# Virtual Private Cloud
# Operation Guide

Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

# Contents

# Operation Guide
# Network topology

Last updated: 2024-01-12 14:33:37

The network topology map displays all VPC resources, so that you can obtain VPC deployments and connections in real time.

## Instructions

1. Log in to the VPC console.
2. Click **Network Topology** on the left sidebar.
3. Select a region and VPC to view the cloud resources contained within the VPC, such as cloud servers, load balancers, cloud databases, and cloud caches, as well as their network topology relationships.
   For instance, as depicted in the diagram below, this VPC comprises two subnets, with the `test6` subnet containing two load balancer instances. This VPC communicates with the Internet via a NAT gateway and public network load balancer, and with the peer VPC via a peering connection.

# Network performance dashboard

Last updated: 2024-01-12 14:33:46

The network performance dashboard displays the private network performance of connections across regions and that between availability zones in a region, allowing you to stay informed about the network performance and better plan your cloud resource distribution.

## Instructions

1. Log in to the VPC console.
2. Click **Network Performance Dashboard** on the left sidebar.

### Private network performance of connections across regions

> **Note**
> The latency data before and after the switching between the source and destination regions may vary due to differences in the underlying transmission lines.

Select a region and click on the inter-region connection line to view the network latency between the chosen regions.

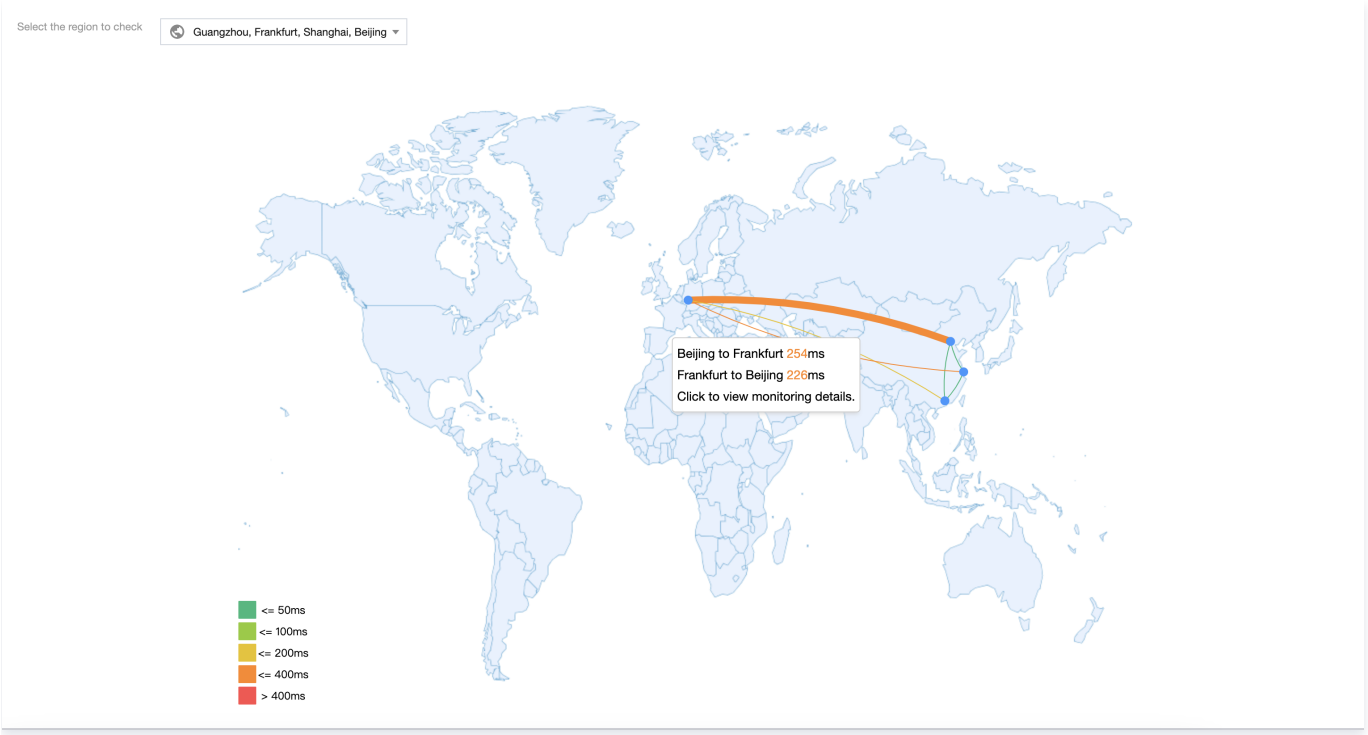


### Private network performance of connections in a region

Select a region to view the network latency between Availability Zones within the region.


**Note**

The latency data of connections in the same availability zone is not available.

# Virtual Private Cloud (VPC) Overview

Last updated: 2024-01-12 14:33:53

You must create a VPC and subnet before using any cloud resources. A VPC is a logically isolated virtual network that you can use exclusively and plan independently on Tencent Cloud. A subnet is a network space in the VPC. VPCs are region-specific, while subnets are availability zone-specific. A VPC must contain at least one subnet. You can create multiple subnets in a VPC to divide it into several networks. Subnets in the same VPC can communicate with each other over a private network by default.
All cloud resources such as CVMs and CLBs in a VPC must be deployed in a subnet.
This document describes the lifecycle of a VPC.
Watch video

## VPC lifecycle

The VPC lifecycle varies with needs, as shown below:

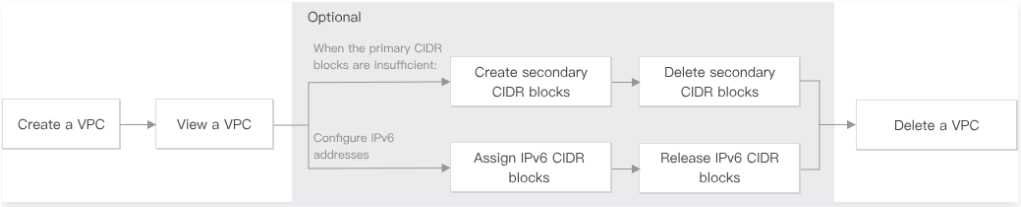


1. Create a VPC: Before creating a VPC, you need to have a network plan in place. The CIDR blocks of the VPC and subnet cannot be modified after creation.
2. Viewing a VPC: You can view the basic information of a VPC, its CCN association, and the resources it contains.
3. (Optional) Choose the operations that apply to your use cases:
   - When the primary CIDR block is insufficient, see Editing IPv4 CIDR Blocks:
     - Creating secondary CIDR blocks: You can create secondary CIDR blocks to meet your actual network demands.
     - Deleting secondary CIDR blocks: You can delete secondary CIDR blocks if you no longer need them.
   - If you need to use IPv6 addresses, see Editing IPv6 CIDR Blocks:
     - Assigning IPv6 CIDR blocks: You can create IPv6 CIDR blocks if you need to use IPv6 addresses.
     - Releasing IPv6 CIDR block: You can release IPv6 CIDR blocks if you no longer need them. After the release of IPv6 CIDR blocks, all IPv6 addresses of cloud resources within the VPC will be released and cannot be recovered.
4. Deleting a VPC: When a VPC is deleted, its subnets and route tables are also deleted.

# Limits

Last updated: 2024-01-12 14:34:03

## Usage Limits

- The IP ranges of VPCs and subnets cannot be modified after creation.
- For each subnet, Tencent Cloud reserves its first two IPs and the last one for IP networking. For example, if the subnet CIDR block is `172.16.0.0/24`, then `172.16.0.0`, `172.16.0.1`, and `172.16.0.255` are reserved by Tencent Cloud.
- When you add a CVM instance to a VPC, the CVM will be randomly assigned a private IP from a specified subnet. After the CVM is created, you can reassign a private IP to it.
- In a VPC, a CVM private IP corresponds to one public IP address.
- Classic network-based CVMs cannot interconnect with cloud resources in the secondary CIDR block.
- A peering connection does not support secondary CIDR blocks.

## Quota Limits

| Resources | Use limits |
| --- | --- |
| Number of VPCs per account per region | 20 |
| Number of subnets per VPC | 100 |
| Number of secondary CIDR blocks per VPC | 5 |

**Note**
If you want to increase the quota, please submit a ticket to apply.

# Creating VPC

Last updated: 2024-01-12 14:34:12

VPCs provide a basis for using Tencent Cloud services. This document describes how to create a VPC in the VPC console.

## Operations Guide

1. Log in to the Virtual Private Cloud Console.
2. At the top of the **Virtual Private Cloud** page, select the region for the VPC and click **Create**.
3. In the **Create VPC** pop-up window, enter the VPC information and initial subnet information.

> **Note**
> The CIDR blocks of the VPC and subnet cannot be modified after creation.

- Tencent Cloud VPC supports CIDR blocks in any of the following IP ranges. If you require internal communication between different VPCs, ensure that the CIDR configurations of both VPCs do not overlap:
  - `10.0.0.0` – `10.255.255.255` (subnet mask range must be between 12 and 28)
  - `172.16.0.0` – `172.31.255.255` (mask range must be between 12 and 28)
  - `192.168.0.0` – `192.168.255.255` (mask range must be between 16 and 28)
- The subnet CIDR must be within or the same as the VPC CIDR.
  For example, if the VPC CIDR block is `10.0.0.0/16`, the subnet CIDR blocks within the VPC can be `10.0.0.0/16`, `10.0.0.0/24`, etc.
- Availability Zone: Subnets have AZ-specific attributes. Please select the AZ for the initial subnet. A VPC can have subnets in different AZs, which communicate with each other over the private network by default.
- Associated Route Table: A subnet must be associated with a route table for traffic forwarding control. The initial subnet is associated with a default route table, which is provided by the system and allows for network communication within the VPC.
- Tags: You can add tags as needed to help manage resource permissions for sub-users and collaborators. This is an optional parameter and can be set as required.

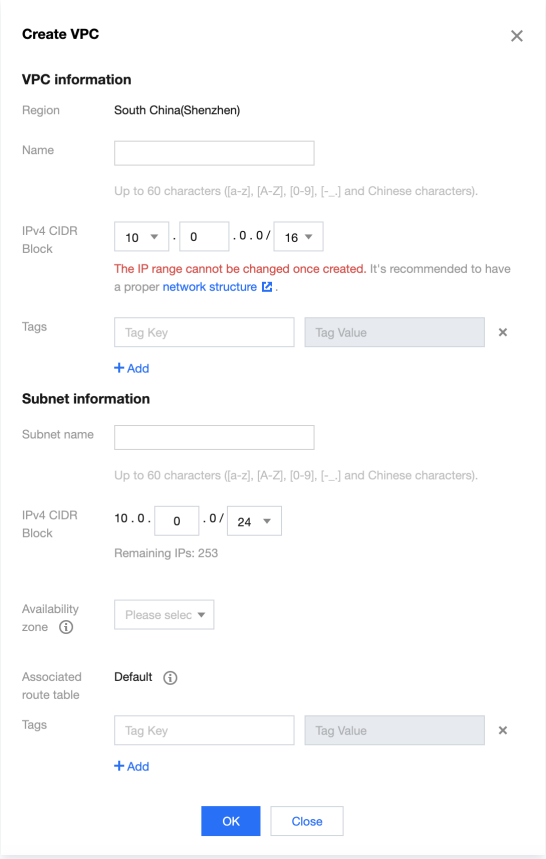

4. After configuring the parameters, click **Confirm** to create the VPC. The successfully created VPC will be displayed in the list, as shown below. The new VPC includes an initial subnet and a default route table.

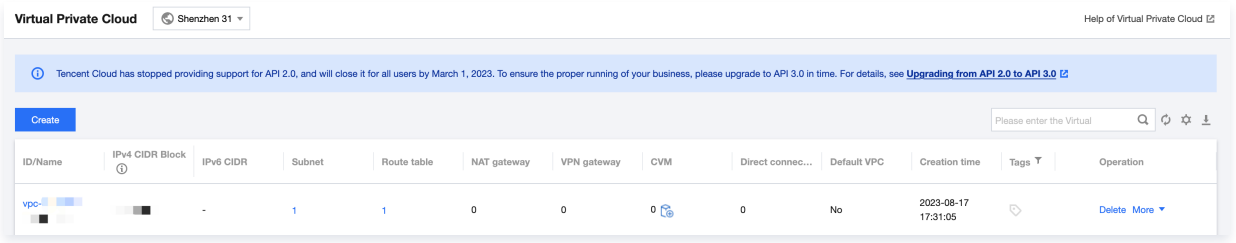


## See Also

Once the VPC and initial subnet are created, you can deploy cloud resources within the VPC, such as Cloud Virtual Machines and Cloud Load Balancers. Click the icon in the red box below to directly access the Cloud Virtual Machine purchase page. For more information, see Quickly Build an IPv4 Virtual Private Cloud.

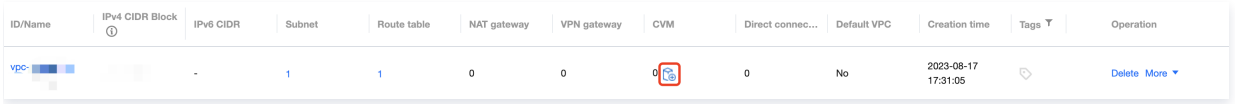


## Related Content

In VPC, there is a default VPC, which means:
When no VPCs have been created in a region, while creating Cloud Virtual Machines, Cloud Load Balancers, or Databases, you can choose to have Tencent Cloud automatically create a default VPC and subnet for you, without the need to create them manually, as shown in the image below.

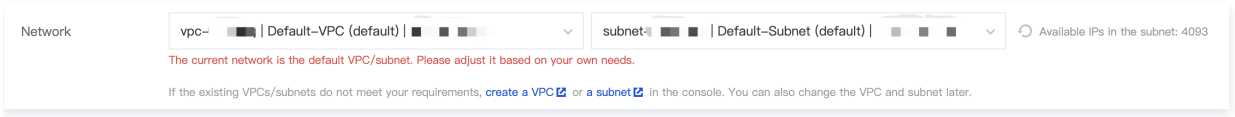


Upon successful instance creation, a default VPC and subnet will also be created. The default VPC and subnet have the same functionality as the ones you create manually and do not occupy your quota in a specific region. You can delete them if they are no longer needed. There can only be one default VPC per region and one default subnet per availability zone.

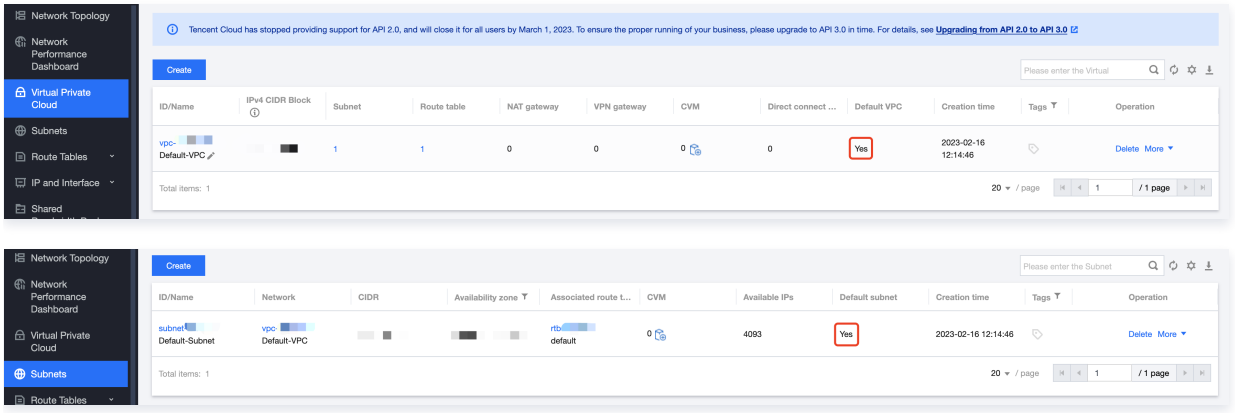

# Viewing VPCs

Last updated: 2024-01-12 14:35:59

You can query all VPC resources via the VPC console, such as cloud resources and connections in a VPC.

## Operations Guide

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page. You can check the information of all VPCs in this region in the VPC list. The following fields are displayed in the VPC list.

| Parameter | Description |
| --- | --- |
| ID/Name | The ID and name of the VPC. The name can be modified. |
| IPv4 CIDR | The IPv4 CIDR block of the VPC. It cannot be modified. |
| IPv6 CIDR | The IPv6 CIDR block of the VPC. This feature is currently in beta. To try it out, please submit an application. |
| Subnet | The number of subnets in the VPC. Click the number shown to access the Subnet page. |
| Route table | The number of route tables in the VPC. Click the number shown to access the Route Table page. |
| NAT gateway | The number of NAT Gateways in the VPC. Click the number shown to access the NAT Gateway page. |
| VPN gateway | The number of VPN gateways in the VPC. Click the number shown to access the VPN Gateway page. |
| CVM | The number of CVMs in the VPC. Click the number shown to access the CVM page. Click the CVM icon to redirect to the CVM purchase page. |
| Direct connect gateway | The number of direct connect gateways in the VPC. Click the number shown to access the Direct Connect Gateway page. |
| Default VPC | If the VPC is a default VPC, this field will be displayed as **Yes**. There can only be a single default VPC and subnet in a given region. With the same features as those being manually created, the default VPC and subnet are created if you choose to automatically create a default VPC and subnet without manual creation when purchasing cloud resources including CVMs. |
| Creation time | The time when the VPC was created. |
| Tags | The number of tags bound to the VPC. You can hover your mouse over the icon to display detailed information. |
| Action | Operations that can be performed on the VPC. Only a VPC without any resources can be deleted. You can click **More** to edit IPv4 CIDR block and IPv6 CIDR block if applicable. |

3. Click the VPC ID to view details, including the basic information, CCN association, and associated resources. Click the number next to a resource to access the resource management page.
4. Return to the VPC list, and click in the top-right corner search box to filter VPCs by different resource attributes.
5. Click the settings icon to customize the list fields.
6. Click the icon to refresh the information displayed in the list.
7. Click the icon to download all the displayed information.

# Editing IPv4 CIDR Blocks

Last updated: 2024-01-12 14:36:05

Each VPC can have one primary CIDR block, which cannot be modified after the VPC creation. When the IPs in the primary CIDR block can not meet your needs, you can create multiple secondary CIDR blocks to expand IP ranges.
You can assign an IP range from the primary or secondary CIDR blocks to a subnet. All subnets of the same VPC are interconnected by default, regardless of whether they belong to the primary or secondary CIDR blocks.

## Usage Limits

- Classic network-based CVMs cannot interconnect with cloud resources in the secondary CIDR block.
- A peering connection does not support secondary CIDR blocks.
- CCN, VPN gateway, and standard direct connect gateway support secondary CIDR blocks. Note the following limits for a direct connect gateway:
  - This feature is unavailable in Finance Cloud regions.
  - Up to 10 secondary CIDR blocks can be propagated.
  - This feature is unavailable to a NAT direct connect gateway.

## Creating secondary CIDR blocks

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. In the VPC list, click **Actions** on the right side of the target VPC, then select **More** > **Edit IPv4 CIDR**.
4. In the pop-up dialog box, click **Add** and enter a secondary CIDR block.

> **Note**
> A secondary CIDR block can overlap with the destination IP range of a custom route. Note that the secondary CIDR block uses a local route, which has a higher priority than that of custom subnet routes.

5. Click **OK**.

## Deleting secondary CIDR blocks

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. In the VPC list, locate the VPC from which you want to delete a secondary CIDR block and select **More** > **Edit IPv4 CIDR block** in the **Operation** column.
4. In the pop-up dialog box, click **Delete** next to the secondary CIDR block.
5. Click **OK**.

# Editing IPv6 CIDR Blocks

Last updated: 2024-01-12 14:36:10

VPC supports IPv6. If your business requires IPv6 communication, you need to first assign an IPv6 CIDR block to your VPC. Then, cloud resources within the VPC can be assigned IPv6 addresses from the IPv6 CIDR block to enable IPv6 communication between the resources. This document describes how to assign an IPv6 CIDR block to a VPC in the console.

> **Note**
> Currently, the IPv6/IPv4 dual-stack feature is in beta test. To try it out, please submit an application.

## Assigning IPv6 CIDR blocks

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. In the VPC list, click **Actions** on the right side of the target VPC, and select **More** > **Edit IPv6 CIDR**.
4. In the pop-up dialog box, click **Get**, and a `/56` IPv6 CIDR block will be assigned to the VPC.
5. Click **OK**.

## Releasing IPv6 CIDR blocks

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. In the VPC list, for a VPC that has already obtained an IPv6 CIDR, select **More** > **Edit IPv6 CIDR** in the **Actions** column on the right.
4. In the pop-up IPv6 CIDR edit dialog box, click **Release**.
5. In the risk warning pop-up window, acknowledge the risks associated with the release operation, confirm the information is correct, and click **Confirm** to complete the release of the IPv6 CIDR block.

# Modifying VPC DNS

Last updated: 2024-01-12 14:36:16

VPC-based CVMs support Dynamic Host Configuration Protocol (DHCP). This document describes how to modify the DHCP parameters, including the DNS address and domain name.

**Notes**
- Dynamic Host Configuration Protocol (DHCP) is a LAN network protocol that defines the standard for transferring configuration information to TCP/IP network servers.
- VPCs created before April 1, 2018 do not support DHCP features. If you cannot modify the DNS address and domain name in the console, it means that your VPC does not support these features.

## Supports and Limits

Changes of the configuration apply to all CVMs in the VPC.
- For newly created CVMs, the modified configurations take effect immediately.
- For existing CVMs, the modified configurations take effect after the CVMs or network services are restarted.

## Instructions

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. Click the VPC ID to access the **Basic information** page.
4. Click the edit icon to modify the DNS and domain name respectively.
   - **DNS**: DNS server addresses.

     **Notes**
     - The Tencent Cloud default DNS addresses are `183.60.83.19` and `183.60.82.98`. If the default DNS addresses are not used, internal services such as Windows activation, NTP, and YUM will be unavailable.
     - Up to four IP addresses can be entered in the **DNS** field. Separate them with commas. Note that some operating systems may not support four DNS addresses.

   - **Domain Name**: CVM hostname suffix, such as `example.com`. You can enter up to 60 characters, or keep the default configuration if you don't have special requirements.

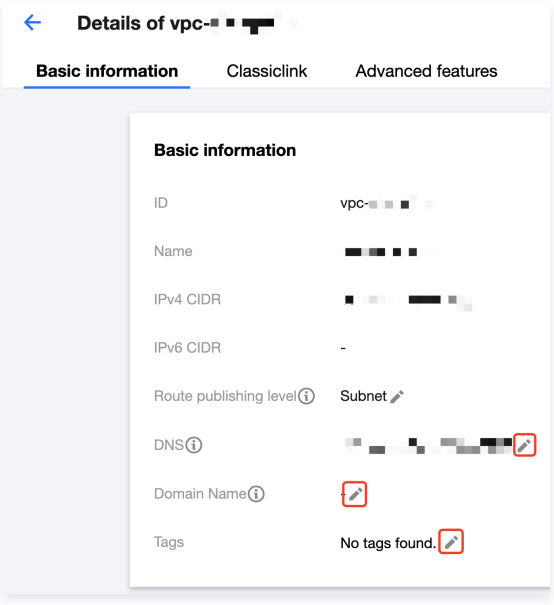

# Modifying VPC Name and Tag

Last updated: 2024-01-12 14:36:21

This document describes how to modify the name, tags, or other information of a VPC.

## Instructions

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. Click the edit icon next to a VPC name to modify it.
4. Click the VPC ID to access the **Basic information** page.
5. Tags are used to identify and manage resources. You can click the edit icon to add or delete tags.

# Associating or Unassociating CCN

Last updated: 2024-01-12 14:36:27

Cloud Connect Network (CCN) can work as a bridge between Tencent Cloud VPCs and between VPCs and local IDCs. It provides you with multipoint private network interconnection. To leverage this CCN feature, you need to first associate VPCs to a CCN. This document describes how to associate a VPC with or disassociate it from a CCN.

## Associating a VPC with a CCN

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. Click the VPC ID to access the **Basic information** page.
4. Click **Associate Now** under the **Associate with CCN** section to open the Associate CCN dialog box.
5. Configure parameters as follows.
   - **Account**: The account to which the CCN belongs. The VPC and CCN can be under the same or different accounts. If you choose **Other accounts**, enter the account ID. The account owner needs to accept the association application within seven days; otherwise, the application will expire. The owner of the CCN bears the network interconnection fee generated by instances connecting to the CCN.
   - **Cloud Connect Network**: Select a CCN ID from the drop-down list if you choose **My account** or enter a CCN ID if you choose **Other accounts**.
6. Click **Confirm** to complete the association process. After the VPC is successfully associated with the CCN instance, the status will be **Connected**.

## Disassociating from a CCN

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. Locate the VPC to be disassociated from the CCN, and click the VPC ID to access the **Basic information** page.
4. Click **Disassociate** under the **Associate with CCN** section.
5. In the risk warning box, fully understand the operational risks, confirm that everything is correct, and click **Disassociate** to complete the operation.

## Related Actions

- Network Instance Interconnection within One Account
- Network Instance Interconnection Crossing Accounts

# Classiclink Overview

Last updated: 2024-01-12 14:36:34

The classiclink feature allows VPC-based resources to communicate with classic network-based CVMs. For example:

- The classic network-based CVMs can communicate with VPC resources such as CVM, TencentDB, private network CLB, Redis/CMEM, and so on.
- VPC resources can only access classic network-based CVMs, but not other resources in the classic network, like TencentDB and CLB.

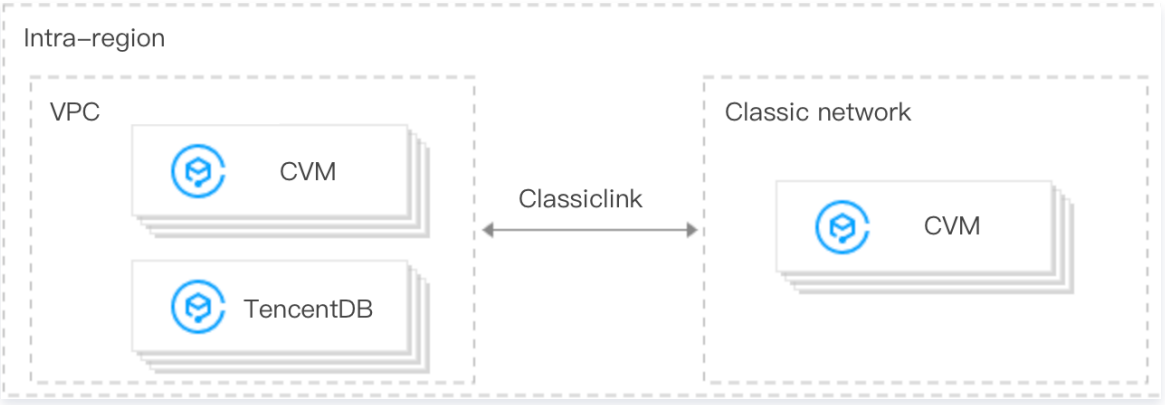


## Usage Limits

- A VPC can only be interconnected with the classic network **in the same region**.
- The VPC IP range must be within `10.0.0.0/16-10.47.0.0/16` (including subsets); otherwise, there may be IP conflicts, which may cause failures while associating and communicating with the classic network-based CVMs.
- A classic network-based CVM can only be associated with one VPC at a time.
- One VPC supports associating with up to 100 classic network-based CVMs.
- After the classic network-based CVMs are associated with a VPC, classic network-based CVMs can only communicate with resources in the primary CIDR block rather than the secondary CIDR block of the VPC.
- CLB instances within a VPC cannot be bound to a classic network-based CVM that interconnects with the same VPC.
- In a classiclink, the CVM traffic can only be routed to private IP addresses within the VPC rather than destinations outside the VPC.

> **Notes**
> The classic network-based CVM cannot access public or private network resources outside the current VPC through network devices such as VPN gateway, direct connect gateway, public gateway, peering connection, and NAT gateway in the current VPC. Likewise, the peer of a VPN gateway, direct connect gateway, and peering connection cannot access the classic network-based CVM.

## Supports and Limits

- Changing the private IP of a classic network-based CVM will invalidate its association with the VPC and cause the configurations to become invalid. To associate them, you need to add a classiclink again in the VPC console.
- The classiclink will not be affected by actions taken regarding the CVM such as isolation due to overdue payment, security isolation, cold migration, failover, configuration modification, and operating system switching.
- The CVM will be automatically disassociated from the VPC if the CVM is returned.

## Documentation

For more information on classiclinks, see Managing Classiclinks.

# Managing Classiclink

Last updated: 2024-01-12 14:36:40

## Creating a classiclink

A classiclink associates classic network-based a CVM with a VPC to enable interconnection between the VPC and the classic network. This allows classic network-based CVMs to communicate with VPC resources.

> **Notes**
> - The private IPs of the associated classic network-based CVMs will be automatically added to the local policy of the VPC's route table. This allows interconnection, without the need to manually modify the routing policy of the VPC.
> - After a classic network-based CVM is associated with a VPC, their security group and network ACL settings will remain effective.

### Instructions

1. Log in to the VPC console.
2. Select the region, and click the ID of the VPC `TomVPC` that needs a classiclink to access the details page.

   > **Notes**
   > In this scenario, the VPC network is `TomVPC`.

3. Click the **Classiclink** tab, and then click **Associate with CVM**.

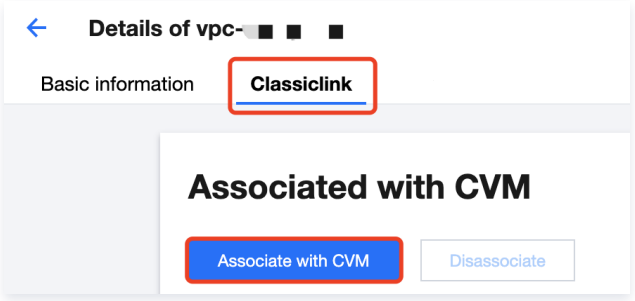


4. In the pop-up window, select the classic network-based CVM that needs to be associated with this VPC, such as `TomCVM`, and click **OK**.

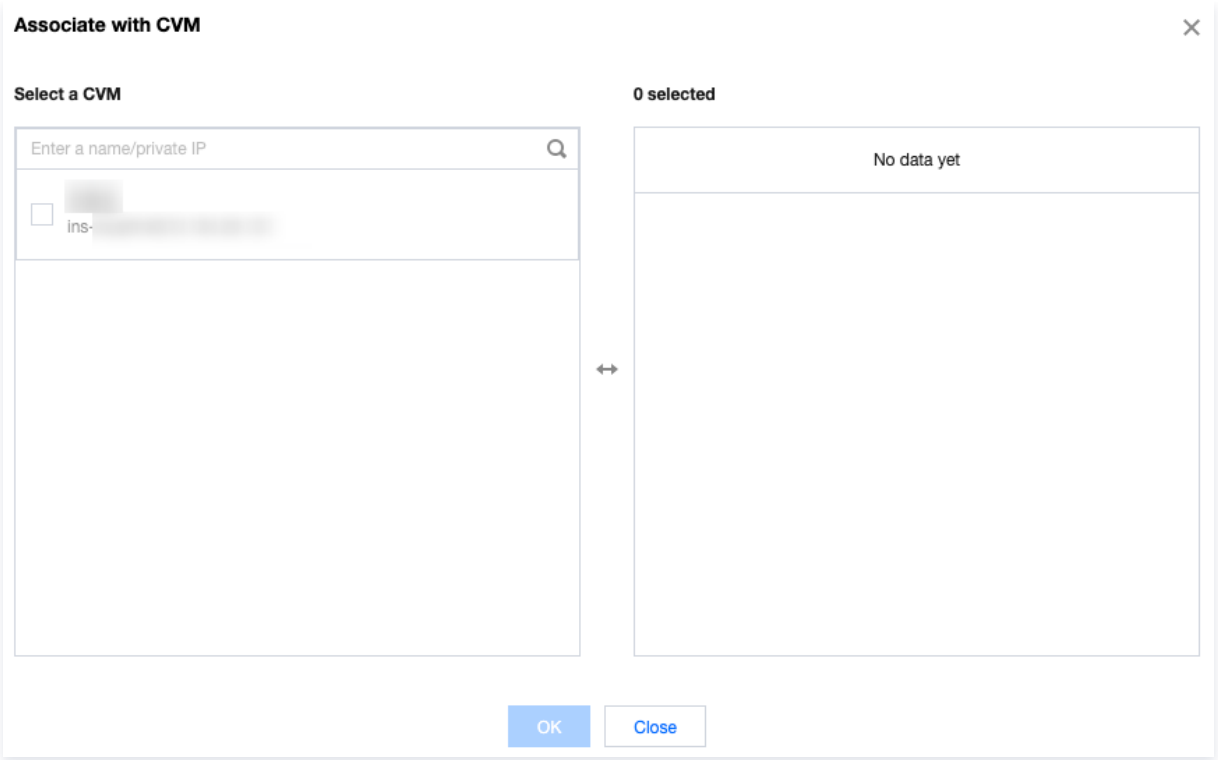


## Viewing classiclinks

You can view the list of classic network-based CVMs that interconnect with a VPC.

