPC 1 is in North China region and VPC 2 and VPC 3 are in East China region. To connect to an on-premises data center, they can use a VPN, as VPC 1 does in Beijing; or a Direct Connect connection, as VPC 2 does in Shanghai. 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 1-2** 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 on-premises data center in Beijing.
- The CIDR block of VPC 2 cannot overlap with the CIDR block of the on-premises data center in Shanghai.
- 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. Different EIPs can use the same shared bandwidth, reducing your bandwidth costs.

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 system 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 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.

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.

**Table 2-1** Accessing the Internet

| Cloud Service | Application Scenario              | Description  | Reference   |
|---------------|-----------------------------------|--|---|
| EIP           | Single ECS accesses the Internet. | <p>An EIP is a static IP address that can be directly accessed through the Internet or provide services accessible from the Internet.</p> <p>An EIP can be bound to an ECS to enable Internet access, or unbound to disable access.</p> <p>Shared bandwidth and shared data packages can be used to lower costs.</p> | <a href="#">Configuring the VPC of ECSs That Access the Internet Using EIPs</a> |

| 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. | <a href="#">Using SNAT to Access the Internet</a><br><a href="#">Using DNAT to Provide Services Accessible from the Internet</a> |
| 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. | <p>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.</p> <p>ELB expands the service capabilities of your applications and improves availability by eliminating single points of failures.</p>   | <a href="#">What Is Elastic Load Balance?</a>  |

## Connecting VPCs

You can connect VPCs using the following cloud services.

**Table 2-2** Connecting VPCs

| 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.   | <a href="#">Creating a VPC Peering Connection with Another VPC in Your Account</a><br><a href="#">Creating a VPC Peering Connection with a VPC in Another Account</a> |
| Cloud Connect | Connect VPCs in different regions.                    | Cloud Connect allows you to connect two VPCs in the same account or in different accounts even if they are in different regions.  | <a href="#">Communication Between VPCs Across Regions</a>   |
| 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. | <a href="#">Connecting to a VPC Through a VPN</a>   |

## Connecting to an On-premises Data Center (IDC)

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

**Table 2-3** Connecting to an IDC

| Cloud Service | Application Scenario                                     | Description   | Reference   |
|---------------|--|---|---|
| VPN           | Use VPN to connect a VPC to a local IDC with a low cost. | VPN uses an encrypted communications tunnel to connect a VPC on the cloud to a local IDC 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. | <a href="#">Connecting to a VPC Through a VPN</a> |

| Cloud Service  | Application Scenario   | Description  | Reference  |
|----------------|--|--|--|
| Direct Connect | Use a physical dedicated connection to connect a VPC to a local IDC. | 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. | <a href="#">Accessing Multiple VPCs Using a Connection</a>                       |
| Cloud Connect  | Connect VPCs in different regions.                                   | Cloud Connect allows the loading of Direct Connect virtual gateways to a Cloud Connect connection, interconnecting an on-premises data center with VPCs across regions.  | <a href="#">Communication Between Data Centers and VPCs in Different Regions</a> |

# 3 Private Network Access

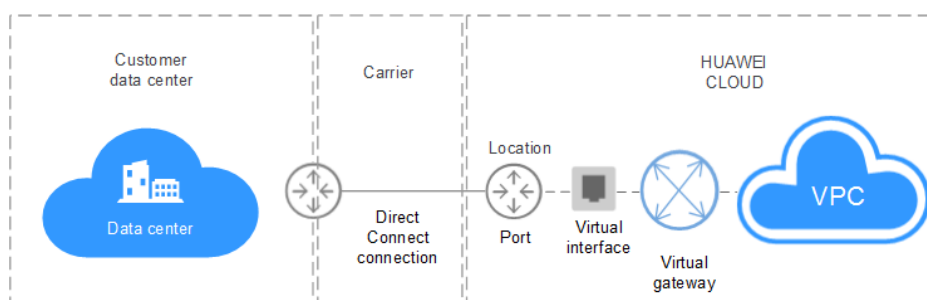
## Connecting to an On-premises Data Center

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 Cloud 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

With Cloud Connect, you can connect VPCs across regions and VPCs and data centers off the cloud.

Cloud Connect allows you to quickly build high-quality networks that are both fast and stable. With Cloud Connect, you can link VPCs across regions and between 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 3-2** 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.

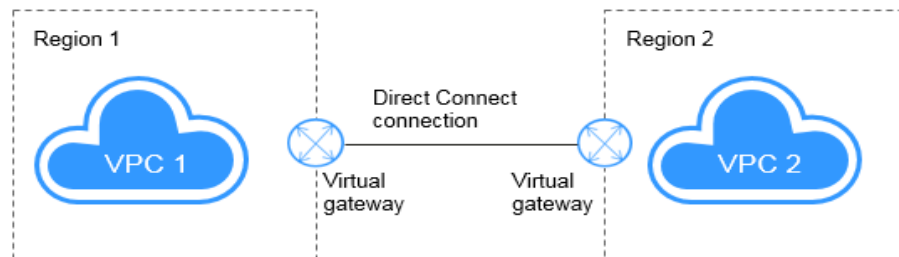
- VPC peering  
You can use VPC peering connections to connect VPCs in the same region.

**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.

**Figure 3-5** Connecting VPCs in different regions with VPN



- **Cloud Connect**  
Cloud Connect allows you to quickly build high-quality networks that are both fast and stable. With Cloud Connect, you can link VPCs across regions and between 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


The public cloud provides EIP, NAT Gateway, and ELB services 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 resources 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.

**Figure 4-1** EIP



- 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.

Figure 4-2 ELB




## Accessing the Internet

- Single ECS accesses the Internet.

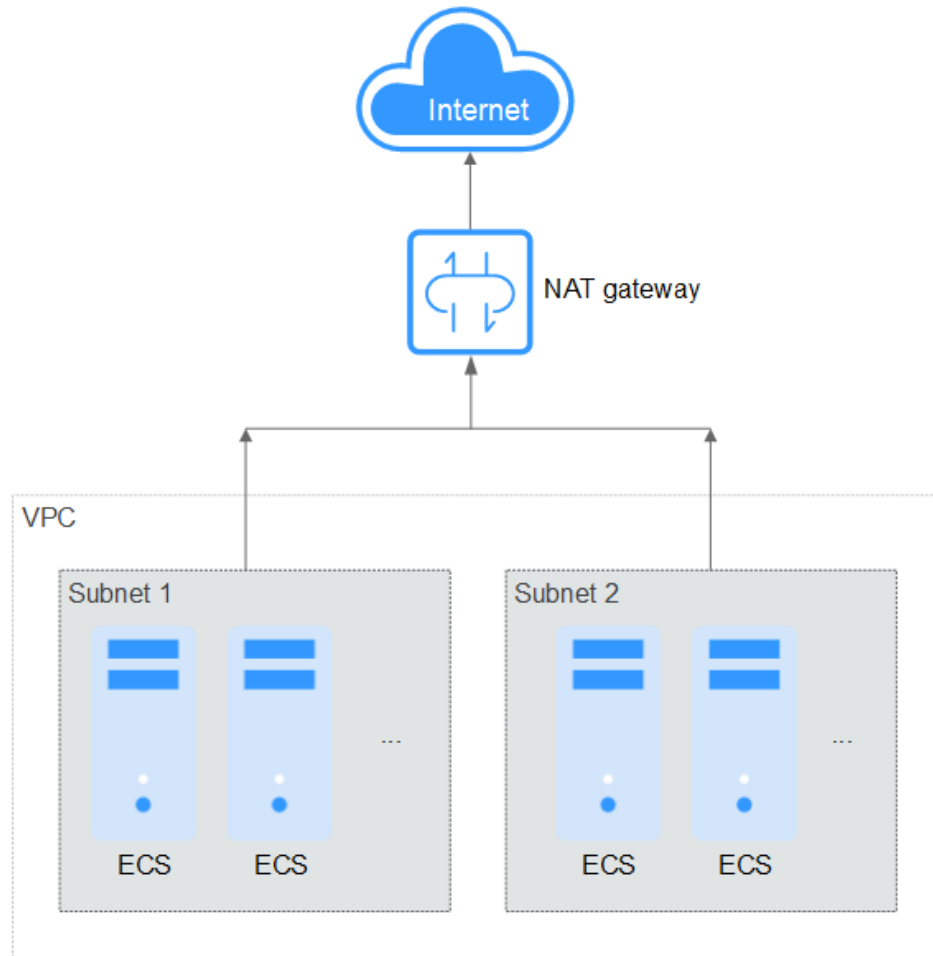
When an ECS needs to access the Internet, you can bind an EIP to the ECS so that the ECS can access the Internet. HUAWEI CLOUD allows your EIP to be billed based on bandwidth usage or amount of traffic. If you do not need to use the EIP, you can flexibly unbind it.

