Cloud Firewall

Product Documentation

Contents

Cloud Firewall 3
- Product Introduction 3
  - Product Overview 3
  - Product Strengths 6
  - Use Cases 8
  - CFW Networking Planning 10
  - Performance Specifications 12
  - Concepts 15
- Operation Guide 18
  - CVM Creation 18
  - Custom Routing Mode Configuration Guide 19
  - Firewall Toggle 21
    - NAT Firewall Toggle 21
      - Feature Description 21
      - Operation Guide 23
      - Creating a NAT Firewall Instance 25
      - Performance Metrics 27
    - VPC Boundary Firewall Switch (Cluster Mode) 28
    - VPC Boundary Firewall Switch (Primary/Secondary Mode) 31
      - Feature Description 31
      - Operation Guide 33
      - Creating Inter-VPC Firewall Instances 47
      - Performance Metrics 58
  - Asset Center 59
  - Traffic Center 60
  - Access Control 61
    - Inter-VPC Rules 61
    - NAT Firewall Rules 65
    - Special Use Cases 69
  - Intrusion Prevention 72
  - CLS 76
  - Log Analysis 78
  - General Settings 79
  - Notification Settings 80
  - Common Tools 81
    - Address Templates 81
    - Policy Backup and Rollback 89
    - Task Scheduler 93
    - Packet Capture Tool 97
    - Simulation Test 100
  - FAQs for Beginners 104
- Ops Guide 106
  - Overview 106
  - CFW Monitoring Metrics and Exception Alarm Handling 107

Product Introduction
Product Overview

What Is CFW?

CFW is a SaaS firewall product developed by the security team based on cloud native strengths. It provides users with inter-VPC and NAT firewall protection and addresses unified cloud access control management and log audit. In addition to the features of traditional firewalls, CFW supports multi-tenancy on the cloud and is an essential network security infrastructure for users to migrate their businesses to the cloud.

Features

Overview

CFW provides users with a unified network access control center. Through the Overview page in the console, tenants can intuitively see all content related to the firewall, including the following modules:

- Asset protection overview: Displays information such as traffic bandwidth statistics, security policy configurations, and log storage statistics.
- Firewall status monitoring: Displays the peak bandwidth of the NAT boundary and between VPCs within the past 7 days.
- Traffic statistics: Display the volume of total traffic between VPCs within the past 24 hours to 6 months.
- Security policy configuration: Displays the number of access control rules configured for the Internet boundary, NAT boundary, or between VPCs, and their remaining quotas.
- Log storage statistics: Display the total log storage capacity, existing capacity, and remaining available capacity.

CFW Instances

- NAT firewall switch: A NAT firewall is a virtualized firewall that works similarly to a NAT gateway. It can provide network address translation and security audit features, such as access control and log retention.

After a NAT firewall instance is activated and created, the system automatically identifies the subnets within the VPC in the selected region. Users only need to locate the subnet to be connected and enable the firewall switch. The system then automatically modifies the subnet route to redirect the Internet traffic of the subnet to the NAT firewall. Users can configure access control lists on the NAT firewall to filter out and manage traffic.

- Inter-VPC firewall switch: The system automatically identifies the quantity, connectivity status, and connection methods of VPCs in the private network of the cloud tenant and displays them through a visual VPC network topology view. If the inter-VPC firewall is used for the first time, you need to enable the unified inter-VPC firewall switch. After the switch is enabled, the system automatically configures sub-firewalls between every two interconnected VPCs for users. Users can enable or disable the sub-firewalls and configure access control rules based on the mapping relationships between VPCs.
- Cross-tenant VPC toggle
  On the basis of the existing Inter-VPC Firewall, protection against cross-tenant peering connections is added, allowing users to deploy the protection solutions with the same experience for multiple VPCs with multiple accounts.

Asset Center

Asset Center allows you to view and manage data and information related to each asset. You can better understand the current status of assets, manage assets, and predict and prevent security incidents by viewing the private network assets and detailed information on all your VPCs.

Traffic Center

Traffic Center displays inter-VPC activities based on the traffic access between VPCs.

Access Control

Access control rules are a collection of security policies, defined in a five-tuple form and arranged in a list. For each data stream passing through internet boundaries, CFW matches the five-tuple information according to the execution order of the rule list. If a matching data flow appears, the corresponding operation is performed on the data stream according to the action of the matched rule, thus meeting the tenant's access control protection requirements for public IP addresses and fulfilling the user's security Ops audit needs through log records.

Intrusion Prevention

CFW automatically identifies unknown risks beyond the access control rules based on its protection mode. It conducts intrusion prevention rule detection on the north-south traffic of public network IP addresses. It also prevents vulnerabilities in CVM from being exposed to the Internet.

Log Audit

CFW provides a 7-day rule hit logging feature, which can record all the traffic and corresponding effective rules of the executed firewall actions, assist security Ops personnel in auditing work, and quickly troubleshoot and repair network connection failures by searching logs. You can also view operation records for the past 30 days to improve the work efficiency of enterprises and network security administrators and reduce management costs.

Log Analysis

In Log Analysis, you can view the detailed traffic logs stored in CFW based on the login account. In addition, Log Analysis supports log search and query based on search statements, and provides report generation and statistical analysis services.

Product Strengths

Easy to Enable and No Deployment Required

CFW uses the software-defined networking (SDN) technology to provide a SaaS firewall on the public cloud, ensuring automatic identification of cloud assets and one-click enabling/disabling of the firewall service. It can be used with simple policy configurations and provides CFW features without deployment costs (when compared with traditional hardware deployment).

High Stability and Reliability, and Smooth Expansion

CFW has a primary-secondary disaster recovery mechanism to ensure stable and reliable performance. At the same time, it fully leverages the strengths of SaaS services to support elastic scaling of bandwidth, assets, and storage, enabling on-demand assignment and smooth scaling.

Unified Management and Control, High Efficiency, and Ease of Use

- CFW provides unified access control and security isolation for NAT firewall and inter-VPC traffic. In addition, it offers a unified access control plane for users.
- NAT firewall access control: Access control rules are configured based on private network assets to perform traffic control and security protection, blocking abnormal external connections.
- VPC boundary access control: Based on the access control rules configured between VPCs within a tenant, CFW can be used to control access behaviors between VPCs, achieving security isolation between VPCs.
- CFW adheres to the rule configuration modes of traditional firewalls and combines the use cases of security Ops personnel, which helps reduce your learning costs and improve usability.

Cybersecurity Classified Protection Compliance

Service and Log Audit

- CFW meets the cybersecurity classified protection requirements, such as boundary protection and access control requirements.
- Rule hit logs: When an access control rule takes effect and is hit, CFW will record the five-tuple information of the hit traffic and provide the corresponding rule to facilitate your security Ops. When a fault occurs, you can rapidly troubleshoot and fix it based on log retention.
- User operation logs: Record the user operation information in CFW, including login operations of all accounts, CFW toggle operations, and operations such as adding, deleting, and editing rules. User operation logs help improve management efficiency and reduce management costs.

Use Cases

Outbound Access Control

Outbound access control is an inside-out scenario. When you need to block access to an external address or domain name from a CVM instance on the cloud, configure outbound access control rules for all public IP addresses that need to be controlled to meet your security needs.

Inbound Access Control

Inbound access control is an outside-in scenario. When you need to prohibit an external address from initiating access to a CVM instance on the cloud, configure inbound access control rules for all public IP addresses that need to be controlled to meet your security needs.

Hit Log Audit

When you need to audit the hits of the access control rules of CFW or trace the effective rules in case of failure, you can use the rule hit log audit feature of CFW for quick log audit and failure recovery, improving your security Ops efficiency and reducing Ops costs.

Operation Log Audit

When you need to manage the operations of CFW, you can use the user operation log feature of CFW to quickly locate the operation time, operation content, and operation account to meet your management needs.

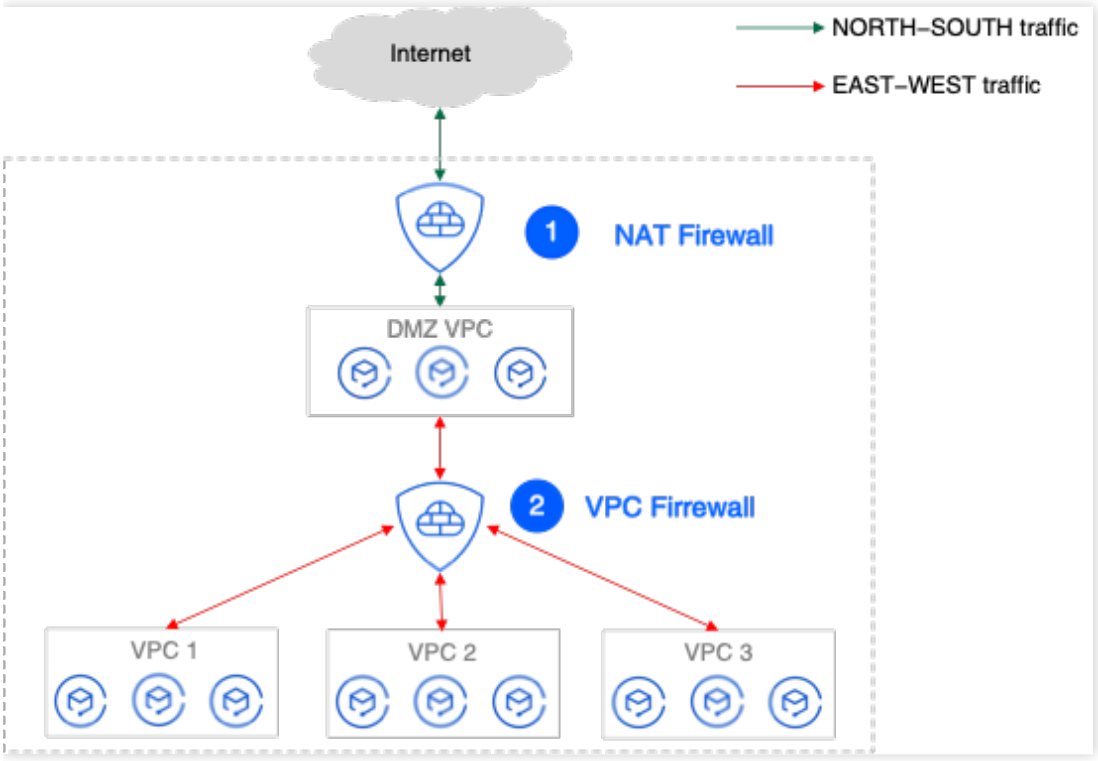
Industry Compliance Requirements

When you are ready to migrate your business to the cloud, you can quickly implement access control and intrusion detection for your business on the cloud by enabling CFW, helping you meet industry compliance requirements and cybersecurity classified protection requirements.

CFW Networking Planning

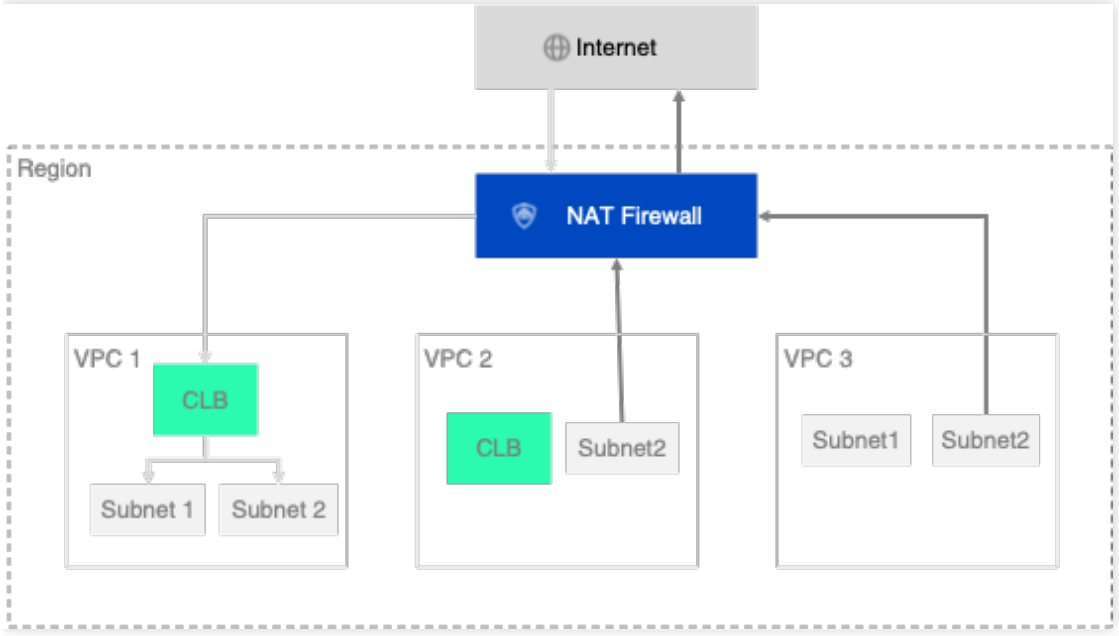
Network Planning Best Practices

CFW supports two types of firewalls: NAT Firewall (NAT_CFW) and Inter-VPC Firewall (VPC_CFW). The network planning shown in the following figure is for your reference.



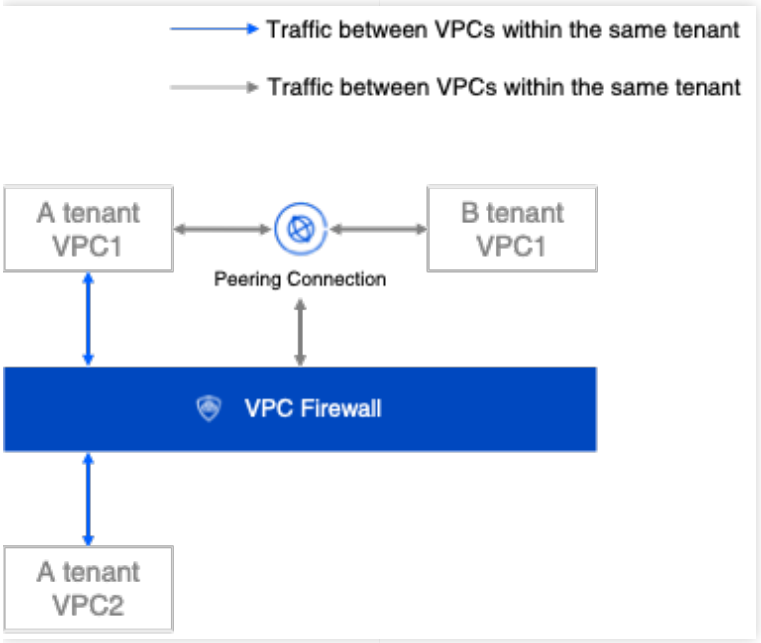
NAT CFW Network Planning Best Practices

NAT CFW is used between the Internet and a VPC.



VPC_CFW Network Planning Best Practices

The VPC firewall is used to protect traffic between VPCs, whether they belong to the same tenant or different tenants.



Performance Specifications

> Note:
>
> CFW runs in a CVM environment. For performance reasons, CVMs of different specifications are bound to different numbers of ENIs. The following describes the specification metrics supported by CFW.

VPC CFW Specification Metrics

- A VPC CFW instance supports a maximum of 10 VPCs.
- A group of VPC CFWs supports a maximum of 10 VPC CFW instances.
- Each tenant has no limit on the number of firewall groups. Each tenant has a total bandwidth limit. You can adjust the total bandwidth limit of a tenant on the cloud platform.
- The performance of a VPC CFW instance is related to the engine CVM configurations. For details, see CFW detection engine performance metrics.

Firewall Specifications	CVM Specifications	Number of Connected VPCs
0M-600M	4c8g	2
600M-2048M	8c16g	4
2048M-4096M	16c32g	6
4096M-6144M	32c64g	10
6144M-10240M	64c128g	10

NAT CFW Specification Metrics

- A NAT CFW instance supports a maximum of 10 VPCs.
- A NAT CFW instance supports a maximum of 8 EIPs.
- A NAT CFW instance supports a maximum of 5 NAT gateways.
- A NAT CFW instance supports a maximum of 6 Gbps outbound bandwidth.
- A NAT CFW instance can be connected to an EIP or a NAT gateway. However, it cannot be connected to them simultaneously.
- Each tenant has no limit on the number of NAT CFW instances. Each tenant has a total bandwidth limit. You can adjust the total bandwidth limit of a tenant on the cloud platform.
- The performance of a NAT CFW instance is related to the engine CVM configurations. For details, see CFW detection engine performance metrics.

Firewall Specifications	CVM Specifications	Maximum Number of Connected VPCs	EIP (New Mode)	NAT Gateway (Access Mode)
0M-600M	4c8g	2	1	1
600M-2048M	8c16g	4	2	2
2048M-4096M	16c32g	6	4	3
4096M-6144M	32c64g	10	8	5

CFW Detection Engine Performance Metrics

Since the maximum configuration of a physical server NIC is 10 Gbps, the maximum traffic (including inbound and outbound traffic) of the CFW cannot exceed 10 GB, and the number of new connections and concurrent connections are also subject to CVM specifications.

- Throughput test results: An 8-core 16 GB CVM model supports 2.33 GB of traffic.
- A 64-core 128 GB CVM model supports 4.8 GB of traffic. Its packet loss rate is 0%, and its blocking rate is 100%. The inbound and

outbound traffic passing through CFW is about 9.6 GB.

Test Scenario	Packet Sending	Network Stream	Sending PPS for eth1	Receiving PPS for eth0	Bandwidth	CVM Configuration	Engine CPU Utilization	Ping Packet Loss Rate (%)	Ping Latency (ms)
One-way packet sending	udp - - 1400	Client	461248	463488	4.8Gbps	S3: 64-core and 128 GB	8.9	0	1.2
		Firewall eth1	461177	462273				/	/
		Firewall eth2	461080	463378				/	/
		Server	0	471740.4				Ping packet blocking rate: 100%	/
Bidirectional packet sending	udp - - 1400	Client	444224	451072	4.8Gbps	S3: 64-core and 128 GB	8.8	0	1.4
		Firewall eth1	461481	461481				/	/
		Firewall eth2	461168	468855				/	/
		Server	460416	462720				Ping packet blocking rate: 100%	/

- Concurrent connection test results for an 8-core 16 GB model:

Scenario	Use Case	Maximum Number of Concurrent Connections
Maximum number of concurrent TCP connections	Peering connection	999999
	Firewall without rules	871881
	Firewall - Adding 5,000 rules with one rule hit	786516
	Adding 5,000 rules with only the last rule hit	837350

- New connection test results for an 8-core 16 GB model:

Scenario	Use Case	Maximum Number of HTTP Requests Per Second
Maximum number of new HTTP connections per second	Blank page	47359
	1 KB page	61530
	2 KB page	58456
	4 KB page	54912
	16 KB page	22179
	21 KB page	17098

Scenario	Use Case	Maximum Number of HTTP Requests Per Second
	64 KB page	5724

Concepts

VPC

A VPC is a logically isolated network space in the Converge Cloud that you can customize. It operates similarly to a traditional network running in an Internet data center (IDC). The resources hosted in the cloud platform VPC include CVM, CLB, and CDB instances.
VPC provides you with:

- Elastic IP: Used for Internet access.
- Peering connection: Supports interconnection between VPCs.
  Through peering connections, VPCs can interconnect with each other on the cloud. Traffic between VPCs is also known as east-west traffic.

Inter-VPC Firewall

Inter-VPC Firewall is a distributed firewall that detects east-west traffic between VPCs.
Inter-VPC Firewall is deployed between two VPCs connected via a peering connection. It supports access control, topology visualization, and log auditing features, without the need for complex routing configurations or image file installations. It provides out-of-the-box and user-exclusive resource configuration.

Access Control

ACL is a collection of traffic filtering rules. All rules that take effect on the same type of traffic form an access control list. Each ACL rule is composed of the rule body and description:

- Rule body: Used to allow or disallow the communication over a certain protocol between a port of one address and a port of another address.
- Access source: Indicates an address that initiates the communication. It is usually an IP address.
- Access destination: Indicates an address that receives the communication. It can be an IP address or a domain name.
- Destination port: Indicates the port number of an access destination address.
- Protocol: Indicates a network protocol used for communication between two parties.
- Policy: If a certain traffic meets the preceding four conditions, it is regarded as the hit traffic, and a

firewall performs the following actions based on this rule policy.

- Allow: Allow access and record traffic logs, but do not record rule hits.
- Observe: Allow the hit traffic to pass through, and record traffic logs and rule hits.
- Block: Not allow the hit traffic to pass through and not record traffic logs, but record rule hits.
- Rule description: Record the purpose of the current access control rule.

> Note:
>
> A reasonable and standardized rule description helps improve the rule's readability, reduce subsequent maintenance costs, and enhance efficiency.

Rule Priority

The execution order is the position of the rule in the list, and the rule with execution order 1 has the highest priority. For each data flow passing through the firewall, the firewall will match the rules in the list in a top-down order:

- If a data flow hits a rule, the firewall executes the corresponding policy and stops matching further rules.
- If a data flow does not hit the current rule, the firewall continues matching the next rule.
- If a data flow hits no rules, the firewall allows the data flow by default.
- In an access control list, the value of the execution priority is an integer ranging from 1 to n. 1 indicates the highest priority, and n indicates the total number of rules, namely, the lowest priority in the current rule list, usually used for wildcard rules.
- Execution priorities adhere to two principles: continuity principle and non-repeatability principle.
- Continuity principle: The execution order should be consecutive positive integers. If the number of current rules is n, the maximum value of the execution order is n.
- Non-repeatability principle: No duplicate execution order values are allowed within the same rule list.
- The continuity and non-repeatability of execution priorities are ensured by the following three points:
- Adding a rule: The execution order is set to n+1.
- Inserting a rule: The execution order is based on the insertion position. All subsequent rules will automatically shift one position backward, with their execution order values increasing by 1.
- Editing a rule: By modifying the execution order, you can move the rule to a new position. The execution order can be adjusted to the target position's execution order value, with a minimum of 1 and a maximum of n. After you enter the target position's execution order value, the rule will be inserted at that position. All subsequent rules will automatically shift one position backward, with their execution order values increasing by 1.

Logs

- Rule hit logs: Rule hits are recorded in rule hit logs, which help Ops personnel with security audit tasks. The rule hit logs display the five-tuple information of allowed, observed, or blocked data flows. You can view the corresponding effective rules through buttons.
- Operation logs: CFW operation logs are classified into user login logs, toggle operation logs, and rule operation logs.
- User login logs: Record the login information of all accounts of the current user.
- Toggle operation logs: Record CFW toggle information.
- Rule operation logs: Record the operations of adding, deleting, and editing access control rules.

Operation Guide
CVM Creation

CVM is an automatically scalable computing service. With the strengths of elastic computing, flexible configuration, stability and reliability, simple management, secure network, and comprehensive protection, it realizes cost control and business protection for pain points such as sudden traffic bursts.
Creating a CVM

1. Choose Cloud Computing & Networking > CVM.
2. Click Create.
3. On the CVM page, choose Customize Configuration > 1: Select Region and Instance Type, select the VPC and its corresponding subnet of the CVM under the Network field, and then click Next: Select Image.
4. Select Public Image, select the CentOS-related image, and then click Next: Select Storage and Bandwidth.
5. Modify the default storage and bandwidth settings as needed and click Next: Configure Security and Host.
6. Select it if a security group that allows all ports already exists. If no available security group exists, you can select Create Security Group.
7. Configure the login password of the host and click Next: Confirm Configuration Information.
8. On the Confirm Configuration Information page, view all the configured information. If everything is correct, you can click Activate to complete the creation.

Custom Routing Mode Configuration Guide

Step 1: Creating an Instance in Custom Routing Mode

See Inter-VPC Firewall Toggle > Creating an Inter-VPC Firewall to create a target instance, where the routing mode is set to custom routing.

Step 2: Configuring a Traffic Steering Route

Check the VPCs at both ends of the peering connection, which are VPC A and VPC B.

1. On the Route Table page, find the route table of VPC A, click the route table ID, and then select the default route table.

2. On the details page, click Add Routing Policy.

3. In the Add Routing Policy pop-up window, add a routing policy with the destination of VPC B IP range and the next hop of a high-availability virtual IP address, and click Create.

4. On the details page, disable the routing toggle for the original peering connection and immediately enable the route you just configured.

   > Note:
   >
   > Changing a route may cause a momentary disconnection in the network. It is recommended to perform this operation during off-peak workload hours.

5. On the Route Table page, find the route table of VPC B and select the default route table.

6. Repeat the previous steps to add a firewall route entry.

Step 3: Verifying Whether the Firewall Is Working Properly

See Log Audit to check whether there are traffic logs.

See Log Audit to check whether the intrusion prevention is normal.

Configure inter-VPC rules and check whether they are hit normally.

Firewall Toggle
NAT Firewall Toggle
Feature Description

Features

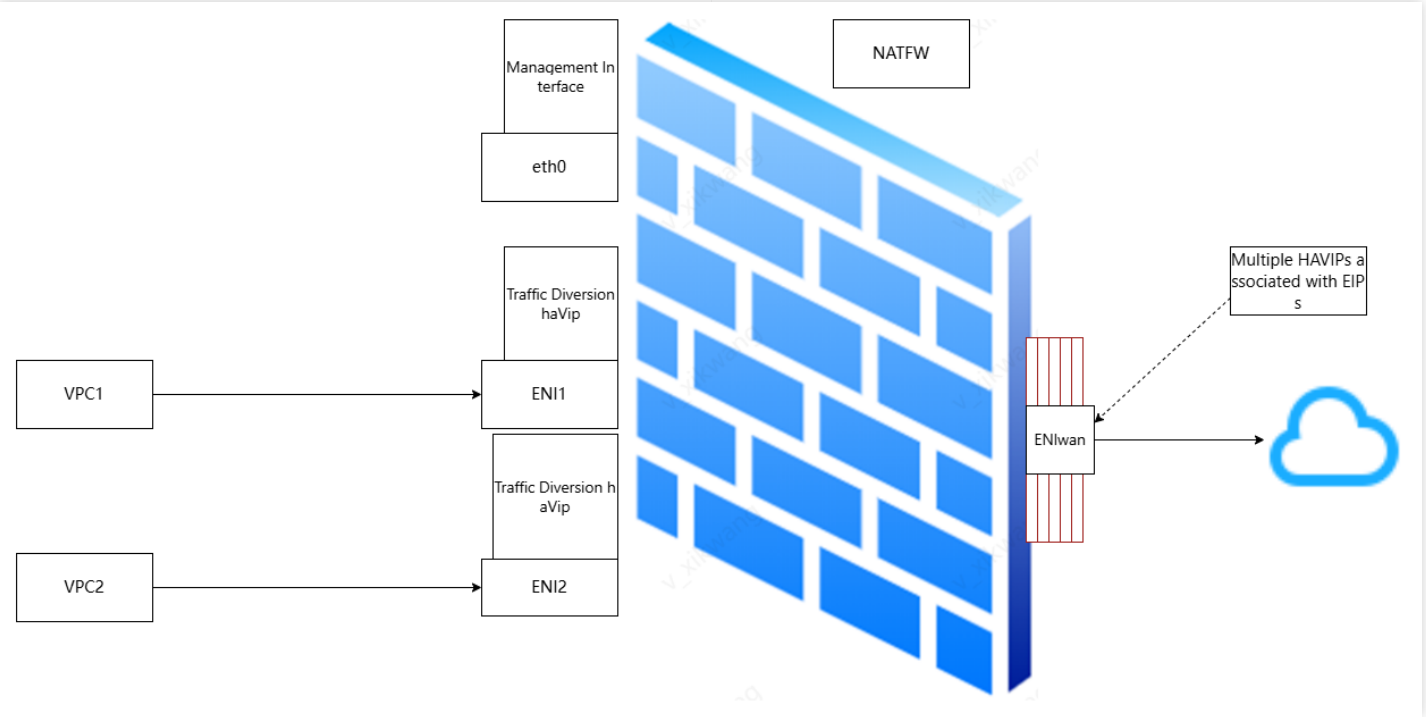
The NAT Firewall toggle supports traffic control and security protection based on private network assets, as well as network traffic forwarding based on SNAT and DNAT.