### Instructions

1. Log in to the VPC console.
2. Select the region, and click the ID of the VPC you want to view to access the details page.
3. Click the **Classiclink** tab to view the list of classic network-based CVMs associated with the VPC.

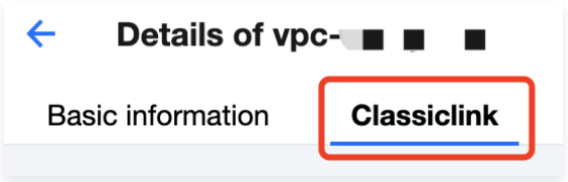


4. Enter a private IP in the top-right corner search box to quickly locate the CVM.

## Deleting a classiclink

This action disassociates a classic network-based CVM from a VPC and terminates their interconnection.

### Instructions

1. Log in to the VPC console.
2. Click the ID of the VPC for which you want to remove a classiclink to access the details page.
3. Click the **Classiclink** tab, locate the CVM to be dissociated in the classic network-based CVM list, and click **Dissociate** in the operation column.

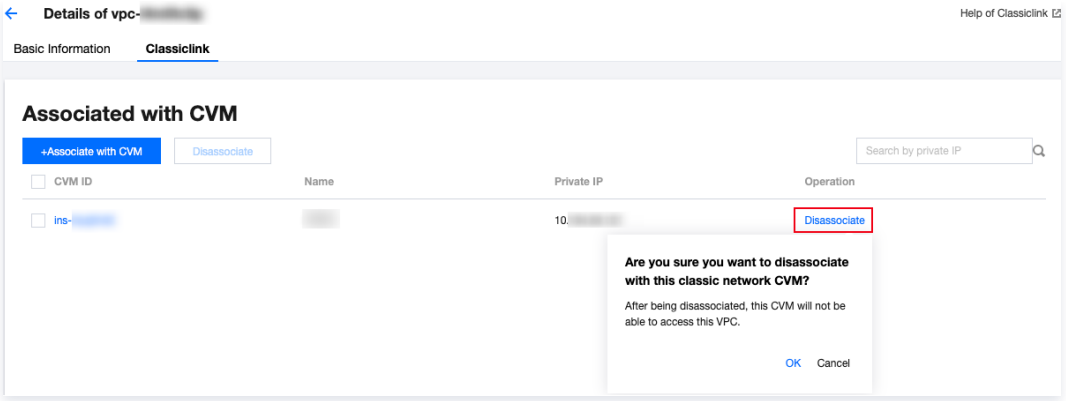


4. Carefully read the notes and click **OK**.
5. To disassociate multiple CVMs, you can select the CVMs to be disassociated and click **Disassociate** above the list.

# Enabling or Disabling Multicast

Last updated: 2024-01-12 14:36:46

This document describes how to enable or disable multicast for VPCs.

## Background

Multicast and broadcast are modes of one-to-many communication, which can save businesses on the network bandwidth and reduce network load through point-to-multipoint efficient data transmission.
In the unicast mode, the initiating server sends data to N servers separately. In the multicast or broadcast mode, the server sends the same data to N servers at one time, which reduces the server resource consumption and also the bandwidth resource of the backbone network.

**Notes**
- The trial period has ended. Please stay tuned for the official launch.
- Currently, multicast and broadcast are supported in the following regions: Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, Nanjing, Hong Kong (China), Singapore, Seoul, Tokyo, Bangkok, Toronto, Silicon Valley, Virginia, Frankfurt, Shenzhen Finance, Shanghai Finance, and Beijing Finance.
- For single-VPC multicast and broadcast, up to 50,000 PPS and 190 Mbps are supported.

- Multicast: Tencent Cloud supports multicast on the VPC dimension.
- Broadcast: Tencent Cloud supports broadcast on the subnet dimension.

## Scenario

Multicast and broadcast are mostly used in the financial and game industries.
- In the financial industry, multicast and broadcast are mainly used for broadcasting services or market data. For example, after obtaining stock prices and other real-time data, brokers can broadcast stock data to multiple clients in real time, effectively reducing network load.
- For the game industry, broadcast and multicast are mainly used for heartbeat holding between multiple servers.

## Instructions

### Enabling multicast

1. Log in to the VPC console.
2. In the list, locate the row of the target VPC for which you want to enable multicast. Click **Enable** under the multicast column and confirm the operation.

### Disabling multicast

1. Log in to the VPC console.
2. In the list, locate the row of the Virtual Private Cloud for which you want to disable multicast. Click **Disable** under the multicast column and confirm the operation.

## Related Actions

For directions about operations regarding broadcast on the subnet dimension, see Enabling or Disabling Broadcast.

# Deleting VPCs

Last updated: 2024-01-12 14:38:14

You can delete VPCs when they are not associated with other resources (peering connections, classiclinks, NAT gateways, VPN gateways, direct connect gateways, CCN instances, and private links). The subnets, route tables, and network ACLs of the VPC must be empty.

> **Note**
> An empty subnet refers to a subnet whose IPs are not used by any resources.

## Instructions

1. Log in to the VPC console.
2. Select the region of the VPC at the top of the **VPC** page.
3. In the VPC list, click **Delete** in the **Operation** column next to the VPC you want to delete, and confirm the action.

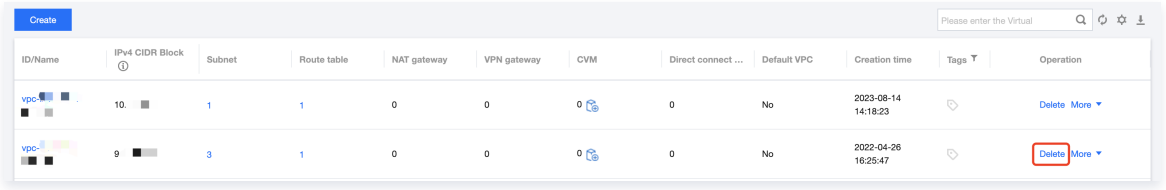

# Subnets
# Creating Subnets

Last updated: 2024-01-12 14:39:12

A subnet is a network space in a VPC where Tencent Cloud resources are deployed. A VPC has at least one subnet. An initial subnet will be created together with a VPC. You can also create more subnets in a VPC that suit your specific requirements to deploy different businesses.
A subnet is specific to an availability zone. A VPC allows subnets in different availability zones, and these subnets can communicate with each other via a private network by default. This document describes how to create a subnet in a VPC.

## Instructions

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar.
3. Select the desired region and Virtual Private Cloud for creating a subnet, then click **Create**.
4. In the **Create a subnet** pop-up dialog box, configure the subnet parameters.

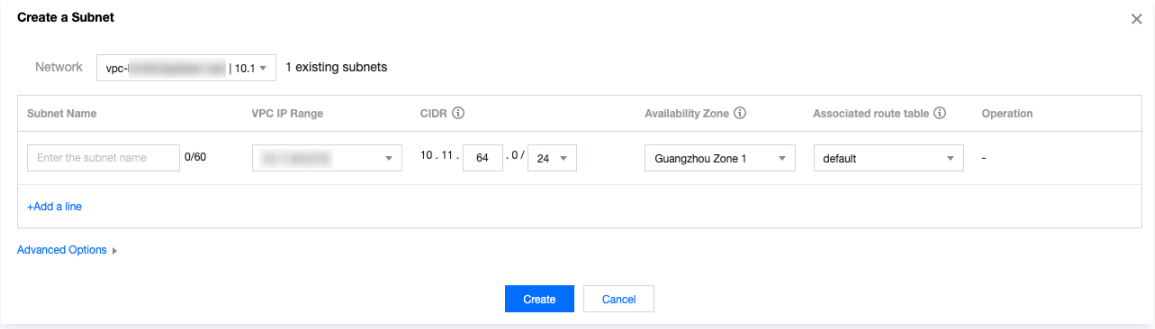


- Network: The VPC where the subnet resides. The VPC selected in Step 3 will be automatically displayed. Alternatively, you can select a VPC from the drop-down list.
- Subnet name: Enter a subnet name within 60 characters.
- VPC IP range: The CIDR block of the selected VPC will be automatically displayed.
- CIDR: Set the CIDR block of the subnet. It must be part of the VPC CIDR block and cannot overlap with the CIDR blocks of the existing subnets under the VPC.

  > **Notes**
  > Plan subnet IP ranges that suit your business scale. A private IP address within the specified subnet will be automatically assigned to the CVM instance you are creating. The primary private IP of a CVM can be modified. For more information, see Modifying Primary Private IPs.

- Availability zone: Select an availability zone for the subnet.
- Associated route table: Select a route table to be associated. The subnet must be associated with a route table to control outbound traffic. The default route table of the VPC will be associated by default to ensure private network interconnection in the VPC. You can also select another route table within the VPC.
- New line: You can click **New line** to create several subnets at a time and click to delete a subnet.
- Advanced options: You can set tags for the subnet to better manage subnet resources. Click **Add** to set multiple tags at a time. You can also click the icon in the operation column to delete a tag.

5. After completing the configurations, click **Create**. Then subnets that have been successfully created will be displayed in the list, as shown below.

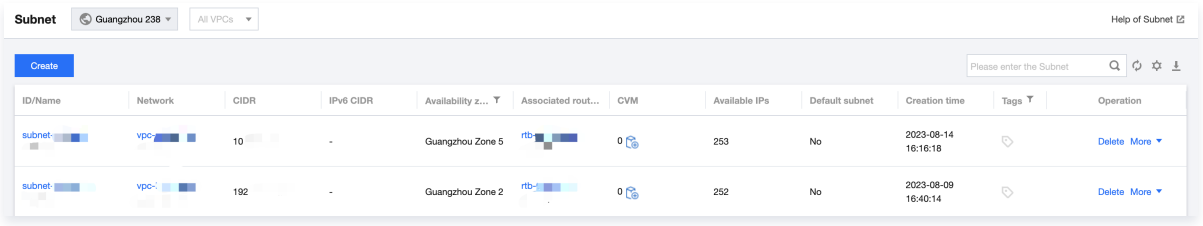


## See Also

Once the subnet is successfully created, you can deploy cloud resources within it, such as Cloud Virtual Machines and Cloud Load Balancers. Click the icon in the red box below to directly access the Cloud Virtual Machine purchase page. For more information, please refer to Quickly Set Up an IPv4 Virtual Private Cloud.

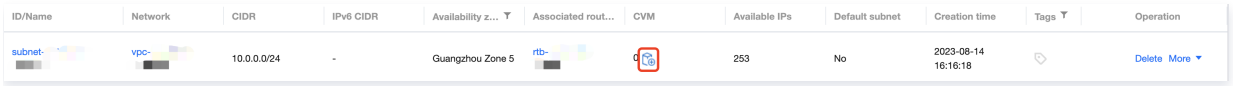

# Viewing Subnet Information

Last updated: 2024-01-12 14:39:18

You can view the resources of all subnets in the VPC on the VPC console, for instance, the cloud resources deployed in the subnet, the route table associated with the subnet, and the ACL rules bound to the subnet.

## Instructions

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar to access the subnet management page.
3. At the top of the **Subnet** page, select the region and VPC to which the subnet belongs. If you keep the default value, namely **All VPCs**, you can view all subnets of all VPCs in the region, as shown below:

> **Notes**
> - Click the VPC ID of the network to which the subnet belongs, or the route table ID of the associated route table to view the detailed information of the corresponding resource.
> - Click the number of CVMs to go to the CVM instance page. If the quantity is 0, click the CVM icon to go to the CVM purchase page.
> - Click the Filter icon next to the **Availability zone** field and select an availability zone to view all subnets in the availability zone.
> - Click the search box on the top right of the list to query subnets by **Subnet ID**, **Subnet Name**, **Tag**, **Keyword** and **IPv4 CIDR Block**.
> - Click the settings icon on the top right to customize the displayed fields.

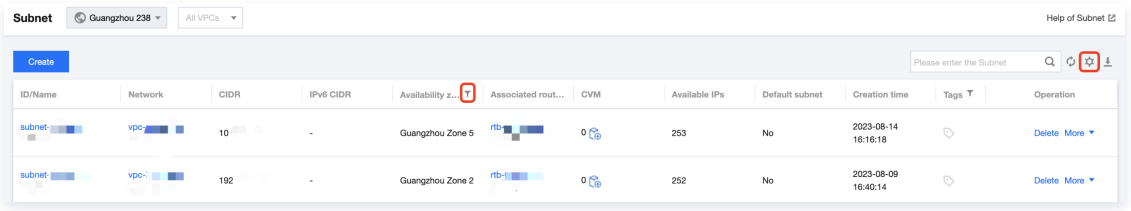


The meaning of the list fields displayed in the interface is as follows:

- **ID/Name**: The subnet ID and name. Each subnet is assigned an ID when it is created, and the subnet name can be modified at any time.
- **Network**: The VPC to which the subnet belongs.
- **CIDR**: The CIDR block of the subnet. The subnet CIDR cannot be modified.
- **Availability zone**: The availability zone where the subnet is located.
- **Associated Route Table**: The route table associated with the subnet.
- **CVM**: The number of CVMs deployed in the subnet.
- **Available IPs**: The number of available IP addresses within the CIDR block of the subnet.
- **Default subnet**: Default subnets are subnets created automatically by Tencent Cloud upon the launch of new CVM instances. Each region has one and only default VPC and subnet.
- **Creation time**: The subnet creation time.
- **Tags**: You can optionally add tags to help you better manage the resource permissions of sub-users and collaborators.
- **Operation**: Available actions. A subnet can be deleted when it's not associated with any resource. You can click **More** > **Change route table** to replace the route table associated with the subnet.

4. Click the subnet ID to view the resource details of the subnet. Switch tabs to view the routing rules and the ACL rules.

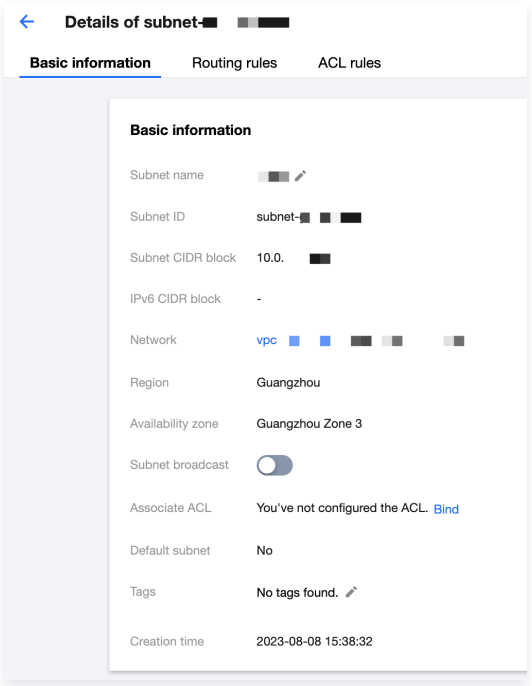

# Acquire and Release IPv6 CIDR blocks

Last updated: 2024-01-12 14:39:24

To use IPv6 for your VPC, you need to configure an IPv6 address range for the network environment first. Then, you can assign IPv6 addresses to resources within the VPC. The whole process is to get a VPC IPv6 CIDR block, and then a subnet IPv6 CIDR block, finally assign IPv6 addresses to cloud resource instances. This document describes how to get and release a subnet IPv6 CIDR block.

> **Note**
> The IPv6/IPv4 dual-stack feature is in beta now. To try it out, please submit a request.

## Obtaining an IPv6 CIDR block

### Preparations

Assign an IPv6 CIDR block to the related VPC of the subnet.

### Instructions

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar. Select a **region** and **VPC**. All subnets in the VPC are listed.
3. Select a subnet, then click on the **More** > **Obtain IPv6 CIDR** in the operation column.
4. Enter the following IP range, from 0 to 255. Ensure that it does not overlap with other subnet's IPv6 CIDR blocks. Click **Confirm**, and the system will allocate a `/64` IPv6 CIDR block for this subnet.

### Documentation

Building Up an IPv6 VPC

## Releasing an IPv6 CIDR block

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar. Select a **region** and **VPC**. All subnets in the VPC are listed.
3. Select a subnet that has already obtained an IPv6 CIDR, **More** > **Release IPv6 CIDR**.
4. In the confirmation pop-up, click **OK**, and the system will reclaim the IPv6 CIDR block of the subnet.

# Changing the Subnet Route Table

Last updated: 2024-01-12 14:39:28

Each subnet must be associated with one route table, which is used to control the outbound traffic direction of the subnet. You can change the subnet's associated route table on the **Subnet** page in the VPC console as needed. If you need to create a route table, see Creating Custom Route Tables.

## Impact on the System

After the route table is replaced by a new one, all instances in the subnet will apply the new route table policy. Please carefully evaluate the business impact.

## Instructions

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar to access the subnet management page.
3. There are two methods to change the route table associated with the subnet.
   - Click **More** > **Change route table** in the operation column on the right of the subnet that needs to change the route table.

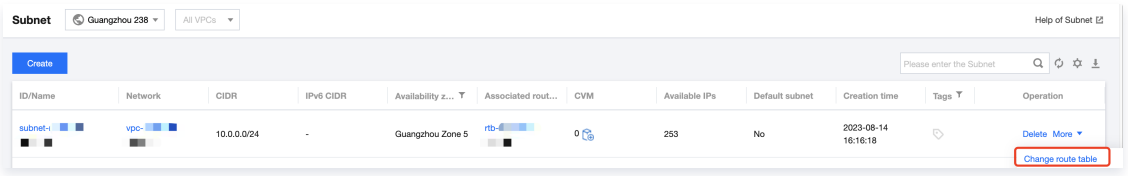


   - Click the ID of the subnet that needs to change the route table to go to the details page, switch to the **Routing rules** tab, and click **Change route table**.

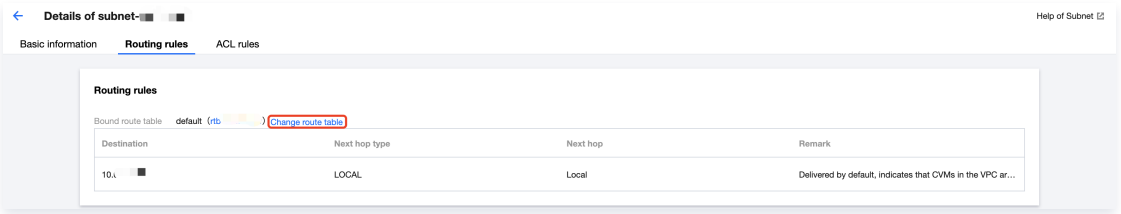


4. In the **Change route table** pop-up window, select a new route table in the drop-down list, confirm the impact on your business, and click **Confirm**.

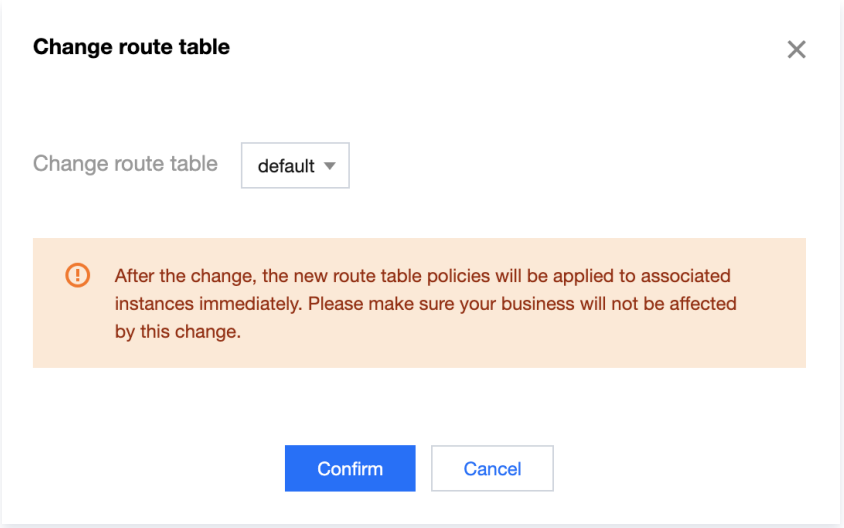

# Managing ACL Rules

Last updated: 2024-01-12 14:39:33

An ACL rule is an optional security layer that operates at the subnet level. It is used to control the inbound and outbound data streams of subnets, which can be accurate to the protocol and port granularity, to achieve fine control of subnet traffic. You can associate the same network ACL with subnets that require the same level of network traffic control.
This document describes how to bind, unbind, and change ACL rules in the VPC console.

## Instructions

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar to access the subnet management page.
3. Click a subnet ID to go to its details page. You can bind, unbind, and change ACL rules on the following tabs:
   - In the **Associated ACL** field under the **Basic Information** tab

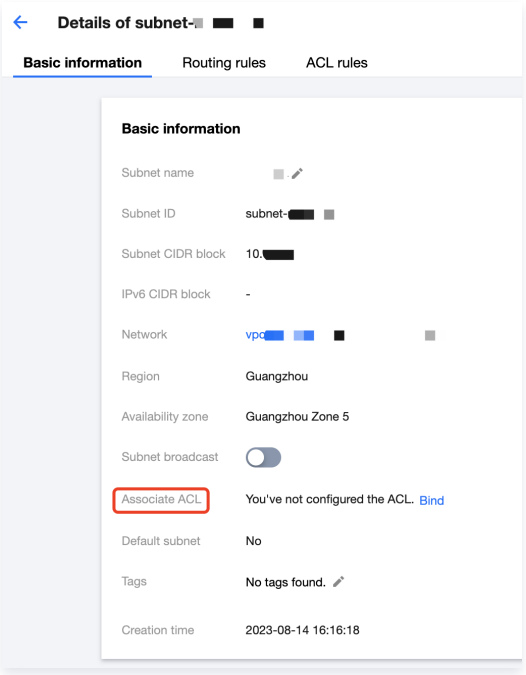


   - Under the **ACL Rules** tab

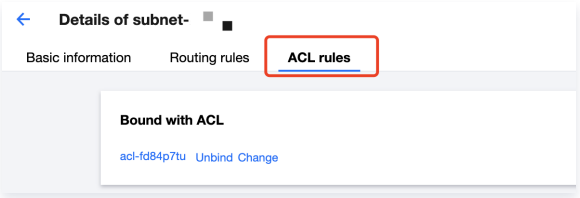


4. Perform the following operations based on the business needs. The following screenshots take the operations in **ACL Rules** as an example.

- If the current subnet is not bound to an ACL rule, you can click **Bind** to select an appropriate ACL rule and click **OK** to complete the binding. The binding will take effect immediately. The inbound and outbound traffic of the subnet is allowed only when the rule is **Allow**.

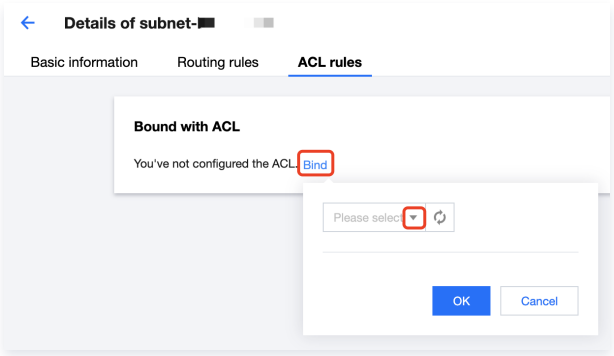


- If the ACL rule bound to the current subnet does not meet network flow requirements, you can click **Change** to change the ACL rule, which will take effect immediately.

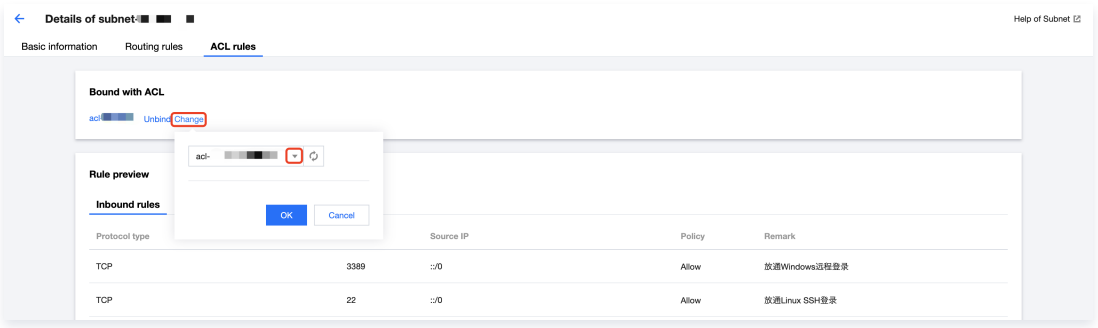


- If the current subnet is bound to an ACL rule, but you no longer need to control the inbound and outbound traffic of the subnet, you can click **Unbind** to unbind the ACL rule. The unbinding will take effect immediately and this will cause the lifting of the ACL rule restriction on the inbound and outbound traffic of the subnet.

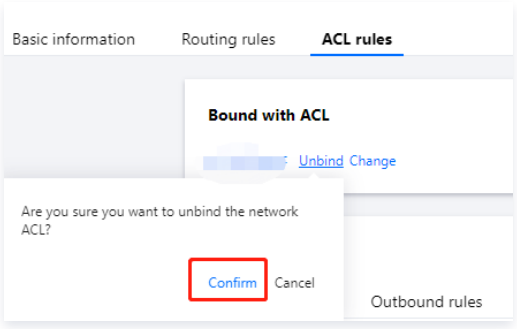

# Enabling or Disabling Broadcast

Last updated: 2024-01-12 14:39:39

## Background

Multicast and broadcast are modes of one-to-many communication, which can save businesses on the network bandwidth and reduce network load through point-to-multipoint efficient data transmission.
In the unicast mode, the initiating server sends data to N servers separately. In the multicast or broadcast mode, the server sends the same data to N servers at one time, which reduces the server resource consumption and also the bandwidth resource of the backbone network.

- Multicast: Tencent Cloud supports multicast on the VPC dimension.
- Broadcast: Tencent Cloud supports broadcast on the subnet dimension.

**Notes**

- The trial period has ended.
- Currently, multicast and broadcast are supported in the following regions: Beijing, Shanghai, Guangzhou, Chengdu, Chongqing, Nanjing, Hong Kong (China), Singapore, Seoul, Tokyo, Bangkok, Toronto, Silicon Valley, Virginia, Frankfurt, Shenzhen Finance, Shanghai Finance, and Beijing Finance.
- For single-VPC multicast and broadcast, up to 50,000 PPS and 190 Mbps are supported.

## Scenario

Multicast and broadcast are mostly used in the financial and game industries.

- In the financial industry, multicast and broadcast are mainly used for broadcast services or market data. For example, after obtaining stock prices and other real-time data, brokers can broadcast stock data to multiple clients in real time, effectively reducing network load.
- For the game industry, broadcast and multicast are mainly used for heartbeat holding between multiple servers.

This document describes how to enable or disable broadcast for subnets.

## Instructions

### Enabling broadcast

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar.
3. In the list, locate the row of the Virtual Private Cloud that requires broadcast functionality, click **Enable** under Subnet Broadcast, and confirm the operation.

### Disabling broadcast

1. Log in to the VPC console.
2. Click **Subnets** on the left sidebar.
3. In the subnet list, locate the subnet that needs to disable broadcast, toggle off **Subnet broadcast**, and confirm the operation.

## Related Actions

For information about VPC-level multicast, see Enabling or Disabling Multicast.

# Deleting a Subnet

Last updated: 2024-01-12 14:39:46

You can delete the subnets that are no longer in use and do not have any IP resources.

> **Note**
> Currently, Tencent Cloud resources that involve IP use in subnets include CVM, private network CLB, ENI, HAVIP, SCF, TKE, and TencentDB (for MySQL, Redis, TDSQL, etc.).

## Instructions

1. Log in to the VPC console.
2. Click **Subnet** on the left sidebar to access the management page.
3. At the top of the list, select the region and VPC that the target subnet belongs to.
4. In the list, locate the row containing the subnet you want to delete, click **Delete** in the operation column, and click **OK**.

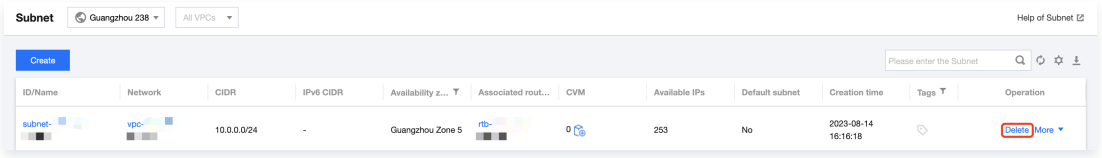

# Route tables
# Overview

Last updated: 2024-01-12 14:40:59

A route table consists of multiple routing policies that control the outbound traffic direction of subnets in the VPC. Each subnet can only be associated with one route table, while each route table can be associated with multiple subnets. You can create multiple route tables for subnets with different traffic routes.
Watch video

## Types

There are two types of route tables: default and custom.

- **Default route table**: When you create a VPC, the system automatically generates a default route table, which will be associated with subnets created later if no custom route table is selected. You cannot delete the default route table, but you can add, delete, and modify routing policies in it.
- **Custom route table**: You can create or delete a custom route table in the VPC. The custom route table can be associated with all subnets to apply the same routing policy. Before deleting the custom route table, you need to first disassociate it from all the subnets.

> **Note**
> You can associate a route table when creating a subnet or change the associated route table after a subnet is created.

## Routing policies

A route table controls traffic routes by using routing policies. A routing policy consists of the destination, next hop type, and next hop:

- **Destination**: Specifies the destination IP range to which you want to forward the traffic. It should be an IP range. If you want to enter a single IP address, set the mask to `32` (for example, `172.16.1.1/32`). The destination cannot be an IP range of the VPC where the route table resides, because the local route already allows private network interconnection in this VPC.

> **Note**
> - If you have deployed Tencent Kubernetes Engine in your VPC, when you create a route table policy for a VPC subnet, the destination IP range cannot be within the VPC IP range or the container IP range.
> - If the container network and VPC routes overlap, traffic will be preferentially forwarded within the container network.

- **Next hop type**: Indicates the egress of data packets for the VPC. The next hop type of VPC supports NAT gateway, peering connection, VPN gateway, direct connect gateway, CVM, and others.
- **Next hop**: Specifies the next hop instance (identified by the next hop ID) to which the traffic is forwarded, such as a NAT gateway in the VPC.

### Routing policy priority

When there are multiple routing policies in a route table, the following routing priority applies, from high to low:

- Traffic within the VPC: Traffic within the VPC is matched first.
- Exact match route (the longest prefix match): When there are multiple routes in the route table that can match the destination IP, the route with the longest (exact) mask is matched to determine the next hop.
- Public IP: If no routing policy is matched, a CVM instance can access the internet through its public IP address.
  **Example:**
  When a subnet is associated with a NAT gateway and the CVM in the subnet has a public IP (or EIP), the CVM accesses the internet through the NAT gateway by default (because the priority of the exact match route is higher than that of the public IP). However, you can set a routing policy to allow the CVM to access the internet by using its public IP address. For details, see Adjusting the Priorities of NAT Gateways and EIPs.

## ECMP

Equal-cost multipath (ECMP) routing means there are multiple equal-cost routes to a single destination. The traditional routing technology only uses one path to transfer packets to the same destination, while the remaining paths are in the standby or invalid status. When the used path fails, it takes time to switch to another path. By contrast, ECMP uses multiple equal-cost routes in the

network environment to increase the transfer bandwidth, balance traffic over multiple routes, and achieve backup with redundant linkages.