**Figure 4-3** EIP



- 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.

A shared bandwidth can be billed by bandwidth or by 95th percentile bandwidth:

- **Bandwidth:** If you use a large number of EIPs and their peak hours are different, use shared bandwidth to greatly reduce costs.
- **95th percentile bandwidth (enhanced):** If your services frequently reach peaks, you can select this option. This ensures that the service system is not affected by the bandwidth limit at service peaks and avoids the cost waste associated with excessive peak bandwidth peaks.

## 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 using bandwidth billed by traffic in a region and buy a shared data package in the same region, the EIPs will use the shared data package.

- 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.
  - A shared data package will not be renewed automatically. If you are uncertain which package to purchase, buy a small one the first time.
  - 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 Access Control

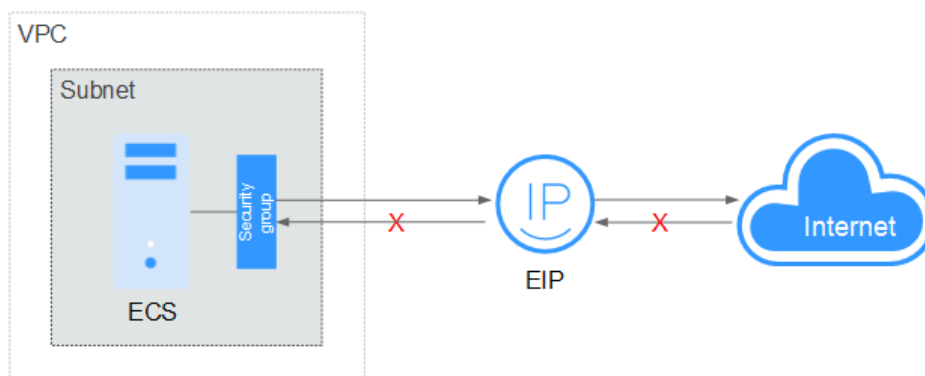
Access control can be managed at the ECS level, the subnet level, or based on services, by using security groups, network ACLs, and whitelists, respectively.

- **Security group: ECS-based access control**  
A security group is a logical group that controls the traffic for one or more ECSs. After a security group is created, you can add rules that control the inbound traffic to ECSs that it contains.
- **Network ACL: Subnet-based access control**  
Network ACLs control traffic in and out of one or more subnets based on priorities. Only packet filtering based on the 5-tuple (protocol, source port, destination port, source IP address, and destination IP address) is supported.
- **Whitelist: Service-based access control**  
Whitelist controls traffic from services (such as ELB, OBS) that use VPC subnet resources.

## Scenario 1: Only Allowing Access to the Internet

An ECS bound with an EIP can access the Internet but cannot be accessed from the Internet with the protection of security group rules.

**Figure 6-1** Only allowing outbound traffic



### Configuration example

Inbound direction of a security group: No rules are added.

Outbound direction of a security group: All protocols and ports are allowed, as shown in [Table 6-1](#).

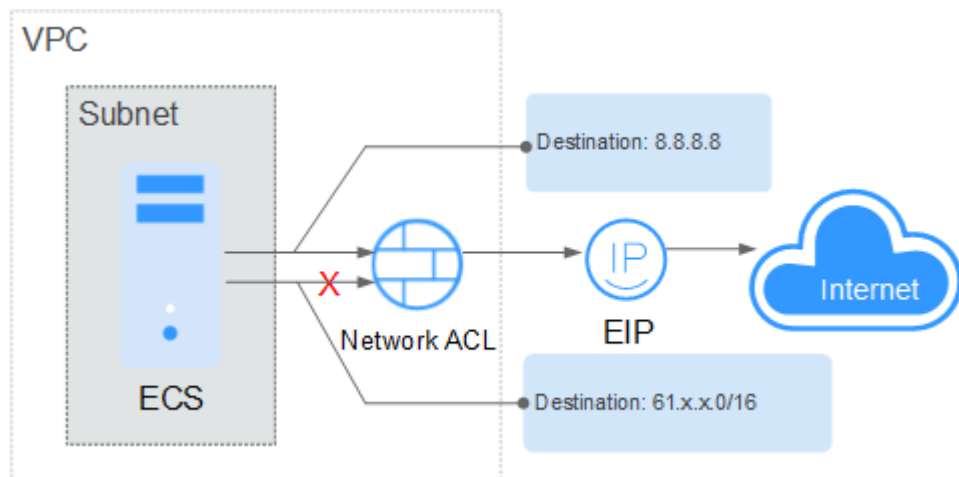
**Table 6-1** Security group rule

| Direction | Protocol / Application | Port | Destination | Description                            |
|-----------|------------------------|------|-------------|--|
| Outbound  | All                    | All  | 0.0.0.0/0   | Allows all outbound traffic. (default) |

## Scenario 2: Denying Access to Specific IP Addresses

Network ACLs can deny the access from all ECSs in a subnet to specific IP addresses. For example, [Figure 6-2](#) shows that ECSs in the subnet are denied to access 61.x.x.0/16.

**Figure 6-2** Denying access to specific IP addresses



### Configuration example

**Table 6-2** Network ACL rule

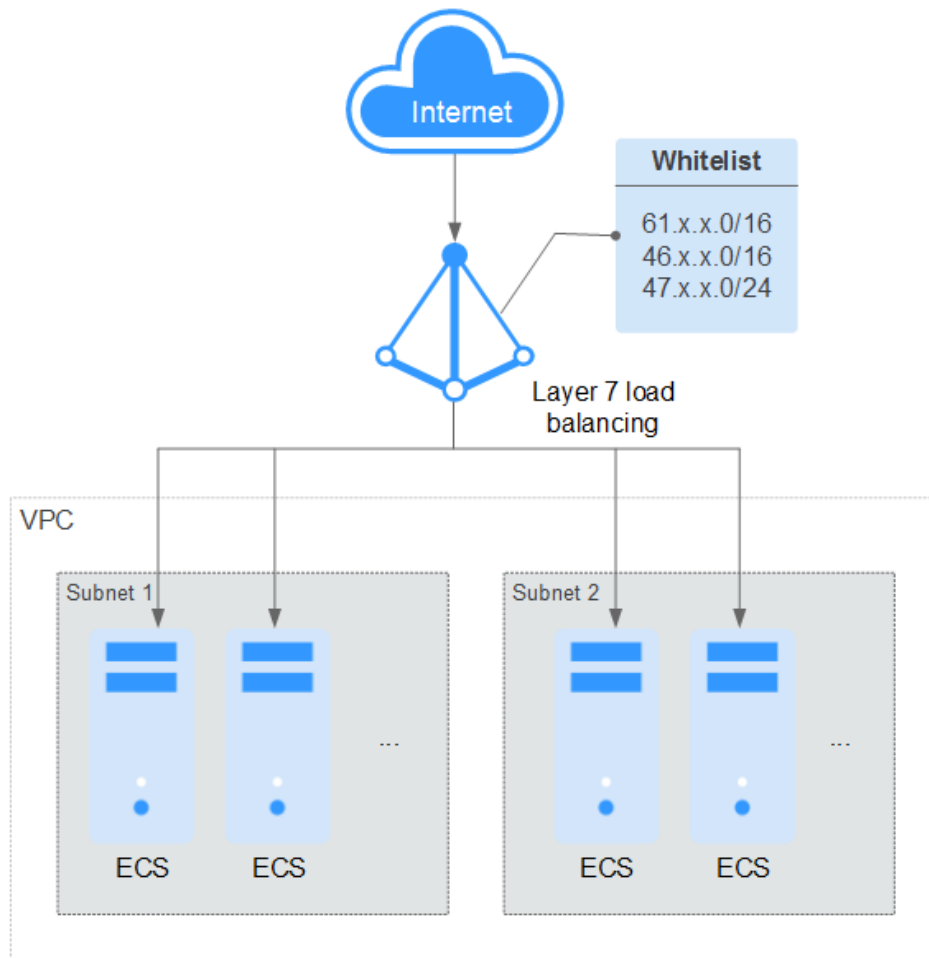
| Direction | Action | Protocol | Source    | Source Port Range | Destination | Destination Port Range | Description                 |
|-----------|--------|----------|-----------|-------------------|-------------|------------------------|-----------------------------|
| Inbound   | Allow  | All      | 0.0.0.0/0 | All               | 0.0.0.0/0   | All                    | Allows all inbound traffic. |

| Direction | Action | Protocol | Source    | Source Port Range | Destination | Destination Port Range | Description                                 |
|-----------|--------|----------|-----------|-------------------|-------------|------------------------|---|
| Inbound   | Deny   | All      | 0.0.0.0/0 | All               | 0.0.0.0/0   | All                    | Denies all inbound traffic. (default)       |
| Outbound  | Deny   | All      | 0.0.0.0/0 | All               | 61.x.x.0/16 | All                    | Denies the outbound traffic to 61.x.x.0/16. |
| Outbound  | Allow  | All      | 0.0.0.0/0 | All               | 0.0.0.0/0   | All                    | Allows all outbound traffic.                |
| Outbound  | Deny   | All      | 0.0.0.0/0 | All               | 0.0.0.0/0   | All                    | Denies all outbound traffic. (default)      |