NAT CFW Use Cases

Working Mode	Prerequisite	Use Case
New mode	The tenant has an EIP that can access the public network.	If there is no NAT gateway in the current region, the new mode allows you to use the built-in NAT features of NAT Firewall to enable specified instances to access the Internet through the firewall.
Access mode	The VPC under the tenant can connect to the Internet through a NAT gateway.	If there is a NAT gateway in the current region or you want to keep the egress IP address of the public network unchanged, the access mode allows for seamless integration of NAT Firewall between the NAT Gateway and CVM instance.

NAT Firewall Routing Principles

Assume that VPC 1 and VPC 2 of a tenant need to connect to CFW. Traffic routing should be implemented first. In this example, a cross-tenant ENI is used to achieve traffic routing. The main implementation principles are as follows:



- The NAT Firewall engine is an overlay instance of the reserved account.
- Each VPC connected to the NAT Firewall will be assigned a firewall subnet under the VPC, and a cross-tenant ENI will be created for the firewall. This ENI serves as a traffic routing NIC of the VPC and will be assigned a high availability virtual IP address from the traffic routing subnet, to achieve high availability.
- If you switch the next hop route to the high availability virtual IP address of the traffic routing NIC, the traffic will be directed to the firewall.
- If a cross-tenant egress network port is assigned from one of the VPCs, for example, the ENIWan port in the figure on the left side, all VPC traffic will access the public network through the ENIWan port. To ensure high availability, you can assign a high availability virtual IP address to the ENIWan port. You can also bind the virtual IP address to an EIP to realize public network access.

> Note:
>
> Traffic that accesses a VPC passes through the subnet routes of the VPC. Based on the principle that larger subnets have higher priorities, traffic is selected and directed to a NAT Firewall. Then, NAT forwarding is performed to access public networks. In new mode, you can bind the ENIWan port to an EIP. In access mode, you can bind the ENIWan port to a NAT gateway.

Operation Guide

Operation Guide

1. Log in to the CFW console. In the left sidebar, choose Firewall Toggle > NAT Firewall Toggle to go to the NAT Firewall Toggle page.
2. On the NAT Firewall Toggle page, click Firewall Instance to go to the firewall instance page. Existing instances are displayed as cards in a list. If there is no firewall instance, you can click Create Instance to add instances.

- The current version only supports a single region.
- The list of available VPCs displays the information of all VPCs under the current account, and you can select multiple VPCs. If the selected VPCs have been connected to other instances, they will be displayed as grayed and cannot be selected.

3. You can perform other auxiliary operations.

- Synchronize Assets: On the Network Topology and Firewall Instance tabs under the Inter-VPC Toggles page, click Synchronize Assets to call the backend API promptly to read and synchronize the asset information of the current user again. This can avoid a situation where the user's asset scale is changed within the polling interval of the backend but has not been synchronized.
- Update Engines: Under the Network Topology or Firewall Instance tab, click Update Engines to immediately update the backend engine. After the update is completed, the Update Engines button becomes gray.

4. On the Firewall Instance tab, click Configure Instance.
5. In the Connect VPC and Public Network IP pop-up window on the right side, click Add Access VPC or Reselect VPC to adjust the VPCs connected to the current firewall instance.
6. On the Port Forwarding tab, create or edit port forwarding rules.
7. On the Egress Binding tab, create or edit egress binding rules.

Network Topology Diagram

CFW provides a visual view to help you quickly organize the access relationship between the VPC and the Internet. The network topology diagram illustrates the access relationship between the Internet and the

VPC. Traffic between NATs with the toggle enabled will be routed to the firewall for filtering and protection.

1. Log in to the CFW console. In the left sidebar, choose Firewall Toggle > NAT Firewall Toggle to go to the NAT Firewall Toggle page.
2. On the NAT Firewall Toggle page, click Network Topology to go to the Network Topology page.
3. Hover the mouse pointer over a VPC node to view its brief information. You can click a VPC node to view its corresponding information.

Description of Abnormal Scenarios for NAT Firewall

- When you configure a firewall instance, at least one EIP needs to be bound. When there is only one bound EIP left, the unbind button is displayed as gray and unclickable.
- When the firewall toggle is disabled, the system will automatically restore the routing policies of all route tables associated with subnets and the port forwarding rules corresponding to subnets. The traffic of all subnets will be recovered to their original paths and will not pass through NAT Firewall.

Creating a NAT Firewall Instance

Choosing Create Firewall Instance > New Mode

1. Log in to the CFW console. In the left sidebar, choose Firewall Toggle > NAT Firewall Toggle. On the NAT Firewall Toggle page, you can perform operations such as creating instances, synchronizing assets, and viewing and monitoring the bandwidth information based on NAT Firewall. Under NAT Firewall Toggle, you can click Create Instance to add instances.

   > Warning:
   >
   > Once a NAT Firewall toggle is enabled, the traffic of the corresponding subnet will pass through the firewall. At that time, access control rules and intrusion prevention will take effect for the firewall, and traffic logs will also be generated.

2. In the Create NAT Firewall pop-up window, you can create a NAT Firewall instance for the current account, configure related fields, select the mode as needed, and then click Next.

3. Under Step 2, select the VPC to be connected to CFW and click Next.

4. Select a traffic routing mode. Three modes are provided: Proprietary IP Range Preferred, Extended IP Range Preferred, and Custom. Then, click Create and wait for the instance creation to complete.

   > Note:
   >
   > CFW will create a /24 subnet in the connected VPC to redirect traffic to the firewall and public network traffic egress. You can select different ways to create subnets.
   >
   > Proprietary IP range preferred: CFW automatically selects an idle subnet IP range in the selected VPC. If the VPC has no subnet quota, an extended IP range in the VPC is used.
   >
   > Extended IP range preferred: CFW will choose an idle extended IP range reserved by the VPC first. This mode does not consume the VPC's subnet quota.
   >
   > Custom: You can only specify a /24 subnet within the CIDR block of the current VPC.

Choosing Create Firewall Instance > Access Mode

1. Choose VPC > NAT Gateway > Create NAT Gateway to configure the VPC to be connected.

2. After the NAT Gateway is created successfully, you can modify the next hop route table of the VPC to

point to the NAT Gateway. At this point, the VPC can access the public network through the NAT Gateway, and the gateway ID is recorded.

3. Log in to the CFW console and click Create NAT Firewall. On the Create NAT Firewall page, configure the instance specification, select the access mode, and then click Next.

4. Check the instance ID to be connected to the NAT gateway and click Next.

5. Select a routing mode based on the VPC assignment information of your project and click Create.

Performance Metrics

Performance Metrics Related to NAT Firewall Access Quantities

Mode	**Machine Specification**	**Maximum Number of Connected VPCs**	**Maximum Number of EIPs**	**NAT Gateway**
New mode	Instance with less than 6 million RPS (4 cores and 8 GB)	8	8	/
	Instance with more than 6 million RPS (over 8 cores and 16 GB)	10	8	/
Access mode	Instance with less than 6 million RPS (4 cores and 8 GB)	4	/	4
	Instance with more than 6 million RPS (over 8 cores and 16 GB)	5	/	5

VPC Boundary Firewall Switch (Cluster Mode)

CFW provides a cluster-mode VPC boundary firewall feature. On the VPC boundary firewall feature page, CFW automatically detects all your Peering Connection instances and helps you configure the CFW feature for the instances.

Connection Dependencies and Risks

Connection Impacts

When a VPC boundary firewall is connected, routing rules are automatically updated to implement traffic routing, which may cause 1–2 seconds of network jitter when the routes take effect. The specific impact scope and mitigation recommendations are as follows:

- Service access: Mutual access between CVM instances across VPCs and between servers within the same VPC will not be interrupted. Users perceive this as brief latency fluctuations.
- PaaS service access: When you enable or disable security protection on a Peering Connection instance, long-lived sessions for accessing PaaS services across VPCs may be interrupted. An automatic reconnect mechanism for the long-lived connections needs to be implemented at the application layer to ensure rapid service recovery.

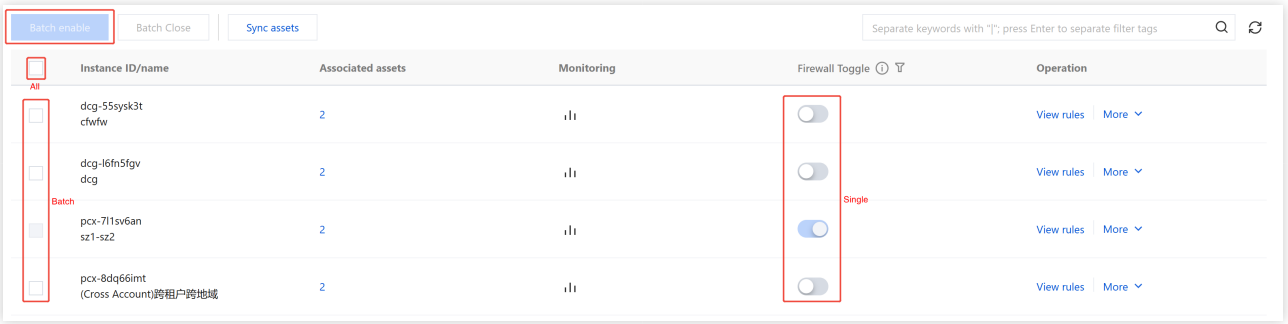
  > Note:
  >
  > The network jitter caused by the connection is a normal technical phenomenon of route updates. It is recommended to enable the firewall feature during off-peak business periods and verify the automatic reconnection capability of critical services in advance.

Connection Restrictions

Currently, inter-VPC traffic across regions cannot be protected. Therefore, this traffic cannot be connected to the VPC boundary firewall.

Firewall Enabling

1. In the left sidebar, choose CFW > VPC Boundary Firewall.
2. On the VPC boundary firewall page, enable protection for a single, multiple, or all Peering Connection instances as needed.
   - To enable protection for a single instance, click the firewall feature in the Actions column of the target Peering Connection instance.
   - To enable protection for instances in batches, select multiple target Peering Connection instances, and then click Batch Enable at the top of the page.
   - To enable protection for all instances, select the Select All checkbox in the list header, and then click Batch Enable at the top of the page.
     After you perform any of the preceding actions, the system displays a confirmation dialog box for enabling the feature.



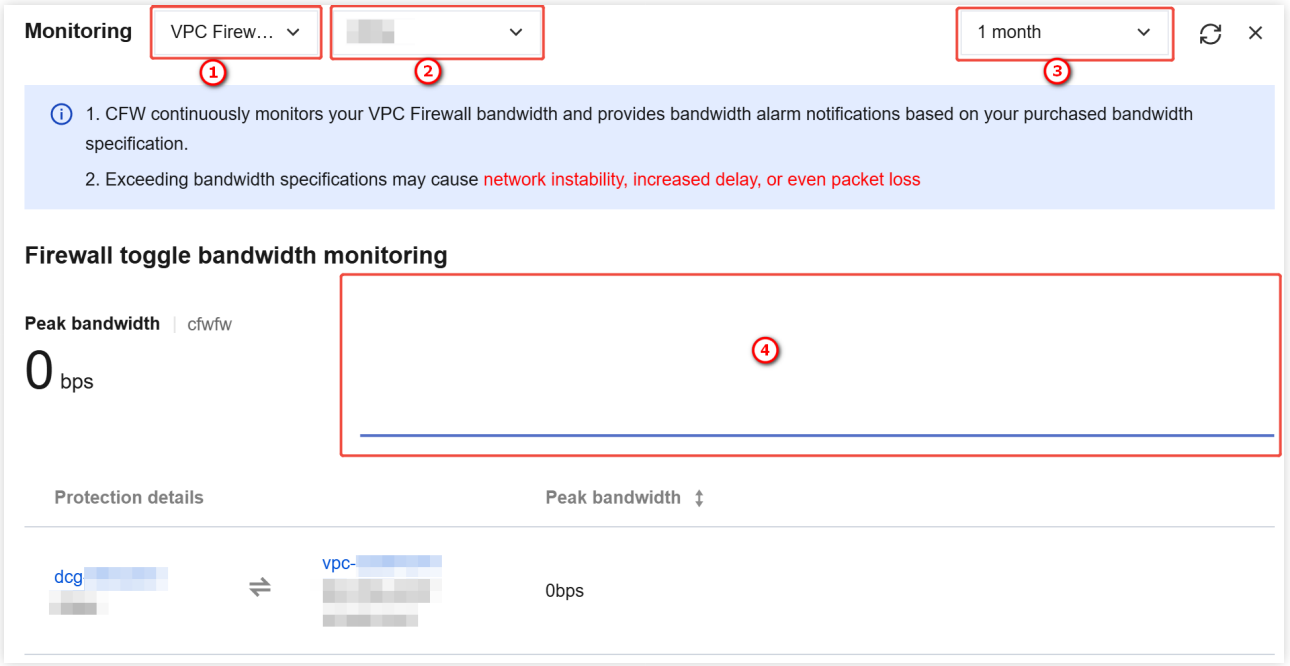
3. In the confirmation dialog box for enabling the feature, select a method for creating routing subnets. Based on your selection, CFW creates a subnet with a /28 CIDR block in the connected VPC to route traffic to the firewall.
   - Proprietary IP range preferred: CFW automatically selects an idle subnet IP range in the selected VPC. If the VPC has no subnet quota, an extended IP range in the VPC is used.
   - Extended IP range preferred: CFW selects an idle extended IP range reserved by the VPC first. This mode does not consume the VPC's subnet quota.
   - Custom: You can specify a /28 CIDR block in the current VPC.
     Example: 192.168.0.0/28
4. Click OK to enable protection.
5. Perform operations such as Disable, View Rules, and View Logs on instances that already have protection enabled in the instance list. The methods for disabling a single instance, instances in batches, and all instances are the same as the corresponding methods for enabling them.

> Note:
>
> - After you enable the firewall, the traffic of the protected object will pass through the firewall. The ACL rules and intrusion prevention feature will take effect, and traffic logs will be recorded.
> - When you disable the firewall, the corresponding traffic routing assets are deleted if no Peering Connection instances associated with the connected VPC are under protection.
> - After you disable the firewall, the traffic of the protected object will not pass through the firewall. The ACL rules and intrusion prevention feature will no longer take effect for it, and traffic logs will no longer be recorded.

Status Monitoring

Status monitoring of the VPC boundary firewall supports bandwidth statistics for each Peering Connection instance.

1. In the upper-right corner of the peak bandwidth trend panel, click View Monitoring.
2. Status Monitoring: On the VPC boundary firewall page ①, specify the instance to monitor based on the Peering Connection instance name ② and the statistical range by selecting a time range ③, and view the monitoring curve ④.



Differences Between Primary/Secondary Mode and Cluster Mode

The following table lists the differences between the primary/secondary mode and the cluster mode editions of the VPC boundary firewall.

Dimension	Primary/Secondary Mode	Cluster Mode
Architecture type	Single-active architecture.	Multi-active architecture.
Upgrade impact	Primary/Secondary switchover causes network jitter.	Zero-downtime upgrade.
Ops complexity	Users are required to manage primary/ secondary instances.	The system automatically maintains nodes and performs upgrades.
High availability	The secondary server automatically takes over when the primary server fails.	Load is distributed across multiple nodes, and a single point of failure has no impact.

VPC Boundary Firewall Switch (Primary/ Secondary Mode) Feature Description

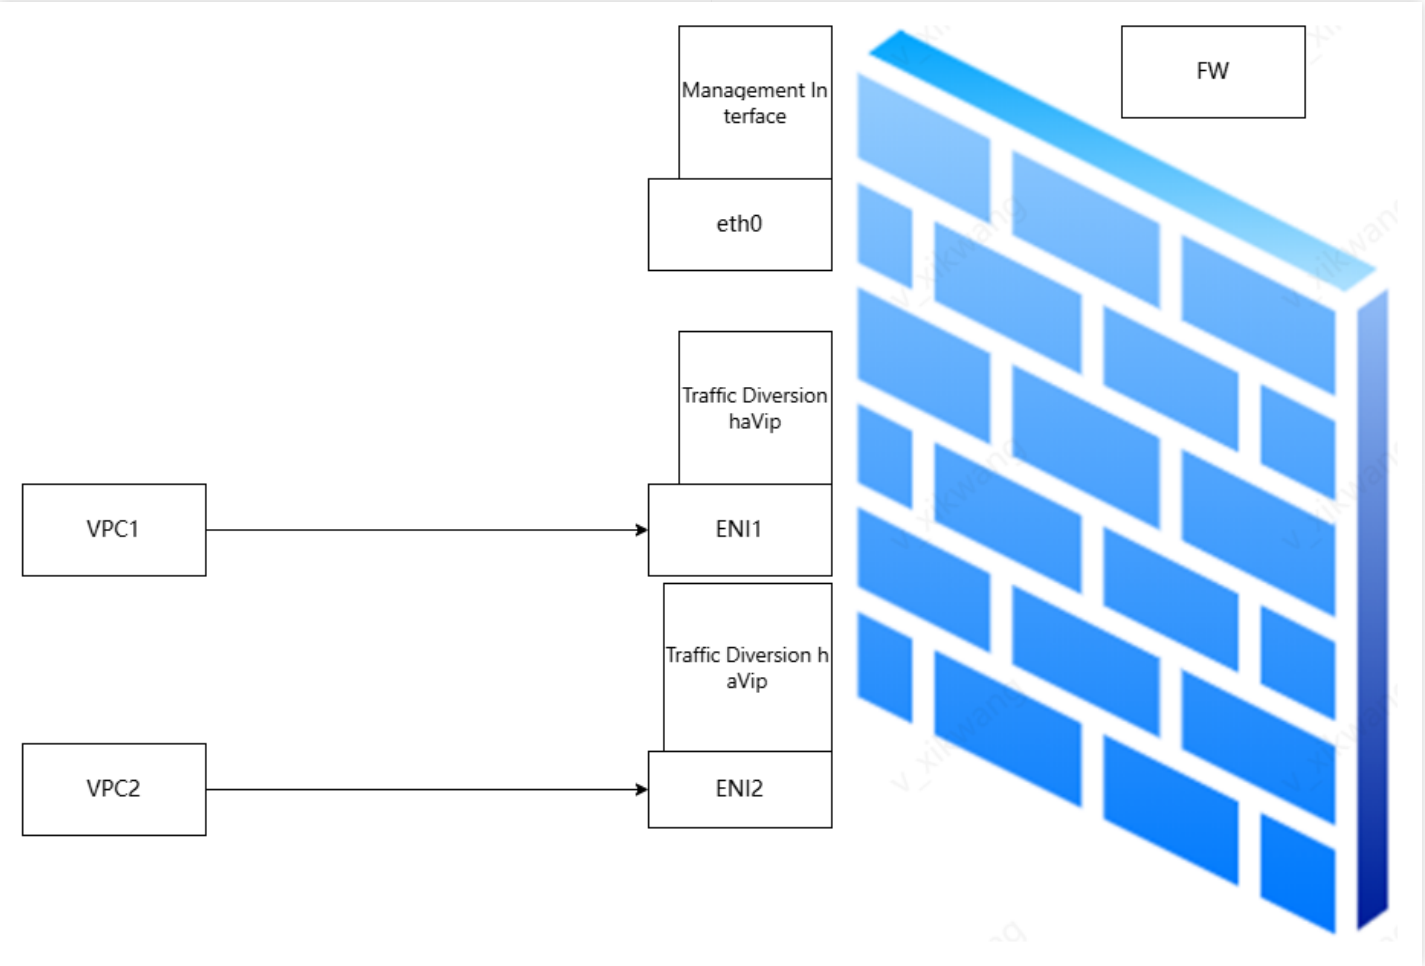
Product Features

Inter-VPC CFW provides traffic control and security protection for private network VPCs and inter-VPC access. It supports cross-region and cross-tenant traffic access control, and offers the IPS feature.

Use Cases of Inter-VPC CFW

Working Mode	Prerequisite	Use Case
Normal mode	A peering connection between VPCs has been established, the route table has been completed, and the VPCs can communicate with each other.	Inter-VPC Firewall can manage inter-VPC access traffic under a single tenant.
Administrator mode	A peering connection between cross-tenant VPCs has been established, the route table has been completed, and cross-tenant VPCs can communicate with each other.	Each tenant is responsible for managing the traffic that other tenants access the VPCs of this tenant, and the administrator mode can be used as the normal mode.
Multi-account mode	A peering connection between cross-tenant VPCs has been established, the route table has been completed, and cross-tenant VPCs can communicate with each other.	CFW needs to be created under multiple accounts across tenants, and traffic passes through two firewalls for access control.

Principles of Traffic Routing for Inter-VPC Firewall



- As shown in the preceding figure, VPC 1 and VPC 2 of the tenant are connected to the firewall. Traffic routing should be implemented first. In this example, a cross-tenant ENI is used to achieve traffic routing. The main implementation principles are as follows:
- FW is an overlay instance of the reserved account.
- Each VPC connected to the inter-VPC firewall will be assigned a firewall subnet, and a cross-tenant ENI will be created for the firewall. This ENI serves as a traffic routing NIC of the VPC and will be assigned a high-availability virtual IP address, to achieve high availability.
- If you switch the next hop route to the high availability virtual IP address of the traffic routing NIC, the traffic will be directed to the firewall.

Operation Guide

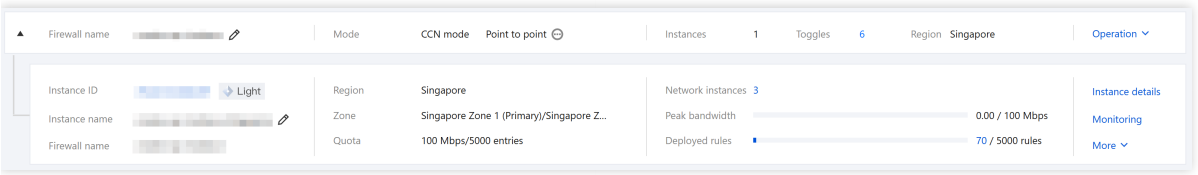
Overview

Scenarios

The CFW toggle provides the Inter-VPC Firewall toggle feature. It establishes a firewall to carry the access traffic between different VPCs and provides access control rules and a log audit system.
This document describes how to create a firewall, view bandwidth usage and specification information, view network topology, and manage firewall toggles on the Inter-VPC Toggles page.

Architecture Overview

Before using this feature in the system, you need to have a general understanding of the composition of Inter-VPC Firewall.
An Inter-VPC Firewall consists of multiple firewall instances, each responsible for connecting to a different VPC and establishing network connections between the VPCs and the firewall.



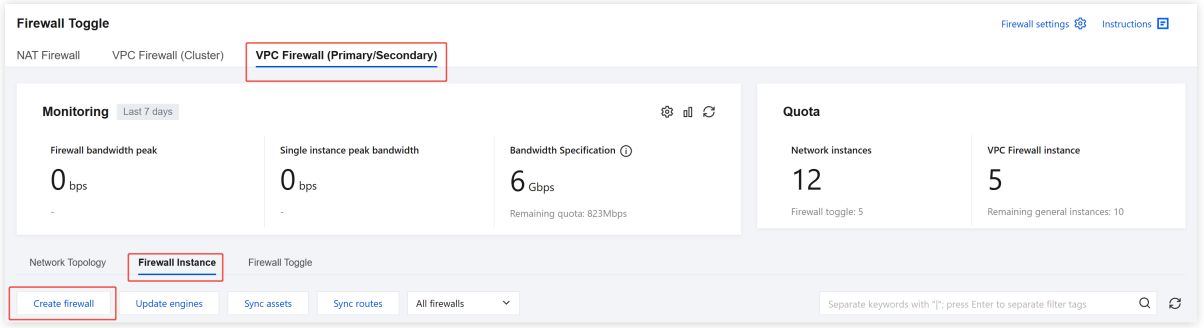
The essence of Inter-VPC Firewall is to route traffic to the firewall by modifying the VPC routing. Whether firewall instances can interconnect with each other depends on whether there are reachable routing paths between the VPCs they are connected to. The firewall does not establish basic network connections. The firewall can establish network connections by modifying the next hop in the route table of the VPC.

Creating an Inter-VPC Firewall

Private Network Mode

1. In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.

2. On the inter-VPC toggles page, click Firewall Instance to go to the firewall instance page, and then click Create Firewall.

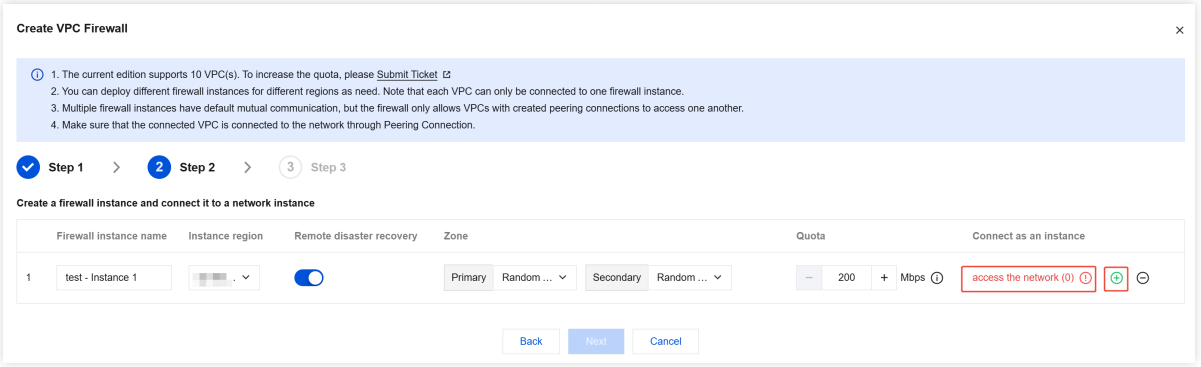


3. In the Create Inter-VPC Firewall dialog box, enter the instance name, select Private Network Mode, and click Next.

> Note:

- Instance Name: custom name defined when you create a firewall instance.
- Private Network Mode: Connect to the firewall via a VPC and modify the VPC route table to implement route steering.