- ECMP with VPC routes of the same type is as detailed below:

| Next Hop Type | Whether ECMP Is Formed with Routes of the Same Type | Maximum Number of Routes Supported by ECMP |
|---|---|---|
| NAT gateway | Yes | N/A |
| CVM public IP | No | N/A |
| CVM | Yes | Eight routes of the same type |
| Peering connection | No | N/A |
| Direct connect gateway | No | N/A |
| CCN | No | N/A |
| High availability virtual IP | Yes | Eight routes of the same type |
| VPN gateway | Yes | Eight routes of the same type |

- ECMP with VPC routes of different types is as detailed below:
  - NAT gateways and CVM instances can form the ECMP.
  - If there is already a self-learning CCN route, when a configured custom route to a direct connect gateway/peering connection is added, CCN and the direct connect gateway/peering connection can form the ECMP.
  - If there is already a custom route for the direct connect gateway/peering connection, and you want to form the ECMP with CCN, please contact our online customer service.

### Scenarios

ECMP is often used to balance the traffic load over gateways with a limited bandwidth. Assume that you need 2,000 Mbps to interconnect your VPC-based and IDC-based businesses, but the current maximum VPN bandwidth is 1,000 Mbps. To achieve the goal, you can create two 1,000-Mbps VPN gateways and two VPN tunnels.

## Primary/secondary routes

Primary and secondary routes refer to two or more paths to the same destination with only one active path. Assume there are two VPC routes to the IDC, that is, paths A and B. All packets are sent to the destination via path A, while path B is invalid or on standby. When path A suffers linkage failures, you can switch to path B to take over traffic from path A, thus ensuring business availability. In this case, paths A and B are called primary and secondary routes.
The next hop type determines the route priority. When adding a routing policy to the VPC route table, you can configure different types of gateways to act as primary and secondary routes to a single destination. Then, the VPC network probe can be used to check the linkage quality and accessibility. After configuring an alarm policy, you can promptly detect any linkage exception and quickly switch between primary and secondary routes to meet the high availability requirements.

**Note**

- VPC does not have the route priority feature by default. This feature is currently in beta test. To try it out, please contact our online customer service.
- The next hop type determines the route priority in the VPC route table. The default route priority sequence from high to low is CCN, direct connect gateway, VPN gateway, and others.
- Currently, you cannot adjust the route priority in the console. If needed, please conta our online customer service.

The following table describes the primary/secondary support of different types of VPC routes:

| Next Hop Type | Support for primary/secondary routes |
|---|---|
| NAT gateway | No |

| | |
|---|---|
| CVM public IP | No |
| CVM | Yes, with CCN, VPN gateway, direct connect gateway, or HAVIP |
| Peering connection (intra-region) | No |
| Peering Connection (Cross-region) | No |
| Direct connect gateway | Yes, with CCN, VPN gateway, HAVIP, or CVM |
| CCN | Yes, with VPN gateway, direct connect gateway, HAVIP, or CVM |
| High availability virtual IP | Yes, with CCN, VPN gateway, direct connect gateway, or CVM |
| VPN gateway | Yes, with CCN, direct connect gateway, HAVIP, or CVM |

## Scenarios

Primary and secondary routes are often used to smoothly forward traffic when a gateway linkage fails, for example:

- VPC-based direct connect gateway (primary) and VPC-based VPN gateway (secondary)
  **Scenario:** Interconnect a Tencent Cloud VPC and an on-premises IDC through a VPC-based direct connect gateway. Meanwhile, create a VPN tunnel through a VPN gateway to act as the secondary communication linkage between the IDC and VPC.

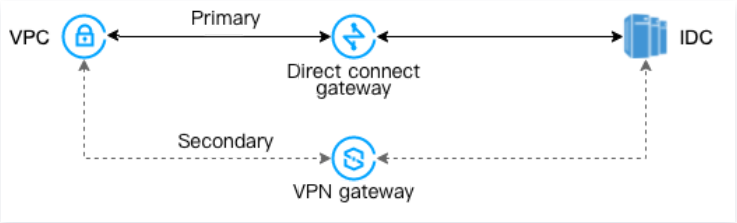


- Direct Connect Gateway (Primary) & VPC-based VPN Gateway (Secondary)
  **Scenario Description:** Users establish communication between the VPC in the cloud and their self-built IDC through a CCN-based direct connect gateway, while also creating a VPN backup channel via a VPN gateway to provide redundancy for the IDC and VPC communication link.

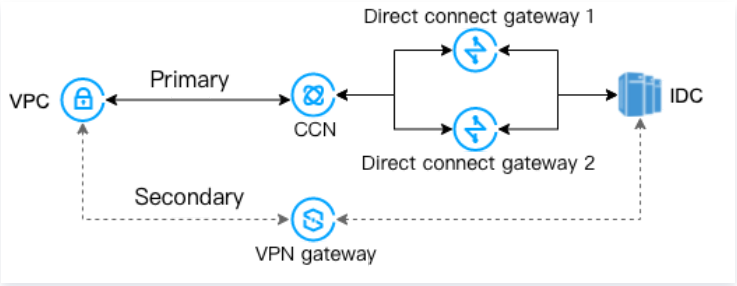

# Limits

Last updated: 2024-01-12 14:41:04

- The default route table for each Virtual Private Cloud cannot be deleted.
- After a VPC is created, a default routing policy will be automatically added to the VPC's route table, indicating that all resources in this VPC are interconnected through the private network. This routing policy cannot be modified or deleted.

| Destination | Next Hop Type | Next Hop |
| --- | --- | --- |
| Local | Local | Local |

- Dynamic routing protocols such as BGP and OSPF are not supported.
- Currently, only the following routes can be published to CCN:

| Next Hop Type | Publishing to CCN by default | Manually publishing or withdrawing | Note |
| --- | --- | --- | --- |
| Local | Yes | No | Assigned by the system. The VPC IP range connecting to CCN will be automatically published to CCN, including primary and secondary CIDR blocks (except for container IP ranges). |
| CVM | No | Yes | A custom route to a CVM. When the IP range is all 0 or the routing policy is disabled, the route cannot be published to CCN. |
| High availability virtual IP | No | Yes | A custom route to an HAVIP. When the IP range is all 0 or the routing policy is disabled, the route cannot be published to CCN. |

**Note**

- A disabled custom route cannot be published to CCN.
- A custom route that has been published to CCN should be withdrawn first before it can be disabled. For more information, see Publishing/Withdrawing a Routing Policy to/from CCN.
- An HAVIP not bound to any CVM cannot be published to CCN. Please retry after binding it to a CVM.

## Quota Limits

| Resources | Limit |
| --- | --- |
| Number of route tables per VPC | 10 |
| Number of route tables associated with each subnet | 1 |
| Number of routing policies per route table | 50 |

# Creating Custom Route Tables

Last updated: 2024-01-12 14:41:35

A route table consisting of multiple routing policies is used to control the outbound traffic of the subnet. There are a default route table and custom route tables. The default route table (local route) allows private network interconnection in the VPC, which cannot be deleted, but can be configured with routing policies the same way as you configure a custom route table. This document describes how to create and configure a custom route table.
Watch video

## Instructions

1. Log in to the VPC console.
2. Click **Route Tables** on the left sidebar to go to the route table management page.
3. Click **Create**.
4. In the pop-up window, enter a name for the route table, select the VPC to which the route table belongs, and configure a routing policy.

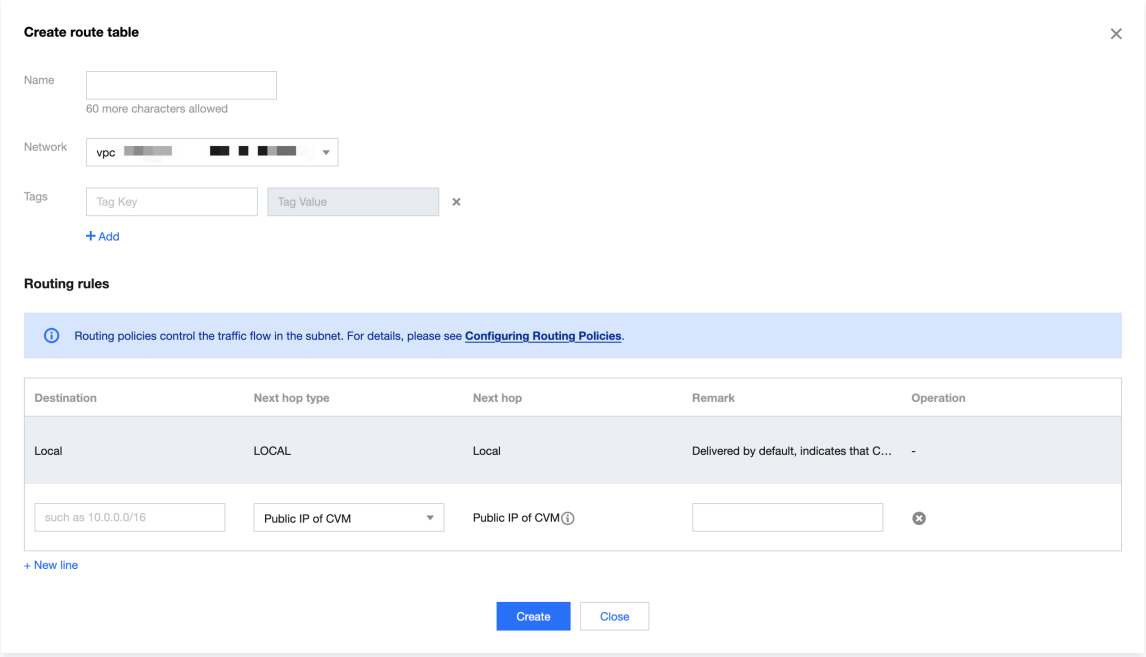


**Note**
Routing policies can be configured when creating a route table or after creating the route table by clicking the route table ID to enter the details page and clicking **Add Routing Policy** to configure. The configuration method is the same in both cases.

Configure a routing policy as instructed below:

| Parameter | Description |
| --- | --- |
| Destination | Refers to the destination IP range to which you want to forward traffic. Configure it as follows:<br>• Enter an IP range. If you want to enter a single IP, set the mask to 32 (for example, `172.16.1.1/32`).<br>• The destination cannot be an IP range of the VPC where the route table resides, because the local route already allows private network interconnection in this VPC.<br>**Note:**<br>When you deploy Tencent Kubernetes Engine in a VPC, the destination in the VPC subnet route table policy cannot be within the VPC CIDR range or include the container network segment. For example, if the VPC CIDR is `172.168.0.0/16` and the container network CIDR is `192.168.0.0/16`, the destination segment in the VPC subnet route table policy cannot be within the `172.168.0.0/16` range or include `192.168.0.0/16`. |

| | |
|---|---|
| Next Hop Type | Refers to the egress of data packets for the VPC. Supported types:<br>• Public NAT Gateway: Traffic directed to a destination IP range is forwarded to a public NAT gateway.<br>• Private NAT gateway: Traffic directed to a destination IP range is forwarded to a private NAT gateway.<br>• Peering connection: Traffic directed to a destination IP range is forwarded to the VPC peer of a peering connection.<br>• Direct connect gateway: Traffic directed to a destination IP range is forwarded to a direct connect gateway.<br>• High availability virtual IP: Traffic directed to a destination IP range is forwarded to an HAVIP.<br>• VPN gateway: Traffic directed to a destination IP range is forwarded to a VPN gateway.<br>• Public IP of CVM: Traffic directed to a destination IP range is forwarded to the public IP (including EIPs) of a CVM instance in the VPC.<br>• CVM: Traffic directed to a destination IP range is forwarded to a CVM instance in the VPC.<br>• CDC local gateway: Tencent Cloud CDC communicates with the customer IDC by using a CDC local gateway. |
| Next Hop | Refers to the next hop instance to which the traffic is redirected, such as a gateway or CVM IP. |
| Remark | Refers to the route description for resource management. This parameter is optional. |
| New line | You can click **New line** to configure multiple routing policies or click the deletion icon in the operation column to delete unwanted routing policies. A custom route table should contain at least one routing policy. |

5. After completing the parameter configuration, click **Create** to finish configuring the route table and policies. Then the route table will be displayed in the list.

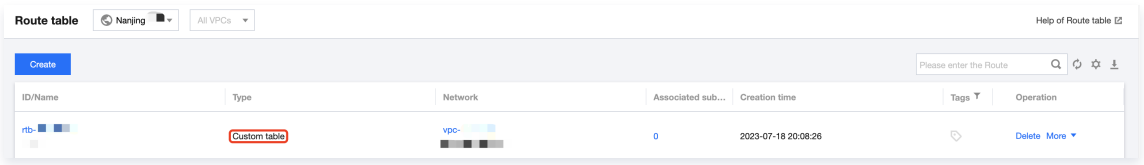


## See Also

Routing policies with **Next hop type** set to **High availability virtual IP or CVM** in a default or custom route table can be manually published to or withdrawn from CCN.

1. Click the route table ID to go to the details page.

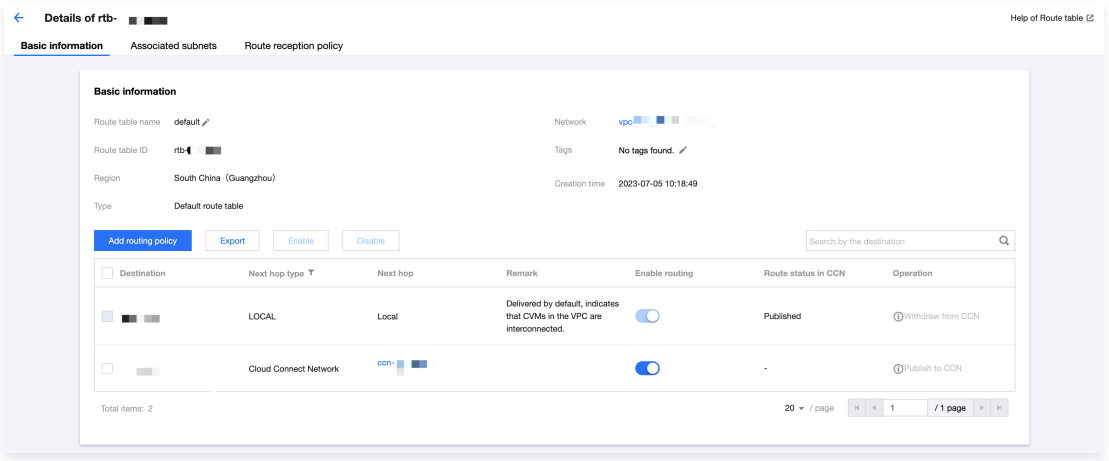


2. Perform the following operations as needed:
   - Click **Publish to CCN** to publish an enabled routing policy to CCN.
   - Click **Withdraw from CCN** to withdraw a custom routing policy that has been published to CCN.
   - Click **Edit** to modify a routing policy.
   - Click **Delete** to delete a disabled routing policy.

# Associating or Disassociating Subnet

Last updated: 2024-01-12 14:41:50

After creating a route table, you need to associate it with a subnet to control the outbound traffic of the subnet. This document describes how to associate a route table with a subnet and disassociate a route table from a subnet.

## Associating a route table with a subnet

1. Log in to VPC console.
2. Click **Route Tables** on the left sidebar to go to the route table management page.
3. There are two methods to associate a route table with a subnet:

- In the list, locate the route table that needs to be associated with a subnet, and click **More** > **Associated subnets** in the operation column.

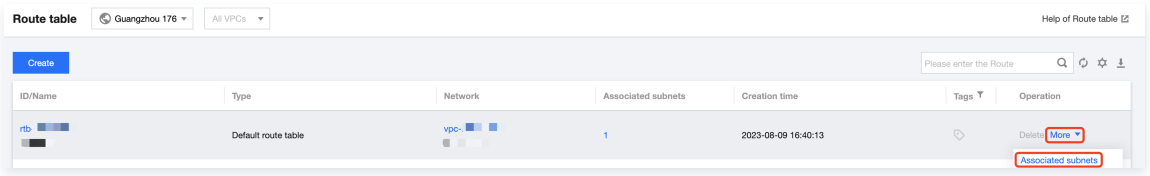


- Click the ID of the route table to go to its details page, select the **Associated subnets** tab, and click **Add associated subnet**.

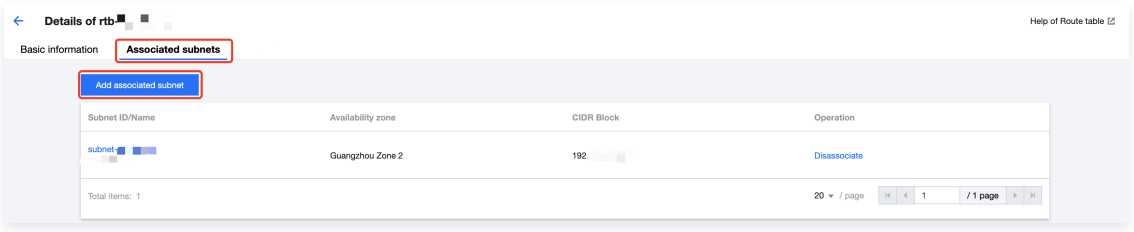


4. In the pop-up window, select one or more subnets to associate (a route table can be associated with multiple subnets at the same time, and you can quickly filter by subnet ID/name). Please evaluate the business impact of the association on the subnet. Confirm the impact, and click **OK**.

> **Note**
> After the route table is associated with the subnet, the original route table associated with the subnet will be replaced with the new one, and the subnet outbound traffic will be executed according to the policies in the new route table. Please carefully evaluate the business impact.

## Disassociating a route table from a subnet

1. Log in to VPC console.
2. Click **Route Tables** on the left sidebar to go to the route table management page.
3. Click the ID of a route table to go to its details page, select the **Associated subnets** tab, and click **Disassociate**.


4. In the pop-up window, select a new route table for the subnet to be disassociated, and click **OK** to complete disassociating the current route table from the subnet. The subnet outbound traffic will be executed according to the policies in the new route table selected for it.

# Managing Routing Policies

Last updated: 2024-01-12 14:41:55

The routing policies in a route table can be managed in real time. For example, you can add, delete, query, and export routing policies, publish routing policies to CCN, withdraw routing policies from CCN, and enable or disable routing policies. This document describes operations related to routing policies.

## Adding a routing policy

1. Log in to the VPC console and go to the Route Tables page.
2. Click the **ID/Name** of the route table to modify to go to its details page.
3. Click **Add Routing Policy**.

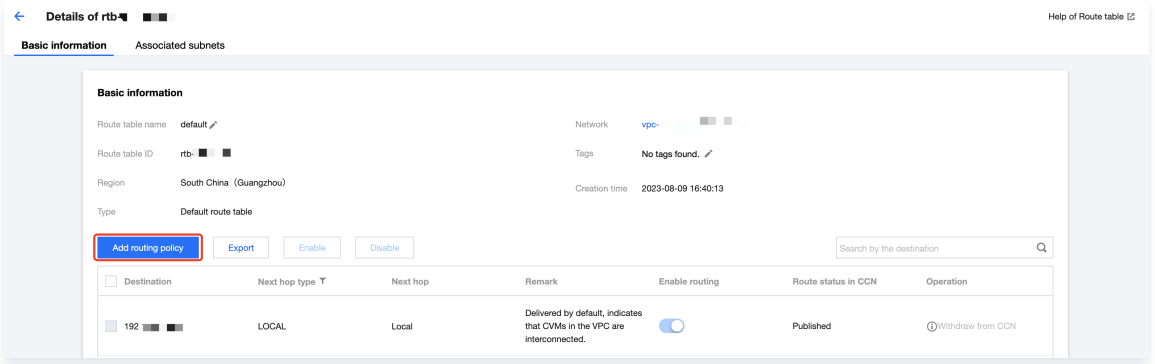


4. In the pop-up window, configure the routing policies.

> **Note**
> When deploying Tencent Kubernetes Engine in a VPC, the destination IP range for the VPC subnet routing policy must not be within the VPC CIDR block or include the container IP range. For example, if the VPC CIDR is 172.168.0.0/16 and the container network CIDR is 192.168.0.0/16, the destination IP range for the VPC subnet routing policy must not be within the 172.168.0.0/16 range or include 192.168.0.0/16.

| Parameter | Note |
| --- | --- |
| Destination | Refers to the destination IP range to which you want to forward traffic. Configure it as follows:<br>• Enter an IP range. If you want to enter a single IP, set the mask to 32 (for example, `172.16.1.1/32`).<br>• The destination cannot be an IP range of the VPC where the route table resides, because the local route already allows private network interconnection in this VPC. |
| Next Hop Type | Refers to the egress of data packets for the VPC. Supported types:<br>• NAT gateway: Traffic directed to a destination IP range is forwarded to a NAT gateway.<br>• Peering connection: Traffic directed to a destination IP range is forwarded to the VPC peer of a peering connection.<br>• Direct connect gateway: Traffic directed to a destination IP range is forwarded to a direct connect gateway.<br>• High availability virtual IP: Traffic directed to a destination IP range is forwarded to an HAVIP.<br>• VPN gateway: Traffic directed to a destination IP range is forwarded to a VPN gateway.<br>• Public IP of CVM: Traffic directed to a destination IP range is forwarded to the public IP (including EIPs) of a CVM instance in the VPC.<br>• CVM: Traffic directed to a destination IP range is forwarded to a CVM instance in the VPC.<br>• CDC local gateway: Tencent Cloud CDC communicates with the customer IDC by using a CDC local gateway. |
| Next Hop | Refers to the next hop instance to which the traffic is redirected, such as a gateway or CVM IP. |
| Remark | Enter the route description for resource management. This parameter is optional. |
| New line | You can click **New line** to configure multiple routing policies or click the deletion icon in the operation column |

to delete unwanted routing policies.

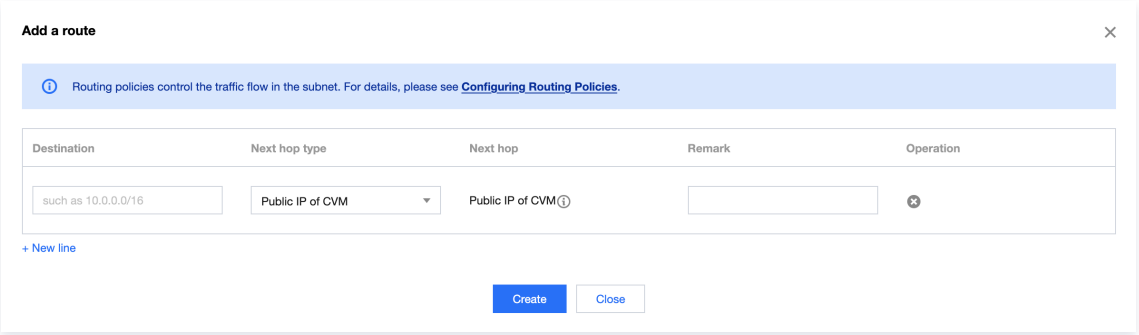


5. Click **Create**.

## Editing a routing policy

1. Log in to the VPC console and go to the Route Tables page.
2. In the list, click the **ID/Name** of the target route table to go to its details page.
3. Click **Edit** on the right side of the routing policy to modify the routing entry.

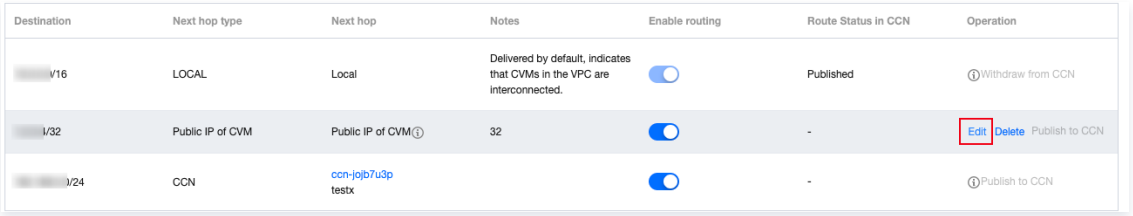


4. Click **OK** to confirm the modification or **Cancel** to cancel the modification.

## Publishing/Withdrawing a routing policy to/from CCN

Routes of a VPC associated with a CCN are published to the CCN by default. For new custom routing policies that are not published, you need to manually publish them. You can also withdraw a routing policy from CCN.
Only routing policies with **Next hop type** set to **High availability virtual IP or CVM** in a default or custom route table can be manually published to or withdrawn from CCN.

### Preparations

The VPC where the HAVIP or CVM resides is associated with a CCN instance.

### Instructions

1. Log in to the VPC console and go to the Route Tables page.
2. Click the **ID/Name** of the route table to modify to go to its details page.
3. Perform the following operations as needed:
- Click **Publish to CCN** to manually publish a custom routing policy to CCN.
- Click **Withdraw from CCN** to withdraw a custom routing policy that has been published to CCN.

> **Note**
> - A disabled routing policy cannot be published to CCN.
> - A routing policy cannot be disabled once being published to CCN.

## Querying and exporting a routing policy

1. Log in to the VPC console and go to the Route Tables page.
2. Click the **ID/Name** of the target route table to go to its details page. On this page, you can view the routing policies in this route table.

3. In the top-right search box, query the routing policies by entering a destination address.

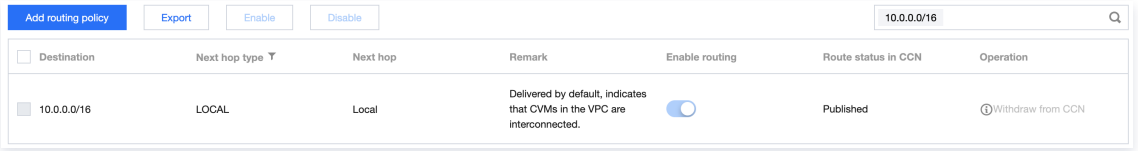


4. Click **Export** to save the search result in a .csv file.

## Enabling/Disabling a routing policy

A custom routing policy can be enabled or disabled.

### Instructions

1. Log in to the VPC console and go to the Route Tables page.
2. Click the **ID/Name** of the target route table to enter its details page. Check the routing policy status:

- indicates that the routing policy is enabled.
- indicates that the routing policy is in a disabled state.

3. Click the icon on the right of an enabled routing policy to disable it.

> **Note**
> Disabling a routing policy may result in business interruption. Please double check before continuing.

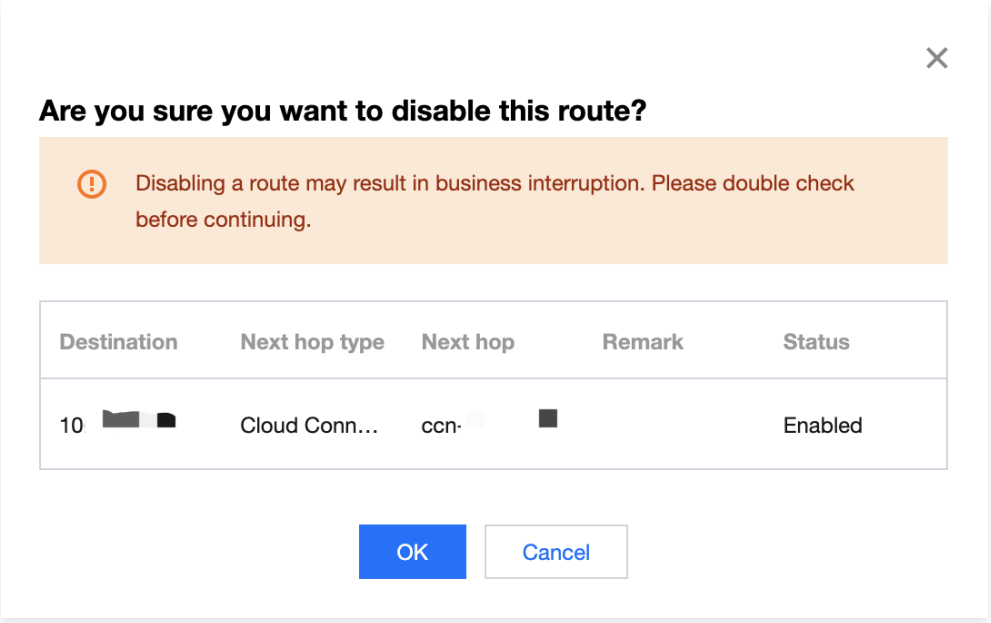


4. Click the icon on the right of a disabled routing policy to enable it.

> **Note**
> The enabled routing policy with the longest mask will be used. This may affect your current business. Please double check before continuing.

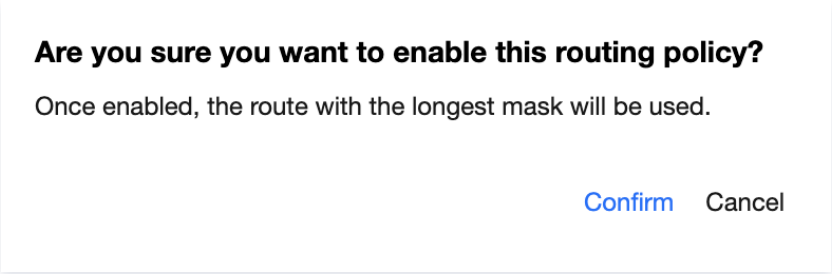


5. You can also enable or disable multiple routing policies at one time. Select multiple routing policies and click **Enable** or **Disable** above the list.

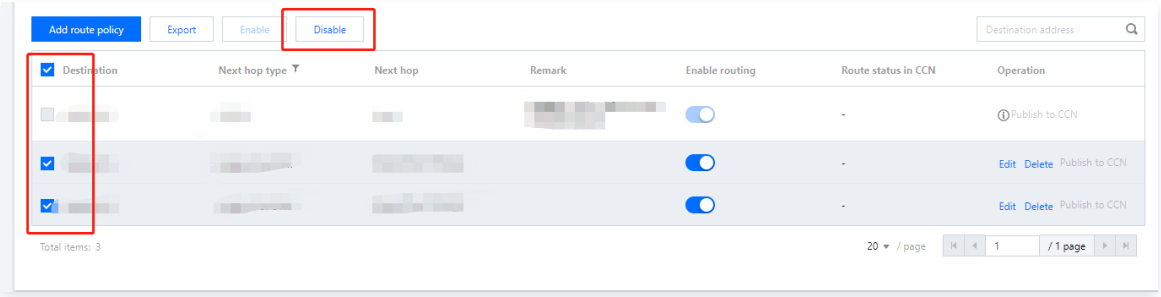


## Deleting a routing policy

You can delete an unwanted routing policy. Only custom routing policies can be deleted.

1. Log in to the VPC console and go to the Route Tables page.
2. Click the **ID/Name** of the route table to modify to go to its details page.
3. Click **Delete** on the right of the routing policy you want to delete.

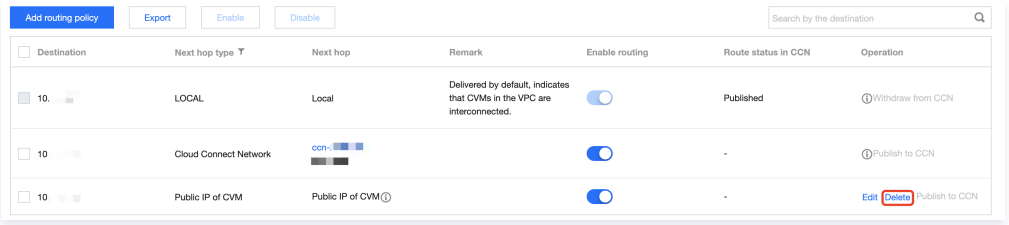


4. Read the notes and click **OK**.

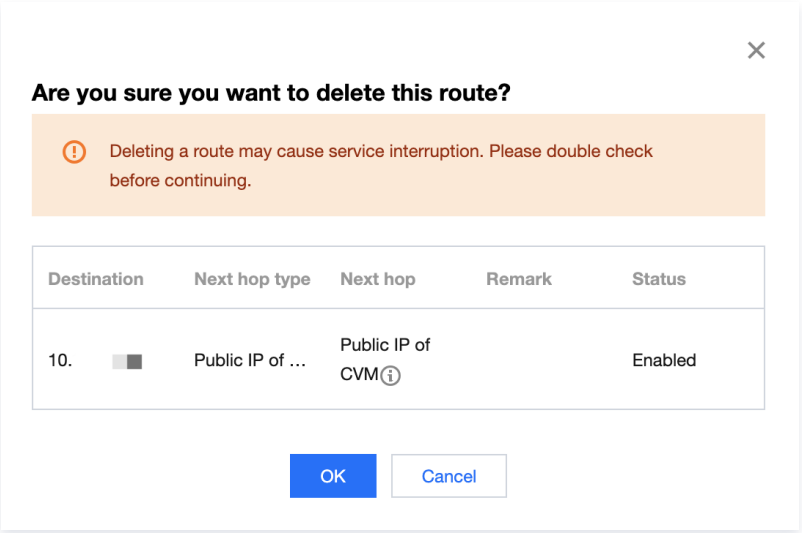

# Deleting a Routing Table

Last updated: 2024-01-12 14:42:01

## Preparations

You can delete a route table that is not associated with any subnet. Only custom route tables can be deleted, while default route tables automatically generated by the system cannot be deleted.