### Scenario 3: Layer 7 Load Balancing

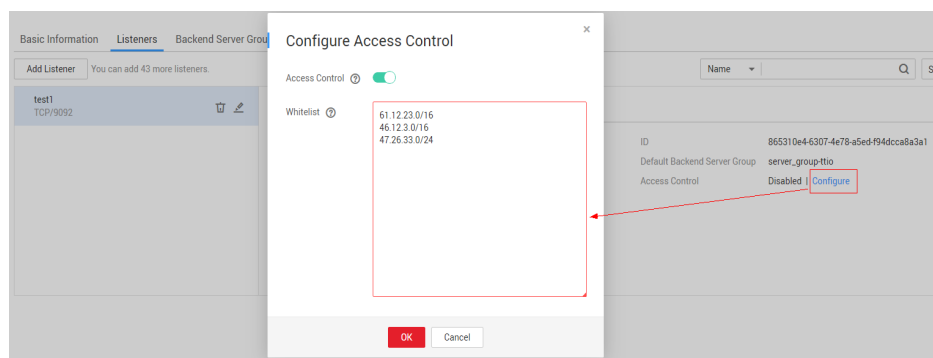
Layer 7 load balancing may be used internally and the access source can be controlled through the whitelist.

Figure 6-3 ELB whitelist



### Configuration example

Figure 6-4 Configuring a whitelist



# 7 Using Third-Party Firewalls When Connecting VPCs

---

## Scenarios

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

This section describes how to use a third party firewall when connecting multiple VPCs in the same region.

## Solution Advantages

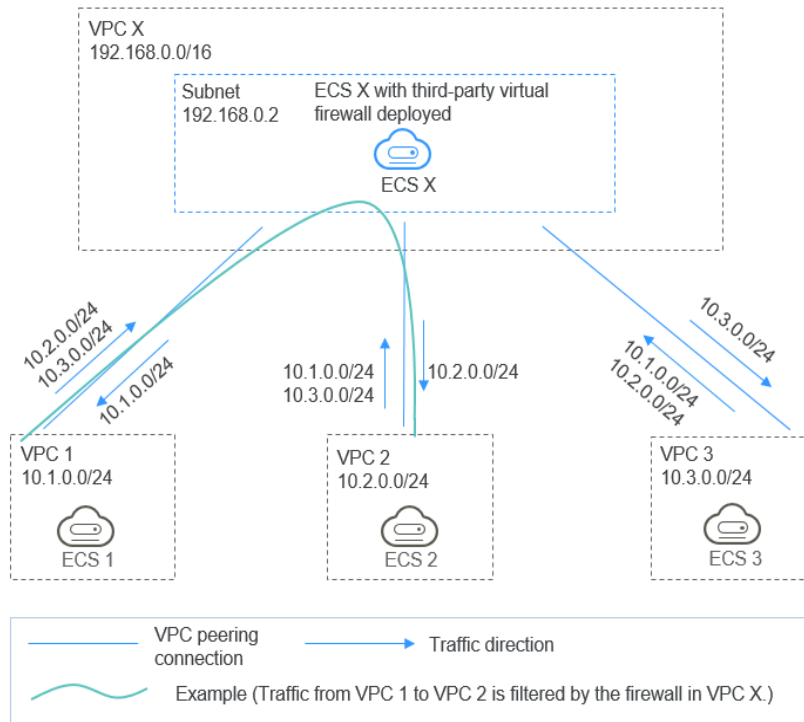
- You can use third-party firewalls.
- You can define security rules as required.

## Typical Topology

Assume that your services are deployed in VPC 1, VPC 2, VPC 3, and VPC X, and you need to use a third-party virtual firewall on the cloud. You can configure the virtual firewall on ECS X in VPC X and use VPC peering connections and configure routes to enable communication between the VPCs. The ECS X with the firewall deployed in VPC X filters incoming and outgoing data based on the firewall rules.

The deployment diagram is as follows:

Figure 7-1 Deployment diagram



## Prerequisites

- The subnet of the ECS with the third-party virtual firewall deployed has been associated with a route table. Ensure that the region you selected allows you to visit the route table module directly from the navigation pane on the left of the network console.
- The subnet CIDR block of VPC X does not overlap with these of VPC 1, VPC 2, and VPC 3. Otherwise, traffic cannot go through firewall on the ECS.

## Procedure

### Step 1 Create VPCs.

Create VPC 1, VPC 2, VPC 3, and VPC X.

For details, see [Creating a VPC](#).

#### NOTE

The CIDR blocks of VPC 1, VPC 2, VPC 3, and VPC X cannot overlap with each other. For example, the CIDR block of VPC 1 is 10.1.0.0/24, VPC 2 is 10.2.0.0/24, VPC 3 is 10.3.0.0/24, and VPC X is 192.168.0.0/16.

### Step 2 Create ECSs.

1. Create ECS 1, ECS 2, ECS 3, and ECS X that belong to VPC 1, VPC 2, VPC 3, and VPC X, respectively.

For details, see [Purchasing an ECS](#).

 **NOTE**

Disable the source/destination check for the ECS X NIC. For details, see [Disabling Source and Destination Check](#).

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

**Step 3 Create VPC peering connections.**

Create VPC peering connections between VPC 1 and VPC X, VPC 2 and VPC X, and VPC 3 and VPC X 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 [Step 6](#).

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 X 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 the ECS X.**

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 ECS X.

For details, see [Binding a Virtual IP Address to an EIP or ECS](#).

**Step 6 Configure routes.**

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

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

- a. Add a route to forward traffic from VPC 1 to VPC X, set the destination of the route to the CIDR block of VPC X, and the next hop of the route to the VPC peering connection between VPC 1 and VPC X.
- b. Add a route to forward traffic from VPC 1 to VPC 2, set the destination of the route to the CIDR block of VPC 2, and the next hop of the route to the VPC peering connection between VPC 1 and VPC X.
- c. Add a route to forward traffic from VPC 1 to VPC 3, set the destination of the route to the CIDR block of VPC 3, and the next hop of the route to the VPC peering connection between VPC 1 and VPC X.

[Figure 7-2](#) is for reference.

**Figure 7-2** Routes in the default route table of VPC 1

| Destination | Next Hop Type        | Next Hop  | Type   | Description                          | Operation       |
|-------------|----------------------|-----------|--------|--------------------------------------|-----------------|
| Local       | Local                | Local     | System | Default route that enables instan... | Modify   Delete |
| 10.2.0.0/24 | VPC peering conne... | VPC1-VPCX | Custom | --                                   | Modify   Delete |
| 10.3.0.0/24 | VPC peering conne... | VPC1-VPCX | Custom | --                                   | Modify   Delete |

2. Add the following routes to the default route table of VPC 2:
  - a. Add a route to forward traffic from VPC 2 to VPC X, set the destination of the route to the CIDR block of VPC X, and the next hop of the route to the VPC peering connection between VPC 2 and VPC X.
  - b. 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 2 and VPC X.
  - c. Add a route to forward traffic from VPC 2 to VPC 3, set the destination of the route to the CIDR block of VPC 3, and the next hop of the route to the VPC peering connection between VPC 2 and VPC X.