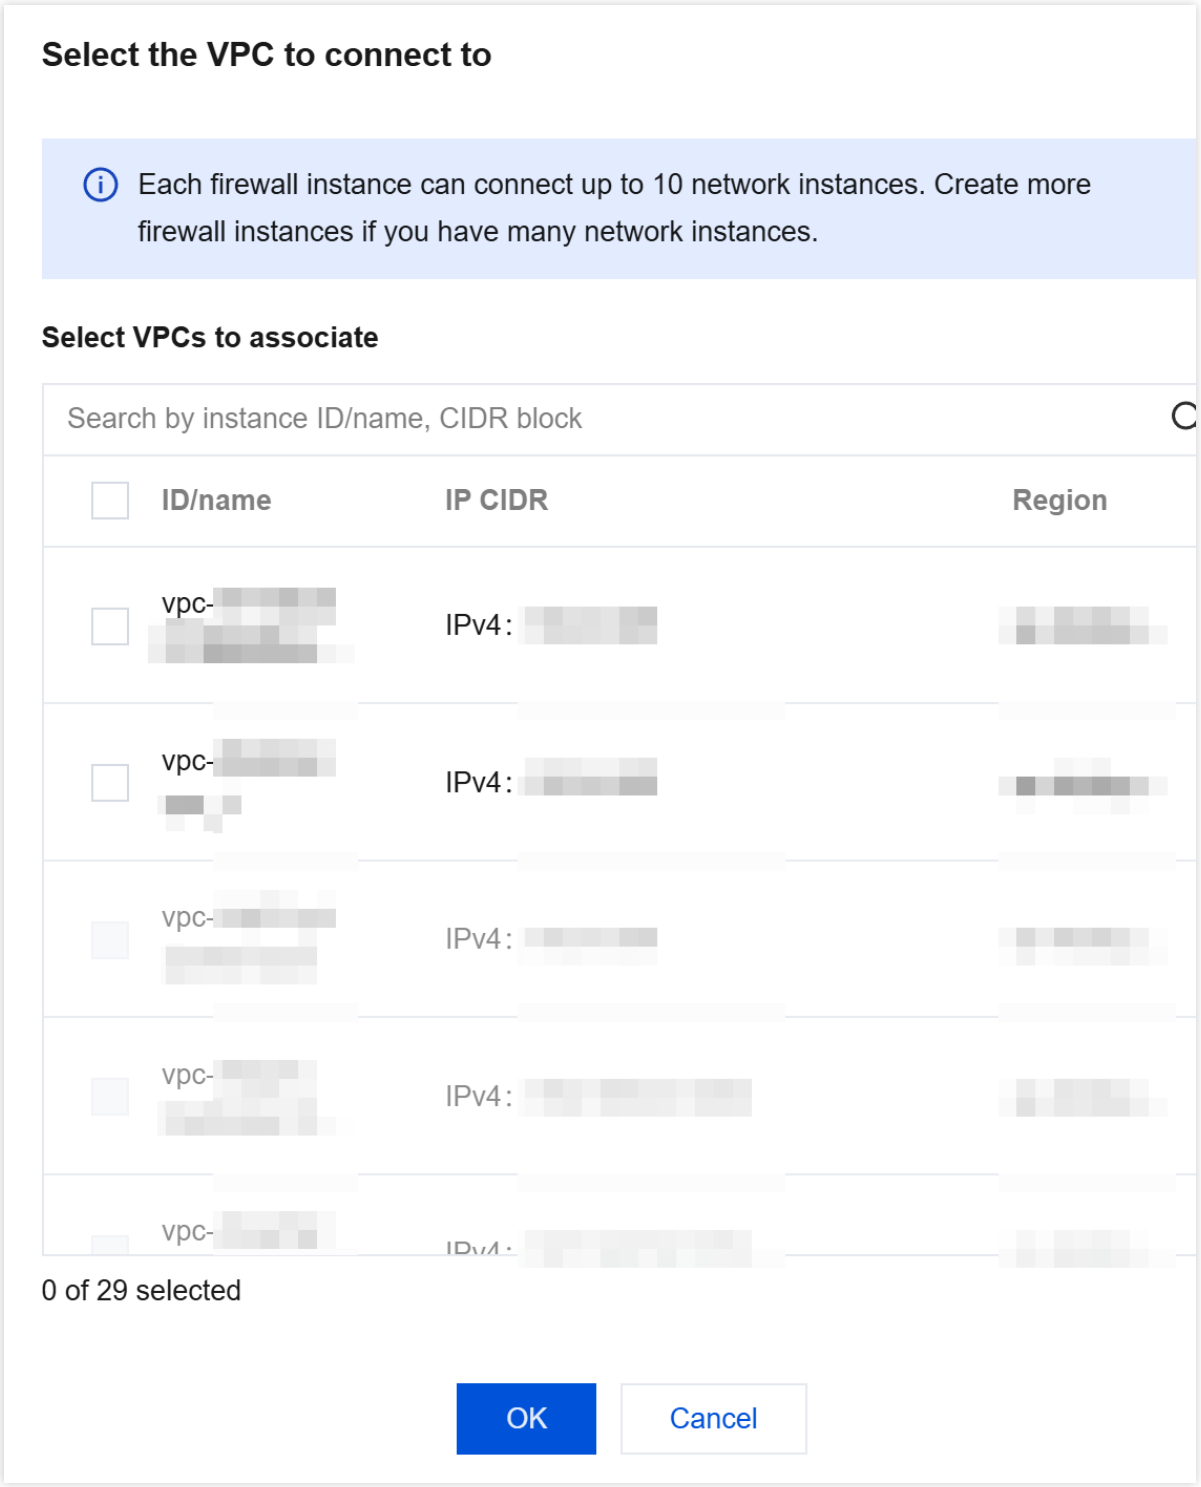
4. Enter the firewall instance name and region, configure disaster recovery information, set the firewall instance bandwidth specifications and connect to the network, and click Next. If the number of instances does not meet your requirements, you can click ➕ on the right to create multiple firewall instances.



- Region: Indicates the region to which the VPC with protection enabled belongs.
- Remote Disaster Recovery: The Inter-VPC Firewall supports remote disaster recovery, which can be enabled by checking the box.
- Availability Zone: Select a suitable AZ based on your requirements.
- Instance Bandwidth: An instance supports throughput of 1 to 20 Gbps (up to 5 Gbps configurable in the console). To configure a bandwidth greater than the configurable limit, submit a ticket or perform a scale-up. If the maximum bandwidth is insufficient, you can create multiple firewall instances to distribute the traffic. However, please note that each firewall instance has its own throughput limit. If there are multiple firewall instances, please ensure that each individual instance stays within its throughput limit.
- Instance Access: Click Access Network, select the required VPC based on the region of the instance to be connected, and click OK.

  Note:

  - A VPC can connect to only one firewall instance.
  - The firewall does not establish basic network connections. Before connecting to the network, make sure that peering connections have been created between VPCs. If no connection is established between VPCs, the connection does not take effect, and no firewall toggle is displayed.
  - Each firewall instance allows up to 10 VPCs in the same region, and you can create multiple firewall instances in the same region. It is recommended to plan the VPCs to be connected by region in advance before creating firewall instances and performing network connection.



5. Configure the method for creating the routing subnet, the firewall VPC, and the routing mode. After confirming the settings are correct, click Create.
    - How to Create the Routing Subnet: CFW will create a /24 subnet in the VPC you connect to, which will be used to redirect traffic to the firewall. You can choose different methods for creating this subnet. The routing subnet cannot be modified after the firewall is created.
        - Preferred Proprietary IP Range: CFW will automatically select an idle subnet IP range in the VPC you chose. If the VPC has no subnet quota, the extended IP range of the selected VPC will be used.
        - Preferred Extended IP Range: CFW will prioritize using the available extended IP range reserved by the VPC.

This mode will not consume the subnet quota of the selected VPC. The extended IP range refers to the auxiliary IP range within the VPC.
    - Custom: You can specify a subnet IP range for the firewall. Note that it must be a /24 subnet, and the custom IP range must fall within the CIDR of the current VPC, for example, 192.168.0.0/24.
- Firewall VPC: It is used to establish network communication between firewall instances. You need to create a new firewall-specific VPC in each region of the selected VPC.
    - Automatic Selection: The firewall will automatically create a /20 VPC that does not conflict with the connected VPC.
    - Custom: Enter a /20 VPC that does not conflict with the planned network, for example, 192.168.1.0/20.
- Routing Mode: This is the traffic routing scheme for the firewall toggle. Different network interconnection methods determine the modes of the firewall toggle and traffic routing. It is recommended to select a mode based on your business network mode.
    - Single-Point Interconnection: This mode is suitable for scenarios where a small number of VPCs need to be connected and the network topology is simple. In this mode, one firewall toggle is generated for each VPC-to-VPC connection.
    - Multi-Point Interconnection: This mode is suitable for scenarios where multiple VPCs need to be connected and the network topology is simple, such as a star network topology. In this mode, one toggle is configured for each VPC, and traffic between VPCs passes through two firewall toggles.
    - Full Interconnection: This mode is suitable for scenarios where many VPCs need to be connected and the network topology is complex, such as a mesh network topology. In this mode, only one firewall toggle is configured to control all VPC routes.
    - Custom Route: In this mode, no firewall toggle exists. You can configure a custom route as guided in Configuring Custom Routes after creating a firewall.

      > Note:
      >
      > When multiple regions are selected, only custom routes are supported. Please refer to the console to verify the availability of specific routing modes.

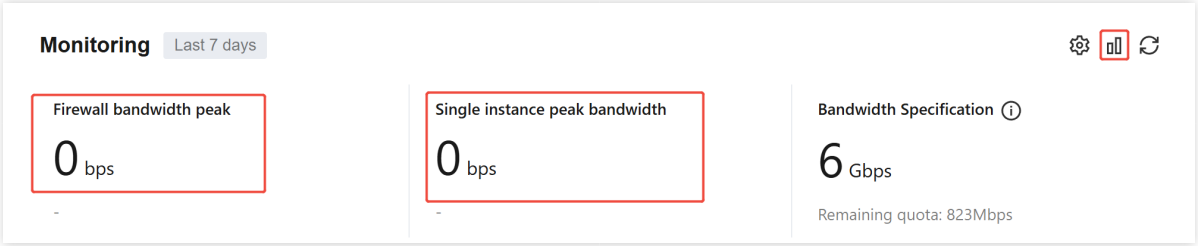
Viewing the Inter-VPC Firewall

Viewing Status Monitoring

The Status Monitoring of Inter-VPC Firewall supports the statistics of the overall firewall bandwidth as well as the bandwidth of each instance. These statistics can be queried through multiple entries.

1. In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.

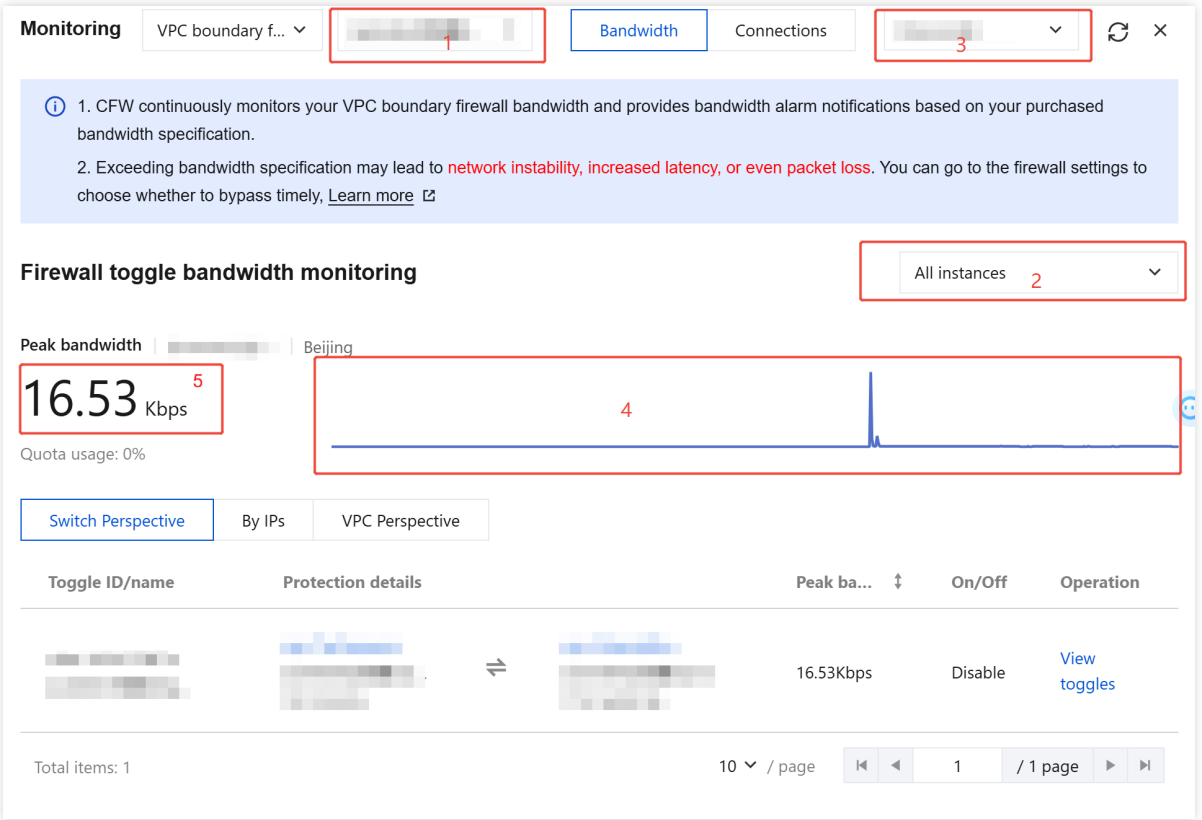
2. On the Inter-VPC Toggles page, open the monitoring panel in one of the following three ways:
    - Click the Monitor icon in the upper right corner of the Status Monitoring section, or click the Firewall Peak Bandwidth or Single Instance Peak Bandwidth panel to quickly access the monitoring panel.



- Click Firewall Instance, select the desired firewall, and then choose Operation > Bandwidth Monitoring to quickly access the monitoring panel.
- Click Firewall Instance, select the desired firewall instance, and click Status Monitoring to quickly access the monitoring panel.

3. In the monitoring panel, you can filter monitoring dimensions based on 1 Firewall Name or 2 Firewall Instance Name. You can also modify statistical dimensions by selecting a 3 Time Range. You can view the 4 Monitoring Curve for bandwidth and check the 5 Peak Bandwidth Data for toggles under the firewall or instance.

> Note:
>
> The minimum statistical granularity varies for different time ranges, which may cause discrepancies with actual peak values. It is recommended to refer to the Cloud Monitor data or the statistics of specific protection units.



4. In the monitoring panel, click View Monitoring Metrics to view all monitoring metric data of the firewall instance, including concurrent connections, inbound and outbound packets in a private network, and inbound and outbound bandwidth in a private network. You can also go to Cloud Monitor to view detailed metrics or configure related alarms.

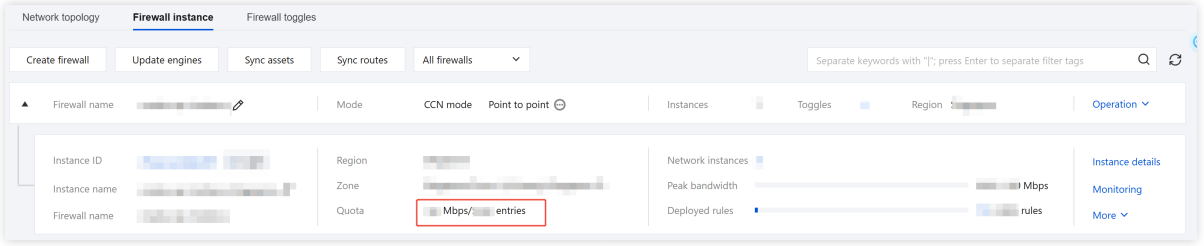
Managing the Inter-VPC Firewall

After the Inter-VPC Firewall is created, you can manage the Inter-VPC Firewall or firewall instances.

Viewing the Overview

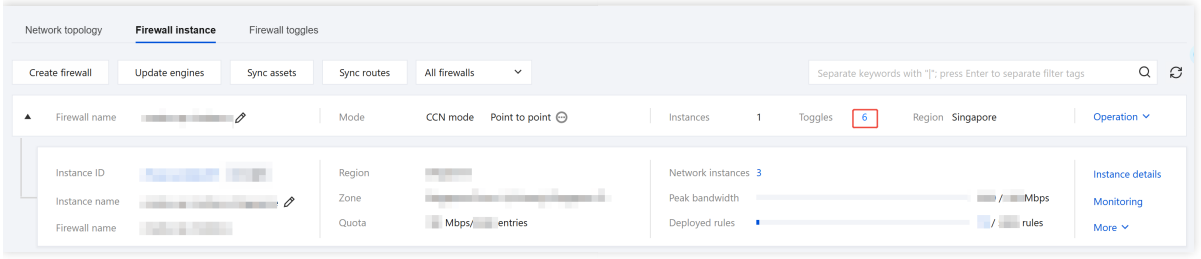
1. In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.
2. On the inter-VPC toggles page, click Firewall Instance to go to the firewall instance page, where you can view the created firewalls and their deployed firewall instance information.

> Note:
>
> Instance specifications: Indicates the maximum bandwidth and the maximum number of rules for the current firewall instance.



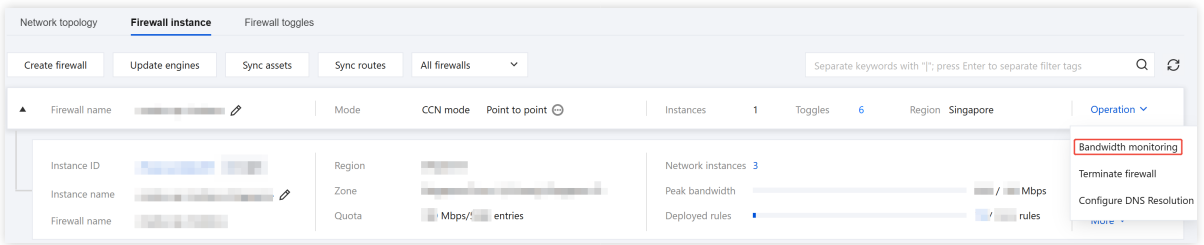
Viewing Associated Toggles

On the firewall instance page, click the number of toggles to go to the firewall toggle page. The toggles associated with the current firewall are automatically filtered.



Configuring Firewall Instances

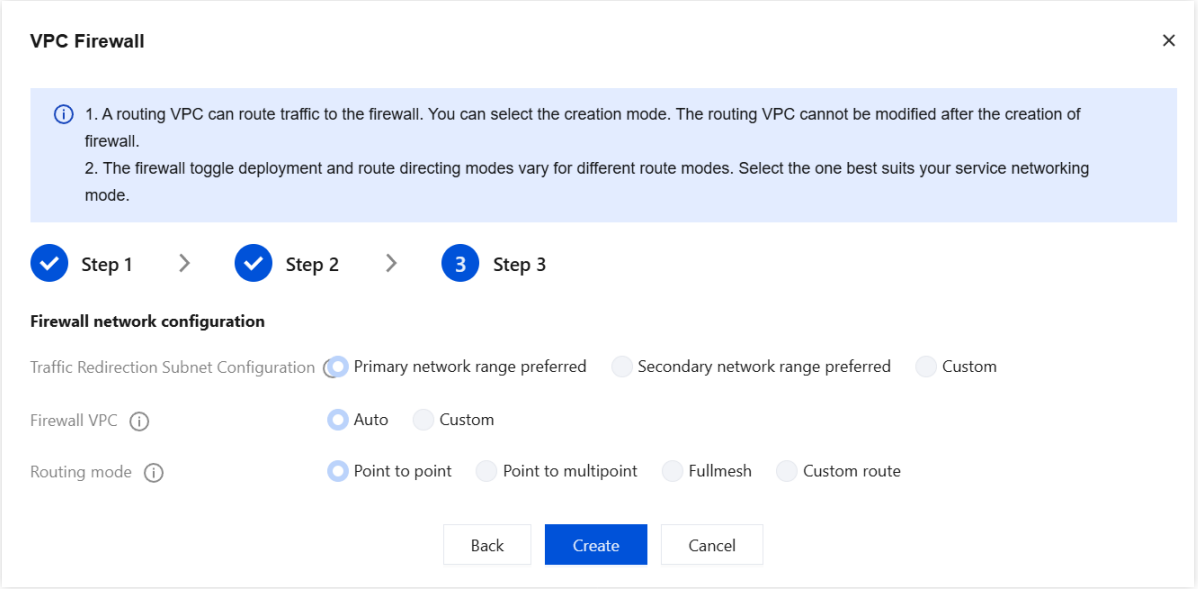
1. On the firewall instance page, hover the mouse pointer over the operation on the right side of the inter-VPC firewall that you want to configure, and then select Configure Instance.



2. Modify the name, specification, and connected network instances that are initially configured when you create the firewall instance, and add more Inter-VPC Firewall instances as needed.

3. When a new VPC is connected to the Inter-VPC Firewall, its network settings configured when you create the Inter-VPC Firewall are used automatically. If you select Custom to create a routing subnet, you also need to manually enter the subnet CIDR of the new VPC.

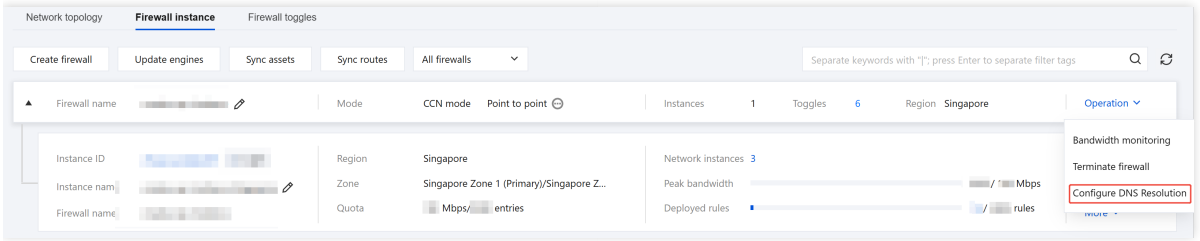
> Note:
>
> The firewall network settings cannot be modified.



Terminating the Inter-VPC Firewall

On the firewall instance page, hover the mouse pointer over the operation on the right side of the inter-VPC firewall that you want to configure, and then select Terminate Firewall. The firewall is terminated after you confirm the operation.

> Note:
>
> If you select Custom Route, make sure that you have restored the route manually.



Managing Inter-VPC Firewall Instances

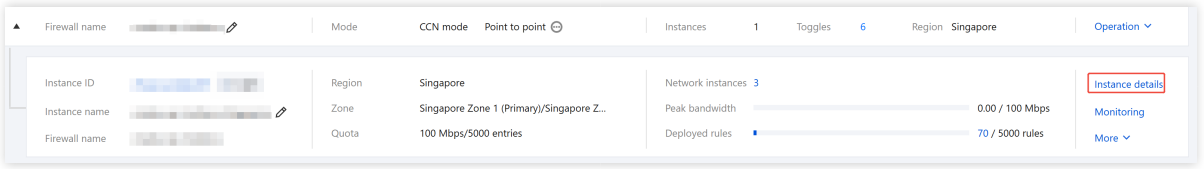
After the Inter-VPC Firewall is created, you can individually manage the Inter-VPC Firewall instances.

- In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.

- On the Inter-VPC Toggles page, click Firewall Instance to go to the firewall instance page.

Viewing Instance Details

1. On the firewall instance page, click Firewall Instance ID or Instance Details on the right.



2. On the Firewall Instance page, view the configuration details of the instance.

Viewing Associated Toggles

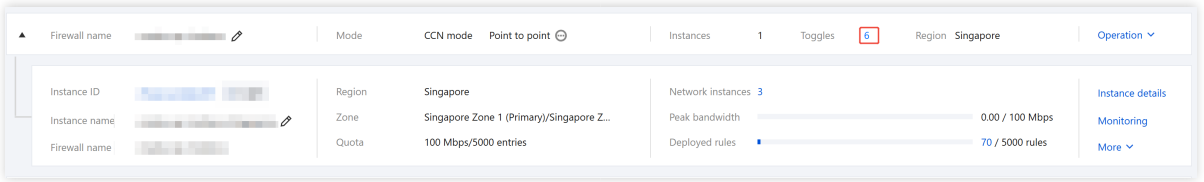
On the firewall instance page, choose More > View Firewall Toggle to go to the firewall toggle page. The toggles associated with the current firewall instance are automatically filtered.

> Note:
>
> A toggle can be associated with multiple firewall instances and controls traffic passing through these firewall instances.

Terminating Inter-VPC Firewall Instances

To terminate any firewall instance, disable all toggles associated with the current Inter-VPC Firewall.

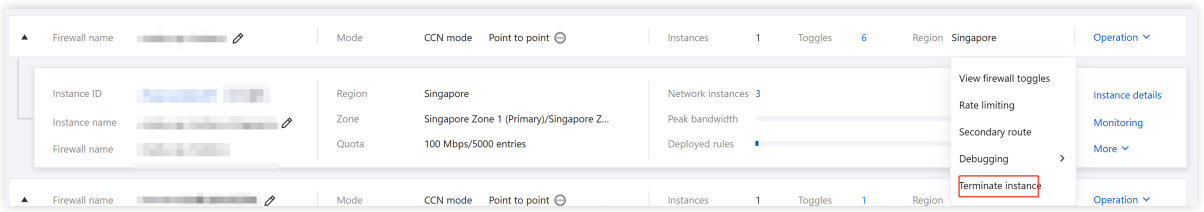
1. On the firewall instance page, click the number of toggles associated with the inter-VPC firewall.



2. Go to the firewall toggle page and click Disable All Toggles for the Current Firewall.

3. On the firewall instance page, choose More > Terminate Instance. The current firewall instance is terminated after you confirm the operation.

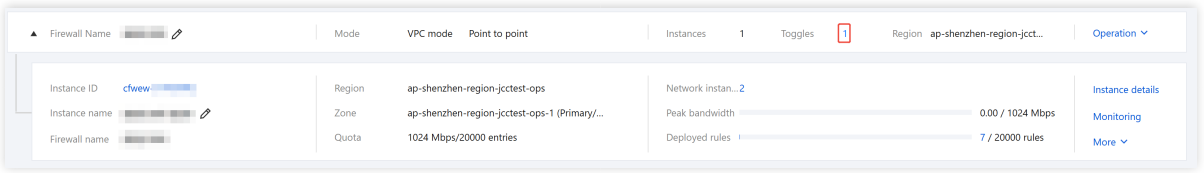
> Note:
>
> After termination, the VPCs will be automatically disconnected from the current instance, and the used quota will be returned.



Changing Associated Instances

To change an associated instance, disable all toggles associated with the current Inter-VPC Firewall.

1. On the firewall instance page, click the number of toggles associated with the inter-VPC firewall.

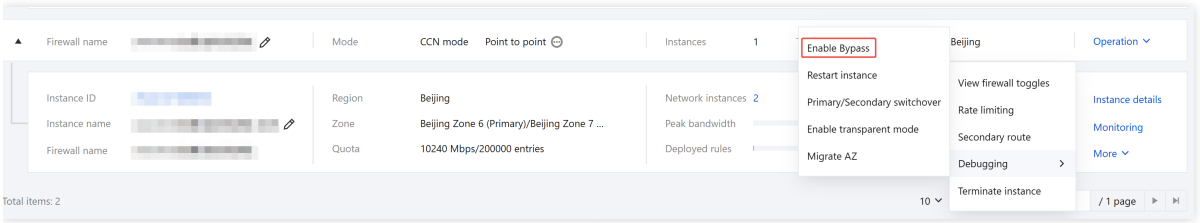


2. Go to the firewall toggle page and click Disable All Toggles for the Current Firewall.
3. On the firewall toggle page, choose More > Change Associated Instance to go to the edit inter-VPC firewall page.
4. In the pop-up window, select the VPC you want to reconnect to and click OK.

ByPass Mode

The firewall can be manually switched to ByPass mode. In ByPass mode, all traffic under the current instance will bypass the firewall, and all firewall configurations will become invalid. It is recommended to use this mode for debugging purposes.
On the firewall instance page, choose More > Enable ByPass (or Disable ByPass).

> Note:
>
> You need to manually disable ByPass Mode after the debugging is completed.



Managing Firewall Toggles

On the Firewall Toggles page, you can control traffic between VPCs through Inter-VPC Firewall toggles. You do not need to adjust the firewall settings when there is an asset change, as CFW can automatically synchronize assets in a short time.

> Note:
>
> Enabling or disabling firewall toggles involves switching networks and routes, which can cause short-term network jitter and momentary disconnection. Therefore, you need to perform this operation at an appropriate time.

Types of Firewall Toggles

There are 4 types of firewall toggles: single-point communication mode, multi-point communication mode, full interconnection mode, and custom route.

- Single-Point Mode: One firewall toggle corresponds to one pair of interconnected VPCs, and one pair of interconnected VPCs corresponds to one peering connection instance.
- Multi-Point Mode: One firewall toggle corresponds to one VPC, and all traffic entering or exiting this VPC is managed by this toggle. Mutual access between one pair of VPCs is independently controlled by their respective toggles.
- Full Interconnection Mode: One firewall toggle corresponds to all interconnected routes of a firewall. When you enable

the toggle, the firewall will control the traffic of all connected VPCs.

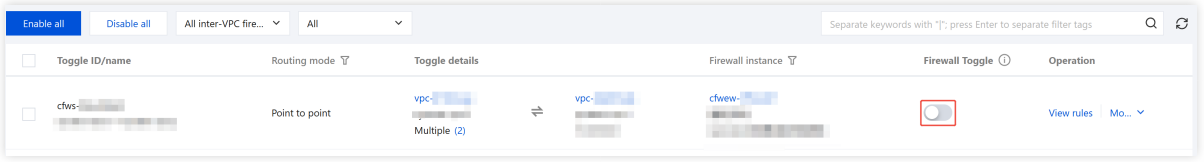
- Custom Route: This toggle is meaningless. If you select this mode, only the connected VPCs are displayed.
  If you change the peering connection instance, the firewall toggle will correspondingly synchronize the changes. To avoid affecting your business, we will only immediately execute changes for toggles that are in the off status. Please ensure that the firewall toggle is disabled when you modify the configuration of the peering connection between VPCs.

Enabling Firewall Toggles

After the toggle is enabled, the system will automatically modify the routing policies of the relevant route tables, routing all traffic between the local and remote networks corresponding to each firewall toggle to Inter-VPC Firewall.

1. On the Inter-VPC Toggles page, click the firewall toggle. You can enable a single firewall, multiple firewalls in batches, or all firewalls at one time.