> **Note**
> Deleting a route table may result in business interruption. Please double check before continuing.

## Instructions

1. Log in to the VPC console and go to the Route Tables page.
2. In the list, locate the route table you want to delete, click **Delete** in the operation column, and confirm the operation.

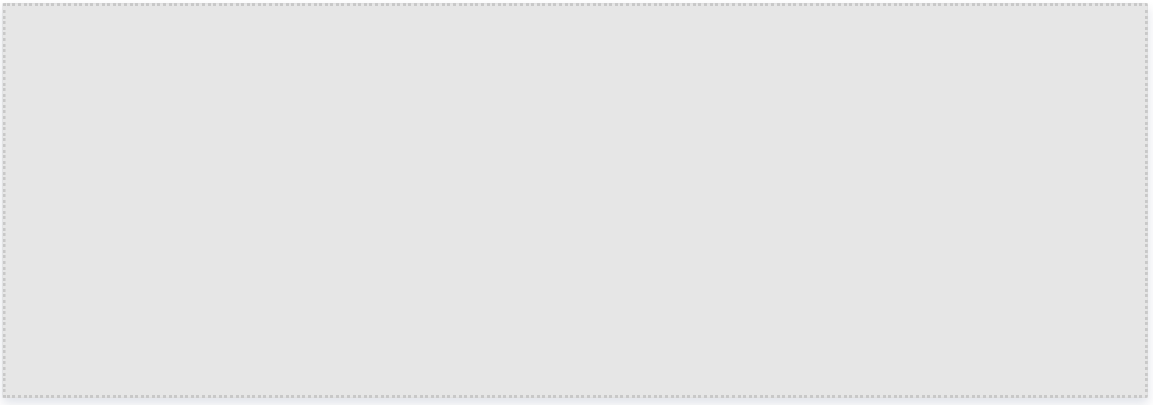

# IPs and ENIs
# Elastic Public IP

Last updated: 2024-01-12 14:43:07

Elastic IP (EIP) is a static IP designed for dynamic cloud computing and a fixed public IP in a certain region. With EIP, you can quickly map an address to another CVM instance or NAT gateway instance in your account to avoid instance failure.
You can keep an EIP in your account until it is released. Unlike a public IP, which can only be released with the CVM, an EIP can be decoupled from the CVM lifecycle and operated independently as a cloud resource. Therefore, if you need to retain a public IP strongly related to your business, you can convert it into an EIP and keep it in your account.
For the common operations of EIP, please see:

- Applying for an EIP
- Binding an EIP to a Cloud Resource
- Building Up an IPv4 VPC

# EIP IPv6

Last updated: 2024-01-12 14:43:13

Elastic IPv6 is the public gateway for IPv6 CVMs, enabling IPv6 CVMs to access the public network. After an IPv6 address is assigned to a CVM, its public network access feature is disabled by default. With Elastic IPv6, you can enable or disable public network access and configure public network bandwidth for each CVM's IPv6 address.
For the common operations of Elastic IPv6, please see:

- Assigning and Releasing an IPv6 CIDR Block for a VPC
- Assigning and Releasing an IPv6 CIDR Block for a Subnet
- Applying for and Releasing an IPv6 Address for an ENI
- Managing IPv6 Public Network
- Building Up an IPv6 VPC

# HAVIPs
# Overview

Last updated: 2025-10-11 15:52:43

A high availability virtual IP (HAVIP) is a private IP address assigned from the VPC subnet CIDR block. It is usually used together with high availability software, such as Keepalived and Windows Server Failover Cluster, to build a highly available primary/secondary cluster.

**Note**
- HAVIP is currently under beta testing. Switching between primary/secondary servers may take 10 seconds. To try it out, please submit a ticket.
- To ensure the high availability of CVMs in a primary/secondary cluster, we recommend assigning CVMs to different hosts using placement groups. For more information about placement groups, see Placement Group.
- The high availability software should support sending ARP messages.

## Features

- You can apply for multiple HAVIP addresses in the console for each VPC.
- You must bind the HAVIP in CVM's configuration file.

## Architecture and principle

Typically, a high availability primary/secondary cluster consists of two servers: an active primary server and a standby secondary server. The two servers share the same virtual IP (VIP). The VIP can only work on one primary server at the same time. When the primary server fails, the secondary server will take over the VIP to continue providing services.

- In traditional physical networks, the primary/secondary status can be negotiated with Keepalived's VRRP protocol. The primary device periodically sends free-of-charge ARP messages to purge the MAC table or terminal ARP table of the uplink switch to trigger the VIP migration to the primary device.
- In a VPC, a high availability primary/secondary cluster can also be implemented by deploying Keepalived on CVMs. However, a CVM instance usually cannot obtain a private IP through ARP announcement due to security reasons such as ARP spoofing. The VIP must be an HAVIP applied from Tencent Cloud, which is subnet-specific. Therefore, an HAVIP can only be bound to a server under the same subnet through announcement.

  **Note**
  Keepalived is a VRRP-based high availability software tool. To use Keepalived, first complete its configuration in the `Keepalived.conf` file.

The following figure shows the HAVIP architecture.

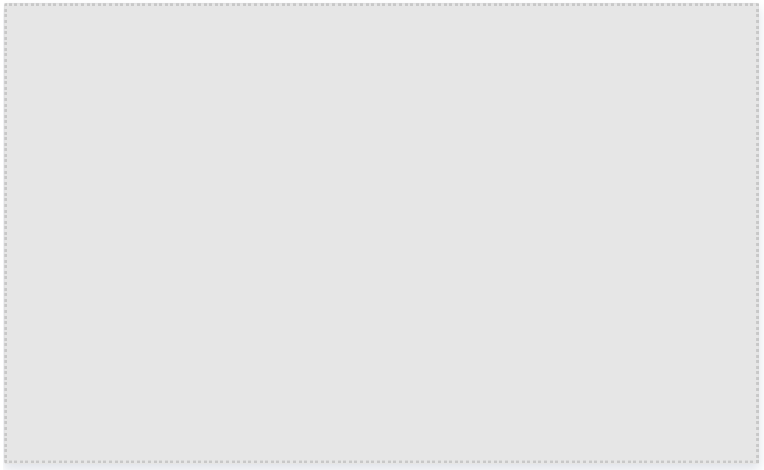

According to the example figure, CVM1 and CVM2 can be built into a high availability primary/secondary cluster with the following steps:

1. Install Keepalived on both CVM1 and CVM2, configure HAVIP as VRRP VIP, and set the priorities of the primary and secondary servers. Larger values represent higher priorities.
2. Keepalived uses the VRRP protocol to compare the initial priorities of CVM1 and CVM2 and determines CVM1 as the primary server due to its higher priority.
3. The Master server sends out ARP messages, announcing the VIP (which is the HAVIP) and updating the VIP and MAC address mapping. At this point, the Master server is the one providing services externally, with the HAVIP as its private network IP. Simultaneously, in the HAVIP console, you can see that the HAVIP is bound to the Master server CVM1.
4. (Optional) Bind an EIP to the HAVIP in the console to implement communication over the public network.
5. The primary server periodically sends VRRP messages to the secondary server. If the primary server fails to send VRRP messages within a certain period, the secondary server will be set as primary and send out ARP update messages that carry its MAC address. In this case, CVM2 becomes the primary server to provide communication services and handle external access requests. You will see the CVM bound to the HAVIP changes to CVM2 in the HAVIP console.

## Common Use Cases

- **Cloud Load Balance HA**
  To deploy Cloud Load Balancers (CLBs), you will generally implement HA between the CLB instances and configure the backend servers as a cluster. Therefore, you must deploy and use an HAVIP as a virtual IP between the two CLB servers.
- **Relational database primary/secondary**
  If Keepalived is used between two databases to build a highly available primary/secondary cluster, use an HAVIP as a virtual IP. For more information, see Building a High Availability Primary/Secondary Cluster with HAVIP + Keepalived .

## FAQs

### Why should I use an HAVIP along with Keepalived in a VPC?

Some public cloud vendors do not support binding a private IP to CVM through ARP announcement due to security reasons such as ARP spoofing. If you directly use a private IP as virtual IP in the `keepalived.conf` file, Keepalived will not be able to update the IP to MAC mapping during the primary/secondary server virtual IP switch. In this case, you have to call an API to switch the IP.
Using Keepalived configuration as an example, the IP configurations are as follows:

```
vrrp_instance VI_1 {
    state BACKUP           #Secondary device
    interface eth0          #ENI name
    virtual_router_id 51
    nopreempt                   #Non-preempt mode
    #preempt_delay 10
    priority 80
    advert_int 1
    authentication {
        auth_type PASS
        auth_pass 1111
    }
    unicast_src_ip 172.17.16.7   # Local private IP
    unicast_peer {
        172.17.16.13           #IP address of the peer device, for example: 10.0.0.1
    }

    virtual_ipaddress {

        172.17.16.3 #Enter the HAVIP address you have applied for in the console.

    }

    garp_master_delay 1
    garp_master_refresh 5

    track_interface {
        eth0
    }
```

```
    track_script {
        checkhaproxy
    }
}
```

If there is no HAVIP, the following section of the configuration file will be invalid.

```
virtual_ipaddress {
   172.17.16.3 #Enter the HAVIP address you have applied for in the console.
}
```

## See Also

- For more information about the use limits of HAVIP, see Limits.
- For more information about the operation guide of HAVIP, see Managing HAVIPs.

# Limits

Last updated: 2024-01-12 14:43:24

## Usage Limits

- The backend CVM can announce its occupation of an HAVIP, but you cannot manually bind HAVIPs to a specified server in the console (the experience is consistent with that of a traditional physical machine).
- The backend RS but not the HAVIP determines whether to migrate based on the configuration file negotiation.
- Only Virtual Private Cloud is supported, while basic networks are not.
- Heartbeat detection must be done by an application on the CVM, but not by the HAVIP, which serves only as a floating private IP address announced by ARP (the experience is consistent with that of a traditional physical machine).
- An HAVIP not bound to any CVM cannot be published to the CCN. Please retry after binding it to a CVM. For more information, see Limits.

## Quota Limits

| Resources | Restrictions |
| --- | --- |
| Number of default HAVIPs per VPC | 10 |

# Managing HAVIP

Last updated: 2025-10-11 15:53:32

This document describes how to create an HAVIP in the VPC console and configure it in third-party software.

> **Note**
> HAVIP is now only available for beta users. To try it out, please submit a ticket.

## Instructions

1. Log in to the Virtual Private Cloud Console, and on the left sidebar, select **IP & Network Interface** > **Highly Available Virtual IP**.
2. Select the target region on the HAVIP management page and click **Apply**.
3. In the pop-up dialog box, configure HAVIP parameters.
   - **Name**: Enter a name for the HAVIP.
   - **Virtual Private Cloud**: Select a VPC where the HAVIP will reside.
   - **Subnet**: Select a subnet for the HAVIP.
   - **IP address**: The IP address of the HAVIP can be automatically assigned or manually specified. If you choose **Automatic assignment**, a subnet IP address will be automatically assigned. If you choose **Custom**, make sure that the entered IP address is within the subnet IP range and is not a reserved IP address of the system. For example, if the subnet IP range is `10.0.0.0/24`, the entered private IP address should be within `10.0.0.2-10.0.0.254`.
4. Click **Confirm**. The successfully created HAVIP will be displayed in the list with a status of **Not Bound to Cloud Virtual Machine**.

## See Also

After the HAVIP is created, you need to manage it in the third-party HA software. An HAVIP is only an operation object and a private IP address that can be bound through the third-party HA software. It cannot be bound to/unbound from CVMs in the Tencent Cloud console. Instead, you need to specify the HAVIP as a floating VIP in the third-party HA software, which then specifies an ENI to be bound to the HAVIP through ARP. The schematic diagram is as follows:

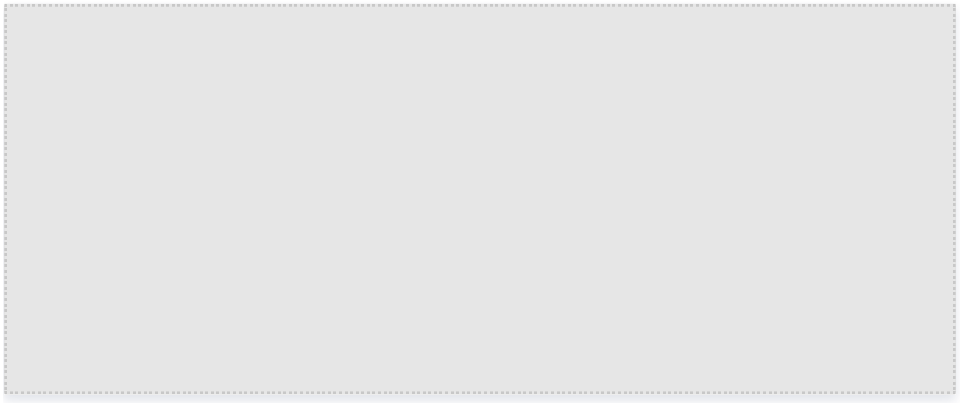

In a traditional physical device environment, all private IP addresses can be bound to ENIs through ARP and specified as floating IP addresses in the HA software by default. In a public cloud environment, a private IP cannot use ARP or be specified as a floating IP address in the HA software. Therefore, you need to follow the same steps as that of the third-party software to specify the HAVIP as a floating IP address instead.

> **Note**
> Common HA software programs include Linux HeartBeat, Keepalived, Pacemaker, and Windows MSCS.

When specifying a VIP in the configuration file of the HA software, you only need to enter the HAVIP that you created. See the example below:

```
vrrp_instance VI_1 {
Select proper parameters for the primary and secondary CVMs.
    state MASTER              #Set the initial status to Secondary.
```

```
    interface eth0              # The ENI such as eth0 used to bind a VIP
    virtual_router_id 51        # Configure the virtual_router_id value for the cluster
        nopreempt                    #Non-preempt mode
        preempt_delay 10            #Set the preempt delay to 10 minutes
    priority 100                #Priority. The larger the value, the higher the priority.
    advert_int 1                #Check interval. The default value is 1 second.
    authentication {            #Authentication
        auth_type PASS           #Authentication method
        auth_pass 1111           #Authentication password
    }
        unicast_src_ip 172.16.16.5 # Set the local device's private IP address
        unicast_peer{
        172.16.16.6                #IP address of the peer device
        }
    virtual_ipaddress {
        172.16.16.12               #Set the "HAVIP" as a floating IP.
    }
    }
```

After configuring the HAVIP in the HA software of the Cloud Virtual Machine, the status of the HAVIP in the console will change to **Bound to Cloud Virtual Machine**.
See the following cases for your configurations: Best Practice – Building HA Primary/Secondary Cluster with HAVIP + Keepalived

## Documentation

Similar to a private IP, an HAVIP can also be bound to or unbound from an EIP in the VPC console. For details, see Binding or Unbinding EIPs.

# Binding or Unbinding EIP

Last updated: 2025-10-11 15:54:03

Similar to a private IP, HAVIP binding can also be configured in the console. Binding an HAVIP refers to EIP operations. You can skip this section if no public network connection is needed.

> **Note**
> HAVIP is now only available for beta users. To try it out, please submit a ticket.

## Binding an EIP

1. Log in to the Virtual Private Cloud Console, and on the left sidebar, select **IP & Network Interface** > **Highly Available Virtual IP**.
2. Select the target region on the HAVIP management page.
3. Select the HAVIP to which you want to bind the EIP, and click **Bind** in the operation column on the right.
4. In the pop-up dialog box, select an EIP to bind.

> **Note**
> - An HAVIP can only be bound with one EIP. If no EIP is available, you must first create an EIP in the console.
> - If the HAVIP is not bounded with any CVM instance, the corresponding EIP will be in **idle** status and incur an idle fee. Please configure the HAVIP correctly and bind it to an instance by referring to the following cases: Best Practices – Building a Highly Available Primary-Secondary Cluster with HAVIP + Keepalived

5. Click **OK** to complete the EIP binding.

## Unbinding an EIP

1. Log in to the Virtual Private Cloud Console, and on the left sidebar, select **IP & Network Interface** > **Highly Available Virtual IP**.
2. Select the target region on the HAVIP management page.
3. Select the HAVIP from which the EIP will be unbound, and click **Unbind EIP** under the operation column on the right.
4. In the pop-up dialog box, read the notes and click **OK** to unbind the EIP.

> **Note**
> - Your public network business may be affected after unbinding the EIP. Please get ready in advance.
> - After being unbound, the EIP will become idle and incur an idle fee. You can directly release unused EIPs to avoid costs.

# Querying HAVIPs

Last updated: 2024-01-12 14:43:39

You can view all HAVIP details in a specific region on the HAVIP console.

> **Note**
> HAVIP is now only available for beta users. To try it out, please submit a ticket.

## Instructions

1. Log in to VPC console.
2. Select **IP and ENI** > **HAVIP** on the left sidebar to enter the HAVIP management page.
3. Select a **region** to view the detailed information of all HAVIPs applied in that region.

   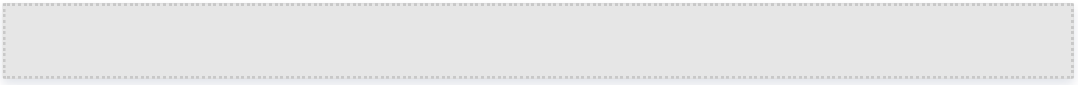

   The following gives a description of the fields:
   - **ID/Name**: An ID is generated automatically when an HAVIP is created. You can set a custom name for the HAVIP. Click the ID to view the basic information of the HAVIP.
   - **Status**: It indicates whether the HAVIP is specified as a floating VIP in the configuration file of the HA software on the CVM. If yes, the status of the HAVIP is **Bound with CVM**; otherwise, the status is **Not bound with CVM**.
   - **Address**: HAVIP address.
   - **Backend ENI**: ENI ID of the CVM bound with the HAVIP. If the HAVIP is not bound with any CVM, this field is **-**.
   - **Server**: ID of the CVM bound with the HAVIP. If the HAVIP is not bound with any CVM, this field is **-**.
   - **EIP**: EIP bound with the HAVIP. If the HAVIP is not bound with any EIP, this field is **-**.
   - **Network**: VPC of the HAVIP.
   - **Subnet**: Subnet of the HAVIP.
   - **Application time**: The time when the HAVIP is applied for.
   - **Operation**: Operations that can be performed on the HAVIP, including binding an EIP, unbinding an EIP, and releasing the HAVIP.
     - **Bind EIP**: Binds an EIP.
     - **Unbind EIP**: Unbinds an EIP.
     - **Release**: Releases the HAVIP.
4. Enter an ID, name, or address in the search box on the right to quickly search for HAVIPs.
5. Click the icon next to the search box to refresh the page.

# Releasing HAVIPs

Last updated: 2024-01-12 14:43:44

This document describes how to release unused HAVIPs in the console.

> **Note**
> HAVIP is now only available for beta users. To try it out, please submit a request.

## Preparations

The HAVIPs to release are **not bound with any CVM**.

> **Note**
> Cancel the **CVM bindings** in the configuration file of the third-party HA software on the CVM.

## Instructions

1. Log in to the VPC console.
2. Select **IP and ENI** > **HAVIP** on the left sidebar. In the HAVIP list, locate the HAVIP to release.
3. Click **Release** in the operation column.
4. In the pop-up window, click **OK** to release the HAVIP.

# Elastic Network Interface

Last updated: 2024-01-12 14:43:49

An Elastic Network Interface (ENI) can be bound to a CVM within a VPC and freely migrated among CVMs. ENIs can help configure and manage networks, as well as develop highly reliable network solutions.
You can bind multiple ENIs to a CVM in the same availability zone (the specific number of ENIs depends on the CVM specifications) to develop a highly available network solution. Additionally, you can bind multiple private IPs to an ENI to enable the deployment of multiple IPs on a single CVM.
For the common operations of ENI, please see:

- Creating an ENI
- Binding to a CVM
- Unbinding from a CVM
- Deleting an ENI
- Applying for a Secondary Private IP
- Releasing a Secondary Private IP
- Binding an EIP
- Unbinding an EIP
- Modifying a Primary Private IP
- Modifying the Subnet of an ENI
- Applying for and Releasing an IPv6 Address

# Bandwidth Package

Last updated: 2024-01-12 14:43:54

Tencent Cloud Bandwidth Package (BWP) is a multi-IP aggregated billing method that significantly reduces public network fees. When public network traffic peaks are distributed across different time periods, BWP enables aggregated bandwidth billing, saving bandwidth costs compared to purchasing bandwidth for each device individually.
Bandwidth packages offer two types of billing modes: Prepaid and Postpaid (Daily Postpaid, Monthly Postpaid), catering to your various business scenarios.

> **Note**
> To use bandwidth packages in the "postpaid – monthly top 5" billing mode, please contact your sales rep for activation.

For common BWP operations, see:

- Viewing the Billable Bandwidth
- Downloading Usage Details
- Changing Billing Mode
- Managing IP Bandwidth Packages
- Managing Device Bandwidth Packages

# Product Overview

Last updated: 2024-01-12 14:44:08

Traffic Package (TP) is a prepaid public network traffic plan that can be used to automatically deduct the traffic consumption of various products. Compared with postpaid traffic, Traffic Package offers lower unit prices and supports unified viewing, analysis, and management of traffic usage, helping you reduce network costs and simplify budget management.
For common traffic package operations, see:

- Creating Traffic Packages
- Viewing Traffic Packages
- Returning Traffic Packages

# Network connection
# NAT Gateway

Last updated: 2024-01-12 14:44:15

NAT Gateway provides IP address translation services, including SNAT and DNAT. After VPCs are connected to a NAT gateway, CVMs deployed in the VPC can access the internet in a secure and fast manner even if they don't have access to the internet on their own.
For the common operations of NAT gateway, please see:

- Getting Started
- Modifying NAT Gateway Configuration
- Managing EIPs of NAT Gateway
- Managing SNAT Rules
- Managing Port Forwarding Rules
- Routing to NAT Gateway

# VPN connection

Last updated: 2024-01-12 14:44:20

VPN Connection is a transfer service based on network tunneling technology that enables connectivity between local IDCs and resources on Tencent Cloud. It can help you quickly build a secure, reliable, and encrypted tunnel on the internet.
For the common operations of VPN Connection, please see:

- VPN Gateway
- Customer Gateway
- VPN Tunnel
- Establishing a Connection Between VPC and IDC (SPD Policy)
- Connecting VPC to IDC (Destination Route)
- Connecting IDC to CCN

# Direct Connect

Last updated: 2024-01-12 14:44:25

Direct Connect provides a fast and secure connection between Tencent Cloud and your local IDC. You can connect Tencent Cloud computing resources in multiple regions with a single connection to implement flexible and reliable hybrid cloud deployment.
For the common operations of Direct Connection, please see:

- Getting Started
- Managing Connections
- Managing Direct Connect Gateways
- Dedicated Tunnels
- Migrating IDC to the Cloud Through CCN

# Cloud Connect Network

Last updated: 2024-01-12 14:44:30

Cloud Connect Network (CCN) is a service that enables private network interconnection between VPCs as well as between VPCs and local IDCs. It supports public + private network multi-point interconnection, route self-learning, linkage prioritization, and fast failure convergence.

For the common CCN operations, please see:

- Connecting Network Instances Under the Same Account
- Connecting Network Instances Across Accounts
- Instance Management
- Route Management
- Bandwidth Management

# Private connection

Last updated: 2024-01-12 14:44:35

Private Link provides Tencent Cloud VPCs with the capability to access other VPCs in the same region through a private network. It allows you to quickly establish access connections between VPCs under the same account or across accounts. Compared with public network services, it can save public bandwidth, enhance security, and greatly simplify the network architecture.
For the common operations of Private Link, please see:

- Getting Started
- Managing Endpoints
- Managing Endpoint Services
- Sharing Services to VPCs in Different Regions
- Sharing Services to VPCs Under Different Accounts

# Security management
# Security Groups
# Security Group Overview

Last updated: 2024-01-12 14:44:43

A security group is a virtual firewall that features stateful data packet filtering. It is used to configure the network access control of CVM, CLB, TencentDB, and other instances while controlling their outbound and inbound traffic. It is an important means of network security isolation. You can learn more about security groups in the following video.
Watch video
You can configure security group rules to allow or reject inbound and outbound traffic of instances within the security group.

## Features

- A security group is a logical group. You can add CVM, ENI, TencentDB, and other instances in the same region with the same network security isolation requirements to the same security group.
- If a security group has no rules, it will reject all traffic by default, and you need to add rules to it to allow traffic.
- Security groups are stateful. Inbound traffic you have allowed is automatically allowed to go out and vice versa.
- You can modify security group rules at any time, and new rules will take effect immediately.

## Usage Limits

For more information on the restrictions and quotas of security groups, please see Use Limits.

## Security Group Rules

### Components

A security group rule consists of:
- **Source** or **Destination**: traffic origin (inbound rules) or target (outbound rules). It can be a single IP address, an IP address range, or a security group. For more information, see Security Group Rules.
- **Protocol:port**: Protocol types such as TCP and UDP.
- **Policy**: Allow or reject.

### Rule priorities

- The rules in a security group are prioritized from top to bottom. The rule at the top of the list has the highest priority and will take effect first, while the rule at the bottom has the lowest priority and will take effect last.
- If there is a rule conflict, the rule with the higher priority will prevail by default.
- When traffic goes in or out of an instance bound to a security group, the security group rules will be matched sequentially from top to bottom. If a rule is matched successfully and takes effect, the subsequent rules will not be matched.

### Multiple Security Groups

An instance can be bound to one or multiple security groups. When it is bound to multiple security groups, the security groups are executed from top to bottom. You can adjust their priorities at any time.

## Security group templates

Tencent Cloud provides the following two security group templates:
- Template of opening all ports: All inbound and outbound traffic will be allowed to pass.
- Template of opening major ports: Port TCP 22 (for Linux SSH login), ports 80 and 443 (for web service), port 3389 (for Windows remote login), and protocol ICMP (for Ping commands) will be open to the public network. In addition, access from the private network (VPC IP ranges) will be allowed.

**Note**
- If these templates cannot meet your needs, you can create custom security groups. For more information, see Creating a Security Group and Use Cases of Security Groups.

- If you need to protect the application layer (HTTP/HTTPS), you can purchase Tencent Cloud Web Application Firewall (WAF), which provides web security at the application layer to defend against web vulnerabilities, malicious crawlers, and CC attacks, protecting your websites and web applications.

## Directions

The following figure shows how to use a security group:

## Security Group Best Practices

### Creating a security group

- We recommend you specify a security group when purchasing a CVM instance via the API; otherwise, the default security group will be used. The default security group cannot be deleted. It adopts the default security rule (that is, allowing all IPv4 addresses), which can be modified as needed after the security group is created.
- If you need to change the instance protection policy, we recommend you modify the existing rules instead of creating a security group.

### Managing rules

- Export and back up the security group rules before you modify them, so you can import and restore them if an error occurs.
- To create multiple security group rules, use a parameter template.

### Associating a security group

- You can add instances with the same protection requirements to the same security group, instead of configuring a separate security group for each instance.
- We recommend you not bind one instance to too many security groups, which may cause rule conflicts and result in network disconnection.

## Security Group and Cloud Firewall

Tencent Cloud Firewall (CFW) is a native Tencent Cloud SaaS firewall that integrates different capabilities, including vulnerability scanning, IPS intrusion blocking, internet-wide threat intelligence, and advanced threat source analysis, making it the traffic security and policy management center in the cloud environment. It also serves as the first security portal for cloud business.
In practical use cases, security groups are typically deployed at the boundaries of cloud products such as CVMs to implement access control between security groups. Tencent Cloud Firewall, on the other hand, is deployed at the boundaries between VPCs or the internet to implement access control between VPCs or from Tencent Cloud to the internet.
In scenarios where security groups cannot meet the requirements, you can use Tencent Cloud Firewall to implement access control:

1. You need to understand the exposure and vulnerabilities of CVM assets on the internet and strengthen protection against network vulnerabilities through the IPS intrusion prevention and virtual patching features.
2. When you need to control the proactive access to the Internet by domain name and enhance the business security.
3. You need to implement access control by region, for example, blocking all IPs outside the Chinese mainland quickly.

# Creating Security Group

Last updated: 2024-01-12 14:44:48

## Scenario

A security group is a virtual firewall for CVM instances. Each CVM instance must belong to at least one security group. Tencent Cloud provides two templates for creating a security group: **Open all ports** and <3>Open ports 22, 80, 443, 3389, and protocol ICMP<3>. For more information, see <5>Overview</5>. You can learn more about how to create security groups in the following video.
Watch video
If you do not want your CVM instance to join the default security group, you can create a security group for your CVM instance in the console.

## Instructions

1. Log in to the security group console.
2. On the security group management page, select a region at the top and click **Create**.
3. In the pop-up dialog box, complete the following configurations:



- **Template**: Based on the services to be deployed for the CVM instance in the security group, select an appropriate template to simplify security group rule configuration, as described in the following table:

| Template | Note | Scenario |
| --- | --- | --- |
| Open all ports | All ports will be open to the public and private networks, which however may incur security risks. | - |
| Open ports 22, 80, 443, 3389, and protocol ICMP | Ports 22, 80, 443, 3389, and protocol ICMP will be open to the public network. In addition, access from the private network will be allowed. | A web service needs to be deployed on instances in the security group. |
| Custom | After creating a security group, you can add security group rules as needed. For more information, see Adding a Security Group Rule. | - |

- **Name**: Set a name for the security group.

- **Associated Project**: The default selection is "Default Project". You can specify another project for easier management in the future.
- **Remark**: Briefly describe the security group to facilitate future management.
- **Advanced options**: You can configure tags for the security group here. By default, no tags are configured. For more information, see Tag Overview.

4. Click **OK** to complete the creation of the security group.
   If you select the **Custom** template when creating a security group, you can click **Add rules now** to add security group rules after the security group is created.

# Adding Security Group Rule

Last updated: 2024-01-12 14:44:52

## Scenario

Security groups are used to determine whether to permit access requests from public or private networks. For security reasons, access denial is adopted for the inbound direction in most cases. If you select the **Open all ports** or **Open ports 22, 80, 443, 3389, and protocol ICMP** template when creating a security group, the system will automatically add security group rules for some communication ports based on the selected template. For more information, see Overview.
This document describes how to add security group rules to allow or forbid CVMs in a security group to access public or private networks.

## Supports and Limits

- Security group rules are divided into IPv4 and IPv6 security group rules.
- **Open all common ports** applies to both IPv4 and IPv6 security group rules.

## Preparations

- You have already created a security group. For operation instructions, see Creating a Security Group.
- You know what public or private network access requests should be permitted or rejected for your CVM instance. For more use cases of security group rule settings, see Use Cases of Security Groups.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region** and locate the security group for which you want to set rules.
3. In the row of the security group, click **Modify rule** in the **Operation** column.
4. On the Security Group Rules page, click **Inbound Rules** and complete the operation in any of the following ways based on your actual needs.

   > **Note**
   > The following instructions use **Add a Rule** as an example.

   - **Open all ports**: this method is ideal if you do not need custom ICMP rules and all traffic goes through ports 20, 21, 22, 80, 443, and 3389 and the ICMP protocol.
   - **Add a Rule**: this method is ideal if you need to use multiple protocols and ports other than those mentioned above.
5. In the **Add Inbound Rule** window that appears, configure the rules.