**Figure 7-3** is for reference.

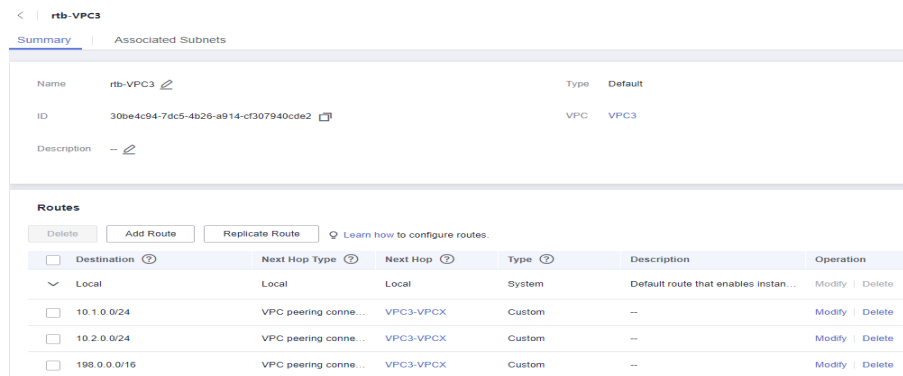
**Figure 7-3** Routes in the default route table of VPC 2

| Destination | Next Hop Type        | Next Hop  | Type   | Description                          | Operation       |
|-------------|----------------------|-----------|--------|--------------------------------------|-----------------|
| Local       | Local                | Local     | System | Default route that enables instan... | Modify   Delete |
| 10.1.0.0/24 | VPC peering conne... | VPC2-VPCX | Custom | --                                   | Modify   Delete |
| 10.3.0.0/24 | VPC peering conne... | VPC2-VPCX | Custom | --                                   | Modify   Delete |

3. Add the following routes to the default route table of VPC 3:
  - a. Add a route to forward traffic from VPC 3 to VPC X, set the destination of the route to the CIDR block of VPC X, and the next hop of the route to the VPC peering connection between VPC 3 and VPC X.
  - b. Add a route to forward traffic from VPC 3 to VPC 2, set the destination of the route to the CIDR block of VPC 2, and the next hop of the route to the VPC peering connection between VPC 3 and VPC X.
  - c. Add a route to forward traffic from VPC 3 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 3 and VPC X.

**Figure 7-4** is for reference.

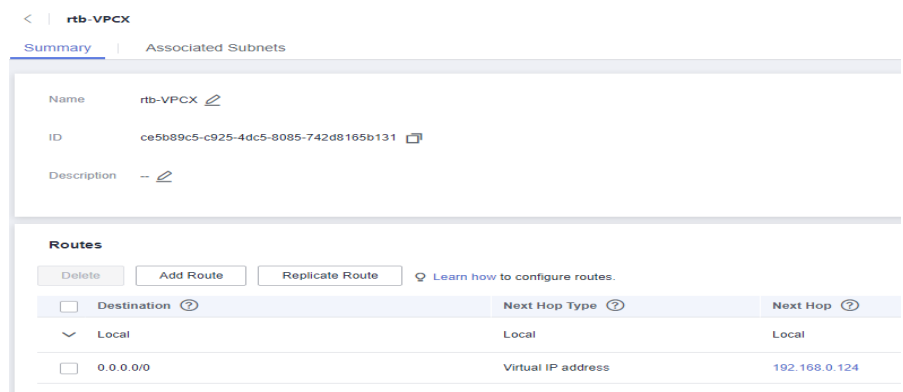
**Figure 7-4** Routes in the default route table of VPC 3



4. Add the following route to the default route table of VPC X:
    - a. Set the destination of the route to 0.0.0.0/0, and the next hop of the route to ECS X.
- 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.

**Figure 7-5** is for reference.

**Figure 7-5** Routes in the default route table of VPC X



5. Add the following routes to the route table of VPC X subnet:
  - a. Add a route to forward traffic from VPC X 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 X.
  - b. Add a route to forward traffic from VPC X to VPC 2, set the destination of the route to the CIDR block of VPC 2, and the next hop of the route to the VPC peering connection between VPC 2 and VPC X.
  - c. Add a route to forward traffic from VPC X to VPC 3, set the destination of the route to the CIDR block of VPC 3, and the next hop of the route to the VPC peering connection between VPC 3 and VPC X.

**Figure 7-6** is for reference.

**Figure 7-6** Routes in the route table of VPC X subnet

< | rtb-VPCX-subnet

Summary | Associated Subnets

Name rtb-VPCX-subnet [✎](#)

ID 8eef4b1f-127a-4285-ad8d-c6ac62bc8284 [📄](#)

Description -- [✎](#)

**Routes**

[Delete](#) [Add Route](#) [Replicate Route](#) [🔗 Learn how to configure routes.](#)

| <input type="checkbox"/> | Destination <a href="#">?</a> | Next Hop Type <a href="#">?</a> | Next Hop <a href="#">?</a> |
|--------------------------|-------------------------------|---------------------------------|----------------------------|
| ▼                        | Local                         | Local                           | Local                      |
| <input type="checkbox"/> | 10.1.0.0/24                   | VPC peering connection          | VPC1-VPCX                  |
| <input type="checkbox"/> | 10.2.0.0/24                   | VPC peering connection          | VPC2-VPCX                  |
| <input type="checkbox"/> | 10.3.0.0/24                   | VPC peering connection          | VPC3-VPCX                  |

----End

## Verification

Log in to ECS 1 and then access ECS 2 from ECS 1. Check whether ECS X can receive packets that are sent from ECS 1 to ECS 2. Check whether the packets pass through and are filtered by the firewall on ECS X.

# 8 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 HUAWEI 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.

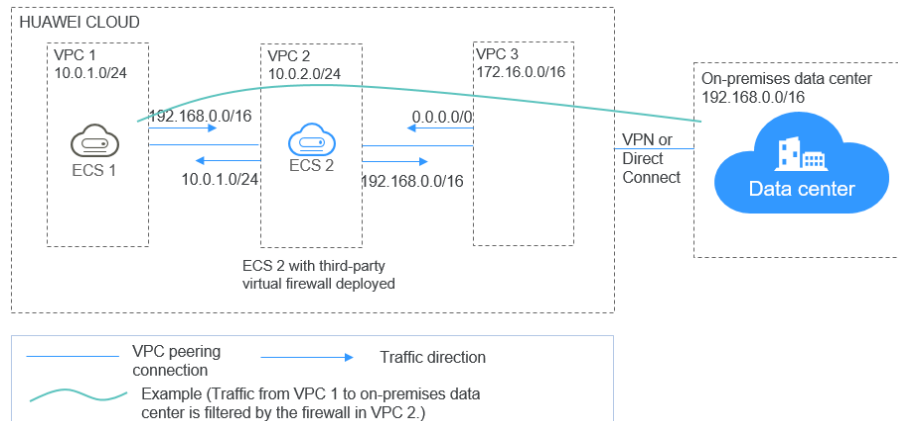
## Solution Advantages

- 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 8-1** 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 [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.

[Figure 8-2](#) is for reference.

**Figure 8-2** Routes in the default route table of VPC 1

| Name                                 | Type     |
|--------------------------------------|----------|
| rtb-VPC1                             | Default  |
| 5e1ee816-ae33-4236-9f23-b4938c70f04f | VPC VPC1 |

| Destination    | Next Hop Type          | Next Hop  | Type   | Description                             |
|----------------|------------------------|-----------|--------|---|
| Local          | Local                  | Local     | System | Default route that enables instance ... |
| 192.168.0.0/16 | VPC peering connection | VPC1-VPC2 | Custom | --                                      |

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.

[Figure 8-3](#) is for reference.

**Figure 8-3** Routes in the default route table of VPC 2

| Name     | Type | Default |
|----------|------|---------|
| rtb-vpc2 | VPC  | VPC2    |

| Destination | Next Hop Type | Next Hop | Type   |
|-------------|---------------|----------|--------|
| Local       | Local         | Local    | System |
| 0.0.0.0/0   | Server        | ecs2     | Custom |

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.

**Figure 8-4** is for reference.

**Figure 8-4** Routes in the route table of VPC 2 subnet

| Name         | Type | Custom Route Table |
|--------------|------|--------------------|
| rtb-vpc2-sub | VPC  | VPC2               |

| Destination    | Next Hop Type          | Next Hop  | Type   | Description                             |
|----------------|------------------------|-----------|--------|---|
| Local          | Local                  | Local     | System | Default route that enables instance ... |
| 10.0.1.0/24    | VPC peering connection | VPC1-VPC2 | Custom | --                                      |
| 192.168.0.0/16 | VPC peering connection | VPC2-VPC3 | Custom | --                                      |

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.

**Figure 8-5** is for reference.

**Figure 8-5** Routes in the default route table of VPC 3

| Name     | Type | Default |
|----------|------|---------|
| rtb-VPC3 | VPC  | VPC3    |

| Destination    | Next Hop Type          | Next Hop    | Type   |
|----------------|------------------------|-------------|--------|
| Local          | Local                  | Local       | System |
| 192.168.0.0/24 | Direct Connect gateway | vgw-0704-01 | System |
| 0.0.0.0/0      | VPC peering connection | VPC2-VPC3   | Custom |

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.