   - Single: Select the required firewall and click Firewall Toggle. A confirmation dialog box will appear.



   - Batch: After you select the firewall toggles, click Batch Enable in the upper left corner. A confirmation dialog box will appear.
   - All: When no items are selected, click Enable All in the upper left corner. A confirmation dialog box will appear.

2. In the confirmation dialog box, click OK to enable protection.

   > Note:
   >
   > - If the peering connection route between VPCs is not correctly configured, the firewall cannot be enabled.
   > - After the firewall toggle is enabled, do not manually change the routes corresponding to the toggle in the VPC console. Otherwise, the firewall will lose the routes, causing network interruption.

Disabling Firewall Toggles

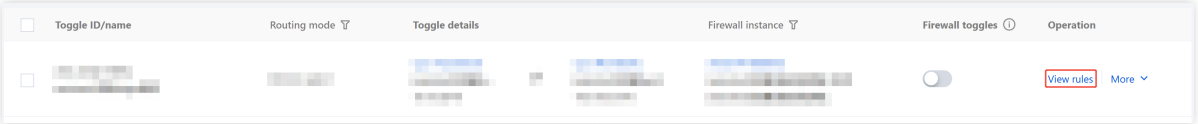
After the toggle is disabled, the system will automatically restore the routing policies of the relevant route tables. All traffic between the local and remote networks corresponding to the firewall will be restored to their original paths and will not pass through Inter-VPC Firewall.

1. On the inter-VPC toggles page, click Firewall Toggle. You can disable a single firewall, multiple firewalls in batches, or all firewalls at a time.
   - Single: Select the required firewall and click the firewall toggle. A confirmation dialog box will appear.
   - Batch: After you select the firewall toggles, click Batch Disable in the upper left corner. A confirmation dialog box will appear.
   - All: When no items are selected, click Disable All in the upper left corner. A confirmation dialog box will appear.
2. In the confirmation dialog box, click OK to disable protection.

   > Note:
   >
   > After the firewall toggle is disabled, you can switch VPC routes as needed. However, you cannot manually enable the firewall routes. Otherwise, it will cause network interruption and firewall toggle failure.

Viewing Rules

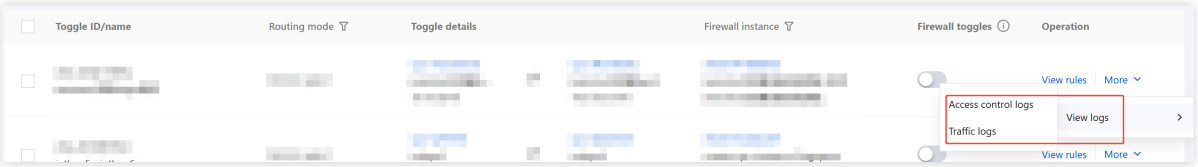
1. On the inter-VPC toggles page, click Firewall Toggle.
2. On the firewall toggle page, select the required firewall toggle and click View Rules.



3. On the inter-VPC rules page, you can view and edit rules. For operation details, see Access Control - Inter-VPC Rules.

Viewing Logs

1. On the inter-VPC toggles page, click Firewall Toggle.
2. On the firewall toggle page, choose More > View Logs. You can then choose to view the traffic access control logs or the traffic logs.

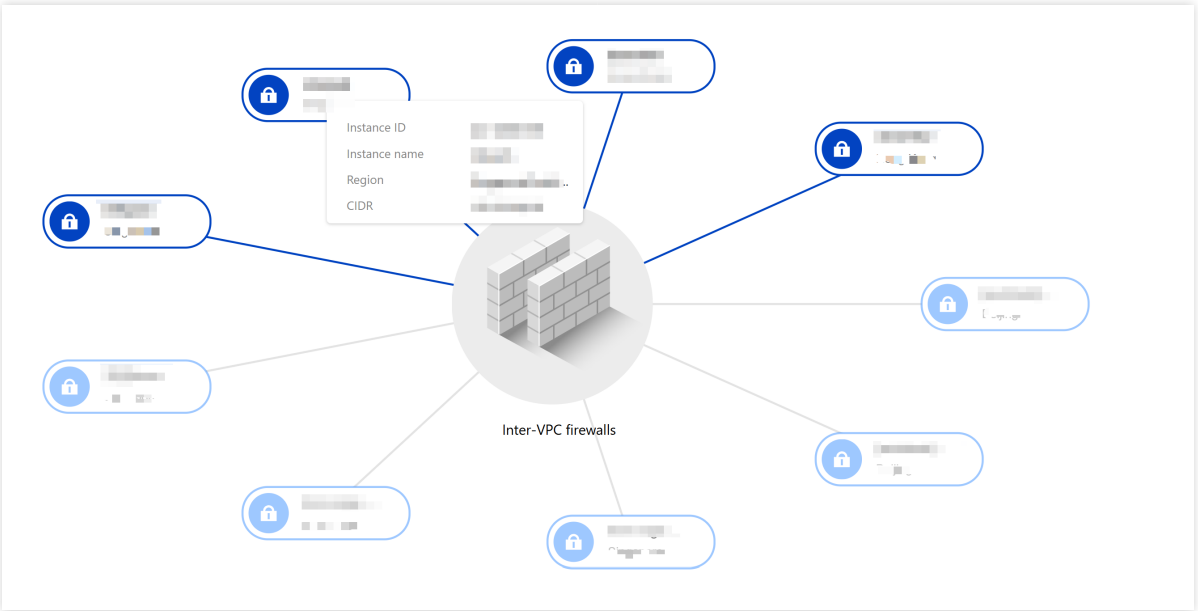


Using Network Topologies

Viewing Network Topologies

CFW provides a visual view to help you quickly organize the access relationship between inter-VPC assets.

1. In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.
2. On the inter-VPC toggles page, click Network Topology to view details of the connected VPC network instances.
3. On the Network Topology page, hover the mouse pointer over a VPC instance to view its details.



4. Click an instance to view its connections to other VPC instances and the firewall toggle status. A dark blue firewall toggle indicator means the toggle is enabled, while a gray one means it is disabled.

5. On the network topology page, click Synchronize Assets to promptly synchronize asset information. Hover the mouse pointer over the engine update area to view the version information.
6. On the network topology page, click the icon in the upper right corner to view the operation guide or refresh the network topology.

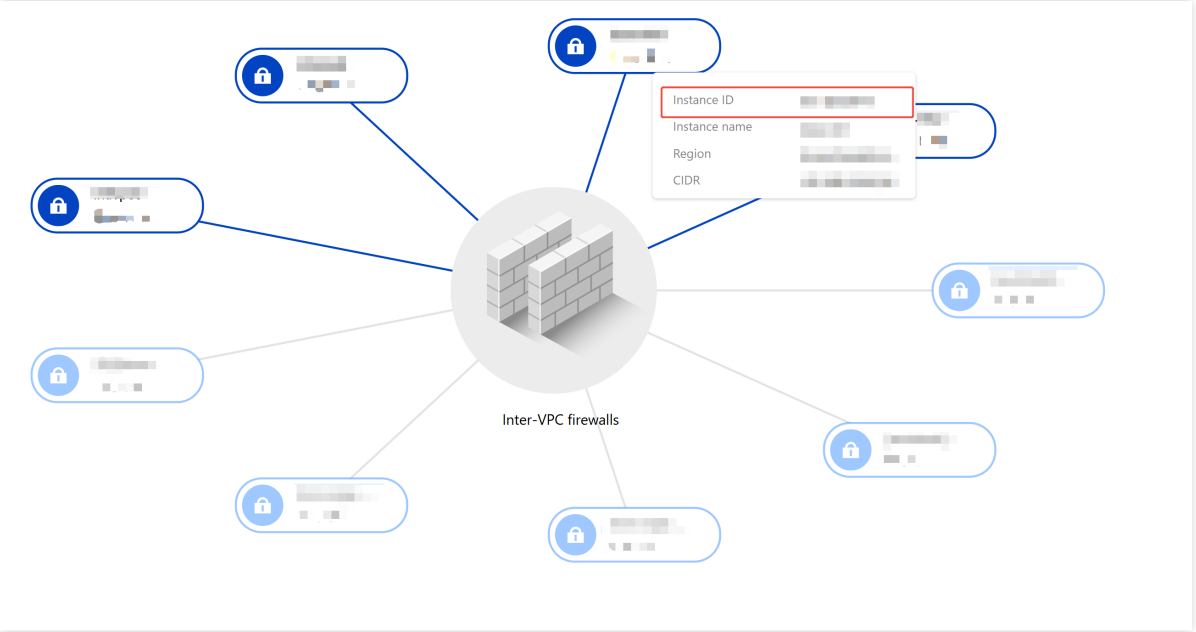
Viewing the Access Relationship Between VPCs Through the VPC View

CFW provides a visual view to help you quickly organize the access relationship between VPCs. In the VPC visual view, each node represents a VPC instance, with Inter-VPC Firewall acting as a centralized device. Each toggle controls different routes. Traffic between VPCs with the toggle enabled will be routed to the firewall for filtering and protection.

1. In the left sidebar, choose Firewall Toggle > Inter-VPC Toggles.

2. On the inter-VPC toggles page, click Network Topology to view details of the connected VPC network instances.

3. Hover the mouse pointer over a VPC node to view its brief information. All VPCs connected to it will also be highlighted. Click VPC ID to go to the VPC details page and view its detailed information.

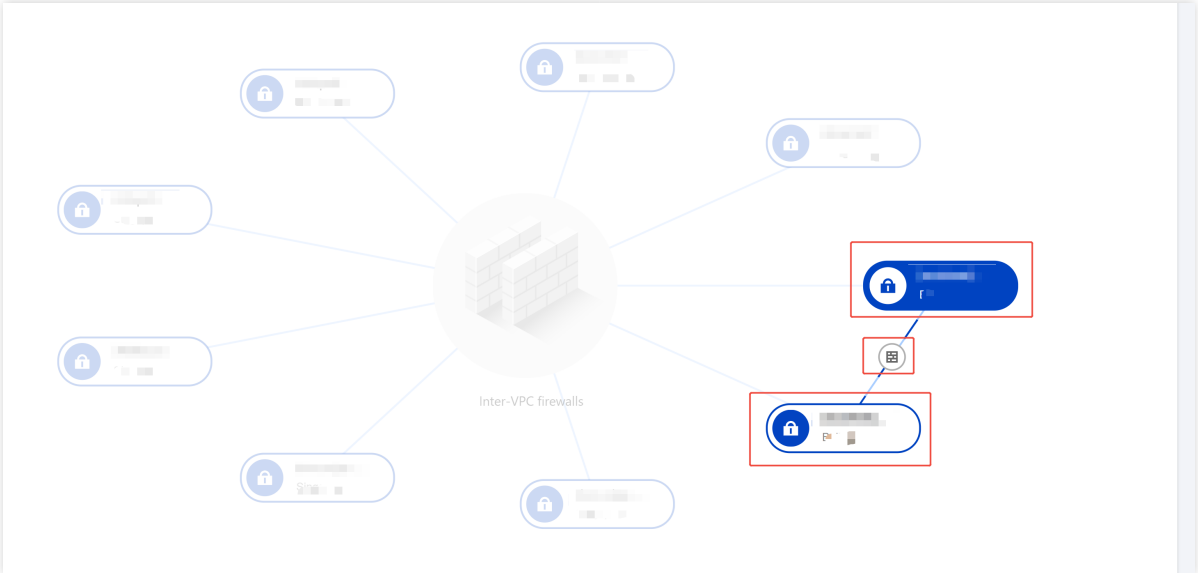


4. Click a VPC node to enter the focus view, which displays the topology centered on the focused VPC.

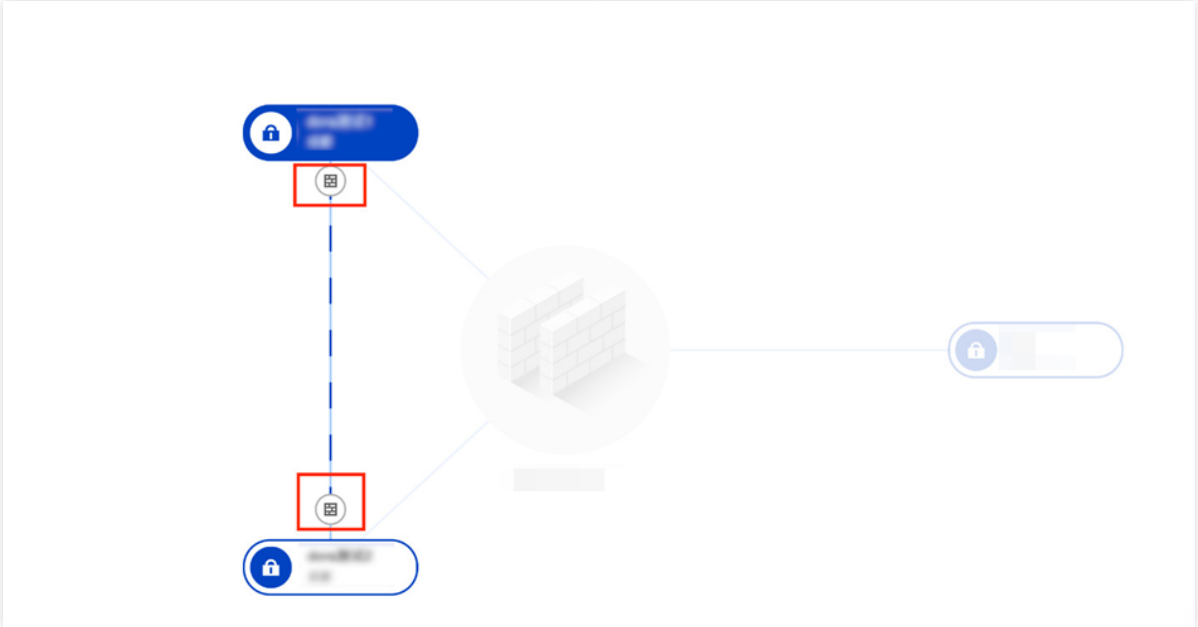
5. Interconnected VPCs are connected via connecting lines. On the lines, you can view or operate the firewall toggle. You can also click Firewall Toggle to directly go to the ACL rule configuration page.

> Note:
>
> - In single-point interconnection mode, only one firewall toggle is displayed for each pair of interconnected VPCs.
> - In multi-point interconnection mode, one toggle is displayed for each VPC.
> - In full interconnection mode, there will be only one toggle for Inter-VPC Firewall.

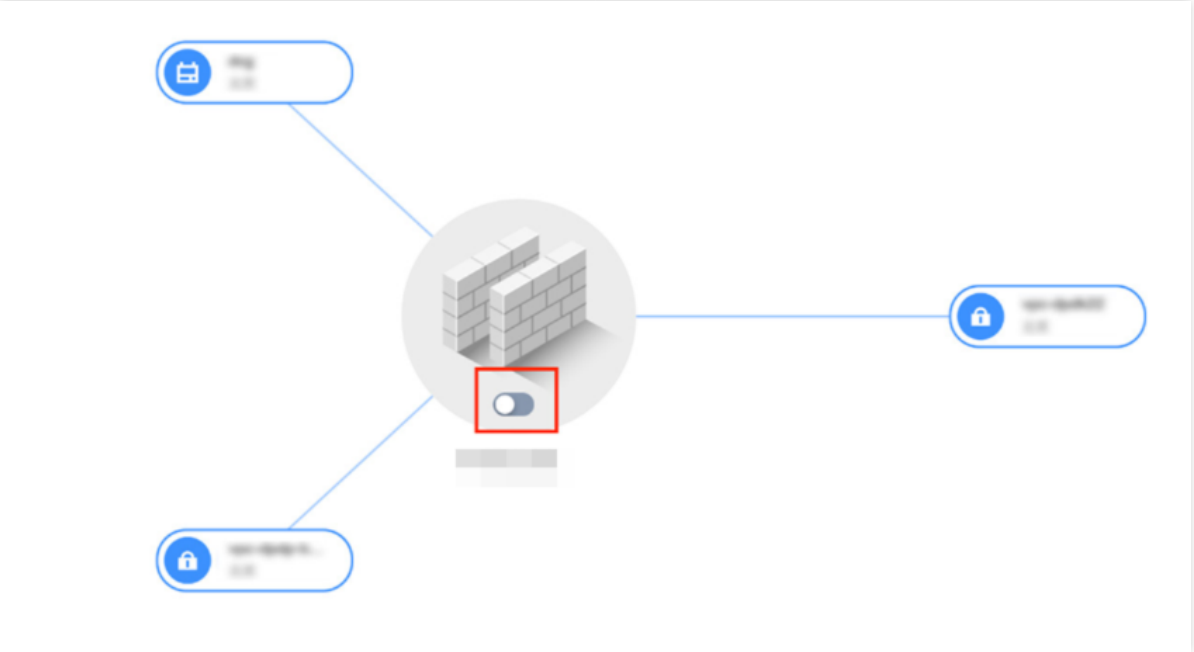
- Single-Point Interconnection Mode:



- Multi-Point Interconnection Mode:



- Full Interconnection Mode:



Description of Abnormal Scenarios for Inter-VPC Firewall

- When you enable or disable the Inter-VPC Firewall toggle, the backend will automatically modify your routing policy, which causes a momentary network disconnection in a very short time. To avoid affecting your workload, you need to operate Inter-VPC Firewall toggles at an appropriate time. If you want to perform batch or frequent operations on toggles, it is recommended to operate at late night when there is little workload traffic.
- The Inter-VPC Firewall toggle is on top of a peering connection between VPCs. If you change or delete the configurations of the peering connection, the firewall toggle will also be automatically changed or deleted. To avoid affecting your workload, CFW can immediately change or delete only the toggles that are disabled.
- If there is no working route between VPCs, the firewall toggle cannot be enabled.
- To configure peering connection routes, see Configuring the Route to Peering Connection.
- If you enable the CFW toggle, do not change the corresponding VPC route table manually in the VPC console. This can invalidate the firewall and disconnect the network, as CFW cannot synchronize the route changes.
- If you disable the CFW toggle, you can change other peering connection routes between VPCs. However, do not enable the route marked with "Firewall". This can cause the firewall toggle to malfunction and disconnect the network.

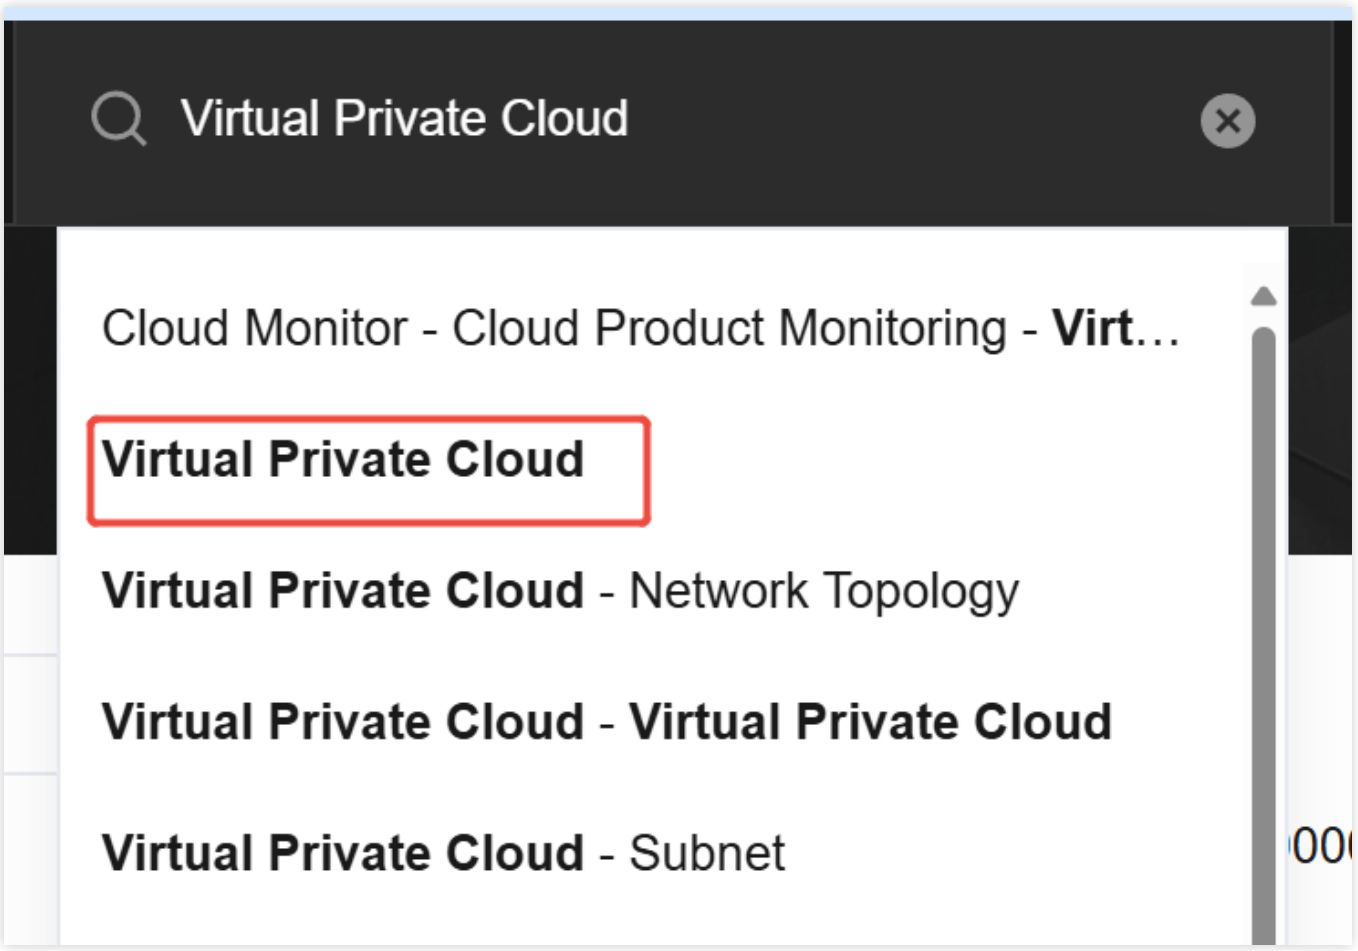
Creating Inter-VPC Firewall Instances

Prerequisites for Inter-VPC Firewall Creation: Establishing a Peering Connection

A VPC is a user-defined and exclusive network space on the cloud that is logically isolated.
A peering connection is a large-bandwidth and high-quality cloud resource interconnection service that supports multi-VPC and heterogeneous network interconnections. The interconnection of a peering connection does not transmit.

Creating a VPC

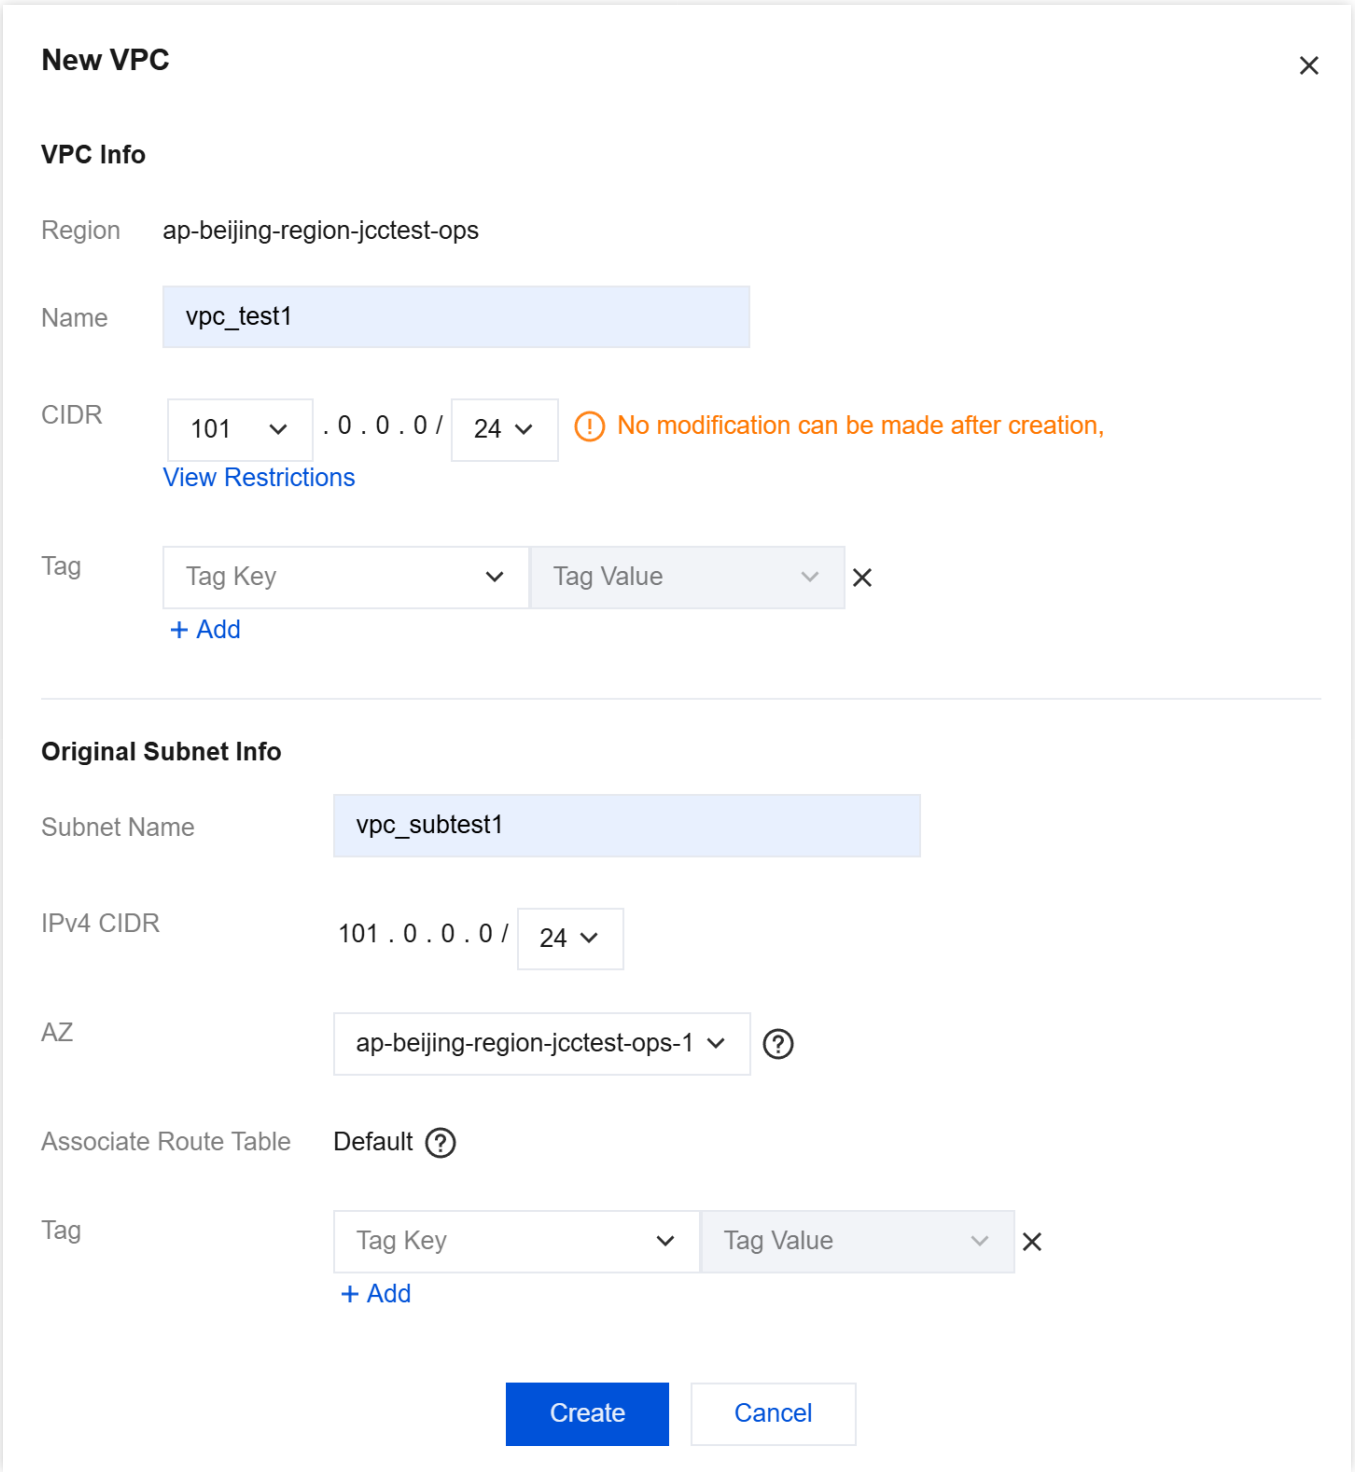
1. Choose Cloud Computing & Networking > VPC.