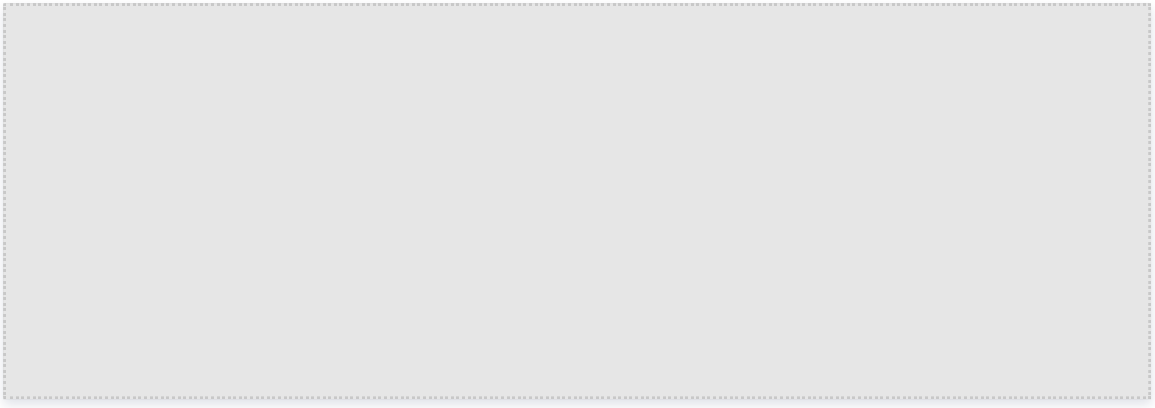

The primary parameters for adding a rule are as follows:

- Type: by default, **Custom** is selected. You can select other types such as **Login Windows CVMs (3389)**, **Login Linux CVMs (22)**, **Ping**, **HTTP (80)**, **HTTPS (443)**, **MySQL (3306)**, and **SQL Server (1433)**.
- **Source** or **Destination**: traffic origin (inbound rules) or target (outbound rules). You can use one of the following to define Source or Destination:

| Source | Notes |
| --- | --- |
| IP Address or CIDR Block | Use CIDR notation to specify the IP address (IPv4, such as 203.0.113.0, 203.0.113.0/24, or 0.0.0.0/0, where 0.0.0.0/0 represents all IPv4 addresses. IPv6, such as FF05::B5, FF05:B5::/60, ::/0, or 0::0/0, where ::/0 or 0::0/0 represents all IPv6 addresses). |
| Parameter Template – IP Address | Refer to the IP address object in the Parameter Template. |
| Parameter Template – IP Address Group | Refer to the IP address group object in the Parameter Template. |
| Security Group | Referencing a security group ID. You can reference the ID of the following security groups:<br>• Current Security Group: The current security group denotes the security group ID associated with the cloud server.<br>• Other Security Group: The other security group denotes another security group ID within the same project in the same region.<br>**Note:**<br>• Referencing a security group ID is an advanced feature. The rules of the referenced security group are not added to the current security group.<br>• If you enter a security group ID in **Source** or **Target** when configuring a security group rule, only the private IP addresses of the CVM instances and the ENIs that are bound to this security group ID are used as the source and destination, excluding public IP addresses. |

**Note**
The "/number" following an IP address represents the subnet mask, where **number** indicates the length of the network portion in the subnet mask. For example, 192.168.0.0/24 represents an IP range, and the subnet mask "/24" indicates that the first 24 bits of 192.168.0.0 are network bits, while the last 8 bits are host bits. Thus, within the 192.168.0.0/24 subnet, the assignable host IP range is from 192.168.0.0 to 192.168.0.255.

- **Protocol:port**: Enter the protocol type and port range. Supported protocol types include TCP, UDP, ICMP, ICMPv6, and GRE. You can reference the protocol ports or protocol port groups in a parameter template.
- **Policy**: **Allow** or **Reject**. **Allow** is selected by default.
  - **Allow**: Allow access requests to the port.
  - **Reject**: Discard data packets directly without returning any response.
- **Remark**: Briefly describe the rule to facilitate future management.

6. Click **OK** to finish adding the inbound rule.
7. On the security group rule page, click **Outbound rules**, and add an outbound rule by referring to Step 4 to Step 7.

# Associating Instance with Security Group

Last updated: 2024-01-12 14:44:58

> **Note**
> Security groups can be associated with CVM, ENI, TencentDB for MySQL, and CLB instance. This document describes how to associate a security group with a CVM instance.

## Scenario

You can configure a security group to control network access of one or more CVM instances. A CVM can be associated with one or more security groups. The following describes how to associate a CVM instance with a security group in the console.

## Preparations

Create a CVM instance.

## Instructions

1. Log in to the Security Group console, and go to the security group management page.
2. On the security group management page, select the related region and locate the target security group.
3. Click **Manage instances** in the **Operation** column to go to the **Associated to** page.
4. On the **Associated to** page, click **Add instance**.
5. In the pop-up dialog box, select the instance to be bound to the security group and click **OK**.

## See Also

- You can view all security groups created in a specific region.
  See Viewing Security Groups.
- If you do not want your CVM instance to belong to a security group, remove it from the security group.
  See Removing Instances from a Security Group.
- Delete security groups when you don't need them anymore. Note that all security group rules in the group are deleted as well.
  See Deleting a Security Group.

# Managing Security Group
# Viewing Security Group

Last updated: 2024-01-12 14:45:06

## Scenario

This document describes how to check security groups under a region.

## Instructions

### Viewing all security groups

1. Log in to the security group console, and go to the security group management page.
2. On the security group management page, select a **region** to view all security groups under that region.

### Searching for specific security groups

You can search for security groups by specifying the group ID/name/tag or using a keyword.

1. Log in to the security group console, and go to the security group management page.
2. On the security group management page, select a **region**.
3. Click the search box, and select a type from the drop-down list.

- **Security group ID**: Input the group ID and click .
- **Security group name**: Input the group name and click .
- **Tag**: Input a tag and click .
- **Keyword**: Input a keyword and click .

## Other operations

To learn more about the syntax for searching for specific security groups, click .

# Removing from Security Group

Last updated: 2024-01-12 14:47:27

## Scenario

You can remove a CVM instance from a security group if necessary.

## Preparations

The CVM instance to be removed is bound to two or more security groups.

## Instructions

1. Log in to the security group console, and go to the security group management page.
2. On the security group management page, select a **region** and locate the security group from which you want to remove the CVM instance.
3. In the row of the security group, click **Manage instances** in the **Operation** column to go to the **Associated to** page.
4. On the **Associated to** page, select the instance you want to remove and click **Remove from the security group**.
5. In the pop-up dialog box, click **OK**.

# Cloning Security Groups

Last updated: 2024-01-12 14:47:33

## Scenario

You may need to clone a security group in the following scenarios:

- You have created a security group named sg-A in region A and want to apply the rules of sg-A to instances in region B. In this case, you can clone sg-A to region B instead of creating a security group in region B.
- Your business needs to execute a new security group rule. In this case, you can clone the original security group as a backup.

## Supports and Limits

- By default, when you clone a security group, only the rules are cloned, not the association with instances.
- You can clone security groups between projects and regions.

## Instructions

1. Log in to the security group console, and go to the security group management page.
2. On the security group management page, select a **region** and locate the security group you want to clone.
3. In the row of the security group you want to clone, click on 'More' > 'Clone' in the operation column.
4. In the pop-up dialog box, select the **target project** and **target region** for the cloning, enter a **new name** for the security group, and click **OK**.

# Deleting Security Group

Last updated: 2024-01-12 14:47:39

## Scenario

If your business no longer needs one or multiple security groups, you can delete them. After you delete a security group, all security group rules in the group will also be deleted.

## Preparations

The security group to be deleted is not associated with any instances. If it is associated with an instance, remove it from the security group first. Otherwise, the security group cannot be deleted. For operation instructions, see Removing Instances from a Security Group.

## Instructions

1. Log in to the security group console, and go to the security group management page.
2. On the security group management page, select a **region** and locate the security group you want to delete.
3. In the row of the security group to be deleted, click **More** > **Delete** in the action column.
4. In the pop-up dialog box, click **OK**.

# Adjusting Security Group Priority

Last updated: 2024-01-12 14:48:23

## Scenario

You can bind a CVM instance to one or more security groups. When a CVM instance is bound to multiple security groups, these security groups are executed based on their priorities (such as 1 and 2). You can adjust the priorities.

## Preparations

The CVM instance is bound to two or more security groups.

## Instructions

1. Log in to the CVM console.
2. On the instance management page, click the ID of the CVM instance to go to the details page.
3. Click the **Security groups** tab to go to the security group management page.
4. In the "Bound Security Groups" section, click on **Sort**.



5. Click the following icon and drag it up or down to adjust the priority of the security group. The higher the position, the higher the priority of the security group.

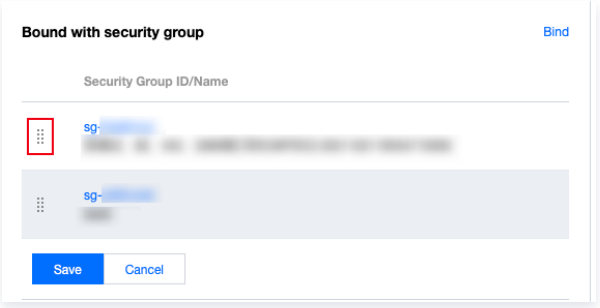


6. After completing the adjustment, click **Save**.

# Managing Security Group Rule
# Viewing Security Group Rule

Last updated: 2024-01-12 14:48:28

## Scenario

After adding a security group rule, you can view its details in the console.

## Preparations

You have created a security group and added a security group rule to the group.
For information on how to create a security group and add a security group rule, see Creating a Security Group and Adding a Security Group Rule.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region** and locate the security group whose rule you want to view.
3. Click the ID/name of the security group to go to the security group rule page.
4. Click the **Inbound rules** or **Outbound rules** tab to view the inbound or outbound rule of the security group.

# Modifying Security Group Rule

Last updated: 2024-01-12 14:48:34

## Scenario

This document describes how to modify a security group rule. Improper configuration of security group rules can incur severe security risks. For example, you can restrict the access to specific ports to prevent risks.

## Preparations

Create a security group and add a security group rule to the group.
For more information, see Creating a Security Group and Adding a Security Group Rule.

## Instructions

1. Log in to the Security Group console and go to the security group management page.
2. Select a **region** and locate the security group to modify.
3. In the row of the security group, click **Modify rule** in the **Operation** column to go to the security group rule page.
4. Click the **Inbound rules** or **Outbound rules** tab of the security group as needed.
5. Locate the security group rule, and click **Edit** in the **Operation** column to modify the rule.

> **Note**
> You don't need to restart the CVM for the rule changes to take effect.

# Deleting Security Group Rule

Last updated: 2024-01-12 14:48:39

## Scenario

If you no longer need a security group rule, you can delete it.

## Preparations

- You have created a security group and added a security group rule to the group.
  For information on how to create a security group and add a security group rule, see Creating a Security Group and Adding a Security Group Rule.
- You have confirmed that your CVM instance does not need to permit or forbid public network access or private network access.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region** and locate the security group whose rule you want to delete.
3. In the row of the security group, click **Modify rule** in the **Operation** column to go to the security group rule page.
4. On the Security Group Rules page, click the **Inbound/Outbound Rules** tab according to the direction (inbound/outbound) of the security group rule you want to delete.
5. Locate the security group rule, and click **Delete** in the **Operation** column.
6. In the pop-up dialog box, click **Confirm**.

# Importing Security Group Rule

Last updated: 2024-01-12 14:48:43

## Scenario

You can quickly create or recover security group rules by importing a rule file.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region** and locate the security group to which you want to import rules.
3. Click the ID/name of the security group to go to the security group rule page.
4. Click the **Inbound rules** or **Outbound rules** tab based on the direction (inbound or outbound) of the rules to be imported.
5. Click **Import rule**.
6. In the pop-up window, select the edited template file for the inbound or outbound rules and click **Import**.

> **Note**
> - If the security group already has rules, click **Append** to add rules in the file before the existing rules.
> - If the security group has no rules, download the template, edit it, and import it.

# Exporting Security Group Rule

Last updated: 2024-01-12 14:48:48

## Scenario

You can export security group rules and save them locally for backup.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region** and locate the security group whose rules you want to export.
3. Click the **name or ID** of the desired security group. The details page of the selected security group appears.
4. Click the **Inbound rules** or **Outbound rules** tab based on the direction (inbound or outbound) of the security group rules to be exported.
5. Click  in the upper right corner to export the rules to a file and save it to your local device.

# Sorting Security Group Rules

Last updated: 2024-01-12 14:48:52

## Scenario

Multiple security group rules can be added to a security group. They take effect in order from top to bottom and can be sorted as needed.

## Preparations

You have created a security group with at least two rules as instructed in Adding a Security Group Rule.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. Select a **region**.
3. Locate the security group with rules to be modified, click the "Security Group ID" or click **Modify Rules** in the operation column to access the security group rules page.
4. On the Security Group Rules page, click **Sort**.
5. Drag the following icon to sort the security group rules. The higher the position, the higher the priority of the security group rule. Click **Save** after adjusting.

# Snapshot Rollback

Last updated: 2024-01-12 14:50:44

If a security group is configured with a snapshot policy, its rules will be backed up according to the configured policy. To roll back its rules, perform a snapshot rollback.

## Preparations

The security group is configured with a snapshot policy, and at least one snapshot backup has been made.

## Instructions

1. Log in to the security group console and go to the security group management page.
2. On the **Security Group** management page, click the ID of the target security group to enter its details page.
3. Click the **Snapshot rollback** tab, which displays all the snapshot records by time. By default, records from the past seven days are displayed. You can specify a query period.

> **Note**
> If there are many backup records, we recommend you limit the time range to three months; otherwise, the system may become slow.

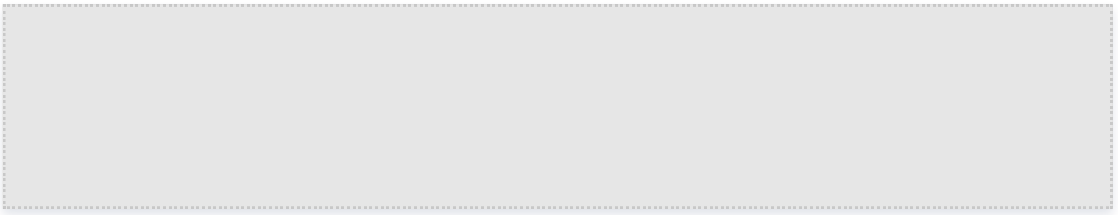

4. Click **Export** to export the inbound and outbound rules saved in this snapshot. Rules are separately in the **Inbound rules** and **Outbound rules** files.
5. Click **Restore** to enter the security group restoration page, which displays the security group rule preview and comparison.

> **Note**
> - **+** indicates newly added entries, while **–** means deleted entries. The comparison result is based on the time period currently selected for comparison preview, during which if rules are changed, the comparison result may be inaccurate.
> - When a security group is restored, if its source contains parameter template rules and nested security groups, the rules within the parameter template and the instances associated with the security group will remain in the current status.
> - Note that when you restore a security group with a snapshot, all current rules are overwritten.

6. Click **OK** to complete the rollback.

# Application Cases of Security Groups

Last updated: 2024-01-12 14:50:49

Security groups are used to manage whether a CVM is accessible. You can configure inbound and outbound rules for security groups to specify whether your CVMs can be accessed by or can access other network resources.
The default inbound and outbound rules for security groups are as follows:

- **To ensure data security, the inbound rule for a security group is a rejection policy that denies remote access from external networks.** To make your CVM accessible to resources in external networks, you need to configure an inbound rule that opens the corresponding ports.
- The outbound rule for a security group specifies whether your CVM can access resources in external networks. If you select **Open all ports** or **Open ports 22, 80, 443, 3389, and protocol ICMP**, the outbound rule opens the ports to external networks. If you select **Custom**, the outbound rule blocks all ports by default, and you need to set the outbound rule to open corresponding ports for your CVM to access resources in external networks.

## Common use cases

This document describes several common use cases of security groups. If any of the following use cases meet your needs, you can set your security groups according to the configurations recommended in that use case.

### Case 1: Remotely connecting to a Linux CVM via SSH

**Case**: You have created a Linux CVM and want to remotely connect to the CVM via SSH.
**Solution**: When adding an inbound rule, select **Login Linux CVMs(22)** for **Type** to open protocol port 22 to allow Linux login via SSH.
You can allow all IP addresses or a specified IP address (or IP address range) as needed by configuring the source IP addresses that can remotely connect to the CVM via SSH.

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Login Linux CVMs(22) | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | TCP:22 | Allow |

### Case 2: Remotely connecting to a Windows CVM via RDP

**Case**: You have created a Windows CVM and want to remotely connect to the CVM via RDP.
**Solution**: When adding an inbound rule, select **Login Windows CVMs(3389)** for **Type** to open protocol port 3389 to allow remote Windows login.
You can allow all IP addresses or a specified IP address (or IP address range) as needed by configuring the source IP addresses that can remotely connect to the CVM via RDP.

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Login Windows CVMs(3389) | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | TCP:3389 | Allow |

### Case 3: Pinging a CVM from the public network

**Case**: You have created a CVM and want to test whether the CVM can normally communicate with other CVMs.
**Solution**: Use the ping program to test. When adding an inbound rule, select **Ping** for **Type** to open the ICMP protocol port to allow other CVMs to access this CVM via ICMP.
You can allow all IP addresses or a specified IP address (or IP address range) as needed by configuring the source IP addresses that can access the CVM via ICMP.

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Ping | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | ICMP | Allow |

### Case 4: Remotely logging in to a CVM via Telnet

**Case:** You want to remotely log in to a CVM via Telnet.
**Solution:** When adding an inbound rule, configure it as follows:

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Custom | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | TCP:23 | Allow |

### Case 5: Authorizing access to a web service via HTTP or HTTPS

**Case:** You have built a website and want to allow users to access your website via HTTP or HTTPS.
**Solution:** When adding an inbound rule, configure it as follows:

- Allow all IP addresses in the public network to access your website

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | HTTP ( 80 ) | 0.0.0.0/0 | TCP:80 | Allow |
| Inbound | HTTPS ( 443 ) | 0.0.0.0/0 | TCP:443 | Allow |

- Allow some IP addresses in the public network to access your website

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | HTTP ( 80 ) | The IP address or IP address range that is allowed to access your website | TCP:80 | Allow |
| Inbound | HTTPS ( 443 ) | The IP address or IP address range that is allowed to access your website | TCP:443 | Allow |

### Case 6: Allowing external IP addresses to access a specified port

**Case:** You have deployed a service and want the specified service port (such as port 1101) to be accessible from external networks.
**Solution:** When adding an inbound rule, select **Custom** for **Type** and open protocol port 1101 to allow access from external networks to the specified service port.
You can allow all IP addresses or a specified IP address (or IP address range) as needed by configuring the source IP addresses that can access the service port.

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Custom | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | TCP:1101 | Allow |

### Case 7: Rejecting access from external IP addresses to a specified port

**Use case:** After deploying your service, you want to prevent external access to a specific service port (e.g., 1102).
**Solution:** When adding an inbound rule, select **Custom** in the **Type** field, configure the 1102 protocol port, and set the **Policy** to **Deny** to reject external access to the specified service port.

| Direction | Type | Source | Protocol port | Policy |
|---|---|---|---|---|
| Inbound | Custom | All IP addresses: 0.0.0.0/0<br>Specified IP address: A specified IP address or IP address range | TCP:1102 | Reject |

### Case 8: Allowing a CVM to access only a specified external IP address

**Case:** You want to allow your CVM to access only a specified external IP address.
**Solution:** Add the two outbound security group rules according to the following configurations.

- Allow the CVM instance to access a specified public IP address.
- Disallow the CVM instance to access any public IP address through any protocol.

> **Note**
> Rules that permit access take priority over those that forbid access.

| Direction | Type | Source | Protocol port | Policy |
| --- | --- | --- | --- | --- |
| Outbound | Custom | The public IP address that the CVM can access | The required protocol and port | Allow |
| Outbound | Custom | 0.0.0.0/0 | ALL | Reject |

### Case 9: Rejecting access from a CVM to a specified external IP address

**Case:** You do not want your CVM to access a specified external IP address.
**Solution:** Add a security group rule according to the following configurations.

| Direction | Type | Source | Protocol port | Policy |
| --- | --- | --- | --- | --- |
| Outbound | Custom | The public IP address that you do not allow your CVM to access | ALL | Reject |

### Case 10: Uploading a file to or downloading a file from a CVM through FTP

**Case:** You want to upload a file to or download a file from a CVM by using an FTP program.
**Solution:** Add a security group rule according to the following configurations.

| Direction | Type | Source | Protocol port | Policy |
| --- | --- | --- | --- | --- |
| Inbound | Custom | 0.0.0.0/0 | TCP:20-21 | Allow |

## Combination of multiple use cases

In an actual use case, you may want to configure multiple security group rules based on service requirements, for example, configuring inbound and outbound rules at the same time. One CVM can be bound to one or multiple security groups. When a CVM is bound to multiple security groups, these security groups are executed from top to bottom. You can adjust their priorities at any time. For more information on rule priorities, see the Rule priorities section in Overview.

# Common Server Ports

Last updated: 2024-01-12 14:50:54

The following are common server ports. For more information on service application ports in Windows, please refer to the official Microsoft documentation ( Overview of Services and Network Port Requirements for Windows ).

| Port | Service | Note |
| --- | --- | --- |
| 21 | FTP | An open FTP server port for uploading and downloading. |
| 22 | SSH | Port 22 is the SSH port. It is used to remotely connect to Linux servers in CLI mode. |
| 25 | SMTP | SMTP server's open port for sending emails. |
| 80 | HTTP | This port is used for web services such as IIS, Apache, and Nginx to provide external access. |
| 110 | POP3 | Port 110 is open for Post Office Protocol 3 (POP3) services. |
| 137、138、139 | NETBIOS Protocol | Ports 137 and 138 are UDP ports for transferring files through My Network Places.<br>Port 139: Connections over port 139 attempt to access the NetBIOS/SMB service. This protocol is used for file and printer sharing on Windows and Samba. |
| 143 | IMAP | Port 143 is mainly used for Internet Message Access Protocol (IMAP) v2, which is a protocol for receiving emails and similar to POP3. |
| 443 | HTTPS | A web browsing port. HTTPS is another type of HTTP that provides encryption and transmission through secure ports. |
| 1433 | SQL Server | Port 1433 is the default port for SQL Server. SQL Server uses two ports: port 1433 for TCP and port 1434 for UDP. Port 1433 is used for SQL Server to provide external services, whereas port 1434 is used to respond to the requester regarding which TCP/IP port is being used by SQL Server. |
| 3306 | MySQL | Port 3306 is the default port for MySQL databases and is used to provide external services. |
| 3389 | Windows Server Remote Desktop Services | Port 3389 is the port for remote desktop services on Windows Server. Through this port, you can connect to a remote server by using the Remote Desktop connection tool. |
| 8080 | Proxy Port | Similar to port 80, port 8080 is used for the WWW proxy service for web browsing. The port number extension ":8080" is often appended to the URL when users visit a website or use a proxy server. In addition, after the Apache Tomcat web server is installed, the default service port is port 8080. |

# Network ACL Rule Overview

Last updated: 2024-01-12 14:51:04

A network Access Control List (ACL) is an optional layer of security that throttles traffic to and from subnets accurate to protocol and port.
Watch video

## Application example

You can associate a network ACL with multiple subnets that require the same level of network traffic control and set inbound and outbound rules to precisely control their inbound and outbound traffic.
For example, when you host a multi-layer web application in a Tencent Cloud VPC instance and create different subnets for the web-layer, logic-layer, and data-layer services, you can use a network ACL to ensure that the web-layer and data-layer subnets cannot access each other, and only the logic-layer subnet can access the web-layer and data-layer subnets.

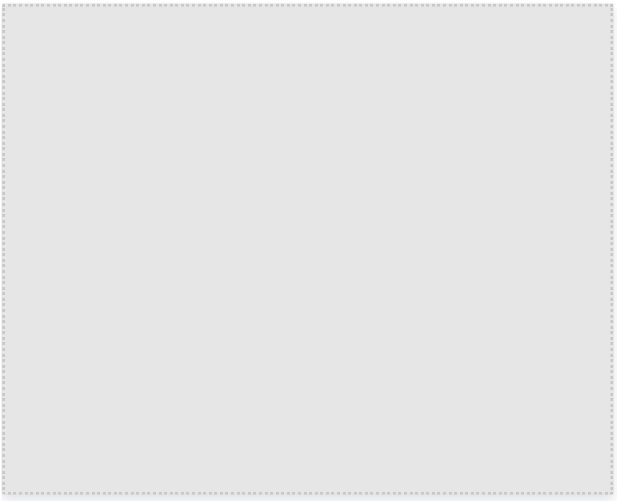

## ACL rules

After you add or delete a rule in a network ACL, the changes you make will be automatically applied to the associated subnets.
You can configure inbound and outbound network ACL rules. Each rule consists of the following elements:

- **Source IP/Destination IP**: The origin or target IP address of the traffic. For inbound rules, enter the source IP; for outbound rules, enter the destination IP. Both source and destination IPs support the following formats:
  - Single IP: Such as "192.168.0.1" or "FF05::B5".
  - CIDR block: Such as "192.168.1.0/24" or "FF05:B5::/60".
  - All IPv4 addresses: "0.0.0.0/0".
  - All IPv6 addresses: "0::0/0" or "::/0".
- **Protocol**: Select a protocol that the ACL rule allows or denies, such as TCP and UDP.
- **Port**: The source or destination port of the traffic. Supported formats include:
  - Single port: Such as "22" or "80".
  - Port range: Such as "1-65535" or "100-20000".
  - All ports: All.
- **Policy**: Select **Allow** or **Refuse**.

## Default Rules

Once created, every network ACL has two default rules that cannot be modified or deleted, with the lowest priority.

- Default inbound rule

| Protocol | Port | Source IP | Policy | Note |
| --- | --- | --- | --- | --- |
| ALL | ALL | 0.0.0.0/0 | Refuse | Denies all inbound traffic. |

- Default outbound rule

| Protocol | Port | Destination IP | Policy | Note |
| --- | --- | --- | --- | --- |
| ALL | ALL | 0.0.0.0/0 | Refuse | Denies all outbound traffic. |

## Rule priorities

- The rules of a network ACL are prioritized from top to bottom. The rule at the top of the list has the highest priority and will take effect first, while the rule at the bottom has the lowest priority and will take effect last.
- If there is a rule conflict, the rule with the higher priority will prevail by default.
- When traffic goes in or out of a subnet that is bound to a network ACL, the network ACL rules will be matched sequentially from top to bottom. If a rule is matched successfully and takes effect, the subsequent rules will not be matched.

### Application example

To allow all source IP addresses to access all ports of CVMs in a subnet associated with a network ACL and deny the source IP address `192.168.200.11/24` of HTTP services to access port 80, add the following two inbound rules to the network ACL:

| Protocol | Port | Source IP | Policy | Note |
| --- | --- | --- | --- | --- |
| HTTP | 80 | 192.168.200.11/24 | Refuse | Denies this IP address of HTTP services to access port 80. |
| ALL | ALL | 0.0.0.0/0 | Allow | Allows all source IP addresses to access all ports. |

## Security group vs. network ACL

| Item | Security Group | Network ACL |
| --- | --- | --- |
| Traffic throttling | Traffic throttling at the instance level, such as CVM and TencentDB | Traffic throttling at the subnet level |
| Policy | Allow or refuse. | Allow or refuse. |
| Stateful/Stateless | Stateful: returned traffic is automatically permitted without being subject to any rules. | Stateless: Returned traffic must be explicitly allowed by rules. |
| Effective time | Security group rules are applied to an instance, such as a CVM or TencentDB instance, only if you specify a security group when creating the instance or associate a security group with the instance after it is created. | ACL rules are automatically applied to all instances, such as CVM and TencentDB instances, in the associated subnet. |
| Rule priorities | If there is a rule conflict, the rule with the higher priority will prevail by default. | If there is a rule conflict, the rule with the higher priority will prevail by default. |

# Limits

Last updated: 2024-01-12 14:51:12

## Usage Limits

- One network ACL can be bound to multiple subnets.
- Network ACLs are stateless. Therefore, you need to set outbound rules and inbound rules separately.
- Network ACLs do not affect the communication among CVM instances in the bound subnets over the private network.

## Quota Limits

| Resources | Restrictions |
| --- | --- |
| Number of network ACLs in each VPC | 50 |
| Number of rules per network ACL | Inbound: 20<br>Outbound: 20 |
| Number of network ACLs associated with each subnet | 1 |

# Managing Network ACLs

Last updated: 2024-01-12 14:51:17

This document describes how to configure network ACLs for access control.
Watch video

## Creating Network ACLs

1. Log in to the VPC console.
2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. Select a region and a VPC above the list and click **Create**.
4. In the pop-up window, enter a name and select a VPC for the ACL, then click **OK**.

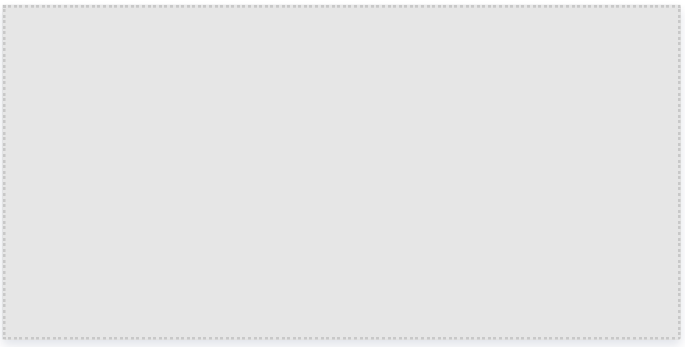

5. In the network ACL list, click the ID of the ACL to go to its details page, where you can add ACL rules and associate ACL rules with subnets.

## Adding Network ACL Rules

1. Log in to the VPC console.
2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. In the network ACL list, locate the ACL you want to modify and click its ID to go to the details page.
4. Click **Outbound rules** or **Inbound rules** > **Edit** > **+ New line**, select a protocol, enter a port and source IP, and select a policy.
   - **Protocol**: Select the protocol that the ACL rule allows or denies, such as TCP and UDP.
   - **Port**: Enter the source port of traffic, which can be a single port or a port range, such as port 80 or ports 90 to 100.
   - **Source IP**: Enter the source IP address or IP range of traffic, which can be an IP or a CIDR block, such as `10.20.3.0` or `10.0.0.2/24`.
   - **Policy**: Select **Allow** or **Refuse**.

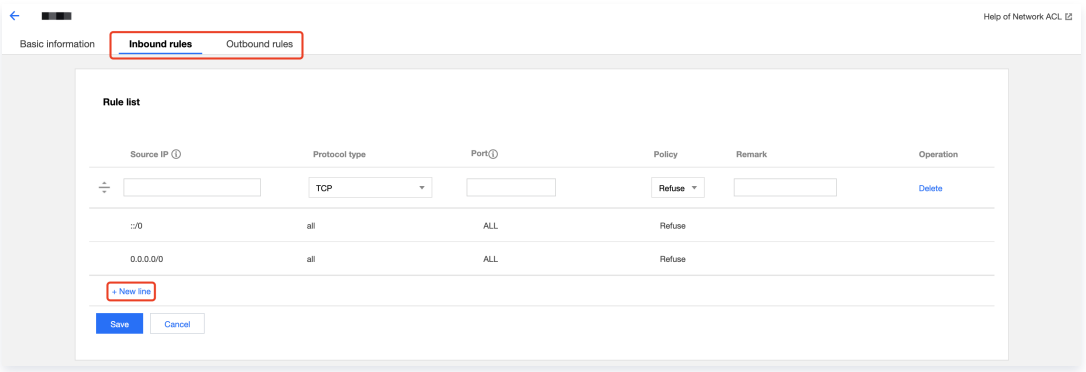


5. Click **Save**.

## Deleting Network ACL Rules

1. Log in to the VPC console.
2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. In the network ACL list, locate the ACL whose rule you want to delete and click its ID to go to the **Basic information** page.

4. Click the **Inbound rules** or **Outbound rules** tab to go to the rule list page.
5. Click **Edit**. The process for deleting inbound rules is the same as for deleting outbound rules. The deletion of inbound rules is used as an example here.

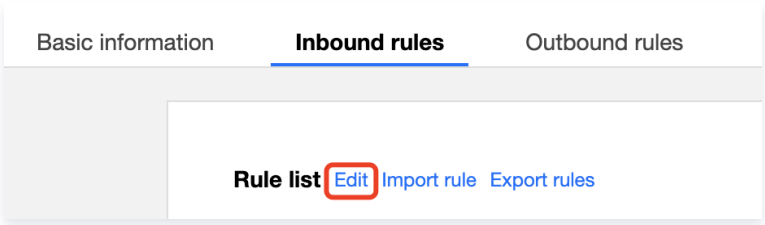


6. In the rule list, locate the row of the rule you want to delete and click **Delete** in the operation column.

> **Note**
> The ACL rule is now grayed out. If you deleted the rule by accident, you can click **Recover the deleted** in the operation column to recover it.