# 9 Deploying Containers That Can Communicate With Each Other on ECSs

---

## Scenarios

If you do not use HUAWEI CLOUD container products, you can deploy containers on HUAWEI CLOUD ECSs and enable containers on different ECSs in the same subnet to communicate with each other.

## Solution Advantages

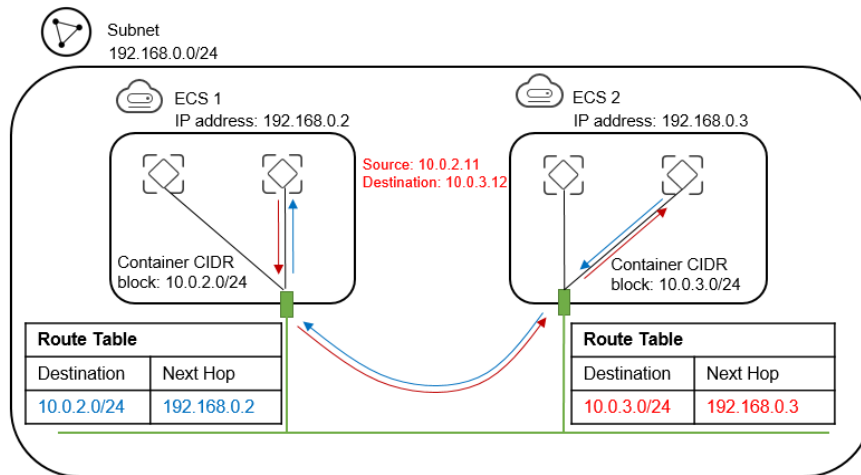
- Containers are deployed on ECSs, use CIDR blocks that are different from those of VPCs to which the ECSs belong, and use routes added to VPC route tables for data forwarding.
- You only need to add routes to the route tables to allow communications among containers, which is flexible and convenient.

## 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 different from those of the VPC subnet to which the ECSs belong. Containers on the same ECS are on the same CIDR block, but containers on different ECSs use 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 located. 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 the next hop of the packets sent to CIDR block 10.0.3.0/24 is 192.168.0.3.

**Figure 9-1** Network topology



## Procedure

### Step 1 Create VPCs.

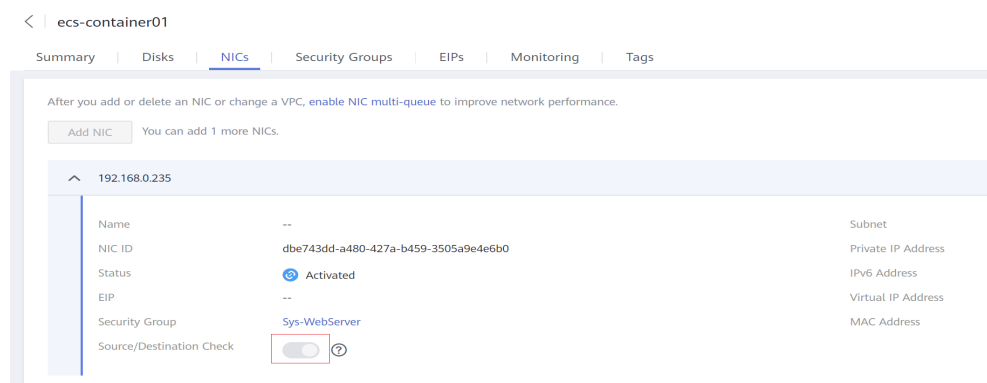
For details, see [Creating a VPC](#).

### 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 9-2](#).

**Figure 9-2** Disabling source/destination check



### Step 3 Deploy containers on an ECS.

Containers on the same ECS must be on the same CIDR block and the CIDR blocks of containers on different ECS 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.

**Figure 9-3** Adding routes

< | vpc1

|                    |                                      |            |                |
|--------------------|--------------------------------------|------------|----------------|
| Name               | vpc1                                 | Status     | Available      |
| ID                 | 3d9e9edc-3bb3-416f-8693-54da2d706865 | CIDR Block | 192.168.0.0/16 |
| Subnets            | 1                                    |            |                |
| Enterprise Project | default                              |            |                |

Subnets | **Route Tables** | Topology | Tags

Custom Route Table

Add Route

| Destination | Next Hop    | Operation     |
|-------------|-------------|---------------|
| 10.0.2.0/24 | 192.168.0.2 | Modify Delete |
| 10.0.3.0/24 | 192.168.0.3 | Modify Delete |

**NOTE**

- By default, a single 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 new route to the VPC route table.