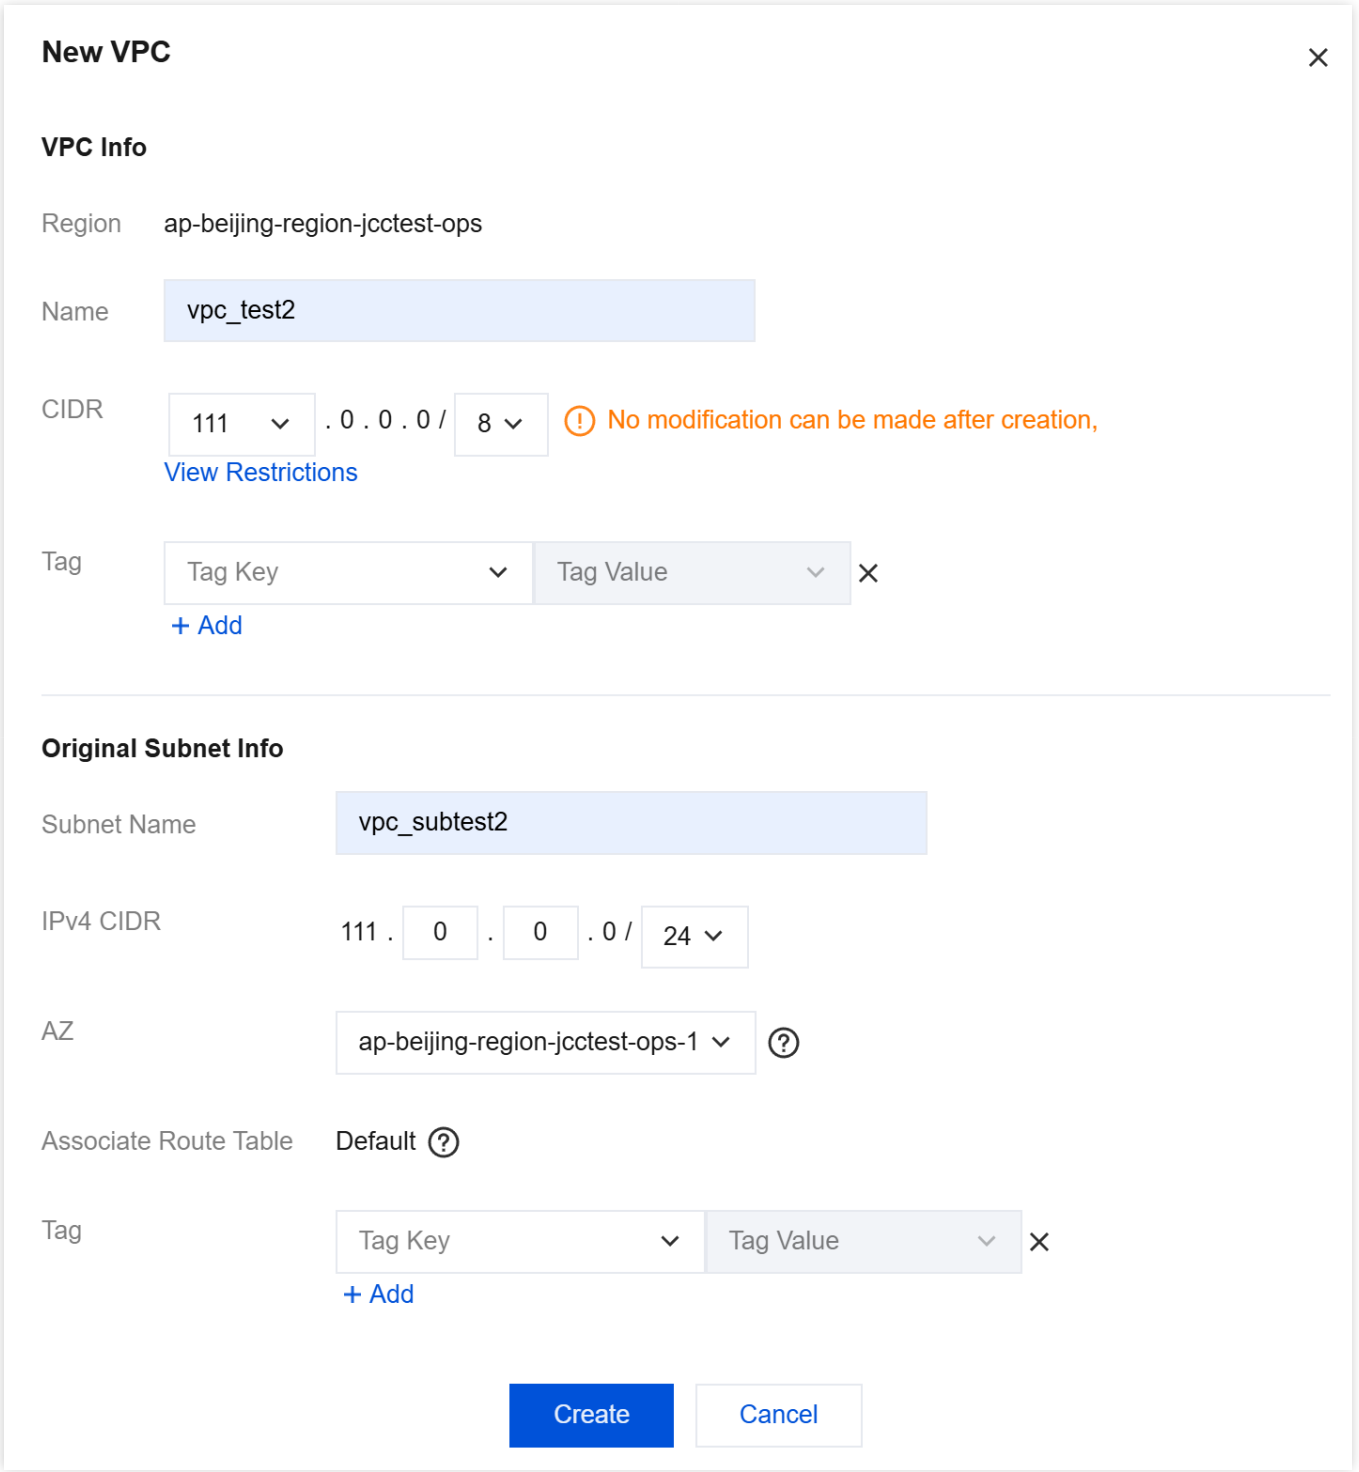


2. Click Create.
3. Enter the VPC name and CIDR block, and the name and CIDR block of the subnet under the VPC, select an

availability zone, and then click Create.

Since a peering connection will be established later, it is recommended that the VPC names here be more semantically readable, such as vpc1 and vpc2.

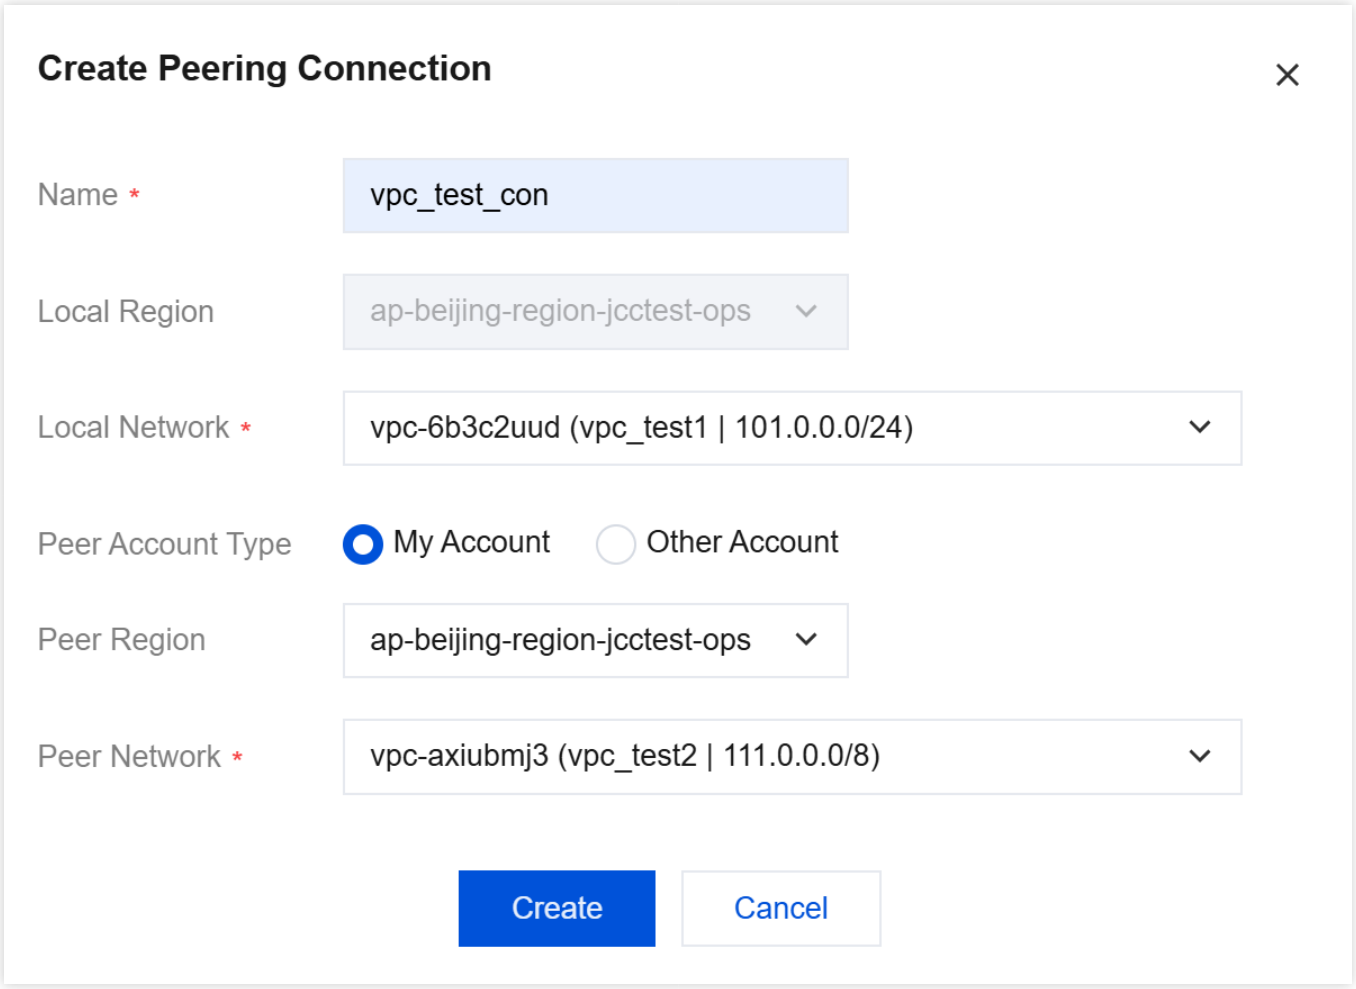




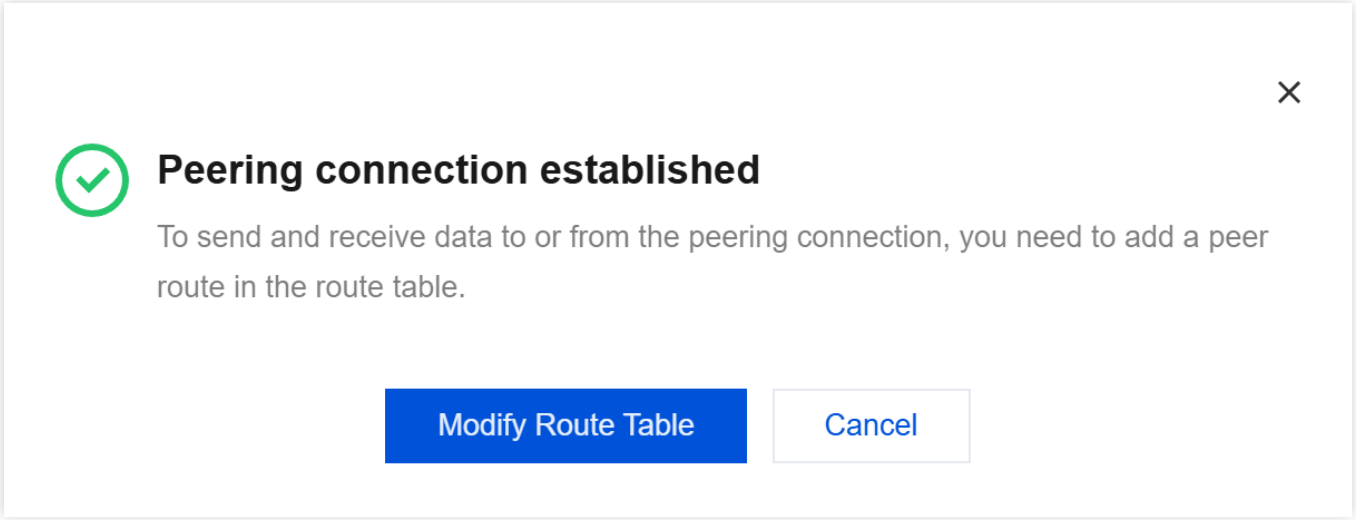
Creating a Peering Connection

1. Choose Peering Connection in the left sidebar.
2. Click Create. In the Create Peering Connection pop-up window, enter the peering connection name and select a region and a VPC.
   Pay attention to the selection of the local network and the peer network here, as it affects the subsequent

use of the UDP protocol in CFW (where the local network corresponds to the client side, and the peer network corresponds to the server side).

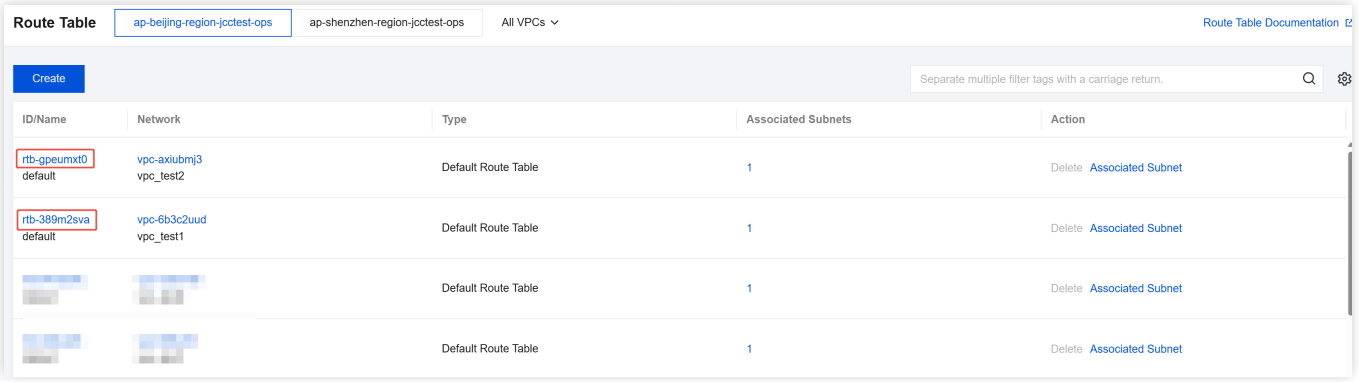


3. After the peering connection is created successfully, a window is displayed, as shown in the following figure. You can go to the Route Table page by clicking Modify Route Table in the pop-up window or through the left sidebar.

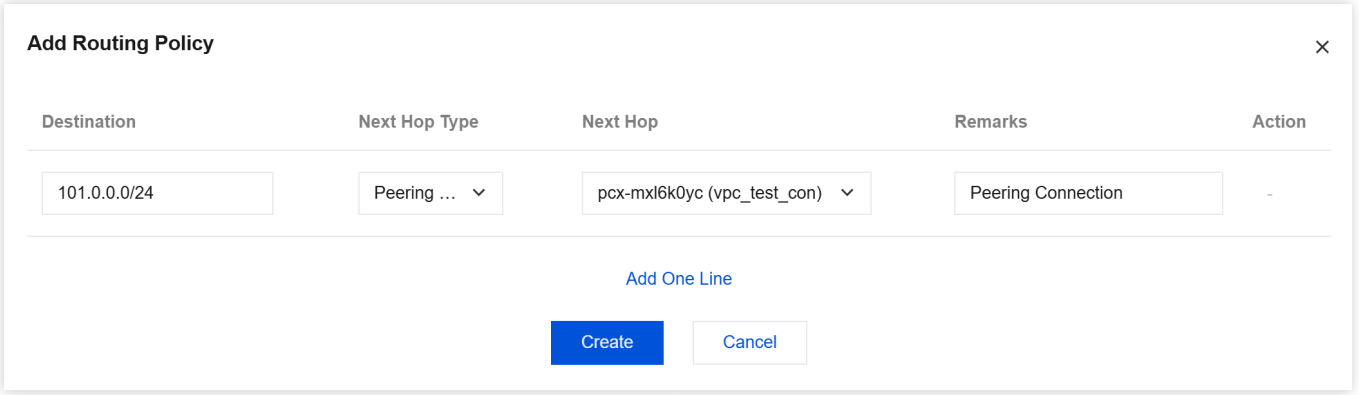


Configuring a Route Table

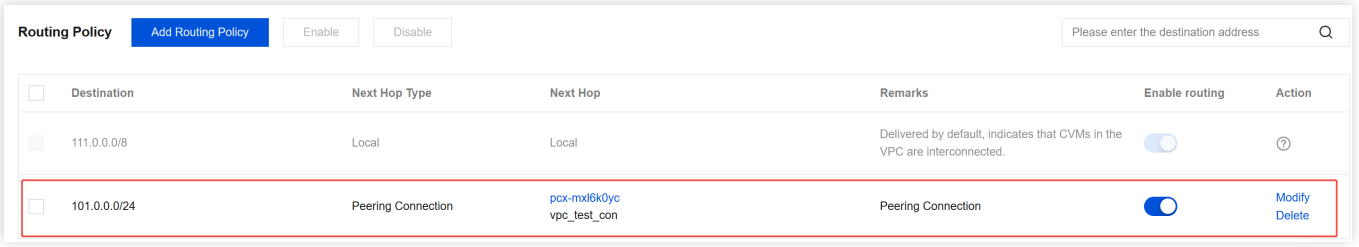
1. On the Route Table page, select the VPCs you created before and configure them. Note that two VPCs need to be configured.



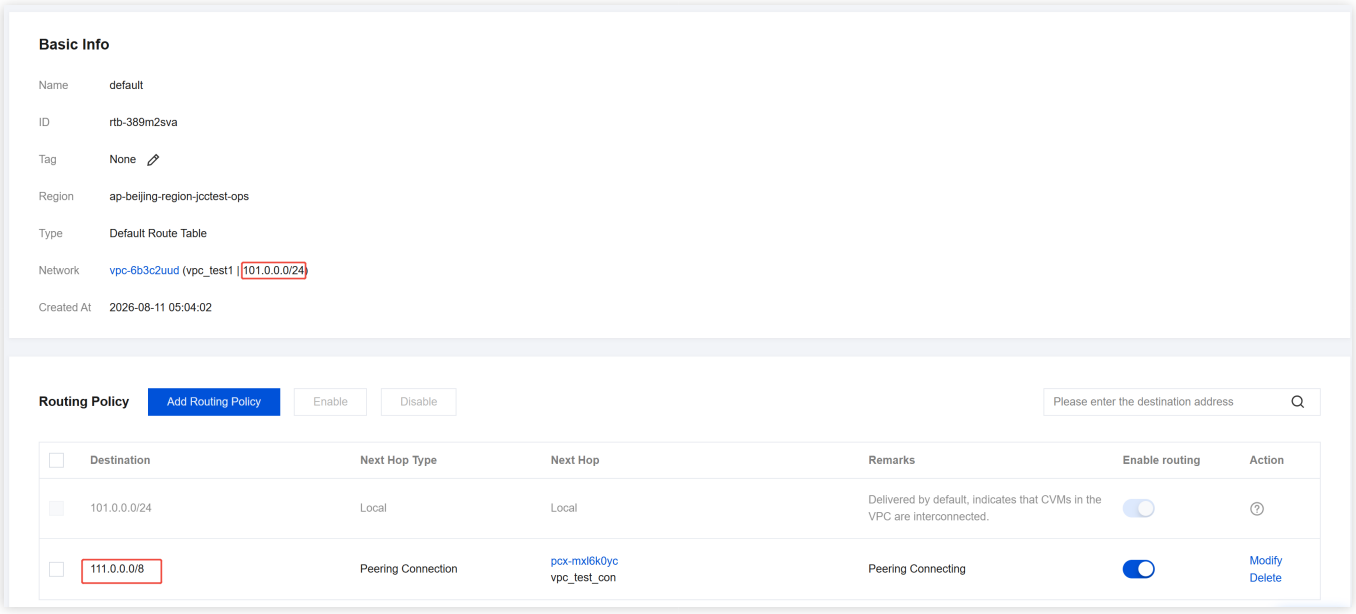
2. Click the blue ID corresponding to a VPC to go to the details page.
3. Click Add Routing Policy and enter the destination (another VPC CIDR), the next hop type (peering connection), and the next hop information. (The previously created peering connection is displayed here by default.) Then, enter the remarks and click Create.



4. After successful creation, the corresponding information is displayed in the list.



5. Perform the preceding steps for another VPC.

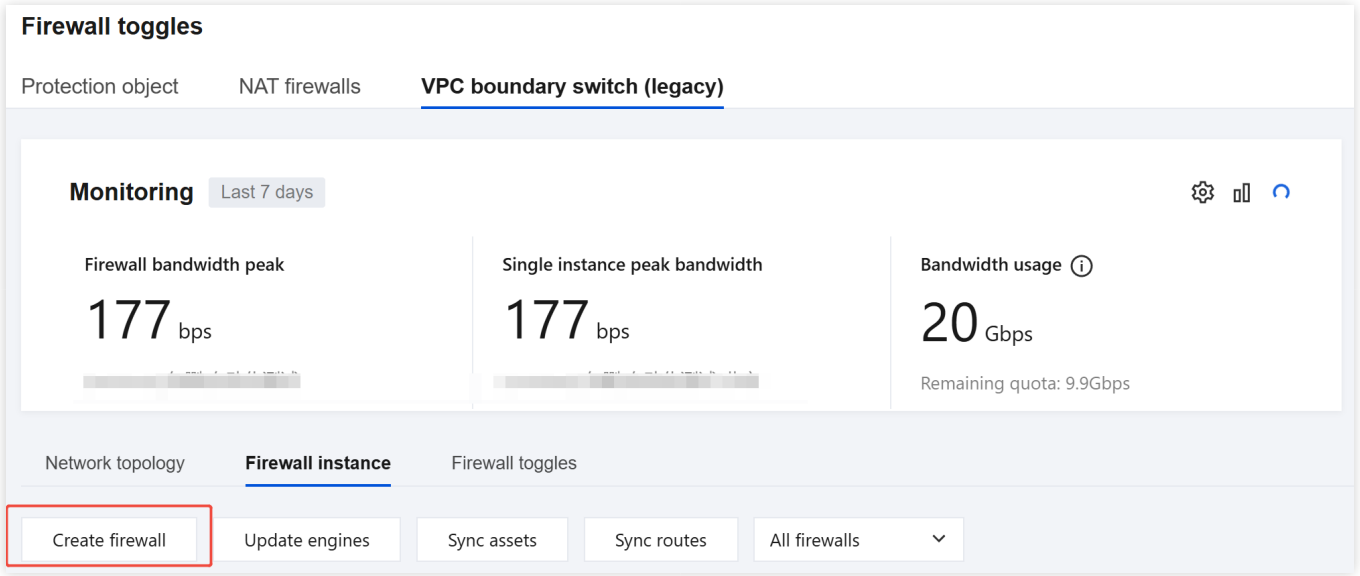


Creating VPC Firewall Instances

> Note:
>
> - In normal mode, you need to select a VPC with a connected peering connection under the local account during access.
> - In administrator mode or multi-account mode, you need to select two VPCs with cross-tenant peering connections during the first access.

1. Create an Inter-VPC Firewall in private network mode. In the left sidebar, choose Firewall Toggle > VPC Boundaries (Primary/Secondary).

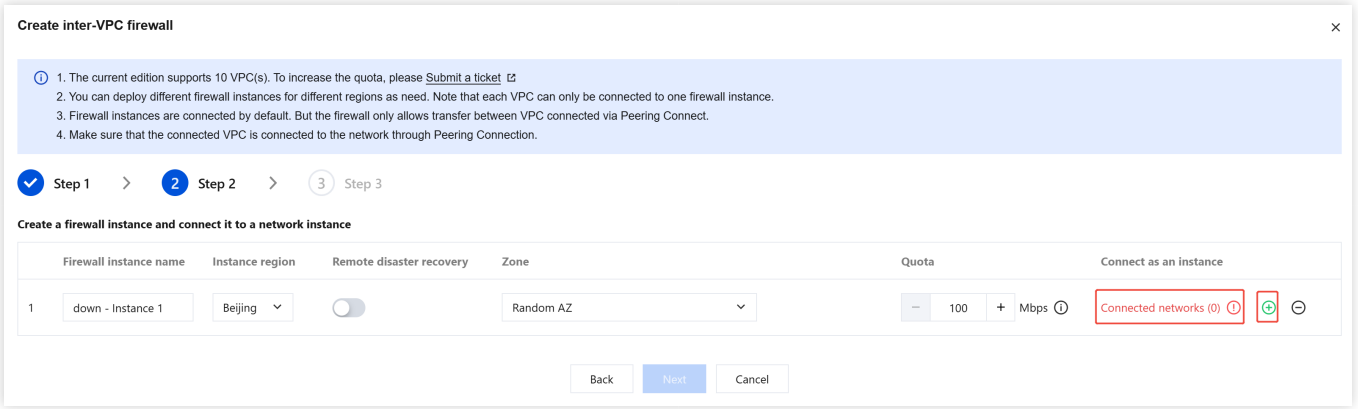
2. On the VPC boundaries (primary/secondary) page, click Firewall Instance to go to the firewall instance page, and then click Create Firewall.