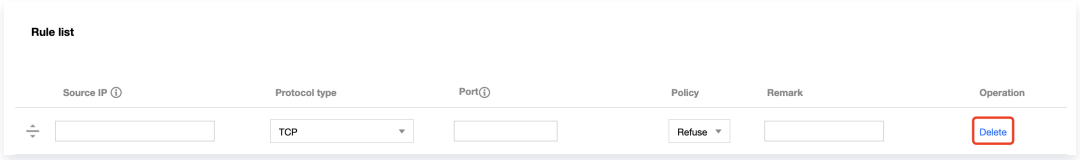


7. Click **Save** to save your operations.

> **Note**
> The deletion or recovery of the ACL rule only takes effect after you save your operations.

## Associating network ACLs with subnets

1. Log in to the VPC console.
2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. In the network ACL list, locate the ACL you want to associate and click its ID to go to the details page.
4. On the **Basic information** page, click **+ Associate** in the **Associated subnets** section.

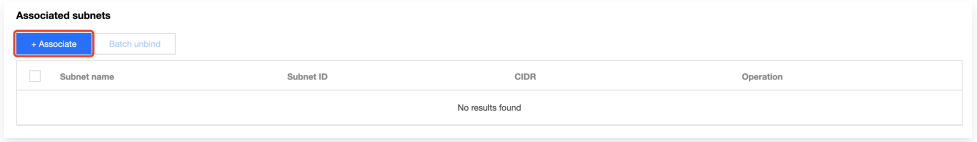

5. In the pop-up window, select subnets to associate and click **OK**.

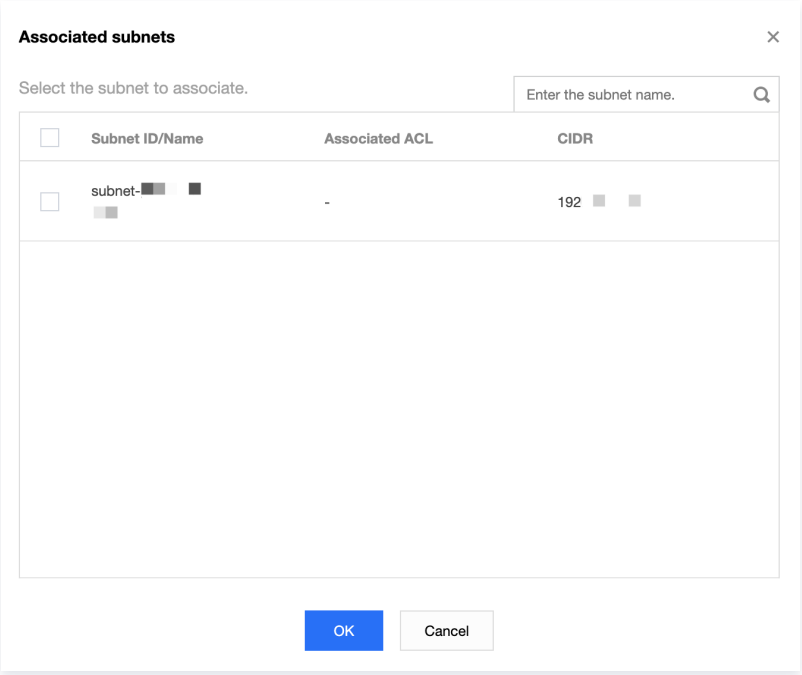


## Disassociating network ACLs from subnets

1. Log in to the VPC console.
2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. In the network ACL list, locate the ACL you want to disassociate and click its ID to go to the details page.
4. There are two methods for disassociating ACLs from subnets:

- Method 1: In the **Associated subnets** section on the **Basic information** page, locate the subnet you want to disassociate and click **Unbind**.


- Method 2: In the **Associated subnets** section on the **Basic information** page, select the subnets you want to disassociate and click **Batch unbind**.


5. In the pop-up window, click **OK**.


## Deleting Network ACLs

1. Log in to the VPC console.

2. Click **Security** > **Network ACL** on the left sidebar to go to the management page.
3. Select a region and a VPC.
4. In the network ACL list, locate the ACL you want to delete, click **Delete**, and confirm the operation. The ACL and all of its rules will be deleted.

> **Note**
> If the **Delete** option is grayed out, such as that for the network ACL `testEg` in the following figure, the network ACL is associated with a subnet. You will need to disassociate it from the subnet first before you can delete it.

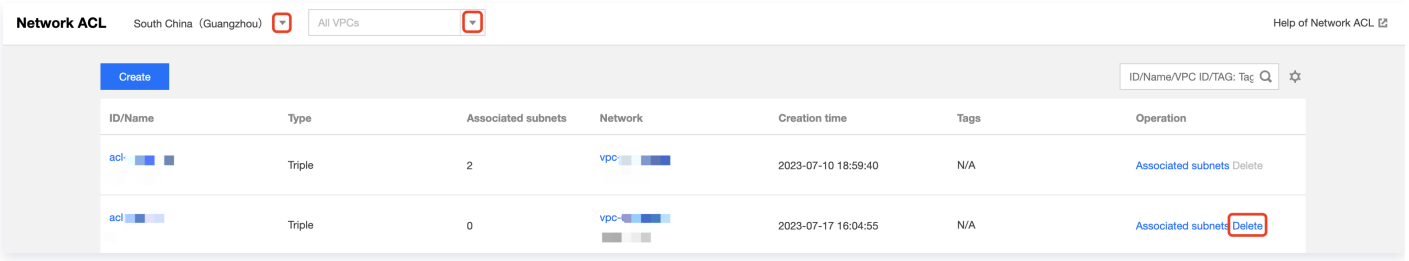

# Parameter Template Overview

Last updated: 2024-01-12 14:51:22

A parameter template is a set of IP address or protocol port parameters. You can save IP addresses or protocol ports with the same needs as a template so that you can directly reference the template as the source/destination IP addresses or protocol ports when adding security group rules. Parameter templates, if properly used, can enhance your efficiency in using security groups.

## Scenarios

Parameter templates are mainly suitable for the following scenarios:

- Manage multiple IP addresses or protocol port groups with the same needs.
- Manage multiple IP addresses or protocol port groups with frequent editing needs.

## Parameter template types

Tencent Cloud supports the following four types of parameter templates:

- IP address: Also known as an IP address object. This template is a set of IP addresses and supports one single IP, CIDR block, and IP range.
- IP address group: Also known as an IP address group object. This template is a set of multiple IP address objects.
- Protocol port: Also known as a protocol port object. This template is a set of protocol ports and supports one single port, multiple ports, port range, and all ports. It supports TCP, UDP, ICMP, and GRE protocols.
- Protocol port group: Also known as a protocol port group object. This template is a set of protocol port objects.

# Limits

Last updated: 2024-01-12 14:51:34

## Usage Limits

- Formats supported by the IP address template are as follows:
  - Single IP address: Such as `10.0.0.1`.
  - Consecutive IP addresses: Such as `10.0.0.1` - `10.0.0.100`.
  - IP range: Such as `10.0.1.0/24`.
- Formats supported by the port template are as follows:
  - Single port: Such as `TCP:80`.
  - Multiple ports: Such as `TCP:80,443`.
  - Consecutive ports: Such as `TCP:3306-20000`.
  - All ports: Such as `TCP:ALL`.

## Quota Limits

| Instance | Upper limit |
| --- | --- |
| IP address objects (ipm) | 1,000 per tenant |
| IP Address Group Objects (IPMG) | 1,000 per tenant |
| Protocol port objects (ppm) | 1,000 per tenant |
| Protocol port group objects (ppmg) | 1,000 per tenant |
| IP address members in an IP address object (ipm) | 20 per tenant |
| IP address object members (ipm) in an IP address group object (ipmg) | 20 per tenant |
| Protocol port members in a protocol port object (ppm) | 20 per tenant |
| Protocol port object members (ppm) in a protocol port group object (ppmg) | 20 per tenant |
| IP address group objects (ipmg) that can reference an IP address object (ipm) | 50 per tenant |
| Protocol port group objects (ppmg) that can reference a protocol port object (ppm) | 50 per tenant |

**Note**
If a parameter template is referenced by a security group, the IPs and ports in the template will be converted to multiple security group rules (up to 2,000).

# Managing Parameter Templates

Last updated: 2024-01-12 14:51:39

This document describes how to create and maintain parameter templates (IP address, IP address group, protocol port, and protocol port group) in the console and how to use them in security groups.

## Creates a parameter template.

### Creating IP address parameter template

Add the IPs with the same needs or frequently edited to this IP address object.

#### Instructions

1. Log in to the Virtual Private Cloud Console.
2. Click **Security** > **Parameter Templates** in the left sidebar to enter the management page.
3. On the **IP Address** tab, click **+Create**.
4. In the pop-up window, enter the name and IP address, then click **Submit**.
   To add multiple IP addresses, separate them with a new line. The supported format is as follows:
   - Single IP: Such as `10.0.0.1` or FF05::B5.
   - CIDR block: such as `10.0.1.0/24` or FF05:B5::/60.
   - Continuous address range: for example, 10.0.0.1-10.0.0.100.

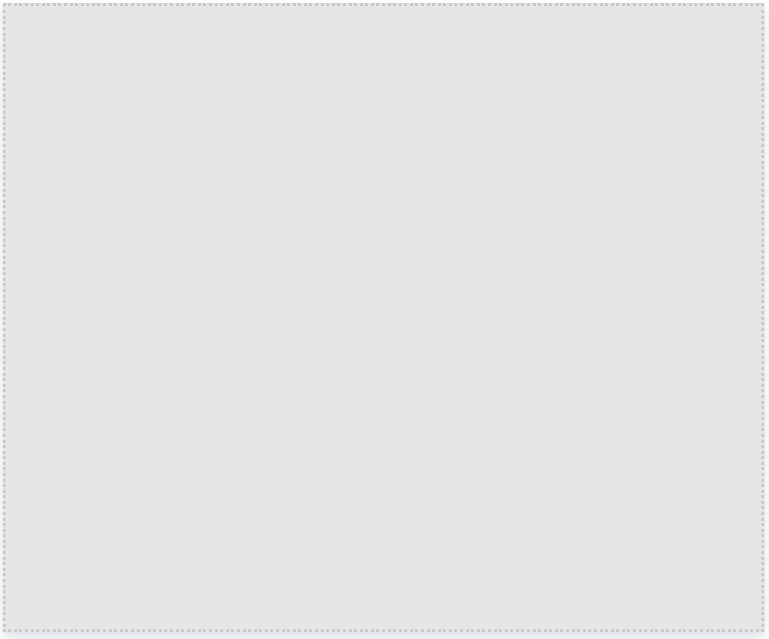

### Creating IP address group parameter template

You can add multiple IP address objects to an IP address group for unified management.

#### Instructions

1. Select the **IP Address Group** tab, enter the management page, and click **+Create**.

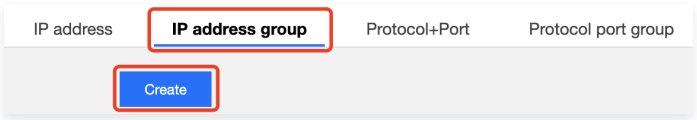


2. In the pop-up window, enter the name, select the IP address object to be added, and click **Submit**.

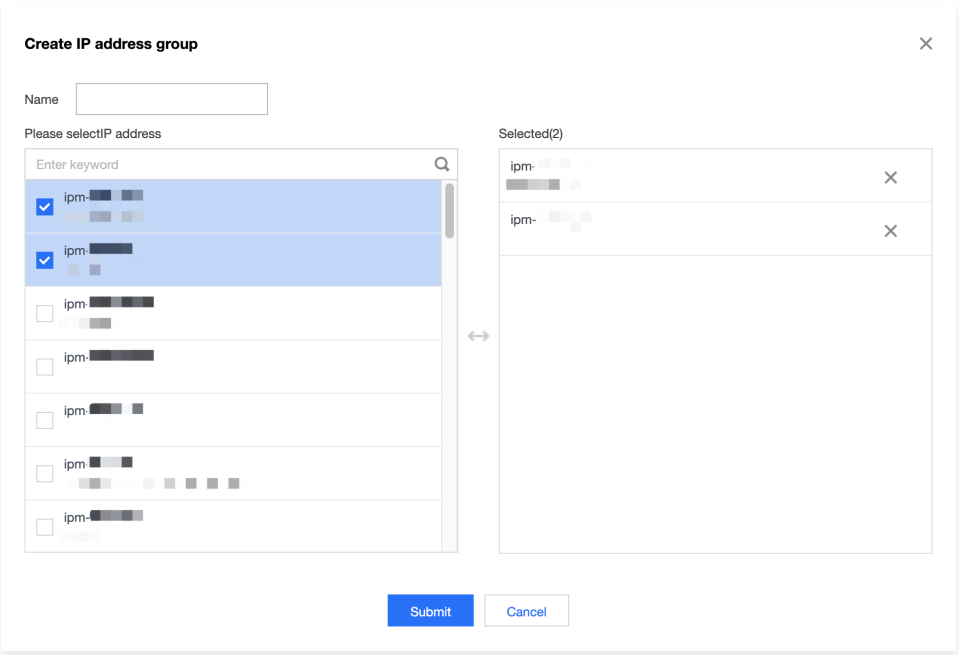


## Creating protocol port parameter template

Create a template and add combinations of protocol and port to it for easy management.

### Instructions

1. Log in to the Virtual Private Cloud Console.
2. Click **Security** > **Parameter Templates** in the left sidebar to enter the management page.
3. Click the **Protocol Port** tab, enter the **Protocol Port** page, and click **+Create**.
4. In the pop-up window, enter the name and protocol port, and click **Submit**.
   To add multiple protocol ports within the specified range, separate them with a newline. The format is as follows:
   - A single port, such as `TCP:80`.
   - Multiple discrete ports, such as `TCP:80,443`.
   - Consecutive ports, such as `TCP:3306-20000`.
   - All ports, such as `TCP:ALL`.

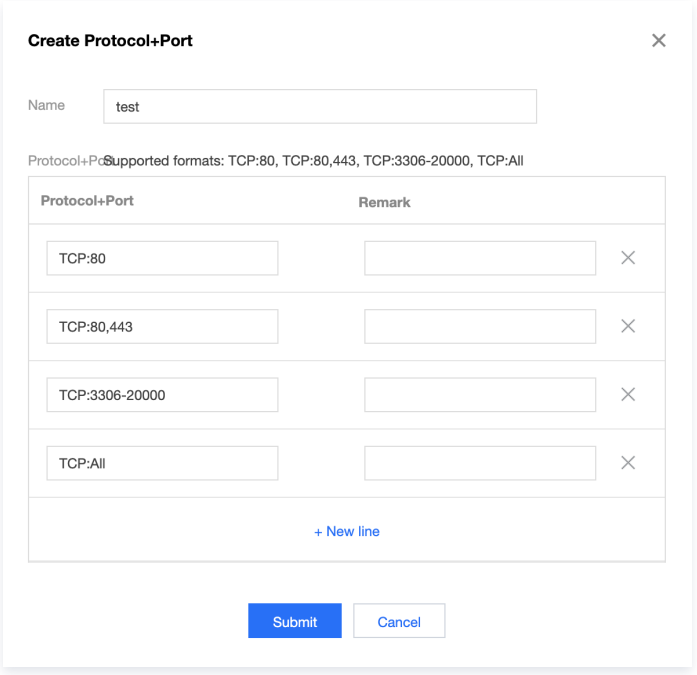

### Creating protocol port group parameter template

You can add multiple created protocol port objects to a protocol port group for unified management.

#### Instructions

1. Select the **Protocol Port Group** tab, enter the management page, and click **+Create**.

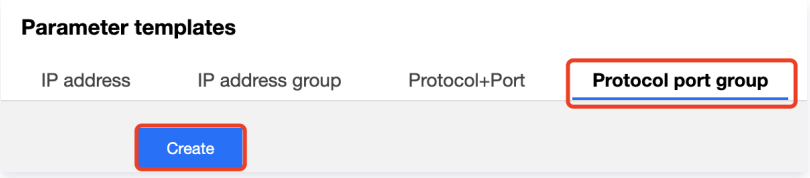


2. In the pop-up window, enter the name and select the required protocol port objects, then click **Submit**.

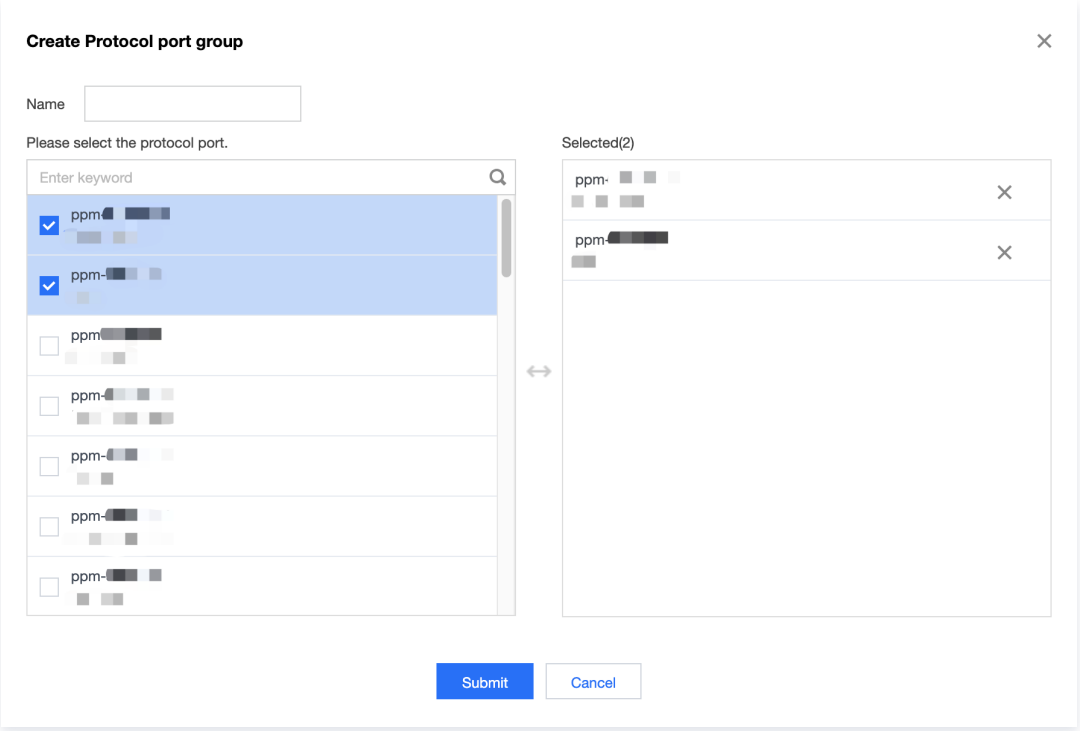


## Modifies a parameter template.

If you need to modify a created parameter template, for example, to add/delete IP addresses or protocol ports, follow the steps below.

#### Instructions

1. Click on the created IP address, IP address group, protocol port, or protocol port group parameter template, and then click **Edit** on the right side. For example, the following image shows modifying an IP address object.

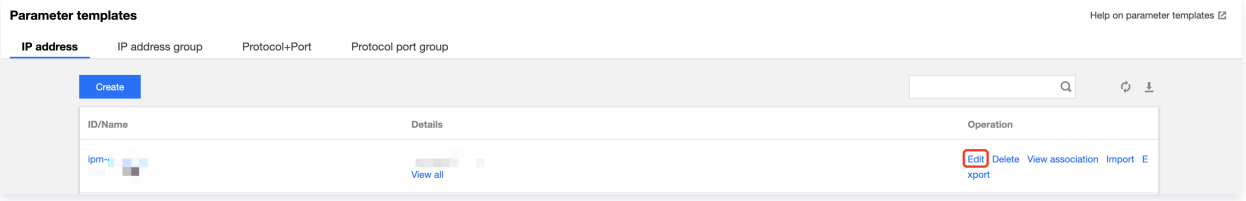


2. In the displayed edit dialog box, modify the corresponding parameters and click **Submit**.

## Deletes a parameter template.

If you no longer use a parameter template, you can delete it. When this template is deleted, all the policy configurations containing it in the security group will be deleted at the same time. Please evaluate and proceed with caution.

#### Instructions

1. Click **Delete** on the right side of the created parameter template.

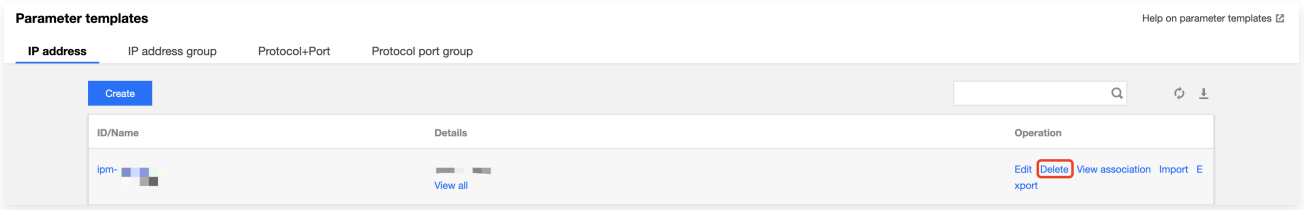


2. After deletion, all policies containing this IP address or protocol port will be deleted. Confirm and continue by clicking **Delete** in the pop-up confirmation box.

## Referencing parameter templates in security groups

When you add a security group rule, you can refer to the templates to add IPs and ports quickly.

### Instructions

1. Log in to the Virtual Private Cloud Console.
2. Click **Security** > **Security Group** in the left directory to enter the management page.
3. In the list, find the security group that needs to import the parameter template and click its ID to enter the details page.
4. In the Inbound/Outbound Rules tab, click **Add Rule**.
5. In the pop-up window, select the "Custom" type, choose the created parameter templates in "Source" and "Protocol Port", and click "Finish". For detailed steps on adding inbound/outbound rules, please refer to Adding Security Group Rules.

> **Note**
> If you need to add a new IP address or protocol port in the future, you only need to add it to the corresponding IP address group or protocol port group, and there is no need to modify the security group rules or create another security group.

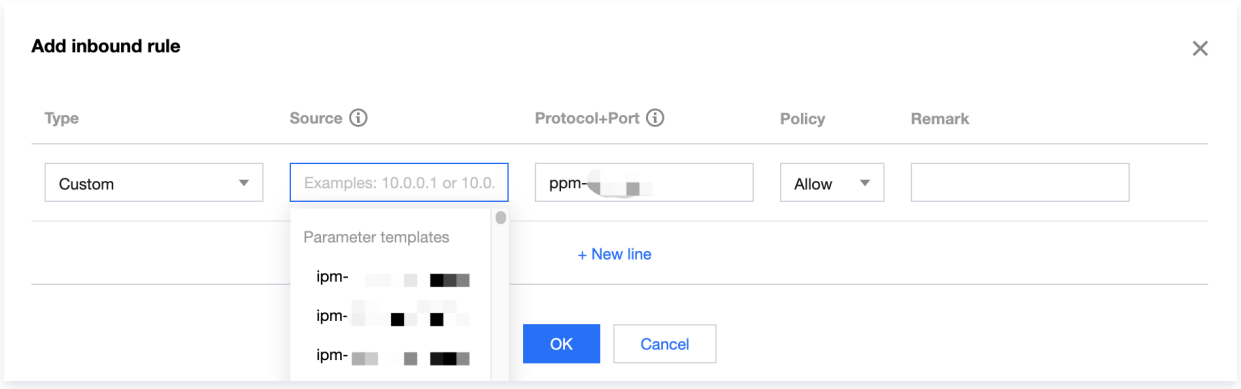


## View associated security groups

You can view all security group instances that import a parameter template in the following steps.

1. Click **View Associations** on the right side of the created parameter template.

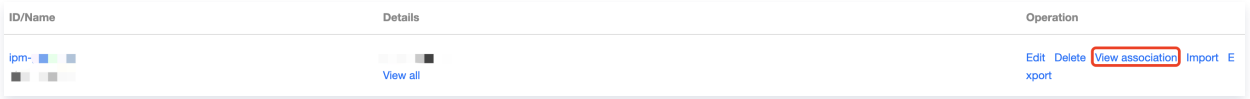

2. The associated security group list that pops up displays all security group instances associated with this parameter template.

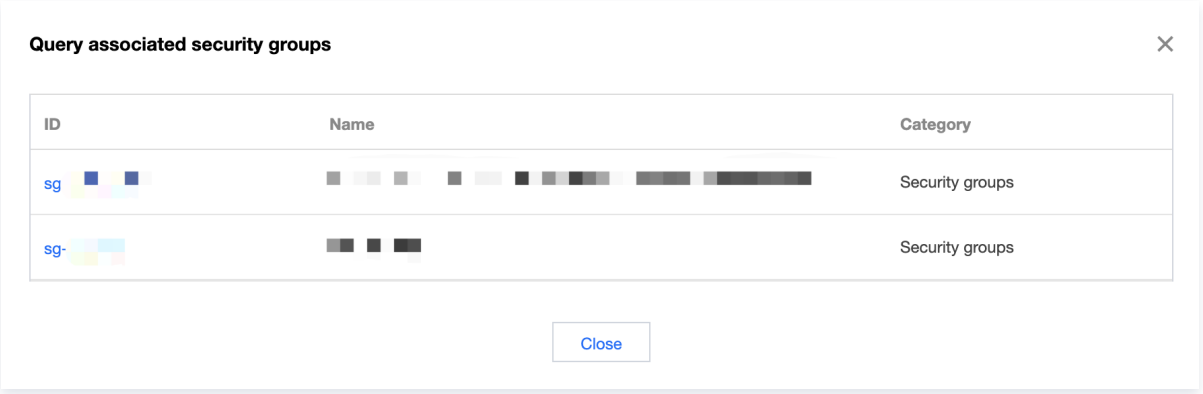


## Importing a parameter template

To batch add parameter template configurations, do the following:

1. Click **Import** on the right side of the created parameter template.
2. Upload local file.

## Exporting a parameter template

To back up the parameter template configuration locally, click **Export** on the right side of the created parameter template.

# Configuration Case

Last updated: 2024-01-12 14:51:47

## Parameter Template Use Cases

Parameter template is an efficient, fast, and easy-to-maintain way to add rules in security groups. For example, when you need to add multiple IP ranges, specified IPs, or protocol ports of multiple types, you can define a parameter template. You can also use the parameter template subsequently to maintain the IP sources and protocol ports in the security group rules.

> **Note**
> All the IP addresses and protocol ports in this document are examples. Please replace them according to your actual business conditions during configuration.

## Example description

Suppose you want to configure the following security group rules and need to update the inbound source IP range and protocol port later:

- Inbound rules:
  - Allowed source IP range: 10.0.0.16-10.0.0.30; protocol ports: TCP:80,443
  - Allowed source CIDR block: 192.168.3.0/24; protocol ports: TCP:3600-15000
- Outbound rule:
  Rejected target IP address: 192.168.10.4; protocol port: TCP:800

## Solution

Because you have the same security group policy for multiple IP ranges and protocol ports, and you need to update the source IP range later, you can use a parameter template to add and maintain security group rules.

### Step 1. Create parameter templates

1. Log in to VPC console.
2. Select **Security** > **Parameter Template** on the left sidebar to go to the management page.
3. On the **IP address** tab, click **Create** to create an IP address parameter template respectively for adding inbound and outbound rules.
4. In the pop-up window, enter the source IP range and click **OK**.

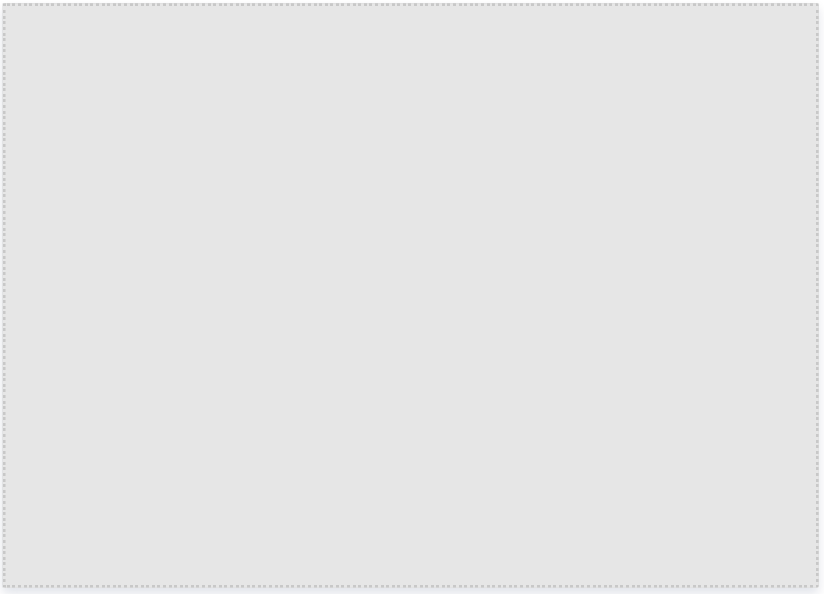

The newly created IP address parameter template is as shown below.

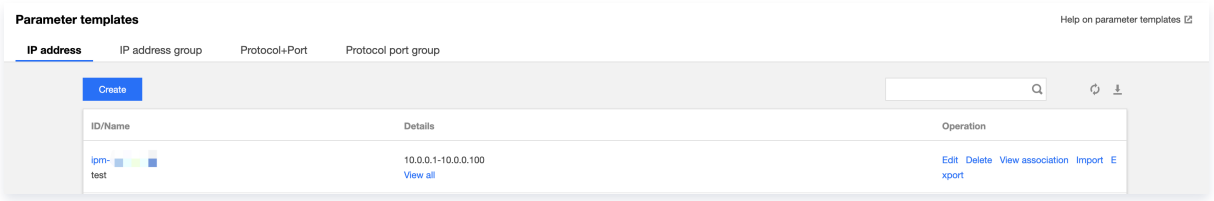


5. On the **Protocol:port** tab, click **Create** to create a protocol port parameter template respectively for adding inbound and outbound rules.

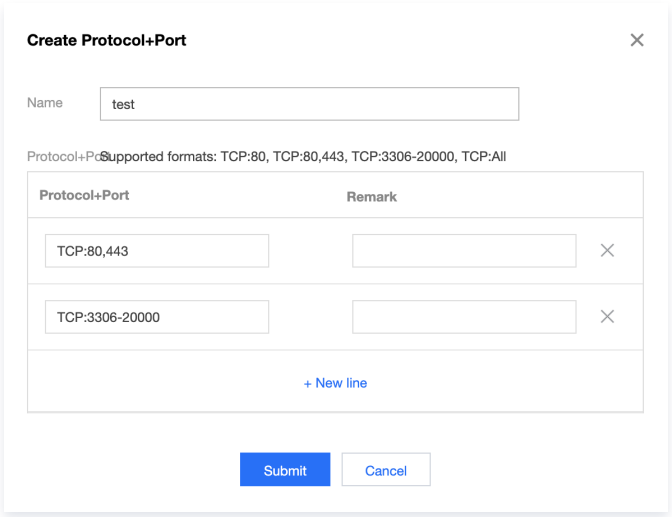


The newly created protocol port parameter template is as shown below:

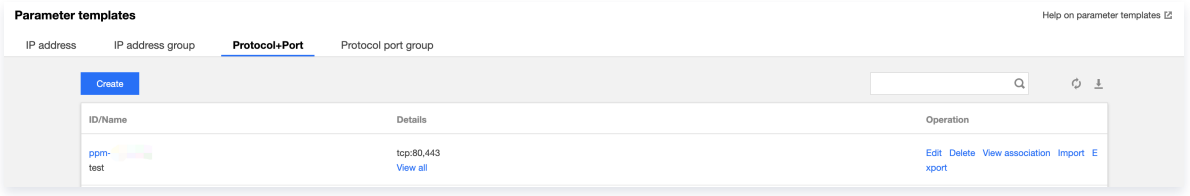


## Step 2. Add a security group rule

1. Log in to VPC console.
2. Select **Security** > **Security Group** on the left sidebar to go to the management page.
3. In the list, find the security group that needs to reference the parameter template and click its ID to enter the details page.
4. On the **Inbound rules** or **Outbound rules** tab, click **Add rule**.
5. In the pop-up window, select the custom type, choose the corresponding IP parameter templates for Source/Destination, and select the corresponding protocol port parameter templates. Click **Finish**.