**----End****Verification**

Use the **ping** command to check whether the containers deployed on two 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 the container on ECS 1 to 10.0.2.0/24, and create a 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=2 ttl=62 time=0.304 ms
64 bytes from 10.0.3.2: seq=3 ttl=62 time=0.319 ms
```

# 10 Creating a L2CG for a Direct Connect Connection to Migrate Services at a Layer 2 Network

---

## Scenario

Layer 2 connection gateways (L2CGs) allow communication between HUAWEI CLOUD and on-premises data centers at a Layer 2 network, while Direct Connect allows communication at Layer 3. L2CGs and switches perform VXLAN tunnel encapsulation.

If your on-premises data center can communicate with HUAWEI CLOUD at Layer 3 through Direct Connect and want to migrate services to the cloud at a Layer 2 network, you can use a L2CG.

Requirements:

- Servers migrated to the cloud can communicate with on-premises servers at Layer 2.
- The Layer 2 network extended to the cloud can communicate with the on-premises data center at Layer 3.
- The Layer 2 network reconstructed at the on-premises data center can communicate with the cloud at Layer 3.

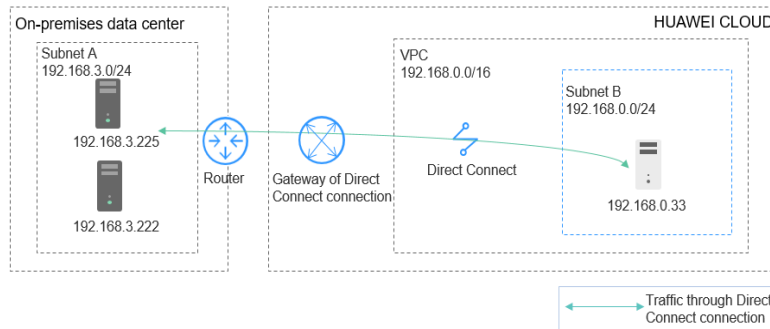
### NOTE

L2CGs are currently available for open beta test in **CN East-Shanghai1** and **CN South-Guangzhou**. You can use this function after obtaining the open beta test permissions.

## Typical Topology

In [Figure 10-1](#), the subnet A in the on-premises data center connects to subnet B on the cloud through Direct Connect.

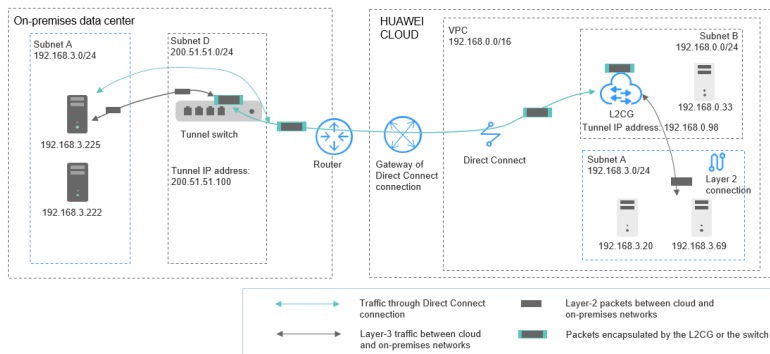
**Figure 10-1** Network topology



If you want to migrate subnet A to the cloud at a Layer 2 network, add a subnet D to the on-premises data center, which will be used as the tunnel network configured on the switch. Add a subnet B on the cloud as the tunnel network on the cloud, and create a L2CG using subnet B. The L2CG will work together with Direct Connect to enable the communication between the cloud and on-premises tunnel networks.

**Figure 10-2** shows the network topology after the reconstruction using L2CG.

**Figure 10-2** Network topology



The subnet planning details are as follows:

**Table 10-1** Subnet planning

| Subnet                                | CIDR Block     |
|---------------------------------------|----------------|
| Subnet A (to be migrated at layer 2)  | 192.168.3.0/24 |
| Subnet B (tunnel subnet on the cloud) | 192.168.0.0/24 |
| Subnet D (on-premises tunnel subnet)  | 200.51.51.0/24 |

## Notes and Constraints

- A maximum of six Layer 2 connections can use the same L2CG to connect cloud and on-premises networks.
- Forwarding unknown unicast, broadcast, and multicast (except VRRP) IP packets from your data center to the cloud is not allowed.
- On-premises servers cannot use VPC peering connections, load balancers, route tables, and NAT gateways on the cloud.
- A VPC can be attached to multiple L2CGs. However, each L2CG can only be attached to one VPC.
- The remote tunnel VNI and tunnel IP address of each Layer 2 connection using the same L2CG must be unique.
- A subnet that has been associated with a Layer 2 connection cannot be used by any other Layer 2 connection or L2CG.
- Each Layer 2 connection of a L2CG requires two IP addresses (interface IP address and tunnel IP address) in the Layer 2 subnet. The two IP addresses must be different from the used IP addresses of your data center.
- Each L2CG gateway requires three IP addresses in the tunnel subnet.

## Prerequisites

The switch of your on-premises data center should support VXLAN and have licenses.

The recommended switch models are as follows:

- Huawei CE58, CE68, CE78, and CE88 series switches support VXLAN. By default, VXLAN is disabled on these switches. To use the VXLAN function, apply for and purchase the license from the switch supplier.
- Huawei CE128 series

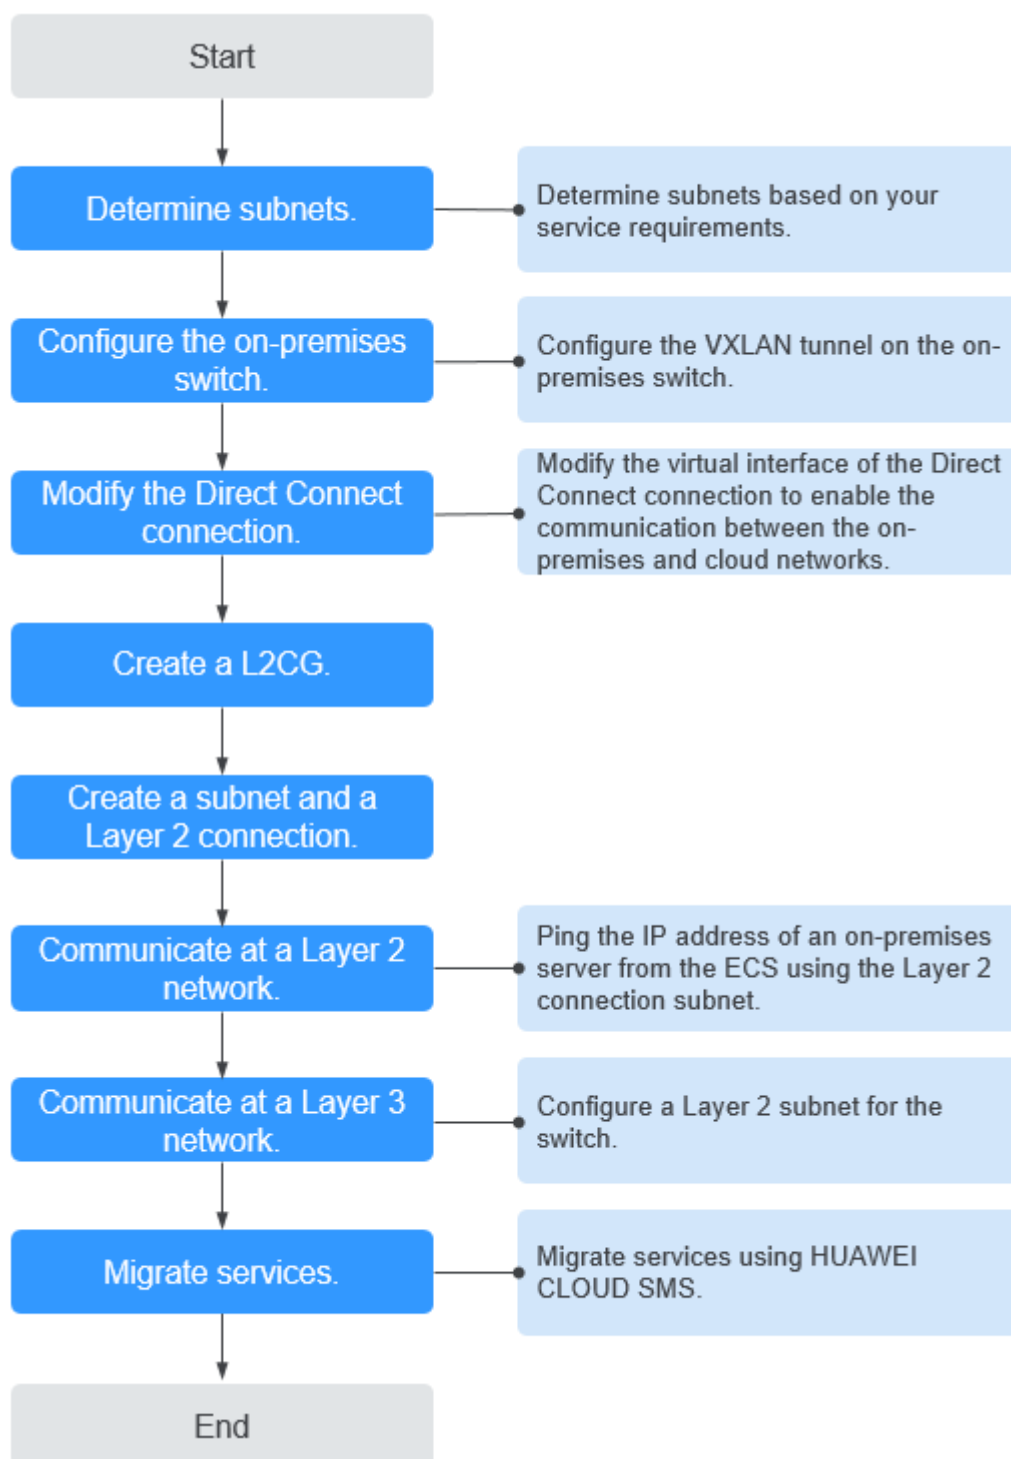
### NOTE

The switch needs to be configured with jumbo frames to allow packets with more than 1500 bytes to pass through.

## Procedure

The overall operation process is as follows:

**Figure 10-3** Operation process



**1. Determine subnets.**

Determine subnets based on your service requirements. [Table 10-1](#) is used as an example.

 NOTE

- The subnets shown in the examples are for demonstration purposes. Adjust them according to actual subnets.
- It is not recommended that the tunnel network be too large. The tunnel IP address is assigned from this tunnel network to establish a VXLAN tunnel with the L2CG on HUAWEI CLOUD. [Figure 10-2](#) shows the example.

**2. Configure the on-premises switch.**

Configure the VXLAN tunnel on the on-premises switch. In this example, subnet D is the tunnel network configured on the switch.

- Source address: Tunnel IP address (192.168.0.98) on the cloud
- Destination address: On-premises tunnel IP address (200.51.51.100)
- Tunnel VNI: 5530

For details about how to configure on-premises switches, see [Configuring a Tunnel Gateway in Your Data Center](#).

**3. Modify the Direct Connect connection.**

[Modify the virtual interface](#) of the Direct Connect connection and add the CIDR block of tunnel subnet D (200.51.51.0/24) to enable the communication between the on-premises and cloud networks.

**4. Create a L2CG.**

[Buy a L2CG](#) and set the following parameters:

- **Tunnel Connection:** Select **Direct Connect**.
- **Connection Gateway:** Select an existing Direct Connect gateway.
- **Tunnel Subnet:** Select subnet B (192.168.0.0/24).
- **Tunnel IP Address:** Specify this parameter value to the local tunnel IP address (192.168.0.98) of the L2CG.

Figure 10-4 Buying a L2CG

Billing Mode: Pay-per-use

Region: cn-hangzhou-Hela-erv

Active AZ: AZ1

Standby AZ: AZ1

Type:

| Basic                          | Standard                      | Enhanced                       |
|--------------------------------|-------------------------------|--------------------------------|
| 10 Gbit/s<br>Maximum Bandwidth | 5 Gbit/s<br>Maximum Bandwidth | 10 Gbit/s<br>Maximum Bandwidth |
| 1,500,000 pps<br>Maximum PPS   | 1,000,000 pps<br>Maximum PPS  | 1,500,000 pps<br>Maximum PPS   |
| 1<br>Connected Subnets         | 3<br>Connected Subnets        | 6<br>Connected Subnets         |

Tunnel Connection: Direct Connect

Connection Gateway: L2CGAuto\_vlan3151

VPC: L2CGAuto\_192\_168\_0\_0\_16...

Tunnel Subnet: L2CGAuto\_192\_168\_0\_0\_24...

Tunnel IP Address: Manually specify (192.168.0.98)

Name: l2cg-e726

Configure ¥0.50/hour

This price is an estimate and may differ from the final price. Pricing details

Next

Click **Next** and then **Submit**. This operation takes 3 to 6 minutes to complete.

5. **Create a subnet and a Layer 2 connection.**

**CAUTION**

After a subnet is created, communication at Layer 3 will be interrupted due to the conflict between the cloud and on-premises routes. Communication at Layer 3 will be restored only after Layer 2 connections are created.

a. Create a Layer 2 connection.

**Create a subnet** (192.168.3.0/24), which corresponds to subnet A (192.168.3.0/24) on the cloud in **Figure 10-2**.

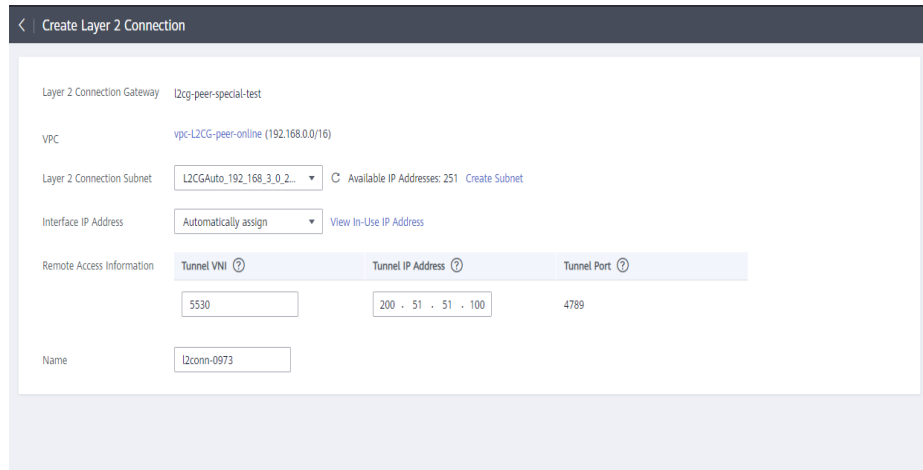
**NOTE**

- Subnets A, B, and D cannot overlap.
- If possible, make the range /28 for subnet D.
- The CIDR block of the VPC on the cloud depends on the number of required L2CGs. Each L2CG needs three IP addresses from the tunnel subnet.

b. **Create a Layer 2 connection.**

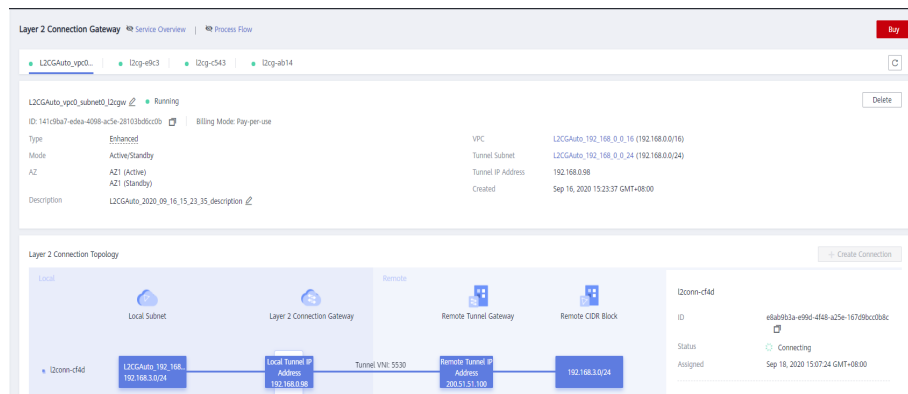
- **Layer 2 Connection Subnet:** Select subnet A (192.168.3.0/24) created in **4.a**.
- **Remote Access Information:** Enter the tunnel VNI (5530) and tunnel IP address (200.51.51.100).

**Figure 10-5** Creating a Layer 2 connection



- c. Click **Create**. If the connection status changes to **Connected**, the layer 2 connection is created successfully.

**Figure 10-6** Layer 2 connection details



**6. Communicate at a Layer 2 network.**

Buy an ECS using the Layer 2 connection subnet A on the cloud, log in to the ECS, and ping the IP address of an on-premises server.

**Figure 10-7** Accessing the on-premises server