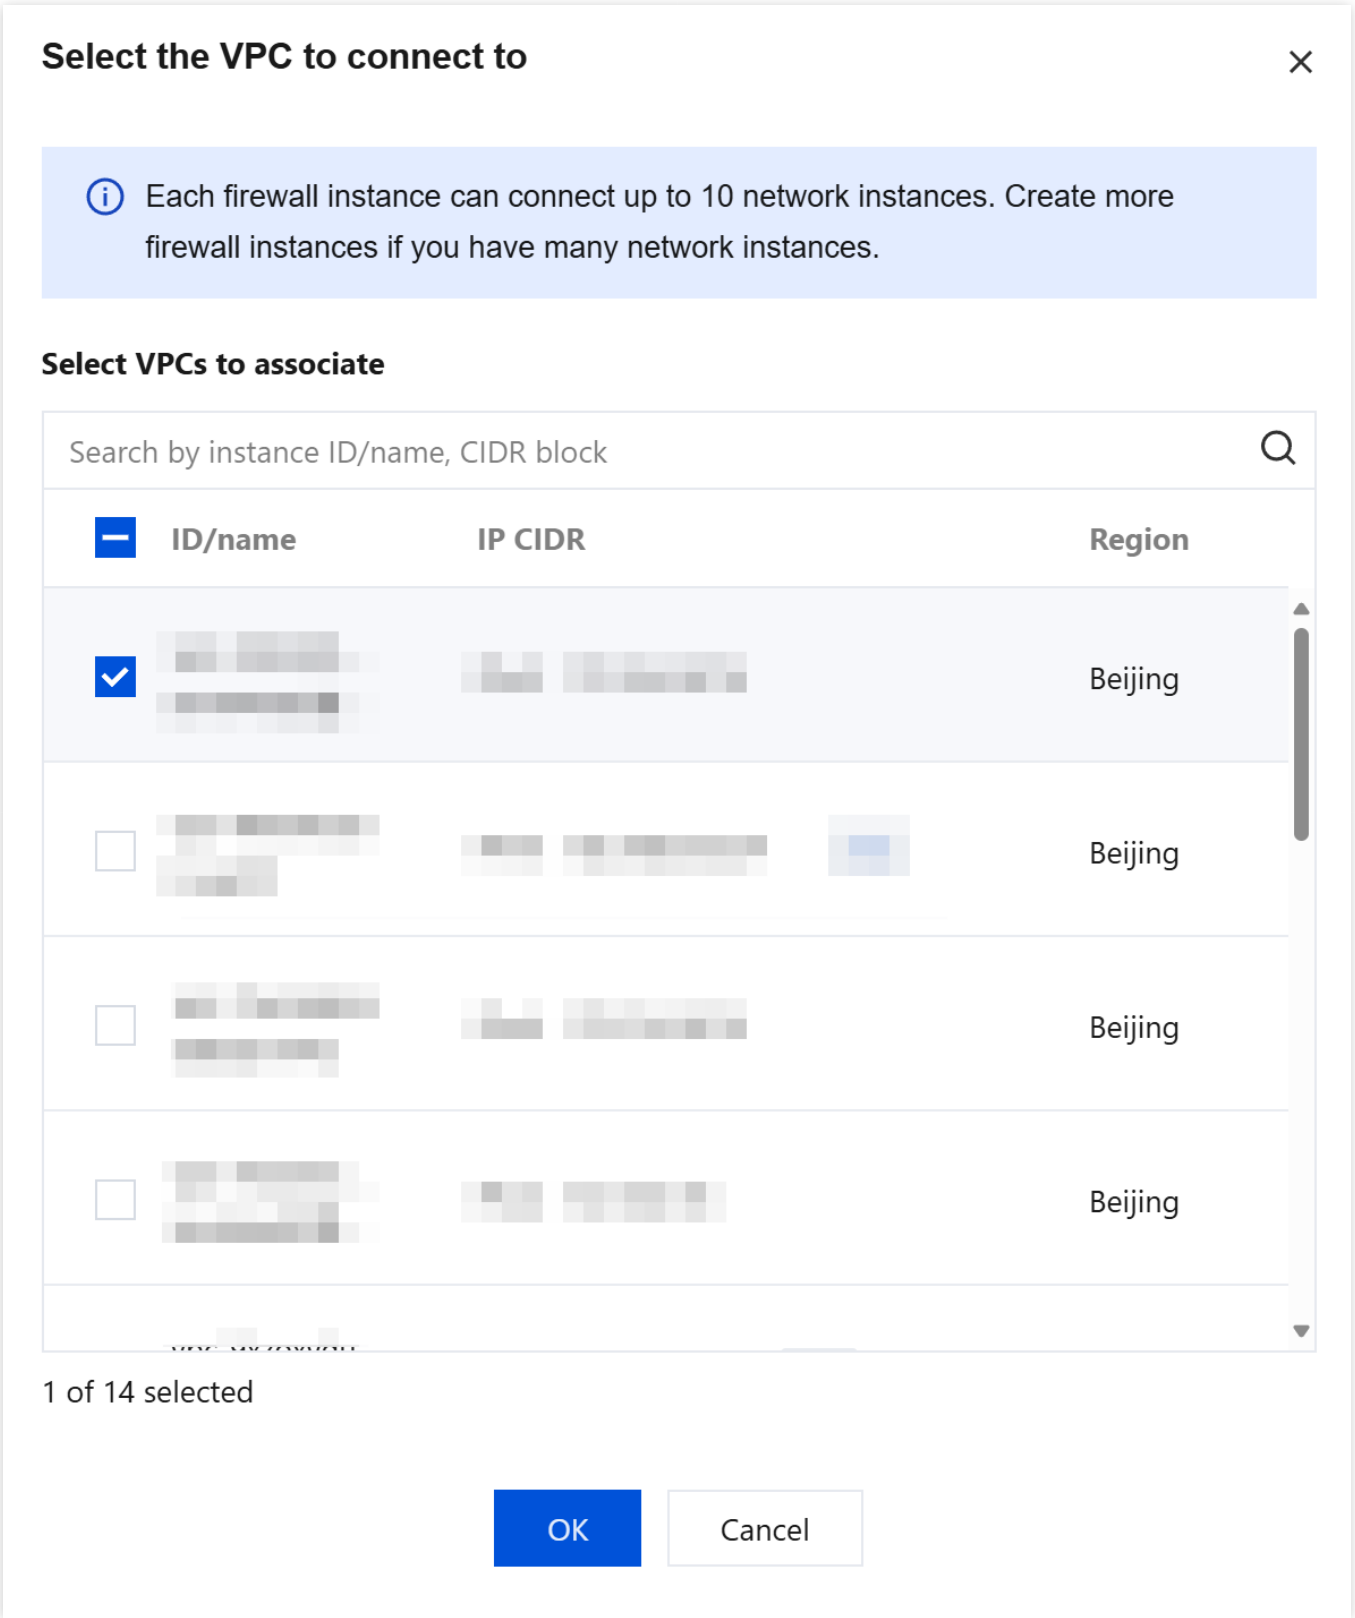


3. In the Create Inter-VPC Firewall pop-up window, enter an instance name, select Private Network Mode, and then click Next.

> Note:
>
> - Instance Name: custom name defined when you create a firewall instance.
> - Set the firewall instance bandwidth specifications and connect to the network, and click Next. If the number of instances does not meet your requirements, you can click ➕ on the right to create multiple firewall instances.
> - Region: Indicates the region to which the VPC with protection enabled belongs.
> - Remote Disaster Recovery: The Inter-VPC Firewall supports remote disaster recovery. After you check it, the primary and secondary engines are automatically deployed in different AZs.
> - Availability Zone: Select a suitable AZ based on your requirements.
> - Instance Bandwidth: An instance supports throughput of 200 Mbps to 10 Gbps (up to 6 Gbps configurable in the console). To configure a bandwidth greater than the configurable limit, submit a ticket or perform a scale-up. If the maximum bandwidth is insufficient, you can create multiple firewall instances to distribute the traffic. However, please note that each firewall instance has its own throughput limit. If there are multiple firewall instances, please ensure that each individual instance stays within its throughput limit.
> - The firewall does not establish basic network connections. Before connecting to the network, make sure that peering connections have been created between VPCs. If no connection is established between VPCs, the connection does not take effect, and no firewall toggle is displayed.
> - Each firewall instance allows up to 10 VPCs in the same region, and you can create multiple firewall instances in the same region. It is recommended to plan the VPCs to be connected by region in advance before creating firewall instances and performing network connection.

4. Click Connected Networks, select VPCs within the region as needed, and then click OK. Note that a VPC can connect to only one firewall instance.



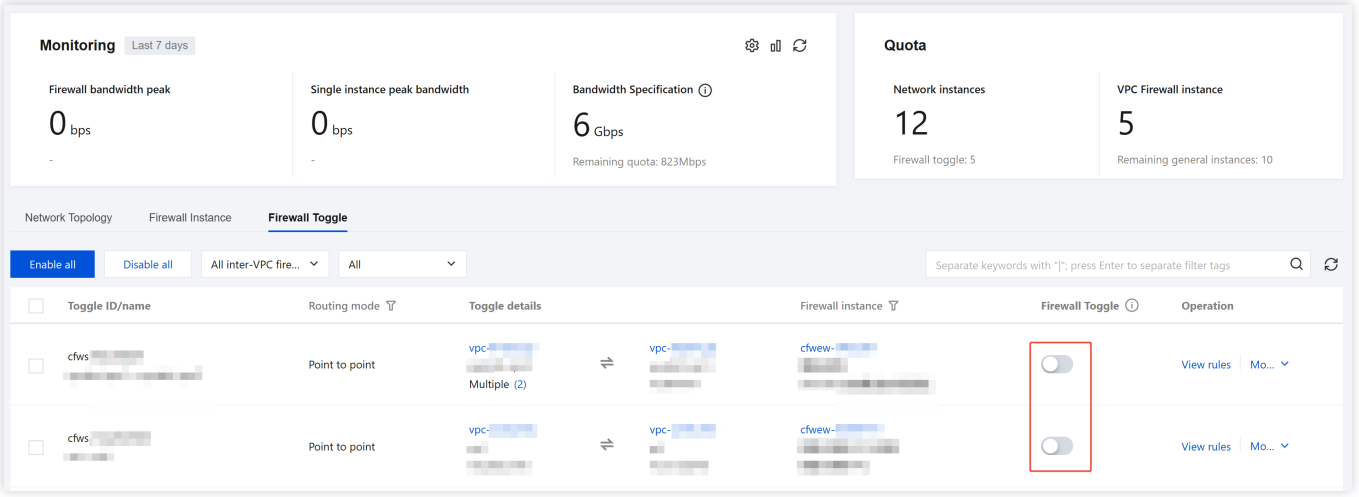


5. Configure How to Create the Routing Subnet, Firewall VPC, and Routing Mode, and click Create after checking these settings.

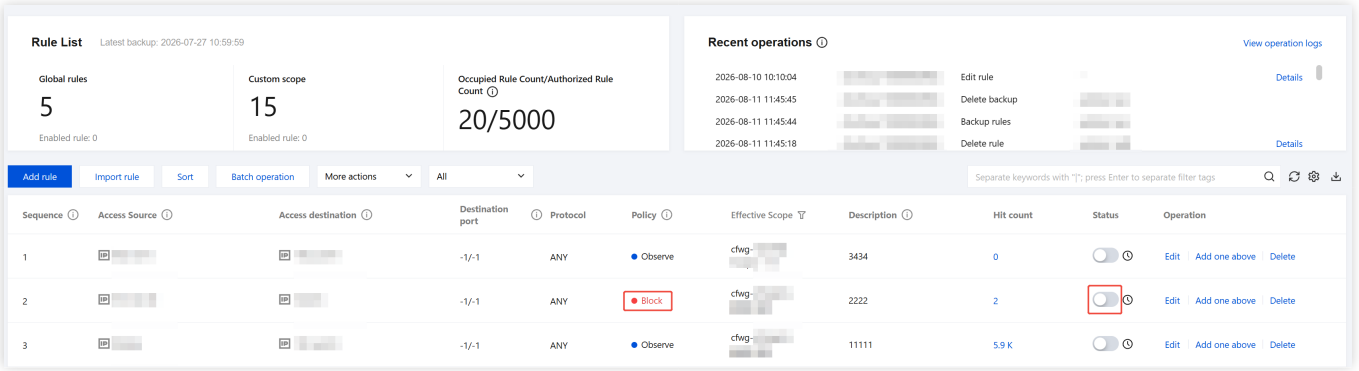
> Note:

> CFW will create a subnet in the VPC you connect to, which will be used to redirect traffic to the firewall. You can choose different methods for creating this subnet. The routing subnet cannot be modified after the firewall is created.
>
> - Preferred Proprietary IP Range: CFW will automatically select an idle subnet IP range in the VPC you chose. If the VPC has no subnet quota, the extended IP range of the selected VPC will be used.
> - Preferred Extended IP Range: CFW will prioritize using the available extended IP range reserved by the VPC. This mode will not consume the subnet quota of the selected VPC. The extended IP range refers to the auxiliary IP range within the VPC.
> - Custom: You can specify a subnet IP range for the firewall. Note that it must be a /24 subnet, and the custom IP range must fall within the CIDR of the current VPC, for example, 192.168.0.0/24.

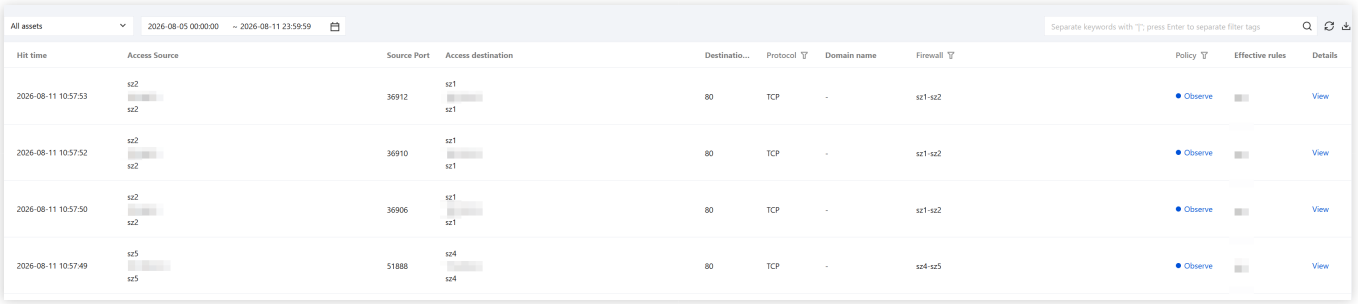
6. Enable the firewall toggle. Once the toggle is enabled, the inter-VPC access traffic will be directed to CFW for filtering and forwarding.



7. Choose Security Policies > Access Control > Add Rules to add access control rules.



8. Choose Cloud Log Service > Access Control Logs to view the generated interception or observation log.


All assets
2026-08-05 00:00:00 ~ 2026-08-11 23:59:59
Separate keywords with "|", press Enter to separate filter tags
Hit time
Access Source
Source Port
Access destination
Destinatio...
Protocol
Domain name
Firewall
Policy
Effective rules
Details
2026-08-11 10:57:53
sz2
sz2
36912
sz1
sz1
80
TCP
-
sz1-sz2
Observe
View
2026-08-11 10:57:52
sz2
sz2
36910
sz1
sz1
80
TCP
-
sz1-sz2
Observe
View
2026-08-11 10:57:50
sz2
sz2
36906
sz1
sz1
80
TCP
-
sz1-sz2
Observe
View
2026-08-11 10:57:49
sz5
sz5
51888
sz4
sz4
80
TCP
-
sz4-sz5
Observe
View

Performance Metrics

Performance Metrics for Inter-VPC Firewall Access Quantities

Bandwidth Performance	Machine Specification	Number of VPC Connections	Intrusion Prevention
0M~600M	4c8g	2	Not supported.
600M~2048M	8c16g	4	Supported.
2048M~4096M	16c32g	6	Supported.
4096~6144M	32c64g	10	Supported.
6144M~10240M	64c128g	10	Supported.

Asset Center

This document describes how to view and manage data and information related to each asset in Asset Center. You can better understand the current status of assets, manage assets, and predict and prevent security incidents by viewing the detailed information on your private network assets, VPCs, and subnets.

Operation Guide

1. Log in to the CFW console. In the left sidebar, click Asset Center to go to the Asset Center page.
2. On the Asset Center page, view server and network asset information and perform related operations. Server assets include both public and private network assets, and network assets include gateways, NICs, VPCs, and subnets.

Server Assets

At the bottom of the Asset Center page, you can click Server Assets to view the details of all server assets under the account, including public and private network assets. You can click the Type button to filter and display information about assets of different types, including asset instance ID/name, IP address, asset type, region, and bandwidth.

Network Assets

At the bottom of the Asset Center page, you can click Network Assets to view detailed information on all network assets under the account. Network assets include four types: gateway, NIC, VPC, and subnet. You can click Type to filter and display assets of different types.

- View the VPC details: In the VPC list, click the ID of the target VPC to jump to the details page of the VPC.
- View additional VPC information: In the Subnet, CVM, or Interconnection VPC column, click the blue field to jump to the corresponding details page.

Traffic Center

The traffic center displays the traffic access between VPCs and contains the Inter-VPC Activities page.

Inter-VPC Activities

You can view the traffic access and protocol distribution between private clouds on the Inter-VPC Activities page.

Operation Steps

1. Log in to the CFW console. In the left sidebar, click Traffic Center and go to the Inter-VPC Activities tab.
2. On the Inter-VPC Activities tab, select a private cloud instance from the filtering box at the top and select a time period.
3. Hover the mouse pointer over the folding line graph for traffic statistics to view the bandwidth rate within a fixed period. You can also hover the mouse pointer over the protocol distribution pie chart on the right side to view the distribution percentage of each protocol.
4. At the bottom of the page, view the specific IP address of Inter-VPC Firewall, and the access destination, access source, access count, and other information on the private cloud. You can also enter the desired access source, access destination, or destination port number in the input box for precise searching.

Access Control
Inter-VPC Rules

An access control list (ACL) is provided for inter-VPC rules. When a rule is created, configure the rule effective scope to provide access control for traffic interconnection between different VPC boundaries. This document describes how to set inter-VPC rules in the CFW console.

Adding Rules

1. Log in to the CFW console. In the left sidebar, choose Access Control > Inter-VPC Rules.

   > Note
   >
   > - Unlike the ACLs of the NAT firewall, the inter-VPC ACLs have no direction limitations.
   > - The local and peer VPCs are equivalent. When configuring rules, you can determine the VPC based on the CIDR block where the access source and destination are located to differentiate the direction.

2. On the Inter-VPC Rules page, click Add Rule to go to the Add Rule page.

3. On the Add Rule page, specify information such as the access source IP address, access destination IP address, destination port number, protocol, policy, and effective scope, and click OK to complete the rule configuration.
   Field description:

- Priority: execution priorities of access control rules, which are independent of those of the rules in the list of the firewall switch. The rule with the highest priority is matched first. If a rule is hit, rules with lower priorities will not be matched. When you modify the priority of a rule, the priorities of all the subsequent rules will increase by 1. When you delete a rule, the priorities of all the subsequent rules will decrease by 1.
- Access Source: It can be an IP address or subnet range in the CIDR block of the local or peer VPC. Note that it cannot be in the same VPC IP range as the access destination. You can also enter 0.0.0.0/0 as the access source for wildcard rules. (It only takes effect within the selected scope.) You can click Details to view details of the access source network instance associated with inter-VPC rules.
- Access Destination: It can be an IP address or subnet range in the CIDR block of the local or peer VPC. Note that it cannot be in the same VPC IP range as the access source. You can also enter 0.0.0.0/0 as the access destination for wildcard rules. (It only takes effect within the selected scope.) You can click Details to view the details of the access destination network instance associated with

inter-VPC rules.

- Destination Port: Single port numbers (such as 80), port ranges (such as 80/80, –1/–1, and 1/65535), and discrete port values separated with commas (such as 80,443,3380/3389) are supported.
- Protocol: Supported protocols include TCP, UDP, ICMP, HTTP, HTTPS, SMTP, SMTPS, DNS, and FTP.
- Policy:
  - Allow: Allow the traffic that hits rules and record the hit count and traffic logs, but not record access control logs.
  - Observe: Allow traffic that hits rules and record the hit count, access control logs, and traffic logs.
  - Block: Block the traffic that hits rules, record the hit count and generate access control logs, and record the complete data packet information of the current request in traffic logs.
- Effective Scope:
  - Global policy: The rule takes effect for all the traffic managed by the inter-VPC firewall.
  - Selecting an inter-VPC firewall switch: The rule takes effect only for inter-VPC traffic managed by the firewall switch.
- Description: the rule description. It can contain a maximum of 50 characters. You can insert special rule settings by using ##. The current version supports #long connection#, which is suitable for the long connection business scenario.
- Inter-VPC Wildcard Rule: supports wildcard IP address ranges. For details, see the Internet Boundary Wildcard Rule in Internet Boundary Rules.

> Note
>
> - The CIDR blocks of the local and peer VPCs for the inter-VPC firewall cannot be the same or overlapped. Otherwise, the firewall cannot be enabled.
> - In inter-VPC access control rules, the access source or access destination can only be an IP address in the CIDR block of the local or peer VPC or a subnet range of it. As the CIDR blocks of the local and peer VPCs cannot be the same, the direction of the traffic controlled by a rule can be distinguished by Access Source or Access Destination.
> - The rule will not take effect if you enter an address other than those in the CIDR block of the local or peer VPC.
> - If Source and Access Destination are set to 0.0.0.0/0, they represent all IP addresses of the VPC.

4. Click Copy in the operation column on the right side to add multiple rules.

> Note
>
> In the Add Inbound Rule pop-up window, one line indicates one rule, and a new rule is added to the end of the list by default. That is, the last rule with the largest priority value is assigned the lowest priority.

Scenario 1: The rule list has been edited locally. Multiple rules need to be added in batches to improve efficiency.

1. Click Copy in the operation column on the right side to add a line of rule below the current position. A maximum of 10 rules can be added at a time.
2. Complete all fields in the form.
3. Before you commit the rules, check whether their execution order meets expectations.
4. Click OK to commit the configured rules.
   Scenario 2: You need to configure multiple rules for an IP address.
5. Edit a line of rule. Edit only the fields that need to be filled in repeatedly.
6. Click Copy in the operation column on the right side to add a line of rule below the current position and automatically copy the content that has been filled in from the previous line. A maximum of 10 rules can be added at a time.
7. Complete other fields in the form and fill in the fields that do not need to be filled in repeatedly.
8. Before you commit the rules, check whether their execution order meets expectations.
9. Click OK to commit the configured rules.
10. After the rules are added, you can view them in the rule list.

Importing Rules

On the Inter-VPC Rules page, click Import Rules to import rules from a local file. You can specify the import location, download the import template, and export existing rules.

Rule Backup and Rollback

For details about rule backup and rollback, see Rule Backup.

> Note
>
> Inter-VPC rules backed up before the revision do not support rollback. If needed, submit a ticket to contact technical support engineers.

Related Information

- If you need to manage inbound and outbound traffic at the Internet boundary in the CFW console, see Internet Boundary Rules.
- If you need to manage inbound and outbound traffic at the NAT boundary in the CFW console, see NAT Firewall Rules.

- If you need to learn about the special use cases of the access control feature of CFW, see Special Use Cases.
- If you encounter issues related to inter-VPC rules, see Inter-VPC Firewall.

NAT Firewall Rules

Access control rules support domain name filtering and traffic filtering based on geographic location requirements. NAT firewall rules provide two access control rule lists: Inbound Rules and Outbound Rules. Inbound Rules: control north-south traffic from outside to inside within the Internet boundary. Outbound Rules: control north-south traffic from inside to outside within the Internet boundary. This section uses Inbound Rules as an example to provide relevant operational instructions. The operations for Outbound Rules are similar.

Operation Guide

1. Log in to the CFW console. In the left sidebar, choose Access Control > NAT Boundary Rule.
2. On the NAT Firewall Rule page, click Inbound Rules to go to the Inbound Rules page.
3. On the Inbound Rules page, you can create access control rules, and view the rule list (including the used quota for inbound and outbound rules and the total quota for the rule list), recent operation records, and access control rule details. The recent operation records display the latest operations performed by users on the rule list.
    - Click Details to view the details of the operation record.
    - Click View Operation Logs to view detailed operation records.
4. Add a rule. On the Inbound Rules page, you can configure rules. The following operations describe how to add an inbound rule.
    i. On the Inbound Rules page, click Add Rule. The Add Inbound Rule pop-up window is displayed.
    ii. In the Add Inbound Rule pop-up window, configure the rule. For Source Type, select IP address, geographic location, cloud vendor, or address template. For Destination Type, select IP address, resource tag, or address template. Check either Execute First or Execute Last based on the rule priority. After you configure Access Source and Access Destination, specify the destination port, set information such as the required protocol and policy, add remarks in the description field, and then click OK to complete the rule configuration.
    Field description:

- Source Type:
    - IP address: any IP address or CIDR block address, such as 10.10.10.10 or 10.10.10.10/24. Multiple entries are allowed and need to be separated by commas.
    - Geographic location: the actual geographic location corresponding to the IP address.
    - Cloud vendor: the service IP address corresponding to each cloud service provider.
    - Address template: a user-defined IP address template.

- Destination Type:
  - IP Address: The access destination of inbound rules applies only to your public IP addresses. If you enter a CIDR block address, the backend will automatically associate all your public IP addresses within the IP address range. Multiple entries are allowed and need to be separated by commas.
  - Resource Tag: the tag of the resource based on which the access destination is selected. The public IP addresses of instances within the tag will match Internet boundary rules.
  - Address Template: Select a user-defined IP address template as the access destination.
- Port Protocol Type: the input mode of the destination port. Bypass firewalls support only manual input of port numbers.
- Rule Priority: the execution priority of a rule. It can be set to Execute First or Execute Last.
- Priority: the priority of access control rules. The execution order of outbound rules and that of inbound rules are independent of each other. The rule with a higher priority is matched first. If a rule is hit, subsequent rules will not be matched. When you modify the priority of a rule, the priorities of all subsequent rules will increase by 1. When you delete a rule, the priorities of all subsequent rules will decrease by 1.
- Access Source: The access source of an inbound rule applies to all public IP addresses. The access source type can be an IP address, geographic location, cloud vendor, or address template.
- Access Destination: The access destination of an inbound rule applies only to all private network assets within the effective scope. It can be an IP address, asset instance, resource tag, address template, or asset group. The access source types and access destination types supported by outbound rules are the exact opposite of those supported by inbound rules.
- Domain Name Rule (Supported for outbound rules): determines whether to allow traffic based on the domain name in the access request. The IP address resolved from the domain name is not verified.
- Domain Name Resolution Rule (Supported for outbound rules): allows the actual IP address requested by the domain name based on the protocol and the region in which the instance resides.
- Destination Port: TCP or UDP rules support a single port number (such as 80), port range with "/" (such as 80/80, -1/-1, and 1/65535), and discrete port values separated with commas (such as 80,443,3380/3389). For ICMP rules, port configuration is not required.
- Protocol: In the current version, the inbound rules support TCP and UDP, and the outbound rules support protocols such as TCP, UDP, ICMP, HTTP, HTTPS, SMTP, SMTPS, DNS, and FTP.
- Policy:
  - Allow: Allow the traffic that hits rules and record the hit count and traffic logs, but not record access control logs.
  - Observe: Allow traffic that hits rules and record the hit count, access control logs, and traffic logs.
  - Block: blocks the traffic that hits rules. The number of hits and access control logs are recorded. Traffic logs record the data packet information of a request in the traffic.

- Effective Scope: the firewall scope to which the current rule applies. It can be set to Global Rules, specific regional instances, or specific firewall instances.
- Description: the rule description. It can contain a maximum of 50 characters. You can insert special rule settings by using ##. The current version supports #long connection#, which is suitable for the long connection business scenario.
- NAT Firewall Wildcard Rule:
- Firewalls provide different rules for wildcard IP addresses, wildcard ports, and wildcard domain names.

Input Field	Input Example	Description
Access source/access destination	0.0.0.0/0	Indicates all IP addresses.
Domain name	*	Indicates all domain names.
Domain name	*.aa.com	Indicates a second-level domain name aa.com starting with *.
Target port	-1/-1	Indicates all ports.
Destination port	1/65535	Indicates all ports.
Target port	80,443,3389	Indicates that it is effective for ports 80, 443, and 3389.
Target port	80/443	Indicates that it is effective for all ports between 80 and 443.
Target port	80/443,3389	Indicates that it is effective for all ports between 80 and 443 as well as port 3389.

5. Click Copy in the operation column on the right side to add multiple rules.

> Note:
>
> In the Add Inbound Rule pop-up window, each row represents a rule. When a rule is added, it is inserted to the end of the list by default. That is, the last rule with the largest priority value is assigned the lowest priority.

- Scenario 1: The rule list has been edited locally. Multiple rules need to be added in batches to improve efficiency.
    i. Click Copy in the operation column on the right side to add a line of rule below the current position. A maximum of 10 rules can be added at a time.
    ii. Complete all fields in the form.

iii. Before you commit the rules, check whether their execution order meets expectations.
iv. Click OK to commit the configured rules.

- Scenario 2: You need to configure multiple rules for an IP address.
  i. Edit a line of rule. Edit only the fields that need to be filled in repeatedly.
  ii. Click Copy in the operation column on the right side to add a line of rule below the current position and automatically copy the content that has been filled in from the previous line. A maximum of 10 rules can be added at a time.
  iii. Complete other fields in the form and fill in the fields that do not need to be filled in repeatedly.
  iv. Before you commit the rules, check whether their execution order meets expectations.
  v. Click OK to commit the configured rules.

6. Import rules: Click Import Rules to import rules from a local file. You can specify the import location and download the import template.

Special Use Cases

This document describes the special use cases of the access control feature of CFW.