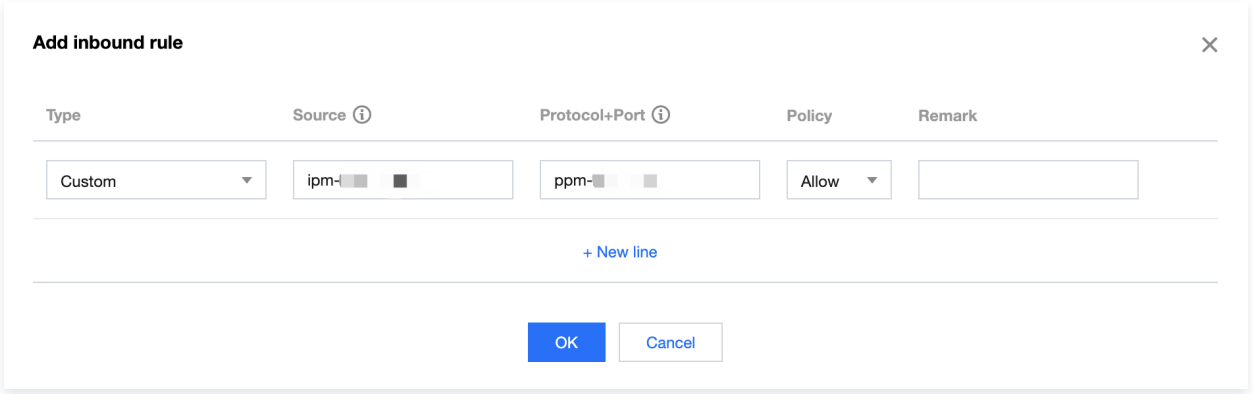

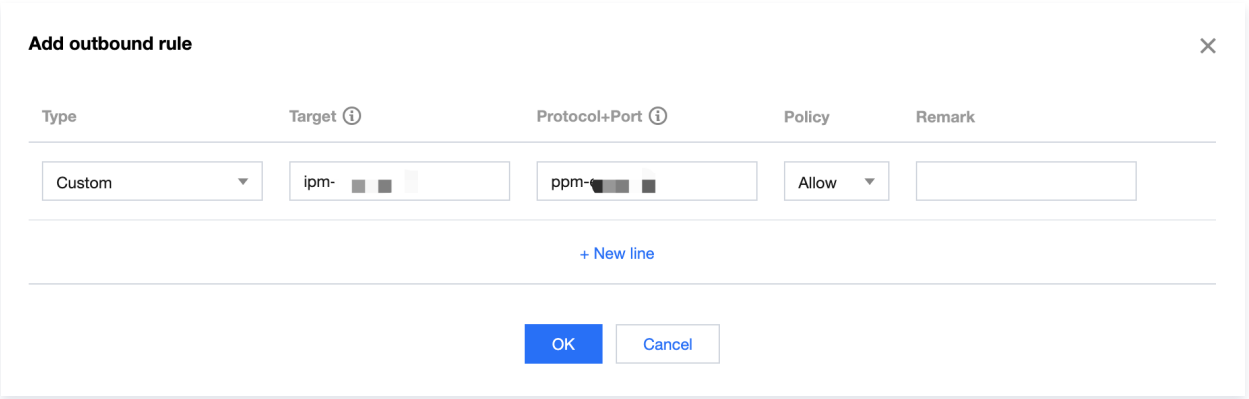


## Step 3. Update the parameter template

Suppose you need to add an inbound rule with the IP source being the `10.0.1.0/27` IP range and the protocol port being `UDP:58`. You can directly update the IP address parameter template `ipm-0ge3ob8e` and the protocol port parameter template `ppm-4ty1ck3i`.

1. On the **IP Address** tab of the parameter template, find the `ipm-0ge3ob8e` parameter template.
2. Click **Edit** on the right.

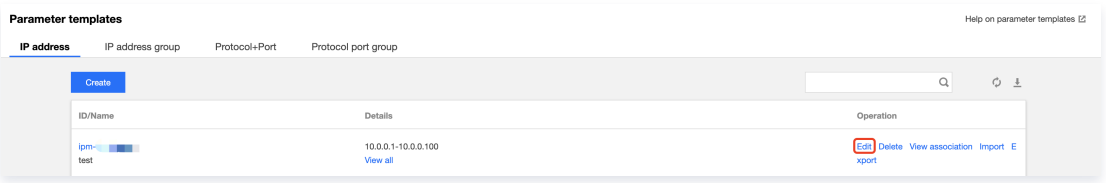


3. In the pop-up window, add the `10.0.1.0/27` IP range in a new line and click **OK**.

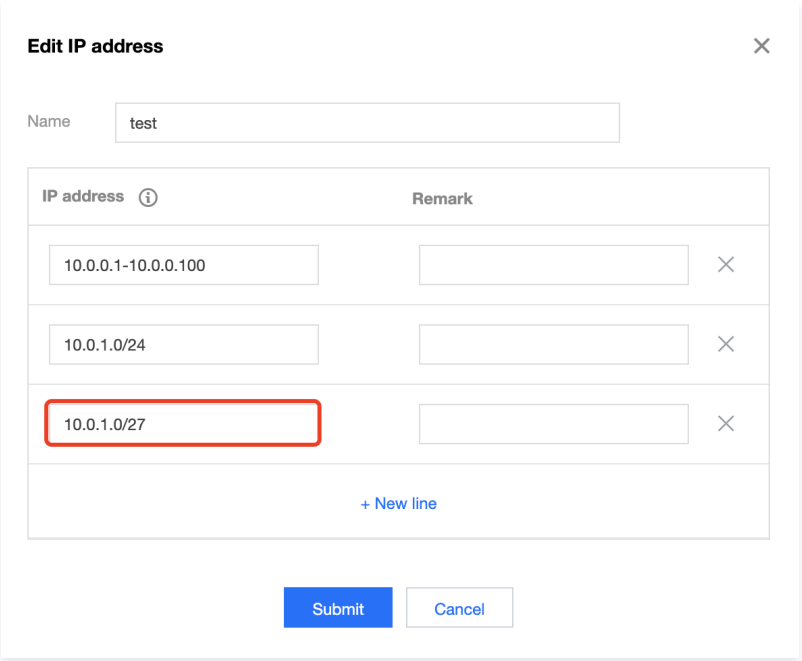


4. On the **Protocol Port** tab of the parameter template, find the `ppm-4ty1ck3i` parameter template.
5. Click **Edit** on the right.

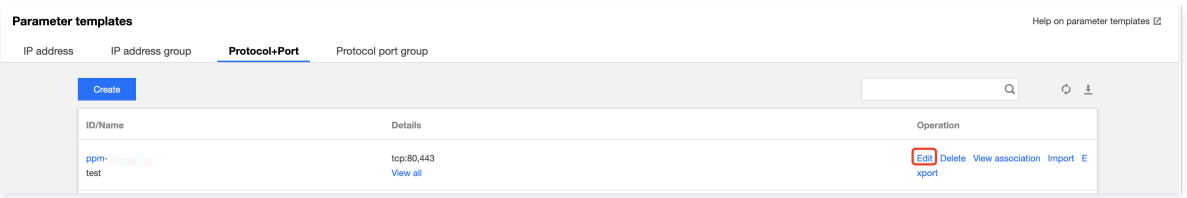


6. In the pop-up window, add the `UDP:58` inbound protocol port in a new line and click **OK**.

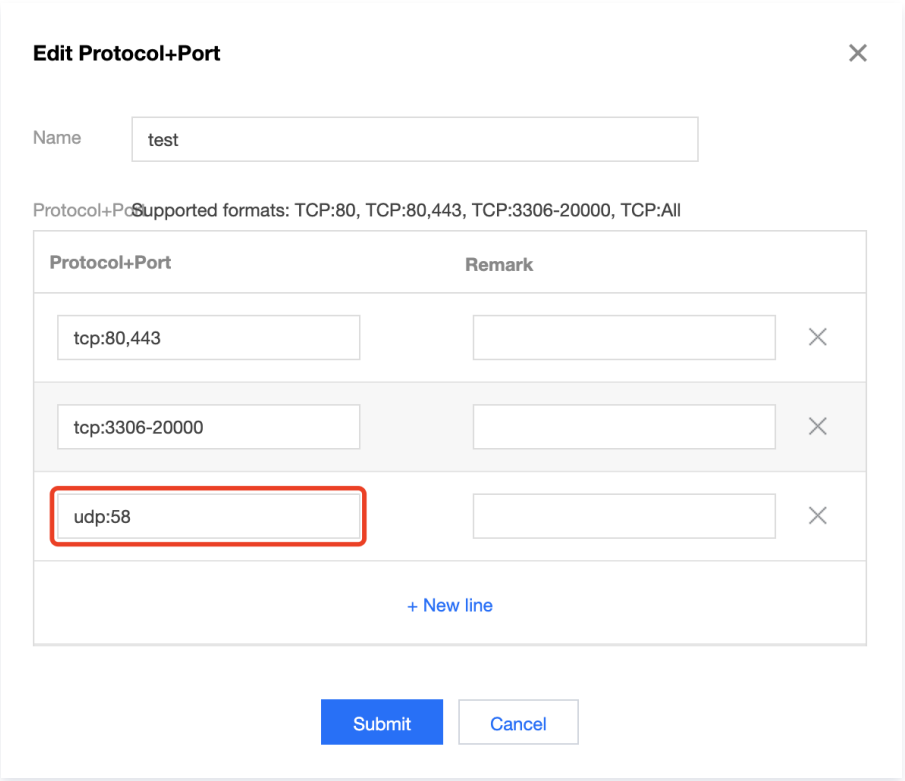

# Diagnostic Tools
# Network Probe

Last updated: 2024-01-12 14:52:01

The Tencent Cloud network probe service is used to monitor the quality of VPC network connections, including latency, packet loss rate, and other key metrics.
Under the hybrid cloud network architecture, you create a network probe in the subnet that needs to communicate with your IDC to monitor the packet loss rate and latency of the probed linkage. The configuration allows you to:

- Monitor the connection quality
- Receive alerts in case of connection failures

## Notes

- The network probe service adopts the ping method with a frequency of 20 pings per minute.
- Up to 50 probes are allowed for each VPC.
- A maximum of 20 subnets under the same VPC can have network probes.

## Creating a network probe

1. Log in to the VPC console.
2. Select **Diagnostic Tools** > **Network Probe** in the left sidebar.
3. Click **Create** at the top of the management page.
4. In the **Create Network Probe** pop-up window, fill in relevant fields.

> **Notes**
> - The network probe route is assigned by the system and cannot be modified.
> - When you switch the route of the subnet, this default route will be removed from the original route table associated with the subnet, and be added to the new route table associated.

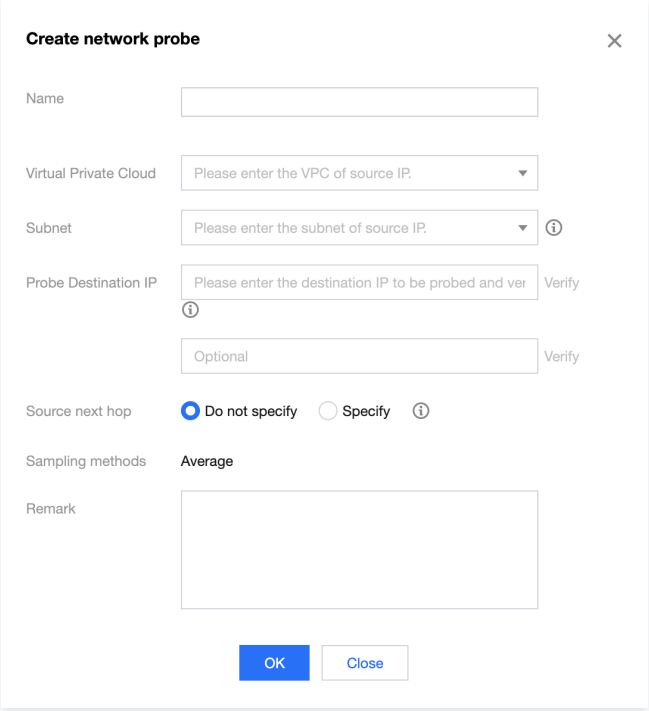


**Field Description:**

| Parameter | Description |
| --- | --- |
| Name | The name of the network probe. |

| | |
|---|---|
| VPCs | The VPC to which the probe source IP belongs. |
| Subnet | The subnet to which the probe source IP belongs. |
| Probe Destination IP | A maximum of two destination IPs are supported for a network probe. Please ensure that you've enabled the ICMP firewall policy for the destination server of the network probe. |
| Source Next Hop | You can choose to **Specify** or **Do Not Specify** the next hop.<br>• If **Do Not Specify** is chosen, no next hop will be selected.<br>**Note:**<br>**Do Not Specify** is now only available to beta users. To try it out, please contact our online customer service.<br>• If you choose **Specify**, you need to select the next hop type and instances. Then, the system automatically adds the corresponding 32-bit route to the subnet-associated route table. Currently, the supported next hop type includes NAT gateway, peering connections, VPN gateway, direct connect gateway, CVM (public gateway), CVM, and Cloud Connect Network.<br>**Note:**<br>If you specify **Cloud Connect Network** as the next hop and the probe destination IPs belong to two VPCs in the CCN, the IP range with the longest mask will be matched and take effect. |

5. (Optional) After completing the fields, click **Verify** next to **Probe Destination IP**.

**Notes**

Skip this step if you do not specify the next hop.

- If the connection is successful, click **OK**.
- If the connection fails, check whether the subnet route is correctly configured and whether the probed device enables network ACL, security group, or other firewalls, which may block the connection. For more information, see Managing Network ACLs and Modifying a Security Group Rule.

## Checking the latency and packet loss of a network probe

1. Log in to the VPC console.
2. Select **Diagnostic Tools** > **Network Probe** in the left sidebar.
3. Click the icon of the target network probe instance to view its latency and packet loss rate.

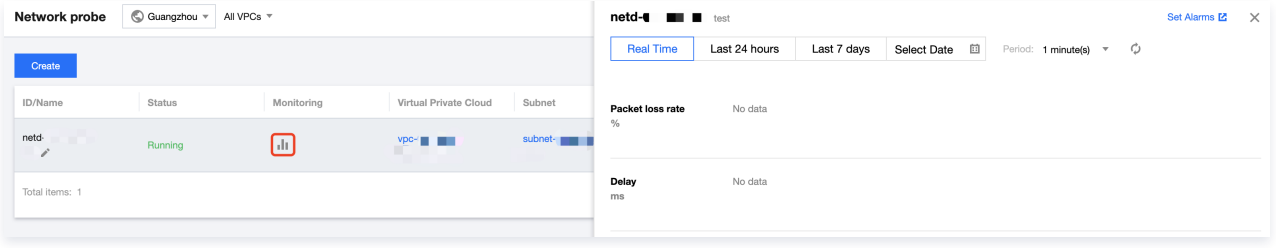


## Modifying a network probe

1. Log in to the VPC console.
2. Select **Diagnostic Tools** > **Network Probe** in the left sidebar.
3. In the network probe instance list, locate the instance you want to modify, and click **Edit** in the operation column.

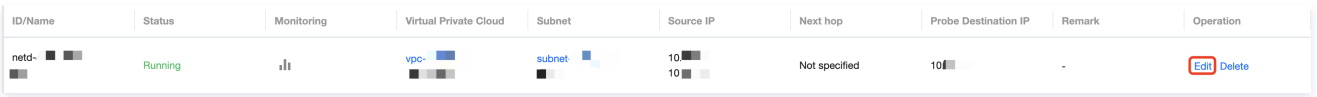


4. In the **Edit network probe** pop-up dialog box, make changes and click **OK**.

**Notes**

- In this example, no next hop is specified.
- If no next hop is specified, the name, probe destination IP, and remark of the network probe can be modified.
- If a next hop is specified, the name, probe destination IP, source next hop, and remark of the network probe can be modified.

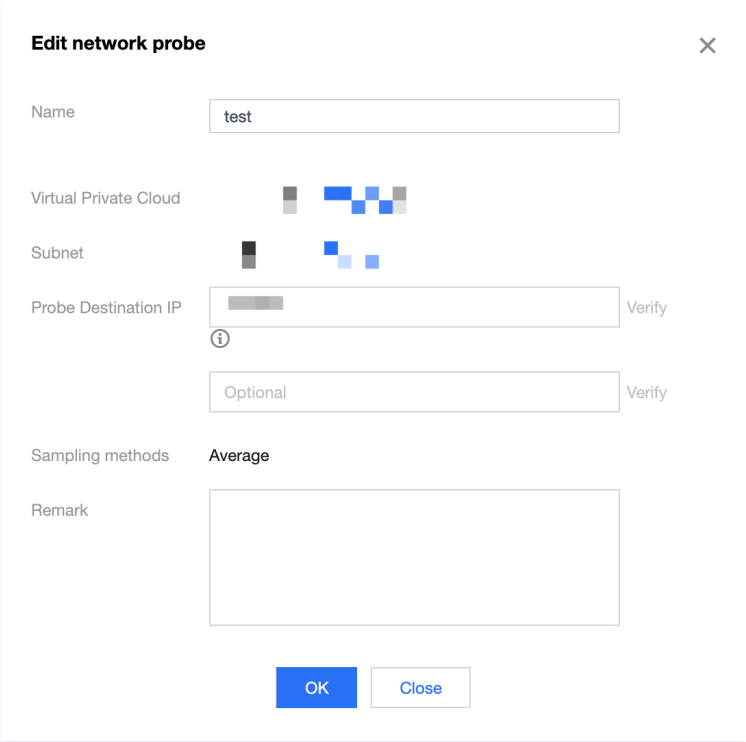


## Deleting a network probe

1. Log in to the VPC console.
2. Select **Diagnostic Tools** > **Network Probe** in the left sidebar.
3. In the network probe instance list, locate the instance you want to delete, and click **Delete** in the operation column.
4. In the pop-up confirmation box, click **Delete** again.

**Notes**

Deleting a network probe also deletes all associated alarm policies and configured routes. Check whether your business will be affected before continuing.


## Configuring an alarm policy

You can configure an alarm policy for the network probe service, so that you can promptly detect any route exception to help switch routes quickly and ensure business availability.

1. Log in to the Observability Platform console and go to the Alarm Policy page.
2. Click **Create Policy**. In the **Create Alarm Policy** pop-up window, enter a policy name, select **VPC / Network Probe** for **Policy Type**, configure the alarm object, alarm trigger condition, and alarm notification, and click **Complete**.

# Instance Port Verification

Last updated: 2024-01-12 14:53:16

The instance port verification feature can help you detect the port accessibility of a security group associated with CVM instances, locate faults, and improve the user experience.
This feature supports the accessibility detection of common ports and custom ports. See below for the common ports.

| Rule | Port | Note |
| --- | --- | --- |
| Inbound rules | ICMP Protocol | Used for transmitting control messages, such as the ping command. ICMP is a control protocol and does not involve port numbers. |
| | TCP:20 | Used to allow uploads and downloads over FTP. |
| | TCP:21 | |
| | TCP:22 | Used to allow Linux SSH login. |
| | TCP:3389 | Used to allow Windows remote login. |
| | TCP:443 | Used to provide website HTTPS service. |
| | TCP:80 | Used to provide website HTTP service. |
| Outbound rules | ALL | Used to allow all outbound traffic for access to external networks. |

## Preparations

You have created a Cloud Virtual Machine instance. For more information, see Instance Creation Guide.

## Operations Guide

1. Log in to the Virtual Private Cloud Console.
2. Click **Diagnostic Tools** > **Instance Port Verification** in the left directory to enter the management page.
3. At the top of the page, select the **Region**, and in the list, locate the row containing the instance you want to verify, then click **One-Click Check**.

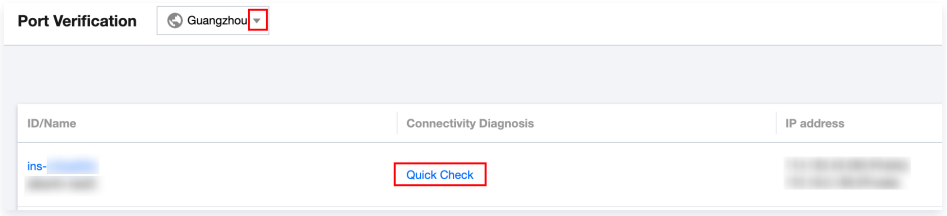


4. You can see the port detection details in the pop-up window. Perform the following operations as needed.
   - If you do not need to detect common ports, you can deselect the corresponding detection entry.
   - If the common ports do not meet your detection requirements, you can enter the desired port number in the custom port section below and click **Save**.
     - Protocol: You can choose between TCP and UDP.
     - Port: Enter the port number you want to verify. Note that the port number must not duplicate the common port numbers.
     - Direction: You can choose between Inbound and Outbound.
     - IP: For inbound direction, please enter the source IP; for outbound direction, please enter the destination IP; for all source or destination addresses, enter ALL.
     - Custom port detection supports up to 15 ports.

   In addition to custom ports, assume that you want to detect a custom port 30 using TCP protocol with the destination IP `10.0.1.12` for the outbound direction, enter the following information in the **Custom port detection** area.

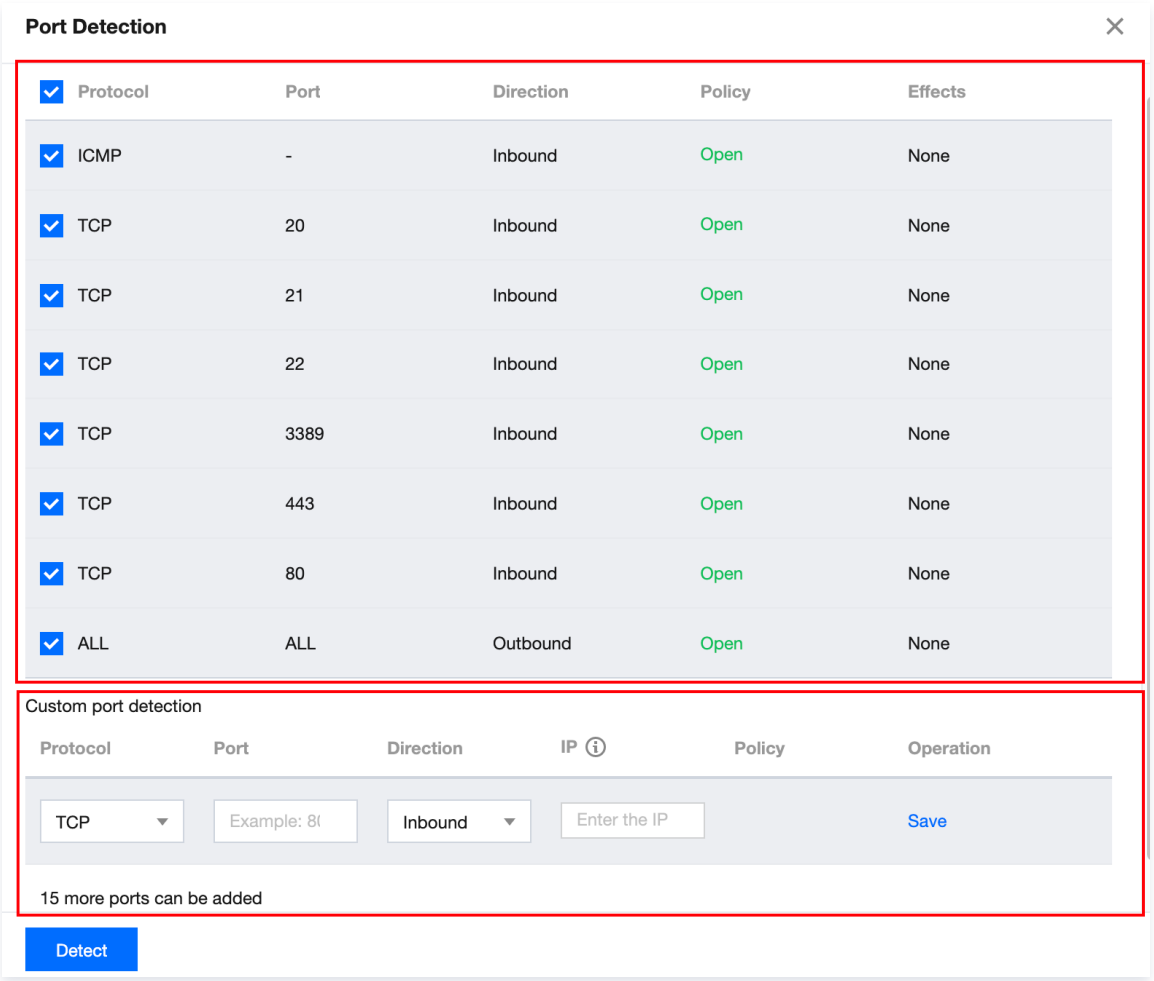


5. After completing the port check settings, click **Start Check**, and the **Policy** column will display the check results.

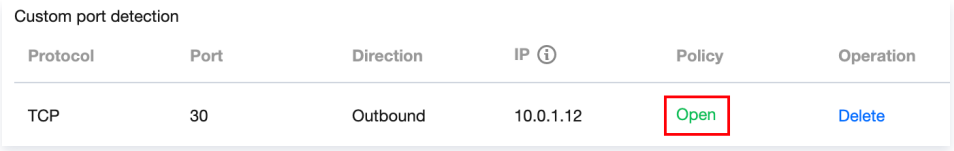


If you have a port policy marked as **not accessible** and you need to open that port (e.g., TCP:22), see the figure below.

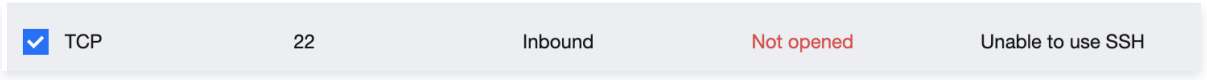


You can go to the Security Group Console and access the security group bound to the instance. Add an inbound rule to allow TCP:22 port access. Based on your actual requirements, you can either select "all" to allow all IP addresses by default in the source or specify an IP address (or IP range), as shown in the image below.

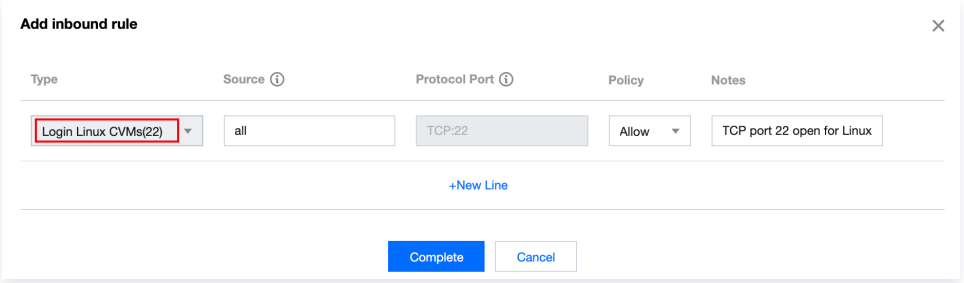


## Relevant Information

- For more information on security groups, see Security Group Overview and Adding Security Group Rules.
- For more information on common server ports, please refer to Common Server Ports.

# Flow logs

Last updated: 2024-01-12 14:53:25

Flow Logs (FL) provide real-time, full-flow, and non-intrusive traffic capture service so you can implement real-time storage and analysis of network traffic, helping you deal with troubleshooting, architecture optimization, security testing, and compliance audit. FL supports the collection of traffic data from ENI, NAT gateways, and cross-region CCN. You can access and display log data through the FL search page or via dashboards.

> **Note**
> The FL service for NAT Gateway and cross-region CCN traffic is currently in beta. To try it out, submit a ticket.

## Common Operations

- Creating Flow Logs
- Topic Configuration
- Advanced Analysis Dashboard
- Creating Logsets and Log Topics
- Deleting Flow Logs
- Viewing Flow Log Records

# Traffic Mirroring Overview

Last updated: 2024-01-12 14:53:30

Traffic mirror is a traffic collection feature that enables you to filter traffic in the specified collection range by different criteria. Then you can copy and forward the filtered traffic to CVM instances in the same VPC. This feature is applicable to use cases such as security auditing, risk monitoring, troubleshooting, and business analysis.

**Notes**

- Using the traffic mirror feature will consume CVM resources such as CPU, memory, and bandwidth. For example, if a network interface has 1 Gbps inbound traffic and 1 Gbps outbound traffic and uses the traffic mirror feature, its application system will need to handle 1 Gbps inbound traffic and 3 Gbps outbound traffic (including 1 Gbps outbound traffic, 1 Gbps mirrored inbound traffic, and 1 Gbps mirrored outbound traffic).
- When using the traffic mirror feature, please pay attention to the traffic forwarding configurations (such as ip_forward) of the related CVMs to prevent traffic loops from affecting business stability.

## Workflow

The following are key components of a traffic mirror, together with its workflow.

- Source: The specified ENIs in the VPC, which support filtering by rules such as the network, collection range, collection type, and traffic filtering.
- Target: The receiving IPs that the collected traffic is copied to.

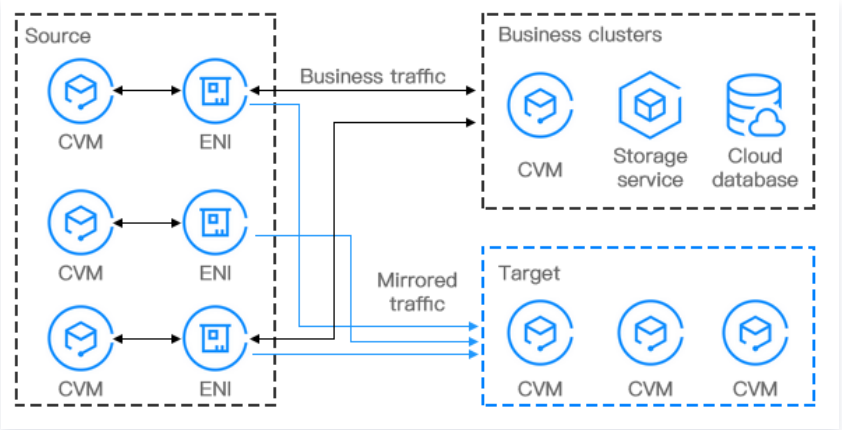


## Use Cases

### Security auditing

A running system may incur unhealthy network traffic or generate an error message due to software exceptions, hardware faults, computer viruses, or improper use. To locate the causes of these issues, you can use the traffic mirror feature to analyze the network messages.

### Intrusion detection

To ensure the confidentiality, integrity, and availability of network system resources, you can use the traffic mirror feature to copy traffic to CVM clusters for real-time analysis.

### Business analysis

Use the traffic mirror feature to clearly and visually present the business traffic mode.

# Limits

Last updated: 2024-01-12 14:53:36

To ensure the proper running of your business, refer to the following limits before using a traffic mirror.

- The traffic mirror feature is in beta now. To try it out, please submit a ticket. Save the link to the Traffic Mirror console for later logins; otherwise, you may need to apply again.
- Using the traffic mirror feature consumes CVM resources such as CPU, memory, and bandwidth.
  The mirrored traffic counts towards the instance bandwidth based on the traffic size and type. For example, if a network interface has 1 Gbps inbound traffic and 1 Gbps outbound traffic, when you use the traffic mirror, the traffic to be processed will be 1 Gbps inbound traffic and 3 Gbps outbound traffic (including 1 Gbps outbound traffic, 1 Gbps mirrored inbound traffic, and 1 Gbps mirrored outbound traffic).
- Flow logs cannot be used to capture traffic mirror data.
- Security group limits:
  - Collection source: Mirrored traffic is not subject to security group policies.
  - Target: It is subject to security group policies.
- Unsupported data:
  - ARP
  - DHCP
  - Instance Metadata Service
  - NTP
  - Windows activation
- Supported source/target models:
  Standard S1, Standard S2, Standard S3, Memory Optimized M1, Memory Optimized M2, Memory Optimized M3, High I/O I1, High I/O I2, Compute Optimized C2, Compute Optimized C3, Compute Enhanced CN3, and Big Data D1.
- Limits on CVM ENIs
  - The upper ENI bandwidth limit of the target CVM must be at least 1/9 of the total ENI bandwidth of all CVM instances in the collection range.
    For example, if there are six S3.6XLARGE48 instances in the collection range of a traffic mirror, and the total ENI bandwidth is 3 Gbps x 6 = 18 Gbps, then the inbound bandwidth at the target must be at least 2 Gbps (18/9 = 2), that is, at least two S3.MEDIUM8 instances or one S3.4XLARGE32 instance.
  - It's recommended to set the number of targets and CVM specifications a little higher than your actual business requirements to keep a buffer. For more information on CVM instance specifications, see Instance Types.