```

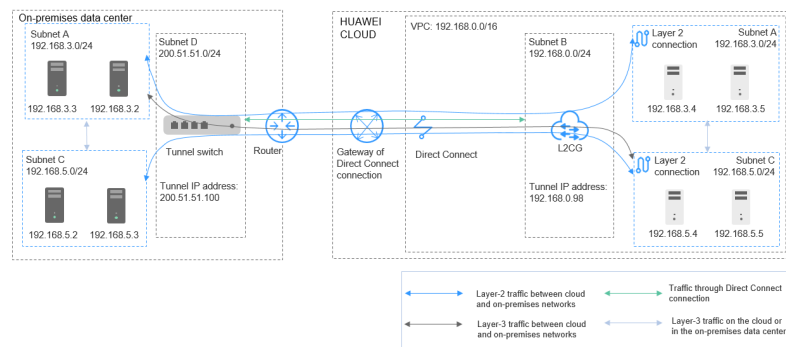
l2cgauto-vpc0-subnet3-az0-pod1-kvm login: root
Password:
Last login: Tue May 5 15:12:15 from 10.173.134.147
##### Notice #####
#
# 1. Please create unique passwords that use a combination of words, #
# numbers, symbols, and both upper-case and lower-case letters. #
# Avoid using simple adjacent keyboard combinations such as #
# "Qwert!234", "QazZwsx", etc. #
#
# 2. Unless necessary, please DO NOT open or use high-risk ports, #
# such as Telnet-23, FTP-20/21, NTP-123(UDP), RDP-3389, #
# SSH/SFTP-22, MySQL-3306, SQL-1433, etc. #
#
# Any questions please contact 4000-955-988 #
#####
pin[192.168.3.28_pod1-kvm ~]#
[192.168.3.28_pod1-kvm ~]#
[192.168.3.28_pod1-kvm ~]#
[192.168.3.28_pod1-kvm ~]#ping 192.168.3.222
PING 192.168.3.222 (192.168.3.222) 56(84) bytes of data:
64 bytes from 192.168.3.222: icmp_seq=1 ttl=64 time=1006 ms
64 bytes from 192.168.3.222: icmp_seq=2 ttl=64 time=6.99 ms
64 bytes from 192.168.3.222: icmp_seq=3 ttl=64 time=2.96 ms
64 bytes from 192.168.3.222: icmp_seq=4 ttl=64 time=2.54 ms
^C
--- 192.168.3.222 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3001ms
rtt min/avg/max/mdev = 2.547/254.669/1006.164/433.879 ms, pipe 2
[192.168.3.28_pod1-kvm ~]#
    