Managing the Execution Priorities of Rules

Managing the Execution Priorities of Inter-VPC Rules

Scenario 1: Quick Sorting of Rules in a List

1. Log in to the CFW console. In the left sidebar, choose Access Control > Inter-VPC rules to go to the Inter-VPC rules page.
2. On the Inter-VPC rules page, click Sort at the top of the list to enter the edit mode.
3. Move the positions and priorities of rules in batches within the current page, and sort the rules by dragging the icons next to them.
4. Click Save when you are done.
   Sorting instructions:

- If the position of a rule is changed after you release the mouse pointer, it indicates that one sort operation has taken place.
- If the position of a rule is not changed after you release the mouse pointer, it indicates that no sort operation has taken place.
- After a sort operation takes place, the Cancel button becomes active.
- You can click Recover once to return the list to the status after the last sort operation.
- If you click Save, you will see the toast message "Sorted successfully" at the top of the page.
- If you click Cancel, the list will return to the initial status, and all sort operations will not take effect.

Scenario 2: Moving a Rule to a Specified Position by Modifying It

When moving a rule over a large range is required, the quick sorting feature may not efficiently perform this operation. Therefore, you can use the editing feature to quickly move the current rule to a specified position. Unlike quick sorting, the editing feature allows you to modify the execution order of one rule at a time.

1. Log in to the CFW console. In the left sidebar, choose Access Control > Inter-VPC rules to go to the Inter-VPC rules page.
2. On the Inter-VPC rules page, find the rule you want to move in the list and determine the new

position.

3. On the right side of the rule, click Edit to enter edit mode.
4. Modify the execution priority to the desired value.
    - Execution priorities adhere to the continuity and non-repeatability principles. When you move a rule by modifying it, you need to ensure that the minimum value of the execution priority is 1 and its maximum value is the total number of rules in the current list.
5. Click Complete and check whether the priorities of rules in the list are as expected.
   When you modify the execution priority of a rule in the list, the positions of all other rules will be automatically adjusted.

Scenario 3: Inserting a Rule to a Specified Position in an Existing List

CFW supports inserting a rule between any two rules. The execution order of the inserted rule is the current insertion position.
Rules are inserted in a pre-insertion manner. If you need to insert a new rule between rules with execution orders "2" and "3":

1. Log in to the CFW console. In the left sidebar, choose Access Control > NAT firewall rules to go to the NAT firewall rules page.
2. On the NAT firewall rules page, find the rule in position 3 in the list and click Insert on the right side of the rule.
3. A rule editing box will be displayed above the rule in position 3.
4. In the rule editing box, enter the rule fields and click Complete to complete the rule insertion operation. After the insertion operation is completed, the execution priority values of all rules after the inserted rule automatically increase by 1.

Checking Whether Rules Take Effect

- Method 1: Check the hit count of a rule in the access control list. If there are hits, the rule has taken effect. If a rule has zero hits, it does not mean that it is incorrectly configured. The rule may simply have no hits or not take effect for the time being.
- Method 2: Choose Log Audit > Access Control Logs in the left sidebar to view access control logs (rule hit logs). If a rule is included in the logs, it indicates that the rule has taken effect.

Operation Locking

- At any one time, only one user is allowed to execute any one of the following operations on a single

access control list with the same AppID (the firewall ID is used for VPCs): Add Rule, Import Rule, Sort, Edit, and Insert.

- When performing operations on a list, you may see the toast message "The list is being edited by others and is locked. Please try again later." at the top of the page. This means that another user is performing operations on the list.
- Operations are locked for 5 minutes and will be automatically unlocked after that time period.

Intrusion Prevention

This document describes how to identify unknown risks beyond access control rules, perform intrusion prevention rule detection on the north-south traffic, and prevent vulnerabilities in Cloud Virtual Machine (CVM) from being exposed to the Internet through the intrusion prevention feature.

Operation Guide

Enabling Threat Intelligence

1. After threat intelligence is enabled, CFW forwards protection traffic to the threat intelligence detection and analysis engine to identify unknown risks beyond access control rules.

- Log in to the CFW console. In the left sidebar, select Intrusion Prevention to go to the Intrusion Prevention page.
- On the Intrusion Prevention page, click the icon next to Threat Intelligence to enable this feature.
- After threat intelligence is enabled, CFW forwards protection traffic to the threat intelligence detection and analysis engine to identify unknown risks beyond access control rules.
- External malicious access: CFW monitors and identifies external access to cloud assets from malicious IP addresses and threat samples, such as malicious scans, brute force cracking, mining trojans, ransomware attacks, and remote control.
- Proactive outgoing access: CFW monitors and identifies proactive outgoing access from cloud assets to external malicious IP addresses or domain names, and determines potential host compromise risks through the comparative analysis of big data provided by threat intelligence.

Enabling Basic Prevention

1. Log in to the CFW console. In the left sidebar, click Intrusion Prevention to go to the Intrusion Prevention page.

2. On the Intrusion Prevention page, find the Basic Prevention module and click View Rules to pop up the Basic Prevention Rules window.

3. In the Basic Prevention Rules pop-up window, view the IPS rule list. You can click any rule ID to view its corresponding description.

4. Click  to enable or disable the rule for NAT and VPC firewalls.

> Note
>
> If the versions of some firewall engines are too early, the prevention rules that rely on later engine versions will not take effect for earlier versions of firewall instances. However, basic prevention rules may still remain effective.

5. In the Current Action column, select the action after a rule is hit. Rule actions take effect only for the assets whose protection mode is the interception mode.

> Note
>
> Custom intrusion prevention rules only take effect for NAT Firewall and VPC boundary firewalls. If the versions of some firewall engines are too early, the custom intrusion prevention rules will not take effect for the corresponding firewalls.

6. For all IPS rules, provide keyword searching, one-click reset, and batch operations.

- Keyword searching: You can query the corresponding rules based on rule attributes, and then configure the toggle status and handling actions.
- One-click reset: You can restore the toggle status of all rules with one click and change a rule action to the default action to make the corresponding rule take effect immediately.
- Batch operations: You can select multiple rules and enable or disable them or change their handling actions in batches.

7. After viewing the rules, click  in the Basic Prevention module to enable this feature.

> Note:
>
> - After Basic Prevention is disabled, all basic prevention rules no longer take effect.
> - In interception mode, malicious behaviors that hit high-confidence rules are automatically intercepted, and security event alarms are generated when other rules are hit.

Enabling Virtual Patching

1. Log in to the CFW console. In the left sidebar, click Intrusion Prevention to go to the Intrusion Prevention page.
2. On the Intrusion Prevention page, find the Virtual Patching module and click View Rules to pop up

the Virtual Patching Rules window.
3. For specific rule operations, see the preceding basic prevention rules.



4. Click in the Virtual Patching module to enable this feature.

Protection Mode Description

1. Log in to the CFW console. In the left sidebar, click Intrusion Prevention to go to the Intrusion Prevention page.
2. On the Intrusion Prevention page, locate the Protection Mode module to configure the protection mode.
3. Protection modes include the observation mode, interception mode, and strict mode.

> Note
>
> The observation mode is used by default.

- If you select Observation Mode, threat intelligence, basic prevention, and virtual patching are all in detection mode. For malicious access or network attack behaviors detected, only alarms are generated, but connections will not be automatically blocked.
- If you select Interception Mode, high-confidence network attacks or malicious access will be automatically intercepted. Threat intelligence supports the automatic interception of outbound malicious access. Basic prevention supports the automatic interception of high-confidence rule alarms. Virtual patching supports the automatic interception of traffic detected as vulnerability exploits.
- If you select Strict Mode, threat intelligence (except the threat intelligence detection for outbound domain names), basic prevention, and virtual patching are all in global interception mode. For any detected alarms, connections are automatically blocked but false positives may be generated. This mode is suitable for important period guarantee or attack and defense scenarios.

Feature Dynamic Description

1. Log in to the CFW console. In the left sidebar, click Intrusion Prevention to go to the Intrusion Prevention page.
2. On the right side of the Intrusion Prevention page, view the feature dynamics and feature descriptions:

- Feature Dynamic: In the Feature Dynamic module, you can view the features of the intrusion

prevention module.

Managing the Blocklist (IPS Blocklist) and Allowlist Policy

1. Log in to the CFW console. In the left sidebar, click Intrusion Prevention to go to the Intrusion Prevention page.
2. On the Intrusion Prevention page, view the blocklist and allowlist policy.

Blocklist

Viewing the Blocklist

Click Blocklist to go to the Blocklist tab.

Disabling the Blocklist

In case of emergency, click  to disable the Enable Blocklist switch to disable the blocklist.

Managing the Effective Time of the Blocklist

When an IP address in the blocklist expires, the IP address is automatically deleted from the list, and subsequent access traffic of the IP address will not be blocked by the firewall. To prevent IP addresses with security risks from being deleted automatically from the blocklist, click Edit in the Operation column on the right side of the list to modify the expiration time of target IP addresses.
Note: For IP addresses in the blocklist, all traffic passing through CFW in the outbound or inbound direction is blocked and recorded in Intrusion Prevention Logs under Log Audit.

Allowlist Policy

Viewing the Allowlist Policy

Click Allowlist Policy to go to the Allowlist Policy tab.

> Note:
>
> The IPS feature is bypassed for IP addresses in the allowlist policy.

Managing the Allowlist Policy Effective Time

When an IP address in the allowlist policy expires, the IP address is automatically deleted from the list, and subsequent access of the IP address will not bypass the IPS feature of the firewall. To prevent trusted IP addresses from being deleted automatically from the allowlist, click Edit in the Operation column on the right side of the list to modify the end time and date of the target IP addresses.

CLS

This document describes how to view CFW-related logs.

Viewing Access Control Logs

1. Log in to the CFW console. In the left sidebar, choose CLS > Access Control Logs to go to the Access Control Logs page.

2. On the Access Control Logs page, you can view the rule hit logs generated by CFW based on the access control rules you configured on the NAT Firewall and Inter-VPC Firewall pages. On the NAT Firewall page, you can also see two rule hit lists generated based on the inbound and outbound directions, respectively.

3. In the rule hit list, click View in the Operation column on the right side.

4. On the Rule Hit Details page, view the hit details of the current rule.

   > Note:
   >
   > - If the rule is deleted after the log generation time, its status is displayed as deleted.
   > - If the rule is edited after the log generation time, its status is displayed as edited.
   > - If the rule has not been deleted or edited since the log generation time, its status is displayed as newly added.

5. To further speed up access control log search and filtering, click the magnifier icon on the right side of the access source or destination to view all rule hit information between the two IP addresses.

6. You can manually download logs and filter them based on different conditions. A maximum of 60,000 records can be downloaded each time.

Viewing Intrusion Prevention Logs

1. Log in to the CFW console. In the left sidebar, choose CLS > Intrusion Prevention Logs to go to the Intrusion Prevention Logs page.
2. On the Intrusion Prevention Logs page, view the details of inbound, outbound, and inter-VPC security

events in the External Intrusion, Host Compromise, and Horizontal Movement lists, respectively.

Viewing Traffic Logs

1. Log in to the CFW console. In the left sidebar, choose CLS > Traffic Logs to go to the Traffic Logs page.
2. On the Traffic Logs page, view the north-south traffic information generated by the NAT firewall based on outbound and inbound directions, as well as the east-west traffic between VPCs.
3. Query and filter logs by asset instance name. To obtain a comprehensive view of traffic based on assets, click All Assets in the upper left corner of the Traffic Logs page, and select an asset instance name in the search dropdown list to filter logs and query all traffic logs of the asset.
4. To further speed up the retrieval and filtering of logs, click the Magnifier button on the right side of the access source or destination to view all traffic information between the two IP addresses.

Viewing Operation Logs

1. Log in to the CFW console. In the left sidebar, choose Log Audit > Operation Logs to go to the Operation Logs page.
2. On the Operation Logs page, view all operations you performed on the security policy page and the firewall toggles page within the current account, as well as their details.

Tag Name	Tag Description
Firewall Switch Logs	Records the firewall switch status and users' operation details on the instance configuration.
Access Control Logs	Records users' operations for adding, editing, and deleting access control rules.
Intrusion Prevention Logs	Records users' operations in the intrusion prevention module.
Login Logs	Records users' account login information.

Log Analysis

Log analysis feature: You can view the detailed traffic logs stored in CFW based on the login account in the past 6 months. In addition, Log Analysis supports log search and query based on search statements, and provides report generation and statistical analysis services. This document describes how to use the log analysis feature.

Operation Steps

1. Log in to the CFW console. In the left sidebar, click Log Analysis to go to the Log Analysis page.
2. To facilitate log retrieval within a specified time range, drag your mouse pointer on the log quantity bar chart in blue to quickly select a time range, and click a bar chart to zoom in on the logs for searching.

- Restore: Click Restore in the upper right corner of the bar chart. The time selector will show the initial time range (today).
- If you switch between time options or log types or enter new keywords in the search boxes to search again, the bar chart will show the initial time (today) of the current time selector.

3. The log list displays the field details of each traffic log from top to down based on the Display Fields module. When the Display Fields module only contains source, the log list displays the field details of each traffic log in the order the fields appear in Hide Fields.

- Show: Hover your mouse pointer over a hidden field and click Show. This field will appear under Display Fields and its details will be displayed in the log list on the right side.
- Hide: Hover your mouse pointer over a displayed field and click Hide. This field will be deleted from Display Fields and its details will not be displayed in the log list on the right side.

General Settings

This document describes how to set access control rules on the General Settings page.

Operation Guide

1. Log in to the CFW console. In the left sidebar, click General Settings to go to the General Settings page.
2. On the General Settings page, configure related parameters for access control rules.

Configuring Access Control Rules

Each time you add, insert, or import a control rule in the access control rule list, you can set the rule's enabling status, which can be either enabled by default or not enabled by default.

Notification Settings

Subscribing to Message Center

1. Log in to the CFW console, and in the left sidebar, click Notification Settings.
2. On the Notification Settings page, configure the firewall push notifications. If you need to adjust to subscription-based push or enable the Message Center, click Switch in the Message Center Subscription Settings.

> Note:
>
> Once the Message Center is enabled, all alarm objects in Notification Settings will become invalid. The actual configuration will follow the Message Center settings.

Configuring Notification Settings

On the Notification Settings page, you can configure common firewall alarms and notifications.

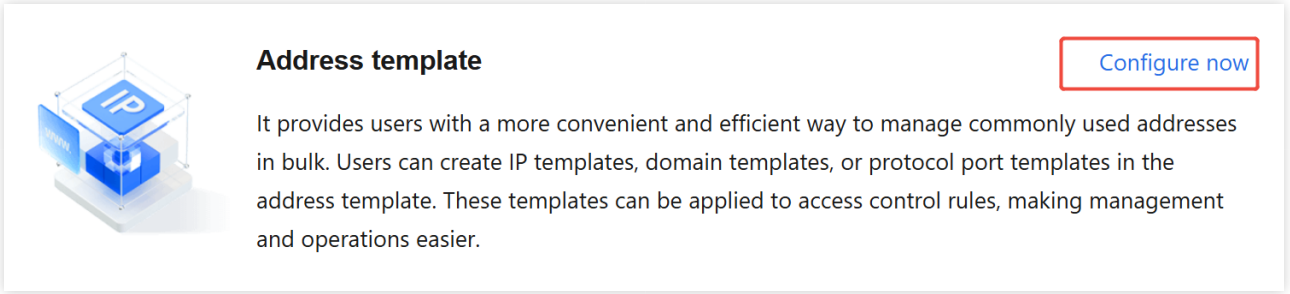
Notification Type	Notification Content	Supported Configuration Item
Bandwidth alarm notification	Three levels of bandwidth alarms are provided. When firewall bandwidth reaches the configured threshold percentage, a notification will be triggered.	It supports tiered alarms by firewall type.
Storage alarm notification	Two levels of storage alarms are provided. When firewall storage reaches the configured threshold percentage, a notification will be triggered.	It supports tiered alarm configuration.
Disaster recovery alarm notification	When your NAT firewall or inter-VPC firewall triggers BYPASS, a notification will be triggered.	-
Automatic task except notification	A notification is sent if anomalies are detected in automatic tasks.	-
System log push notification	When new system logs are generated, notifications are sent according to your configured settings.	It supports triggering notifications by system log event type.

Common Tools
Address Templates

Address templates provide users with a more convenient and efficient way to manage common addresses in batches. Users can create IP address templates, domain name templates, or protocol port templates in the address templates. These templates can be applied to access control rules to facilitate management and Ops. This section describes how to manage addresses in batches by using address templates. Users can also match the existing templates with relevant access control rules for easier management.

Feature Entry

1. Log in to the CFW console. In the left sidebar, click Common Tools.
2. On the Common Tools page, click Go to Configuration under Address Template.

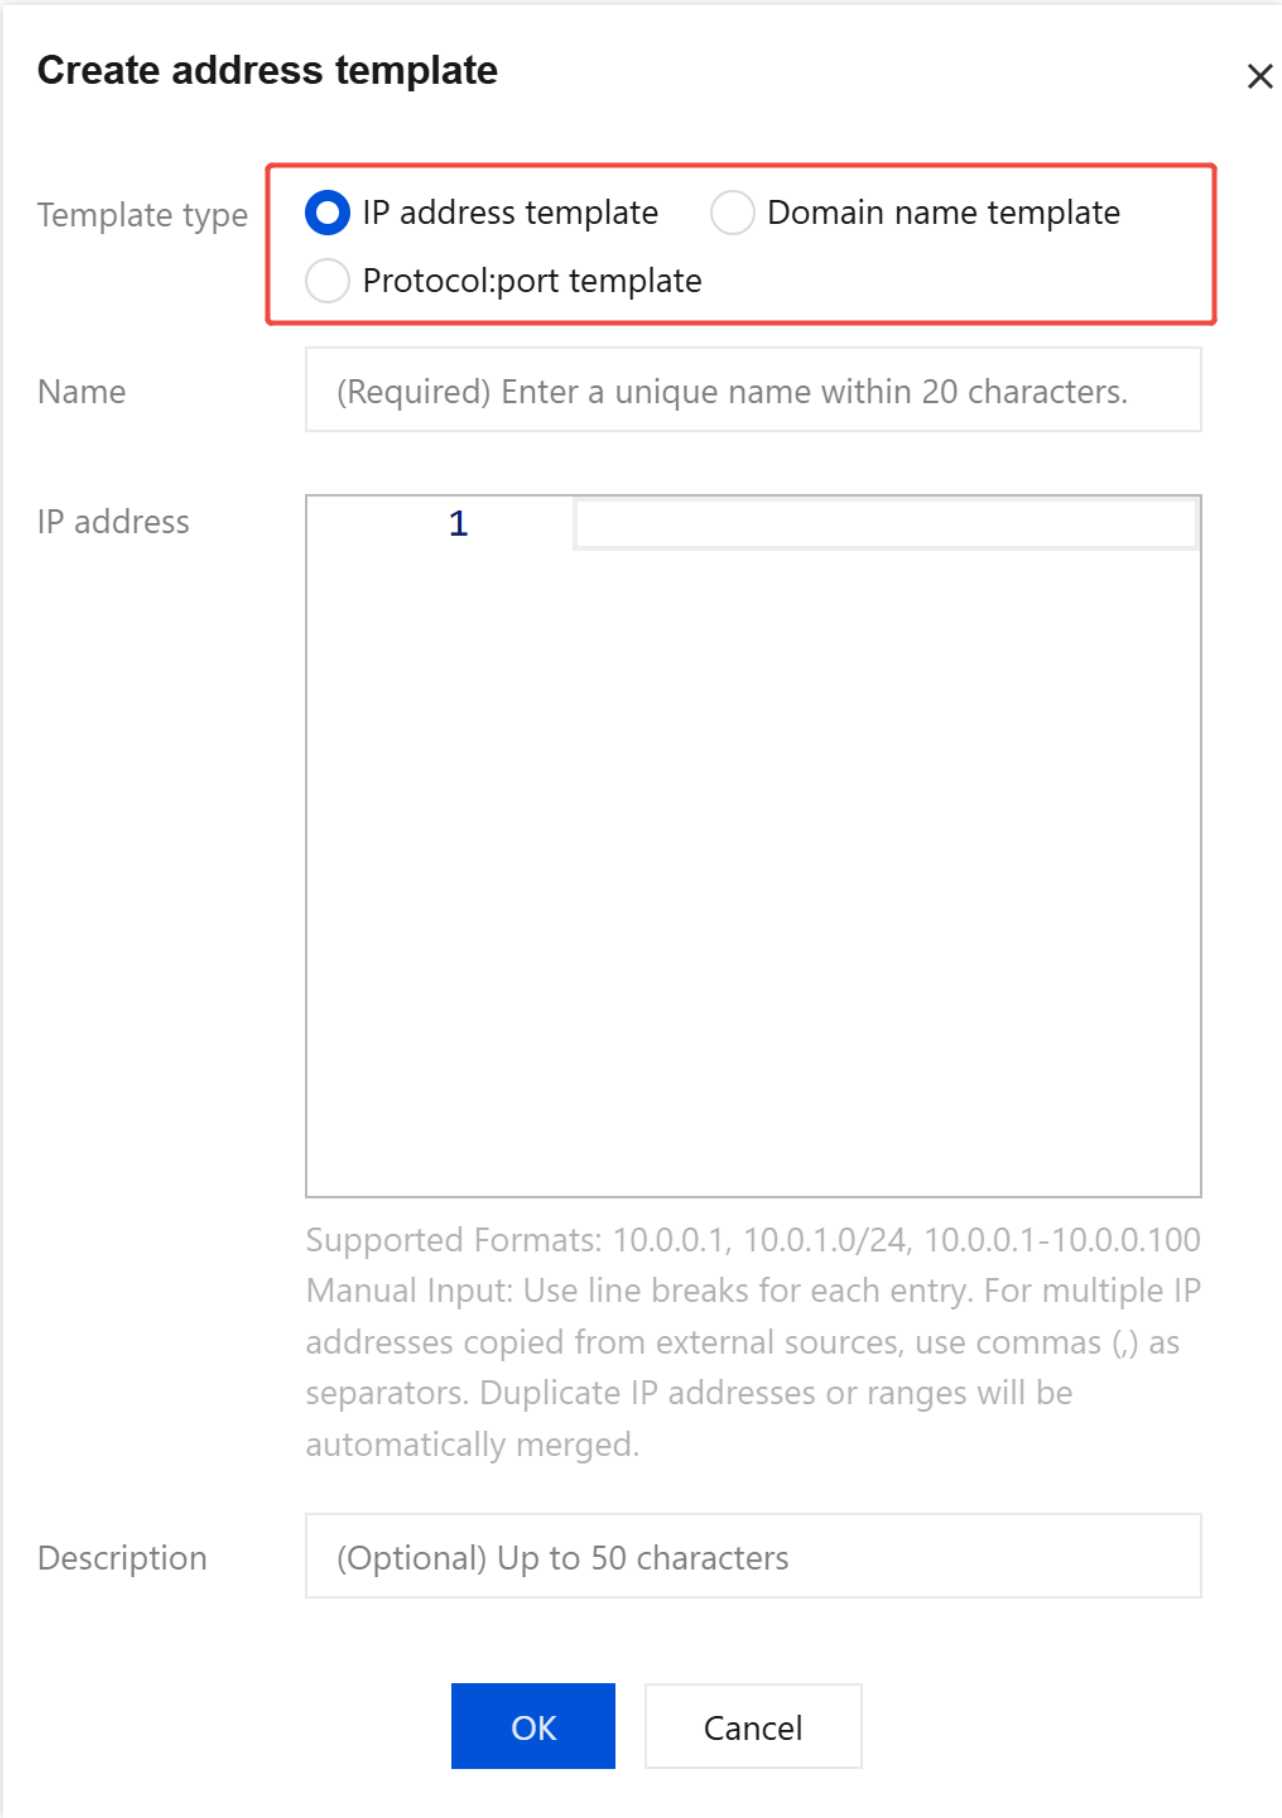


3. On the Address Templates page, users can create and edit templates. When adding rules on the Access Control page, they can directly invoke the templates to reduce the operation frequency.

Creating a Template

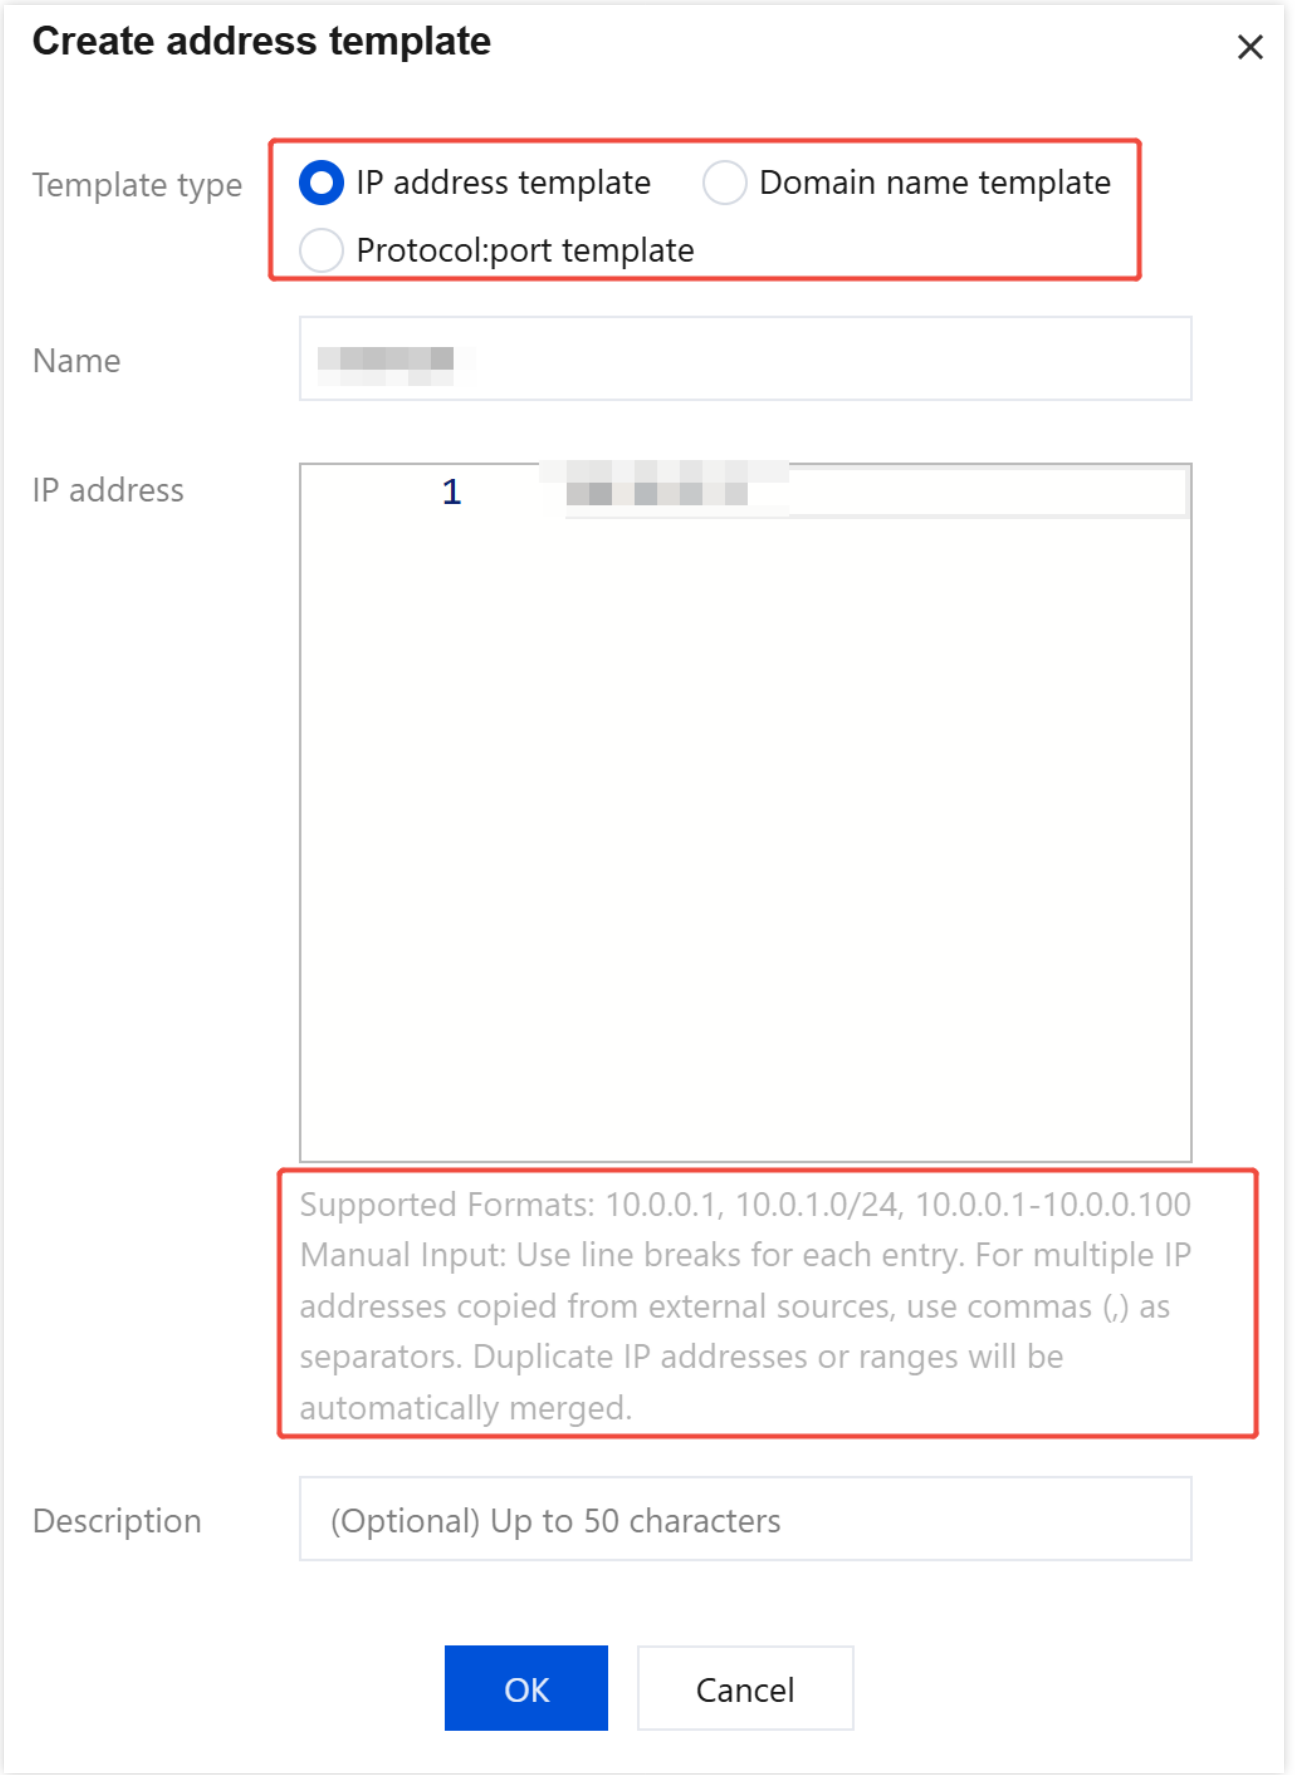
1. On the Address Templates page, click Create Template.