# Creating Traffic Mirror

Last updated: 2024-01-12 14:53:40

Traffic mirror is a traffic collection feature that enables you to filter traffic from the specified ENI by using quintuple and other rules. Then you can copy and forward the filtered traffic to CVM instances in the same VPC. This feature is applicable to use cases including security auditing, risk monitoring, troubleshooting, and business analysis. This document describes how to create a traffic mirror.

> **Note**
> The traffic mirror feature is currently in beta test. To try it out, submit a ticket for application. Save the link to the Traffic Mirror console for later login; otherwise, you may need to apply again.

## Preparations

Make sure that the source IP and target IP are in the same VPC and that the source IP has a route table pointing to the target IP.

## Instructions

### Step 1. Create a traffic mirror source

1. Log in to the VPC console.
2. Click **Diagnostic Tools** > **Traffic Mirror** on the left sidebar and select the target region.
3. Click **Create**.

> **Note**
> Up to five traffic mirrors can be created in a VPC.

4. In the pop-up window, configure as follows:
   - Enter a name for the traffic mirror (up to 60 characters).
   - Select a network.
   - Select **ENI** for **Collection range**. That means to collect all traffic in the VPC, excluding the traffic of the ENI that is bound to the receiving IPs. If you select this option, you need to select a specific ENI.

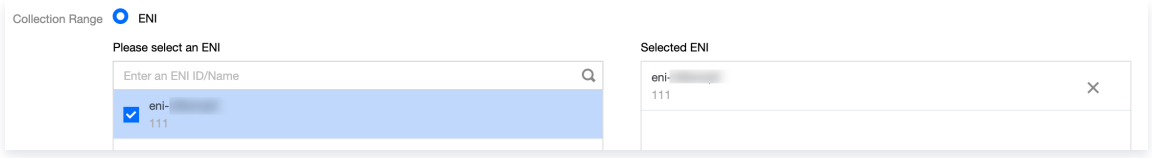


   - Set **Collection type**: Select a traffic direction as needed. There are three options: **All traffic**, **Traffic out**, and **Traffic in**.
   - Set **Traffic filtering**: Select a method to filter out unnecessary traffic and keep the mirror small and lightweight.
     - **N/A**: All traffic configured will be collected.
     - Five-tuple: Collect traffic that meets the five-tuple conditions. After selecting "Five-tuple", you need to set the "Protocol", "Source IP range", "Destination IP range", "Source port", and "Destination port". To add more filtering conditions, click "Add". Multiple filtering conditions are related by an "AND" relationship.

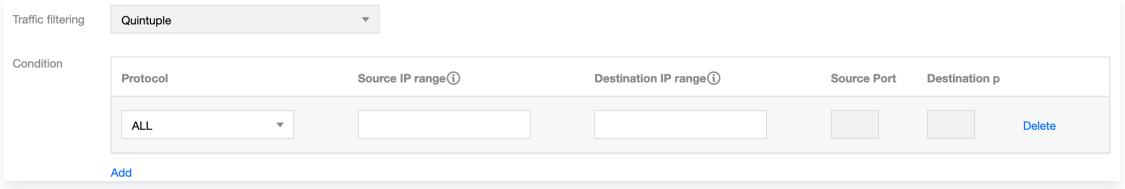


     - **The next hop is the NAT gateway**: Collect traffic whose next hop address is the NAT gateway. After selecting this option, select a specific NAT gateway next to **Condition**.
5. After completing the configuration, click **Next**.

### Step 2. Create a traffic mirror target

1. Set the following fields of traffic receiving configurations:

- **Target type**: Select the target ENI to receive the forwarded traffic.

> **Note**
> - At least one target ENI needs to be selected.
> - Traffic to the target ENI from inside the VPC will not be collected.

- **Balancing method**: Select one of the following method.
  - **Evenly distribute traffic**: All traffic is distributed among all target ENIs evenly.
  - **Hash by ENI**: Traffic from the same ENI is always forwarded to a fixed target ENI.

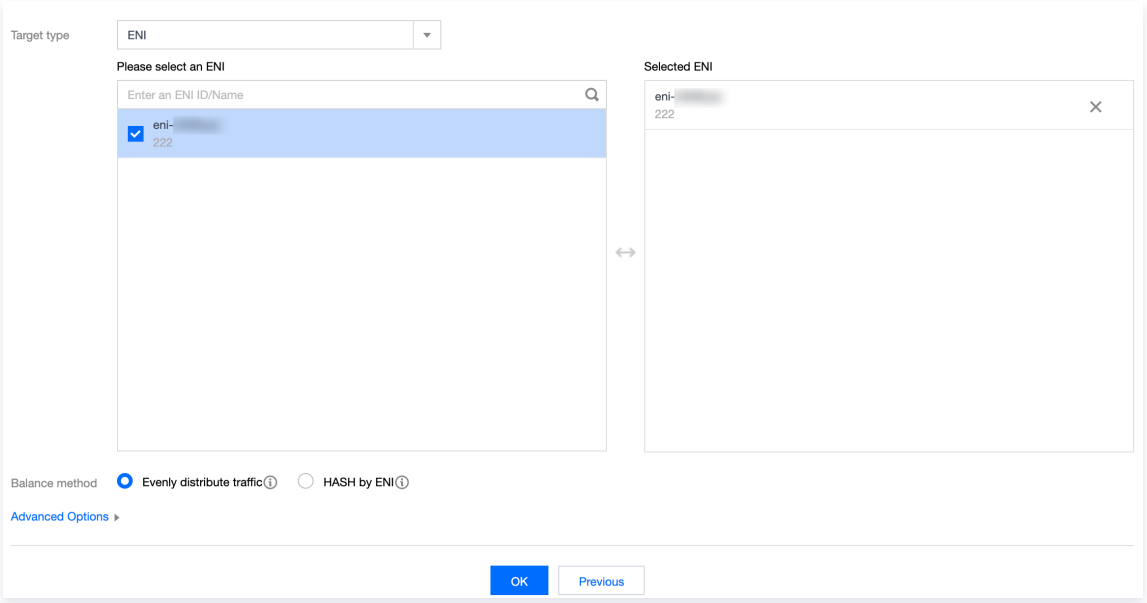


2. Click **OK**.

## Result validation

> **Note**
> This document takes creating a traffic mirror that collects the outbound traffic of the 10.0.0.14 ENI accessing the www.qq.com website as an example.

1. Return to the **Traffic mirroring** page. If the created traffic mirror is displayed in the list with **Collect traffic** enabled, it has been created successfully.

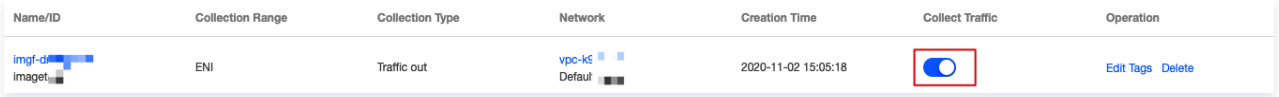


2. Perform the following steps to verify whether the collected traffic is mirrored to the receiving IP.
   2.1 Generate the ENI traffic. For example, you can log in to the source CVM and run the "ping **public IP**" command.
   **Source data:**

```
[root@VM-0-14-centos ~]# ping www.qq.com
PING https.qq.com (58.250.137.36) 56(84) bytes of data.
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=1 ttl=56 time=4.55 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=2 ttl=56 time=4.61 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=3 ttl=56 time=4.61 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=4 ttl=56 time=4.62 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=5 ttl=56 time=4.58 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=6 ttl=56 time=4.57 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=7 ttl=56 time=4.57 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=8 ttl=56 time=4.59 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=9 ttl=56 time=4.58 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=10 ttl=56 time=4.60 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=11 ttl=56 time=4.57 ms
64 bytes from 58.250.137.36 (58.250.137.36): icmp_seq=12 ttl=56 time=4.62 ms
^C
--- https.qq.com ping statistics ---
12 packets transmitted, 12 received, 0% packet loss, time 27ms
rtt min/avg/max/mdev = 4.548/4.588/4.619/0.065 ms
```

2.2 Log in to the receiving Cloud Virtual Machine and execute the following command to capture data and save it as a ".cap" or ".pcap" file. In this example, we will use ".pcap".

```
tcpdump -i eth0 -w capture-2020-10-27.pcap    #Enter the actual filename.
```

**Destination packets:**

```
[root@VM-0-11-centos ~]# tcpdump -i eth0 -w capture-2020-10-27.pcap
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 262144 bytes
^C721 packets captured
735 packets received by filter
0 packets dropped by kernel
[root@VM-0-11-centos ~]# ls
capture-2020-10-27.pcap
```

2.3 Use a terminal simulator (such as SecureCRT) to log in to the destination CVM and export the file saved in Step ii.

```
sz -bye capture-2020-10-27.pcap
```

2.4 Use a packet parser (such as Wireshark) to get the data from the downloaded "capture-2020-10-27.pcap" file. In this sample, 12 mirrored packets of the source CVM instance are obtained from the destination CVM instance.

**Packet verification:**

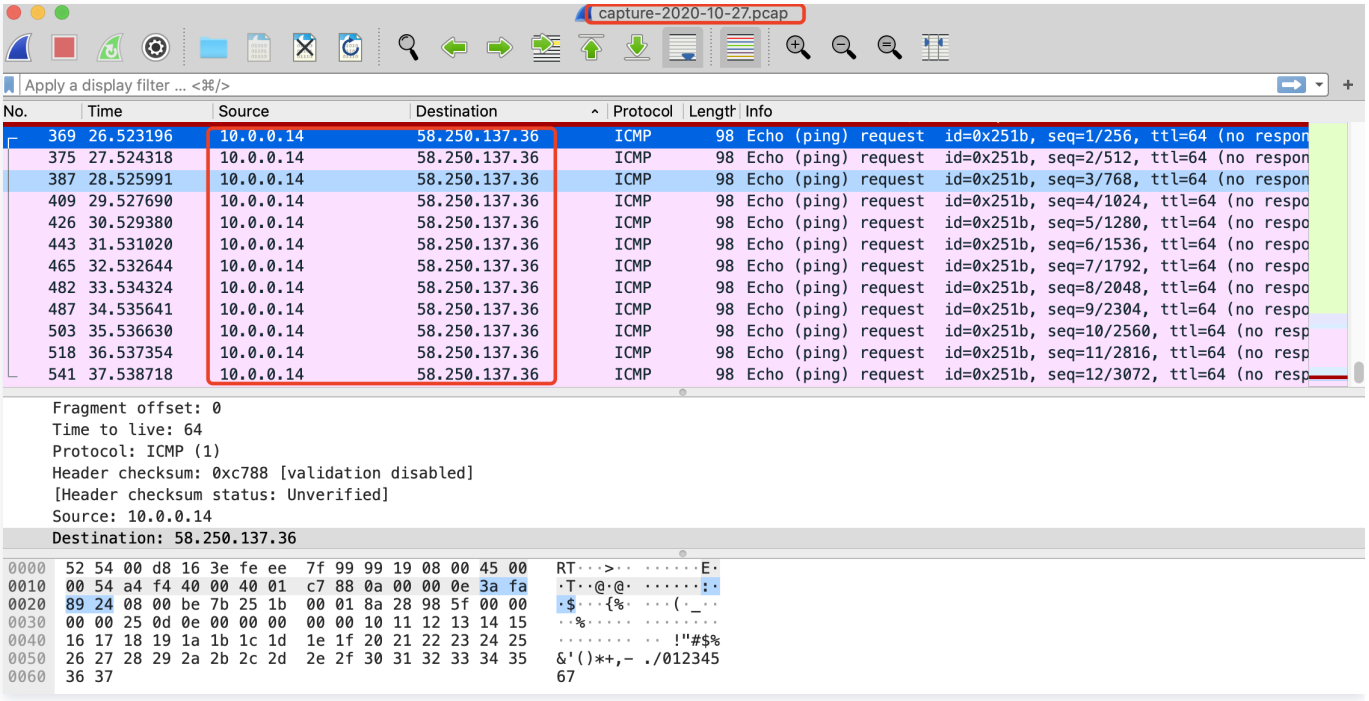


3. If an abnormal packet is obtained or it's unable to obtain packets, please contact us or seek online consultation.

## Subsequent operations

- Enabling or disabling the traffic mirror
- Modifying the traffic mirror
- Adding tags
- Deleting the traffic mirror

# Managing Traffic Mirror

Last updated: 2024-01-12 14:53:46

After a traffic mirror is created, you can enable, disable, modify, or delete it or add tags to it in the console.

## Enabling or disabling a traffic mirror

A newly created traffic mirror is enabled by default. You can follow the steps below to disable it and then enable it again.

1. Log in to the VPC console.
2. Click **Diagnostic Tools** > **Traffic Mirror** on the left sidebar and select the target region.
3. Locate the traffic mirror you want to manage, and disable or enable it under the **Collect traffic** column.

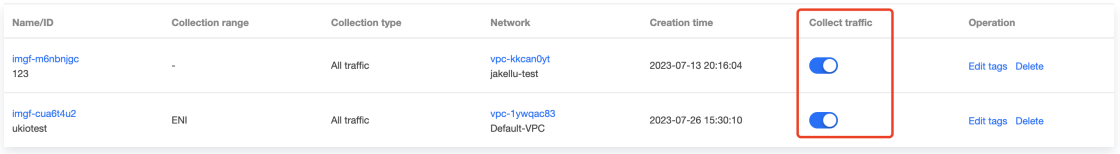


## Modifying a traffic mirror

To modify an existing traffic mirror, follow the steps below:

1. Log in to the VPC console.
2. Click **Diagnostic Tools** > **Traffic Mirror** on the left sidebar and select the target region.
3. Click the ID of the traffic mirror to be modified.
4. Edit the items you want to modify.
   - Edit traffic collection configurations.
     a. Click **Edit** in the top-right corner of the **Traffic collection configurations** module.
     b. In the pop-up window, modify parameters such as **Collection range**, **Collection type**, and **Traffic filtering**. Then click **OK**.
   - Edit traffic receiving configurations.
     a. Click **Edit** in the top-right corner of the **Traffic receiving configurations** module.
     b. In the pop-up window, modify the **Target type** and **Balancing method** parameters, then click **Confirm**.

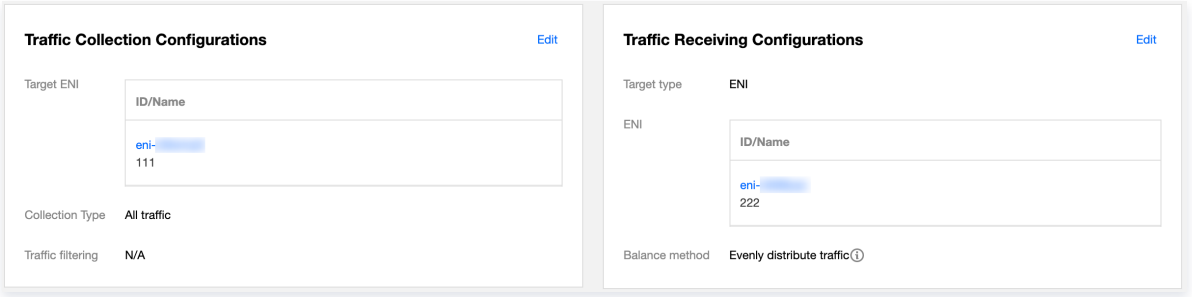


## Adding tags

Tags are used to identify and organize Tencent Cloud resources. Each tag contains a tag key and a tag value. Adding tags to a traffic mirror makes it easy to filter and manage traffic mirror resources.

1. Log in to the VPC console.
2. Click **Diagnostic Tools** > **Traffic Mirror** on the left sidebar and select the target region.
3. In the traffic mirror list, click **Edit Tags** under the **Actions** column to the right of the target traffic mirror.
4. In the pop-up dialog box, configure as follows:
   4.1 For **Tag key**, enter the key name or select from the drop-down list.
   4.2 For **Tag value**, enter the key value.

> **Note**
> A tag key may have none or many tag values.

   4.3 (Optional) Click **Add** and configure **Tag key** and **Tag value** to add a tag.

4.4 Click **OK**.

## Searching for a traffic mirror

1. In the upper right corner of the **Traffic mirroring** page, click 🔍 and select a filter. Three filters are available, as shown in the following figure.

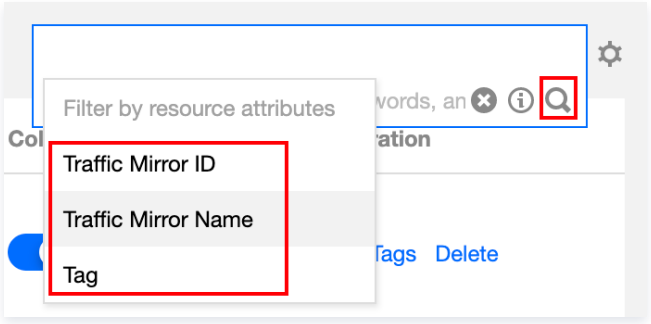


2. Enter a keyword in the edit box and click 🔍.

> **Note**
> Separate keywords with vertical bars (|).

## Deleting a traffic mirror

1. Log in to the VPC console.
2. Click **Diagnostic Tools** > **Traffic Mirror** on the left sidebar and select the target region.
3. In the traffic mirror list, click **Delete** under the **Actions** column on the right side of the target traffic mirror, and confirm the operation.

# Snapshot Policy Overview

Last updated: 2024-01-12 14:54:44

A snapshot allows you to back up the associated object data according to the configured backup policy. Currently, it can be associated with a security group, so that you can back up the outbound and inbound rules of the associated security group. This feature is an auxiliary feature, which does not affect operations on security groups.

## Scenarios

If you need to frequently update your security group rules, we recommend you configure a snapshot policy for the security group and promptly back up the security group rules. When the newly modified rules are abnormal, you can use the snapshot rollback feature to roll back to the original rules, thus ensuring business availability.

## Usage Limits

| Resources | Quota |
| --- | --- |
| Maximum number of snapshot policies that can be created by a user | 5 |
| Number of time points that can be set in each scheduled snapshot policy | 5 |
| Number of snapshot policies that can be associated with an object (security group) | 1 |
| Number of objects (security groups) that can be associated with a snapshot policy | 50 |
| Maximum retention period of a scheduled snapshot policy | 365 days |
| Backup change frequency | Five times per ten seconds |

# Creating Snapshot Policy

Last updated: 2024-01-12 14:54:49

When you need to back up security group rules to meet subsequent business needs or roll back to the original rules due to new rule exceptions, you can configure a snapshot policy.

> **Note**
> Authorize the COS service: As snapshot records are stored in a COS bucket, you need to perform read and write operations on COS. If you have not authorized the COS service when creating the snapshot policy, the system will automatically pop up a window for you to perform authorization as prompted. After performing the authorization, you can refresh the page to go to the snapshot policy page. No more authorization is required after that.

## Instructions

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. Click **Create** to enter the New Snapshot Policy page.
3. In the **Create snapshot policy** pop-up window, configure the parameters.

| Category | Note |
| --- | --- |
| Name | Customize the snapshot policy name, which can contain digits, letters, and special symbols. |
| Action | Two backup modes are supported: whenever updated and scheduled:<br>• **Whenever updated**: This mode refers to triggering a backup each time an operation is performed on the security group rules.<br>> **Note:** Currently, you can perform up to five operations per ten seconds. More frequent operations will not be recorded.<br>• **Scheduled**: Backups are performed at fixed time points. |
| Backup time | This parameter is displayed only when you select **Scheduled** for **Action**. The date ranges from Monday to Sunday, and the time can be accurate to the second. Up to five backup times can be added, and at least one is retained.<br>> **Note:** The backup operation is affected by the data volume. When the data volume is large, there may be some deviation between the selected time point and the actual time point. |
| Retention period | You can customize the retention period of backup records, after which the records will be deleted automatically. The maximum value is 365 days. |
| Create a COS bucket | • **Yes**: Create a new COS bucket.<br>• **No**: Use an existing COS bucket. |
| COS bucket region | • If you choose to create a new COS bucket, please select a region and enter a name for the bucket. The name cannot be changed once set. The name can only contain lowercase letters, numbers, and hyphens (-), and the total number of characters in the bucket name must not exceed 60.<br>• If you choose to use an existing COS bucket, please select the region and bucket name, which cannot be changed once selected.<br>> **Note:** The backup information is stored in a COS bucket. After the bucket is deleted, the snapshot information cannot be queried or restored. |

| | |
|---|---|
| COS bucket name | A COS bucket name is in the format of "custom name – developer app ID". A COS bucket name can contain up to 60 characters, supporting lowercase letters, digits, and hyphens (-). A COS bucket name cannot be changed once set. |

4. Click **OK**.

## See Also

Associating, Disassociating, and Querying Security Groups

# Associating, Disassociating, and Querying Security Group

Last updated: 2024-01-12 14:54:58

After creating a snapshot policy, you can associate it with security groups. Security groups added to the snapshot list will be backed up according to the snapshot policy and can be disassociated when snapshot backup is no longer needed. This document describes how to associate/disassociate security groups with/from a snapshot policy and how to view the associated security groups.

## Preparations

- You have created a snapshot policy.
- You have prepared security groups to be associated.

## Associating security groups

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, click the snapshot policy ID to enter the details page.
3. Click **Associate Security Group**.
4. In the **Associate security group** pop-up window, select a **region** and click the arrow icon on the right of the security groups to be associated in the **Please select** list. The selected security groups are displayed in the **Selected** list on the right. Click **OK**.

> **Notes**
> - A snapshot policy is not specific to a region and can be associated with security group instances in all regions. However, only security groups in the same region can be associated with the policy at a time. To associate a snapshot policy with security groups in different regions, perform multiple association operations.
> - A security group can be associated with only one snapshot policy.

## Disassociating security groups

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, click the snapshot policy ID to enter the details page.
3. Click **Unbind** on the right side of the security group to be unbound.
4. In the **Unbind security group** pop-up window, confirm the unbinding information and click **OK** to complete the unbinding process. After unbinding, the security group rules will no longer be backed up, but existing backup records will not be deleted.
5. (Optional) To disassociate multiple security groups at a time, select them, click **Batch unbind** above the list, and click **OK** in the pop-up window.

> **Notes**
> Only security groups in the same region can be disassociated at a time.

## Querying security groups

To query the security groups associated with a snapshot policy, follow the instructions below.

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, click the snapshot policy ID to enter the details page.
3. In the **Associate security group** section, you can view all the security groups associated with the snapshot policy.
4. Click the filter icon next to the region to filter security groups by region. Click the settings icon in the top-right corner to customize the list fields.
5. Click the **Security Group ID** to enter the security group details page.

# Enabling and Disabling Snapshot Policy

Last updated: 2024-01-12 14:55:03

A successfully created snapshot policy is enabled by default. This document describes how to disable and enable it again when needed. Disabling it will stop generating snapshot backups of all its associated security groups without deleting the original backup information.

## Disabling the policy

After the policy is disabled, no more backups will be performed but the original backup information will not be deleted.

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, click the snapshot policy ID to access the details page. The blue toggle button in the figure indicates that the policy is enabled.
3. Click the button. In the **Disable snapshot policy** pop-up window, confirm the information and click **OK**.

## Enabling the policy

After the policy is disabled, you can enable it again as instructed below. Then the associated security group rules will continue to be backed up according to the previously configured policy.

1. Click the **Enable policy** toggle button.
2. In the **Enable snapshot policy** pop-up window, click **OK**.

# Modifying Snapshot Policy

Last updated: 2024-01-12 14:55:09

You can modify the name, retention period, and backup time of a snapshot policy as instructed below.

## Instructions

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the Snapshot Policy page, click **Modify Policy** on the right side of the policy you wish to edit.
3. In the **Modify snapshot policy** pop-up window, make changes as needed.
   - For **Operational Backup**, you can modify the policy name and snapshot retention period.
   - For **Scheduled Backup**, you can modify the policy name, backup time, and snapshot retention period.
4. Click **OK**.

# Querying Snapshot Policy

Last updated: 2024-01-12 14:55:14

The **Snapshot Policy** page displays the details of all created snapshot policies, including policy name, COS bucket, backup policy, retention period, creation time, policy status, and related executable operations.

## Instructions

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, you can view the list of all created snapshot policies.
   - Click a policy ID to view the basic information and associated security groups of the policy.
   - Click a **COS bucket** name to view the details of the bucket.
   - Click the icon in the top-right corner to customize the list fields.
   - Click the icon in the top-right corner to refresh the displayed information on the page.

# Deleting Snapshot Policy

Last updated: 2024-01-12 14:55:19

This document describes how to delete a snapshot policy if you no longer need to generate snapshot backups of security group rules and want to delete the backup records of all security group rules associated with the snapshot policy.

## Instructions

1. Log in to the VPC console and select **Diagnostic Tools** > **Snapshot Policy** on the left sidebar.
2. On the **Snapshot Policy** page, click **Delete** on the right of the snapshot policy you want to delete.
3. In the pop-up window, confirm the information and click **OK**.

> **Note**
> After the deletion, all security group rules associated with the snapshot policy will no longer be backed up, and the existing backup records will be deleted.

# Alarming and Monitoring

Last updated: 2024-01-12 14:55:23

By configuring alarm policies, you can monitor the status of resources on a VPC, such as NAT gateway, VPN gateway, direct connect gateway, and EIP, so as to discover the abnormal running of cloud resources in time and locate and solve problems ASAP.

## Alarming and Monitoring

1. Log in to the Observability Platform console.
2. Select **Alarm Configuration** > **Alarm Policy** on the left sidebar to go to the policy management page.
3. Click **Create**, enter the alarm policy name, select the policy type for the Virtual Private Cloud resource you want to configure, such as **Virtual Private Cloud** > **Elastic Public IP**, and then configure the alarm rules and notifications based on your specific needs.
4. Click **Complete**. You can see the set alarm policy in the alarm policy list.

> **Note**
> To delete an alarm policy, you need to first unbind all resources from it.

5. When an alarm is triggered, you will receive the alarm notification through the selected alarm channel (SMS, email, Message Center, and so on).

Alarm policy configurations for different cloud resources are detailed below:

- Direct Connect: Configuring Alarm Policies
- CCN: Configuring Alarm Policies
- NAT Gateway: Setting Alarms
- Peering Connection: Configuring Alarm Policies
- VPN Connection: Setting Alarms

## Viewing monitoring information

You can view the monitoring information of the corresponding cloud resources in the VPC console to help you troubleshoot network failures. See:

- Direct Connect: Viewing Monitoring Data
- CCN: View Monitoring Data
- NAT Gateway: Viewing Monitoring Information
- Peering Connection: Viewing Monitoring Data of Network Traffic Over a Cross-region Peering Connection
- VPN Connection: Viewing Monitoring Data
- EIP: Viewing Monitoring Data

# API 2.0 to API 3.0 Transition Guide

Last updated: 2024-01-12 14:55:45

VPC APIs have been upgraded from v2.0 to v3.0. Technical support is no longer provided for APIs v2.0 due to their high access latency and complexity. APIs v2.0 were deprecated on **January 1, 2023**, so we recommend that you upgrade to APIs v3.0 as soon as possible to avoid impact on your business.
You can refer to the comparison tables below, where you can find the new APIs you need to upgrade and complete the upgrade accordingly.

- v2.0 Documentation: API Overview
- v3.0 Documentation: API Overview

## API v2.0 and v3.0 comparison tables

### VPCs

| Feature | API 2.0 | API 3.0 | Remarks |
| --- | --- | --- | --- |
| Creates a VPC | CreateVpc | CreateVpc | - |
| Deletes a VPC | DeleteVpc | DeleteVpc | - |
| Changes the name of a VPC | ModifyVpcAttribute | ModifyVpcAttribute | - |
| Queries the list of VPCs | DescribeVpcEx | DescribeVpcs | - |
| Binds a VIP to a CVM within a VPC | AssociateVip | / | The feature of this API has been replaced with that of Elastic Network Interface (ENI), so we don't recommend this API. |
| Creates a classiclink between a VPC and a classic network device | AttachClassicLinkVpc | AttachClassicLinkVpc | - |
| Deletes a classiclink between a VPC and a classic network device | DetachClassicLinkVpc | DetachClassicLinkVpc | - |
| Queries the classiclinks between a VPC and classic network devices | DescribeVpcClassicLink | DescribeClassicLinkInstances | - |

### Subnet

| Feature | API 2.0 | API 3.0 |
| --- | --- | --- |
| Creates a subnet | CreateSubnet | CreateSubnet |
| Deletes a subnet | DeleteSubnet | DeleteSubnet |
| Changes the name of a subnet | ModifySubnetAttribute | ModifySubnetAttribute |
| Queries the list of subnets | DescribeSubnetEx | DescribeSubnets |
| Queries subnet details | DescribeSubnet | DescribeSubnets |

### Route table

| Feature | API 2.0 | API 3.0 |
| --- | --- | --- |
| Creates a route table | CreateRouteTable | CreateRouteTable |
| Deletes a route table | DeleteRouteTable | DeleteRouteTable |
| Modifies a route table | ModifyRouteTableAttribute | ModifyRouteTableAttribute |

| | | |
|---|---|---|
| Queries route tables | DescribeRouteTable | DescribeRouteTables |
| Changes the route table associated with a subnet | AssociateRouteTable | ReplaceRouteTableAssociation |
| Adds a routing policy | CreateRoute | CreateRoutes |
| Deletes a routing policy | DeleteRoute | DeleteRoutes |

## Network ACL

| Feature | API 2.0 | API 3.0 |
|---|---|---|
| Creates a VPC network ACL | CreateNetworkAcl | CreateNetworkAcl |
| Deletes a network ACL | DeleteNetworkAcl | DeleteNetworkAcl |
| Changes the name of a network ACL | ModifyNetworkAcl | ModifyNetworkAclAttribute |
| Queries the list of network ACLs | DescribeNetworkAcl | DescribeNetworkAcls |
| Sets network ACL rules | ModifyNetworkAclEntry | ModifyNetworkAclEntries |
| Binds a network ACL to a subnet | CreateSubnetAclRule | AssociateNetworkAclSubnets |
| Unbinds a network ACL from a subnet | DeteleSubnetAclRule | DisassociateNetworkAclSubnets |

## ENI

| Feature | API 2.0 | API 3.0 |
|---|---|---|
| Creates an ENI | CreateNetworkInterface | CreateNetworkInterface |
| Deletes an ENI | DeleteNetworkInterface | DeleteNetworkInterface |
| Queries ENI information | DescribeNetworkInterfaces | DescribeNetworkInterfaces |
| Applies for a private IP for an ENI | AssignPrivateIpAddresses | AssignPrivateIpAddresses |
| Returns the private IPs of an ENI | UnassignPrivateIpAddresses | UnassignPrivateIpAddresses |
| Binds an ENI to a CVM | AttachNetworkInterface | AttachNetworkInterface |
| Unbinds an ENI from a CVM | DetachNetworkInterface | DetachNetworkInterface |
| Migrates an ENI | MigrateNetworkInterface | MigrateNetworkInterface |
| Migrates a private IP | MigratePrivateIpAddress | MigratePrivateIpAddress |

## Security Group

| Feature | API 2.0 | API 3.0 | Remarks |
|---|---|---|---|
| Queries the security groups associated with an instance | DescribeInstancesOfSecurityGroup | DescribeInstances | `security-group-id` is carried in the `Filter` parameter. |
| Queries security group rules | DescribeSecurityGroupPolicys/DescribeSecurityGroupPolicy | DescribeSecurityGroupPolicies | – |
| Deletes a security group | DeleteSecurityGroup | DeleteSecurityGroup | – |
| Queries security groups | DescribeSecurityGroupEx/DescribeSecurityGroups | DescribeSecurityGroups | – |
| Replaces a single security group rule | ModifySingleSecurityGroupPolicy | ReplaceSecurityGroupPolicy | – |

| | | | |
|---|---|---|---|
| Modifies the outbound and inbound rules of a security group | ModifySecurityGroupPolicy | ModifySecurityGroupPolicies | – |
| Queries the security groups associated with an ENI | DescribeNetworkInterfacesOfSecurityGroup | DescribeNetworkInterfaces | `groups.security-group-id` is carried in the `Filter` parameter. |
| Adds a security group rule | CreateSecurityGroupPolicy | CreateSecurityGroupPolicies | – |

## Others

| Feature | API 2.0 | API 3.0 |
|---|---|---|
| Queries VPC async task execution results | DescribeVpcTaskResult | DescribeVpcTaskResult |
| Queries account attributes | DescribeAccountVpcAttributes | DescribeAccountAttributes |