```

**7. Communicate at a Layer 3 network.**

Two Layer 2 connections need to be created using the L2CG to implement Layer 3 communication between the cloud and on-premises networks. **Figure 10-8** shows the network topology.

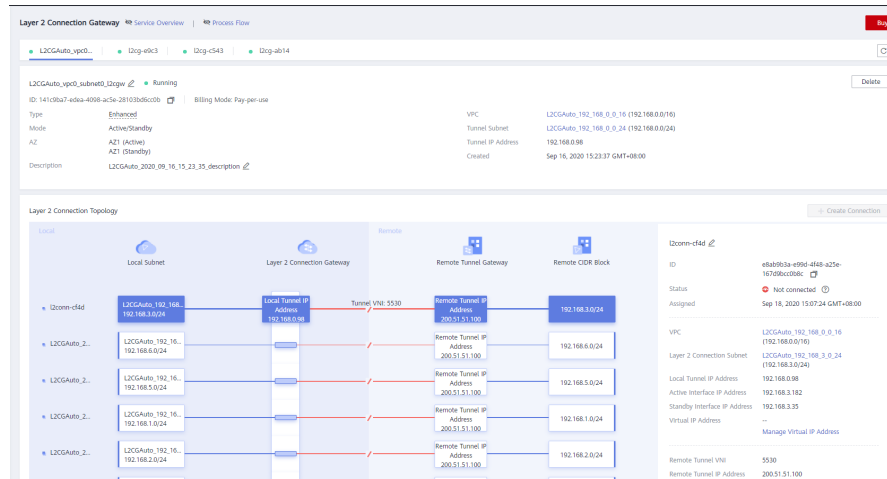
In addition to the Layer 2 connection created in **4.b**, you need to create another Layer 2 connection.

**Figure 10-8** Communication at a Layer 3 network




- Add a Layer 2 subnet on the switch and divert the traffic of the Layer 2 subnet to the new tunnel. (A new tunnel can have the same IP address as an existing tunnel but their tunnel numbers must be different.) For details about how to configure a switch, see **2**.
- Create subnet C (192.168.5.0/24).
- Create a Layer 2 connection. A maximum of six Layer 2 connections can be created. After the creation is successful, the page shown in **Figure 10-9** is displayed.

Figure 10-9 Layer 2 connection details



After the creation is successful, the cloud and on-premises networks can communicate at Layer 3.

Figure 10-10 Communication at Layer 3



8. Migrate services.

You can use [Server Migration Service \(SMS\)](#) to migrate services.

Common Questions

- If the subnets to be connected at Layer 2 are not on the same network, the VPC of the L2CG must support the multiple CIDR block. In this case, you need to use a tool to create subnets across CIDR blocks. If you need help, [submit a service ticket](#).
- If the system IP addresses 192.168.1.253 and 192.168.0.254 on the cloud have been used, but servers with these IP addresses need to be migrated to the cloud, you need to use a tool to change the system IP addresses on the cloud. If you need help, [submit a service ticket](#).

# 11 Building Highly Available Web Server Clusters with Keepalived

## Scenario

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:

**Table 11-1** Data planning


| No. | Item       | Quantity | 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<br>ecs-HA2: 192.168.0.20 |
| 3   | EIP        | 1        | 122.xx.xx.189                                  |

| No. | Item               | Quantity | 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 11-1** Network topology

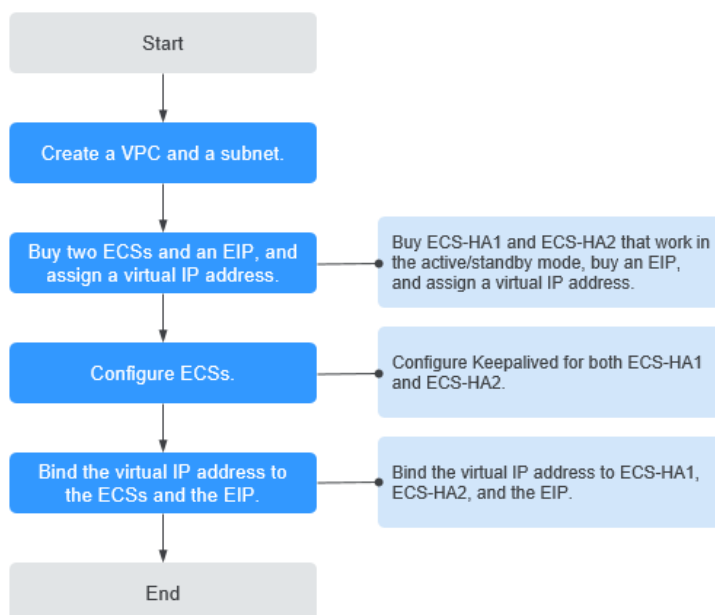


**NOTE**

- 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 11-2** Operation process**Step 1 Create a VPC and a subnet.**

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

Set required parameters as prompted based on [Table 11-2](#).

**Table 11-2** Parameter configurations

| Parameter                  | Example Value  |
|----------------------------|----------------|
| Name (of the VPC)          | vpc-HA         |
| 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 **Computing**, click **Elastic Cloud Server**.
  - c. Click **Buy ECS**.
  - d. On the **Buy ECS** page, set parameters as prompted. For details, see [Table 11-1](#).

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

 **NOTE**

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

2. Buy an EIP.
  - a. Log in to the management console.
  - b. Click **Service List**. Under **Network**, click **Elastic IP**.
  - c. Click **Buy EIP** and set parameters as prompted. For details, see [Table 11-1](#).
3. Assign a virtual IP address.
  - a. Log in to the management console.
  - b. Click **Service List**. Under **Network**, 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 **Computing**, 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:
  - 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.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;
}
}
}
```

- 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 11-3** ECS-HA1 access verification



- 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
}
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
}
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 **Computing**, 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 **Computing**, 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 11-4** ECS-HA2 verification result



# 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 101
advert_int 1
authentication {
auth_type PASS
auth_pass 1111
}
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
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 **Network**, 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.
  - g. **Bind the virtual IP address to 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 **Network**, 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 11-5](#), the virtual IP address has been bound to the eth0 NIC of ecs-HA1.


**Figure 11-5** 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
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    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
    link/ether fa:16:3e:a2:c5:72 brd ff: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. 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 11-6](#) is displayed, the access is normal.

Figure 11-6 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 11-7 Virtual IP address of ecs-HA2

```
root@ecs-ha2 ~# ip addr show
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    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
    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 04958sec preferred_lft 04958sec
    inet 192.168.0.100/32 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::f816:3eff:fe79:321/64 scope link
        valid_lft forever preferred_lft forever
root@ecs-ha2 ~#
```

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 11-8](#) is displayed, the access is normal.

Figure 11-8 ecs-HA2 access verification



## Welcome to ECS-HA2

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

## Scenario

Finance and securities enterprises have high security requirements when planning cloud networks. Access to servers is often controlled based on IP addresses. To simplify security group rule configuration and provide refined security control, you can use IP address groups in case of the following scenarios:

- A security group has more than 40 rules.
- The direction, type, protocol, and port of security group rules are the same except the address.

## Constraints

- IP address groups are currently available in three regions: **CN North-Beijing4**, **CN South-Guangzhou**, and **CN South-Guiyang1**.
- An IP address group can contain a maximum of 20 IP addresses or IP address ranges.

## Prerequisites

You have created one or more security groups for access control.

## Typical Case

For example, you plan to configure the following rules for security group A.


| Direction | Type | Protocol | Port Range | Source/Destination        |
|-----------|------|----------|------------|---------------------------|
| Inbound   | IPv4 | TCP      | 22122      | Source: 11.19.255.64/30   |
| Inbound   | IPv4 | TCP      | 22122      | Source: 113.31.128.252/30 |

| Direction | Type | Protocol | Port Range | Source/Destination        |
|-----------|------|----------|------------|---------------------------|
| Inbound   | IPv4 | TCP      | 22122      | Source: 113.31.138.0/25   |
| Inbound   | IPv4 | TCP      | 22122      | Source: 183.232.25.208/28 |

The four inbound rules have the same port, type, and protocol but different source IP addresses. In this case, you can use an IP address group to reconfigure the security group rules.

## Procedure

### Create an IP address group.

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. Under **Network**, click **Virtual Private Cloud**.
4. In the navigation pane on the left, choose **Access Control > IP Address Groups**.
5. Click **Create IP Address Group**.
6. Set the parameters.
  - **Name:** ipGroup-A
  - **IP Address:**
    - 11.19.255.64/30
    - 113.31.128.252/30
    - 113.31.138.0/25
    - 183.232.25.208/28

**Figure 12-1** Creating an IP address group

**Create IP Address Group**

\* Name: ipGroup-A

\* IP Address: 11.19.255.64/30  
113.31.128.252/30  
113.31.138.0/25  
183.232.25.208/28

Description: ipGroup-A

OK Cancel

7. Click **OK**.

#### Configure a security group rule.

8. In the navigation pane on the left, choose **Access Control > Security Groups**.
9. Locate security group A and click **Manage Rule** in the **Operation** column.
10. Under **Inbound Rules**, click **Add Rule**.
11. Set the parameters.
  - **Protocol & Port: TCP and 22122**
  - **Type: IPv4**
  - **Source: ipGroup-A**

**Figure 12-2** Configuring a security group rule

**Add Inbound Rule** [Learn more](#) about security group configuration.

Inbound rules allow incoming traffic to instances associated with the security group.

Security Group: A

You can import multiple rules in a batch.

| Protocol & Port | Type | Source  | Description | Operation |
|-----------------|------|---|-------------|-----------|
| TCP<br>22122    | IPv4 | IP address group<br>ipGroup-A(a4494ab9-8074-4ee7-...) |             | Operation |

Add Rule

OK Cancel

12. Click **OK**.

#### Delete old security group rules.

13. Delete four old security group rules after the configured security group rule takes effect.