2. In the New Address Template pop-up window, select Template Type.
    - IP Address Template: contains only IP addresses or CIDR blocks. IP address ranges are supported.
    - Domain Template: contains only domain name-type addresses. Wildcard domain names are supported.
    - Protocol Port Template: a combination of protocols and ports, which are categorized by Layer-4 and Layer-7 protocols. The port can be a single port, a combination of ports, or a port range.


Create address template
Template type
IP address template
Domain name template
Protocol:port template
Name
(Required) Enter a unique name within 20 characters.
IP address
1
Supported Formats: 10.0.0.1, 10.0.1.0/24, 10.0.0.1-10.0.0.100
Manual Input: Use line breaks for each entry. For multiple IP addresses copied from external sources, use commas (,) as separators. Duplicate IP addresses or ranges will be automatically merged.
Description
(Optional) Up to 50 characters
OK
Cancel

3. Enter the address template name, IP address, and description, and click OK to create an address template.

> Note
>
> - When you enter a single address, press Enter to wrap to a new line.
> - When you enter multiple addresses at a time, make sure that they are separated by commas before pasting to facilitate system recognition.
> - Duplicate addresses will be automatically merged by the backend.
> - A single address template allows you to add a maximum of 500 addresses at a time. For details about the filling specifications and content, see the related description in the console.


Create address template
Template type
IP address template
Domain name template
Protocol:port template
Name
IP address
1
Supported Formats: 10.0.0.1, 10.0.1.0/24, 10.0.0.1-10.0.0.100
Manual Input: Use line breaks for each entry. For multiple IP addresses copied from external sources, use commas (,) as separators. Duplicate IP addresses or ranges will be automatically merged.
Description
(Optional) Up to 50 characters
OK
Cancel

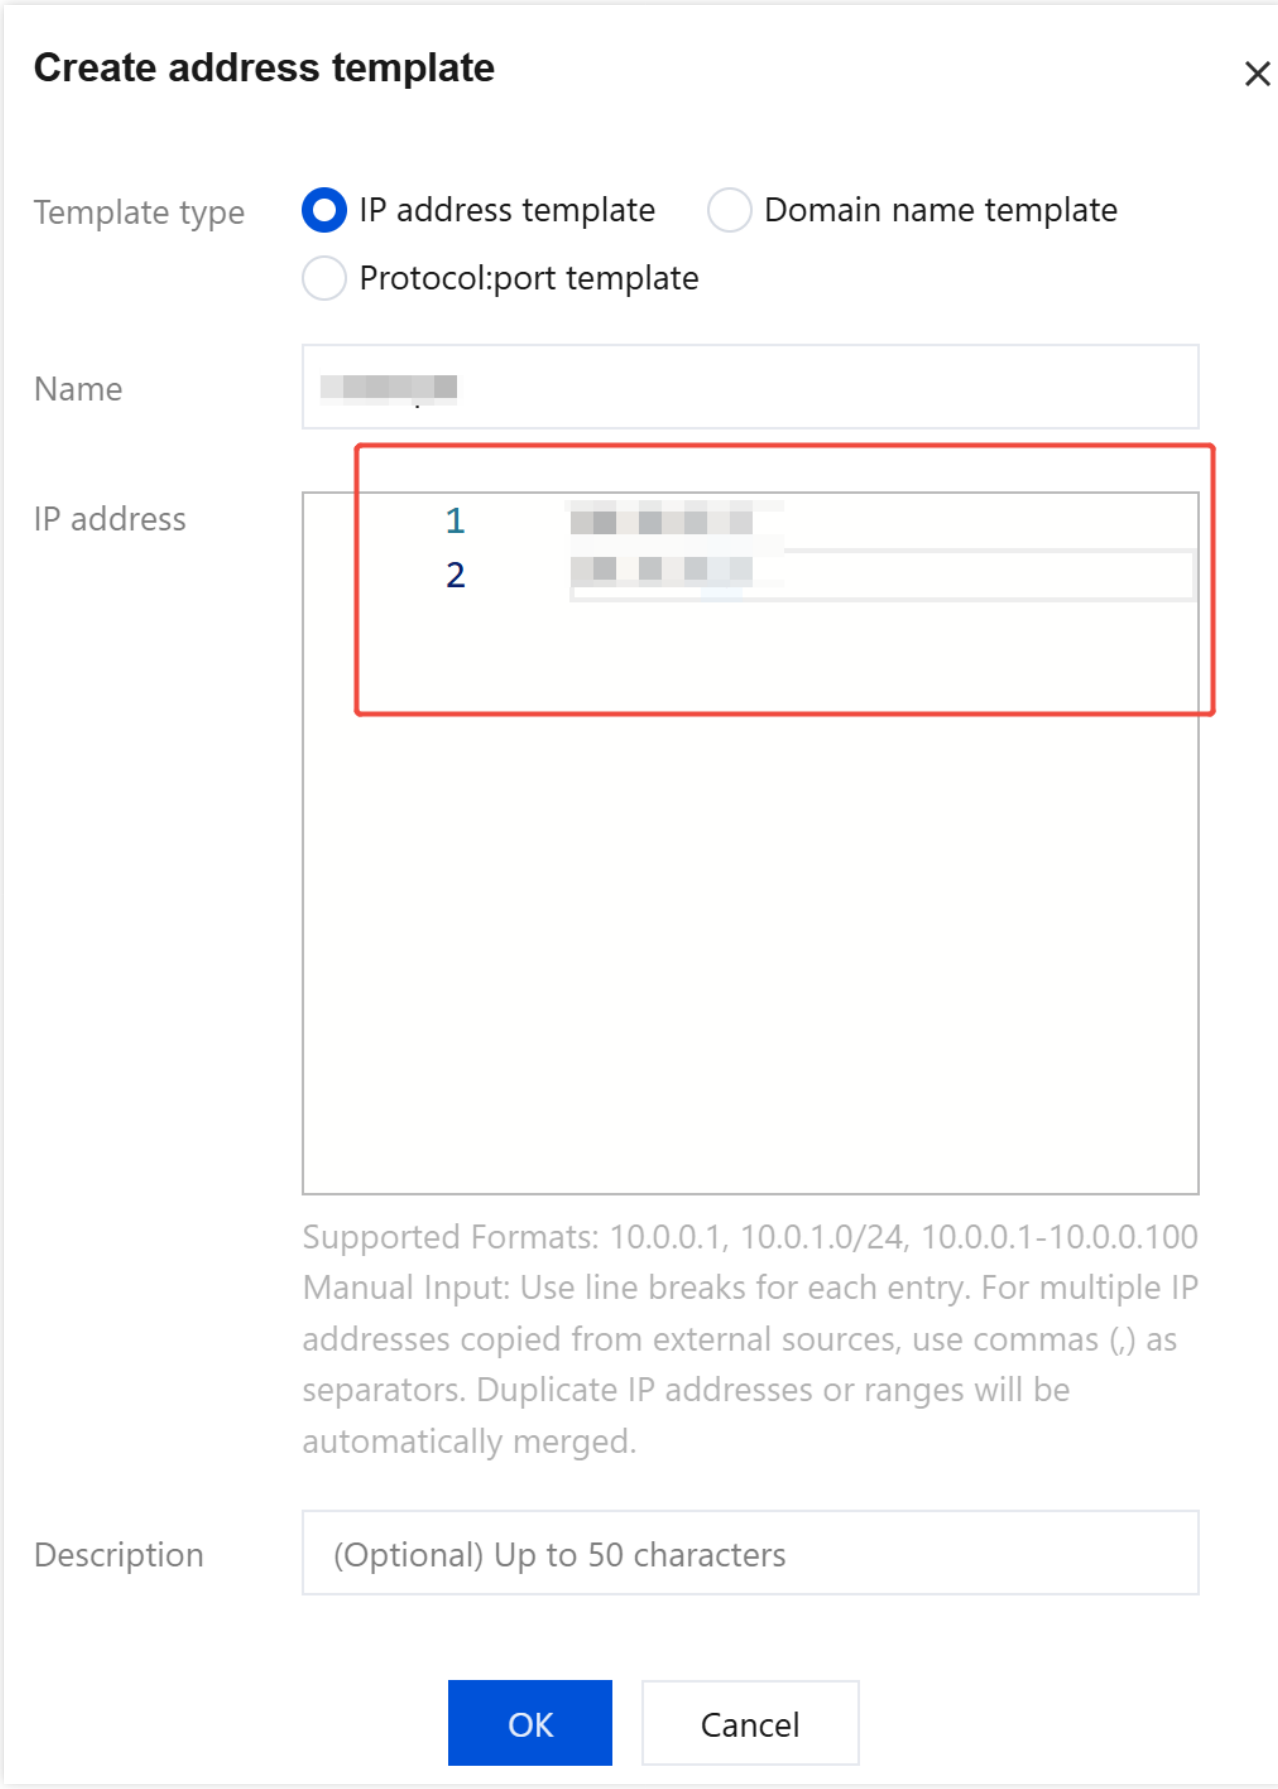
Editing a Template

After creating an address template, users can view it in the address template list on the left. Users can add, delete, and modify addresses of the corresponding category in the existing template list.

1. On the Address Templates page, click Edit in the operation column on the right of the target address template.

2. In the Edit Address Template pop-up window, select the IP address you want to modify or delete, or search for the IP address you want to modify in the search box to perform the operation. Then, click OK to save the modification.

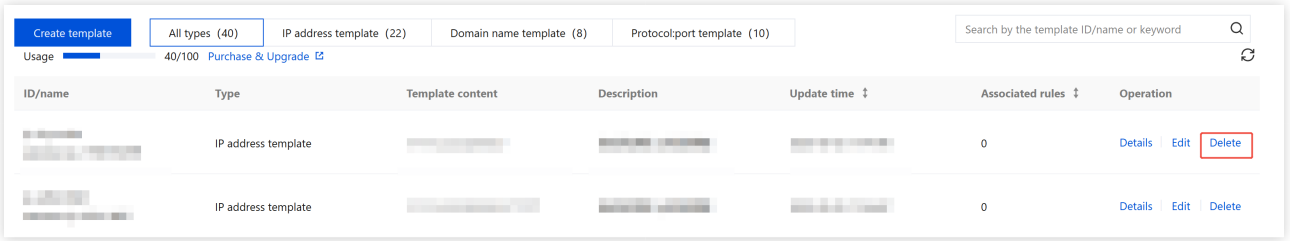
   > Note:
   >
   > Template type cannot be edited. If modification is required, delete the template and create a new template.


Create address template
Template type
IP address template
Domain name template
Protocol:port template
Name
IP address
1
2
Supported Formats: 10.0.0.1, 10.0.1.0/24, 10.0.0.1-10.0.0.100
Manual Input: Use line breaks for each entry. For multiple IP addresses copied from external sources, use commas (,) as separators. Duplicate IP addresses or ranges will be automatically merged.
Description
(Optional) Up to 50 characters
OK
Cancel

Deleting a Template

1. On the Address Templates page, click Delete in the operation column on the right of the target address template.

> Note:
>
> - If associated rules exist, go to the Access Control page to delete the associated rules
> - The deletion operation cannot be undone. It is recommended to back up the current template before deletion.



2. In the Confirm Deletion pop-up window, click OK to delete the rule.

Using an Address Template

On the Access Control page, click Add Rule or locate the rule to be modified and click Edit on the right, then select to use an address template.

> Note:
>
> - Access rules configured with an address template apply to all addresses within the template.
> - The number of address templates is limited, so please plan and use them reasonably.

IP Address Template or Domain Name Template

In access control rules, you can select Address Template in Access Source or Access Destination. The dropdown menu will list supported IP address templates or domain name templates that can be configured.

> Note:
>
> The selection of template types is related to the protocol. Please fill in according to the prompts of the corresponding access control rule.

Protocol Port Template

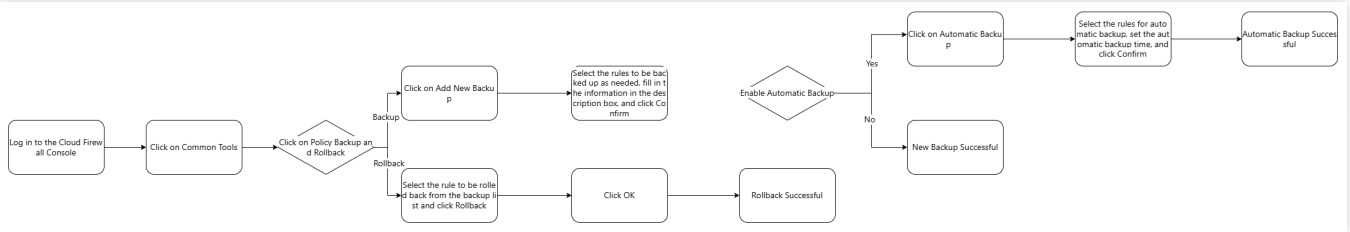
In access control rules, you can select Protocol Port Template in Port Protocol Type. The dropdown menu will list supported protocol port templates that can be configured.

> Note:
>
> The selection of template types is related to the configuration format of the source or destination. Please fill in according to the prompts of the corresponding access control rule.

Policy Backup and Rollback

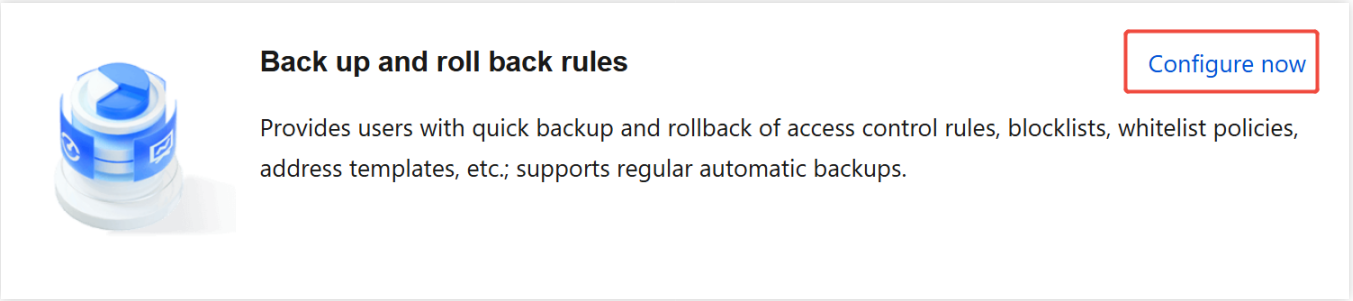
To handle possible misoperations, policy change requirements, or emergency fault recovery scenarios, policy backup and rollback features are provided to ensure your network protection policies remain stable and reliable. With policy backup and rollback, you can quickly undo recent modifications and restore the policy configuration of NAT firewalls or inter-VPC firewalls to a previously saved status.

Flowchart

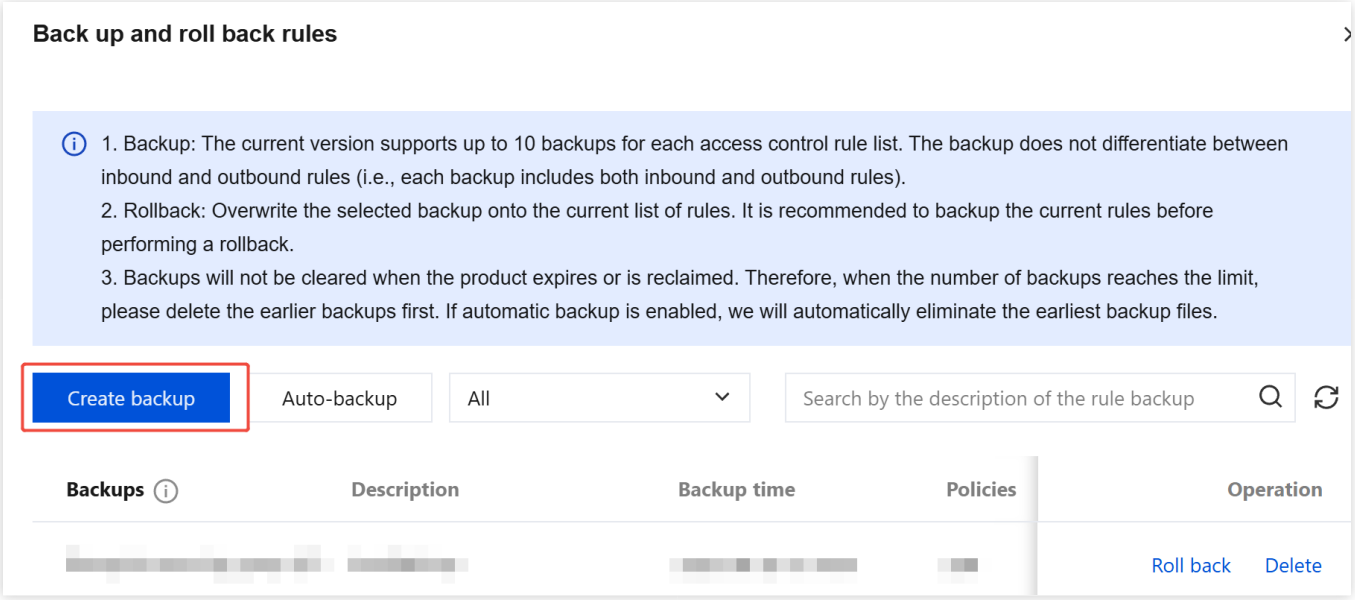


Policy Backup

1. Log in to the CFW console. In the left sidebar, click Common Tools.
2. On the Common Tools page, click Go to Configuration under Policy Backup and Rollback.



3. On the Policy Backup and Rollback page, click New Backup.



4. Expand the backup list box, and the system will display the list of rules currently available for backup. You can make selections based on your needs. Meanwhile, to facilitate later identification of the backup purpose, please enter a clear and concise description information in the Description field.
5. After selecting the rules and filling in the description information, click OK. The system will immediately initiate the backup process. After the backup is completed, a 'Backup successful' prompt will appear, indicating that the backup has been completed and securely stored.
6. The system can also schedule Automatic Backup for your selected rules as needed. Click Automatic Backup, and the system will list the rules available for backup. Select the rules for which you want to enable automatic backup configuration and set an appropriate time. After selecting the rules and time and verifying the information is correct, click OK, and the system will start the automatic backup task according to your settings.

Policy Rollback

> Note:
>
> Performing a rollback with a selected policy backup will overwrite the corresponding policy list. Existing policies will be deleted. To ensure data security, it is recommended to back up the current list before rollback.

1. Log in to the CFW console. In the left sidebar, click Common Tools.

2. On the Common Tools page, click Go to Configuration under Policy Backup and Rollback.

3. On the Policy Backup and Rollback page, the pop-up window includes all backed-up rule resources, descriptions, and backup times.

4. In the dropdown list box, locate the rules that require rollback and click Rollback.

5. Click OK, and the system will immediately start overwriting the current rule list with the backed-up rules.

Use Cases

Scenario Type		Policy Details
Policy Backup	Regular backup	Backing up firewall policies periodically helps retain different stages of network protection statuses. This allows for quick reference when reviewing policy evolution, conducting compliance audits, or comparing policies across different periods. For example, during an enterprise's quarterly internal network security review, regular backups can clearly present the direction and focus of firewall policy adjustments between quarters.
	Before the major business change	Before launching new business systems, expanding network areas, or making large-scale server architecture adjustments, it is essential to perform a policy backup. If the new policy causes issues such as disrupting existing business access, introducing network latency, or exposing security vulnerabilities, the backup can be quickly restored to avoid prolonged business interruption. For example, before the "Double 11" shopping festival, e-commerce enterprises scale out servers and adjust firewall policies. If they back up in advance, even if unexpected issues arise with the new policies, they can quickly roll back to ensure smooth online transactions during the festival.
	Personnel handover scenario	When members of the network security management team are replaced, backing up policies in advance helps reduce delays and risks caused by personnel changes. Backup policies can also serve as a reference baseline for new team members when making subsequent adjustments, ensuring consistency and rationality in policy modifications.
Policy Rollback	After misoperation	If an error occurs, such as mistakenly deleting critical rules or misconfiguring access permissions, resulting in business disruption or security risks, the policy rollback feature can immediately recover policies to the correct status before the error, ensuring normal business operations. For example, if an administrator mistakenly deletes active key business access rules while cleaning up expired rules, rolling back to the previous day's backup policy can quickly resolve the issue.
	After a failed policy adjustment test	When enterprises optimize firewall policies due to business development or new security threats, testing may reveal incompatibility with the existing network architecture, inability to block new attacks, or performance degradation. Using rollback restores the stable status before the adjustment, allowing re-evaluation of the test plan and avoiding losses from failed adjustments. For example, a financial enterprise tests a new intrusion detection policy but finds that a high false-positive rate impacts efficiency. Rolling back to the original policy allows further research and improvements.

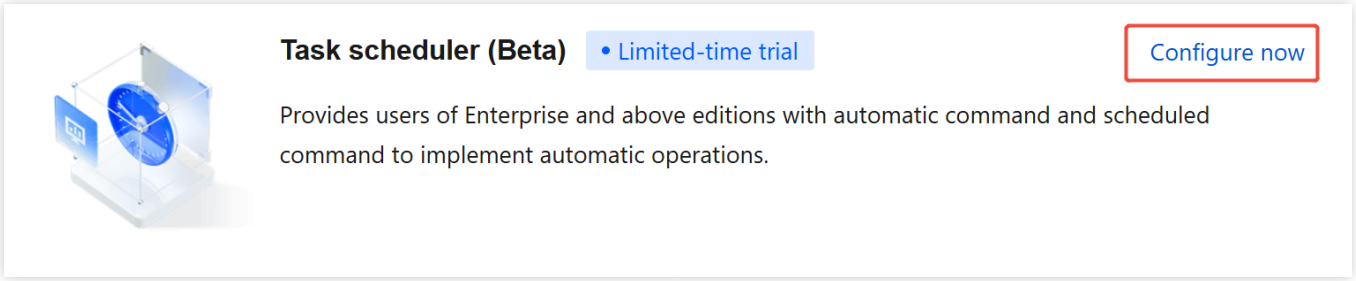
Scenario Type	Policy Details
After a network attack or failure	If an enterprise network suffers malicious attacks that tamper with firewall policies, or if system failures (such as hard disk damage or software crashes) cause policy files to be lost or corrupted, the rollback feature helps quickly recover the previously reliable policy configuration, minimizing losses and gaining time for troubleshooting and repair. For example, during a midnight DDoS attack, firewall policies are maliciously modified. Rolling back to the previous day's backup quickly recovers protection while technical personnel respond to trace the source and fix the issue.

Task Scheduler

Task Scheduler is an Ops tool that provides users with automated instructions and scheduled instructions, allowing users to orchestrate instructions and achieve automated operations.

Feature Entry

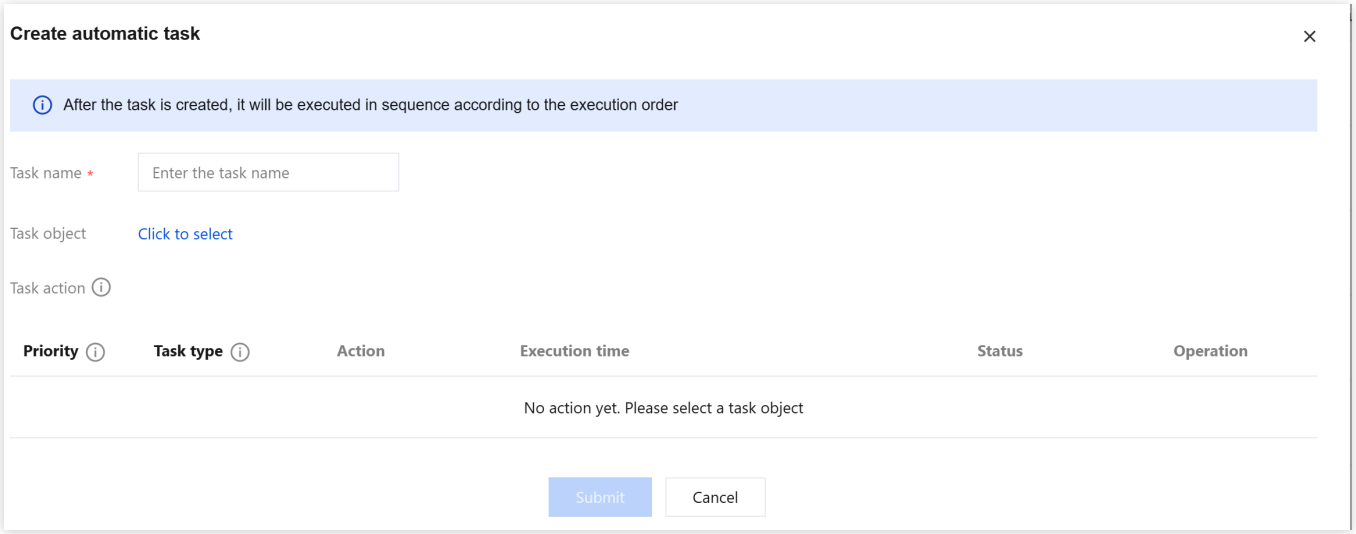
1. Log in to the CFW console. In the left sidebar, click Common Tools.
2. On the Common Tools page, click Go to Configuration under Task Scheduler.



3. On the Task Scheduler page, users can create and orchestrate automatic tasks to achieve features such as scheduled enabling/disabling of access control rules and scheduled ignoring of categorized alarms.

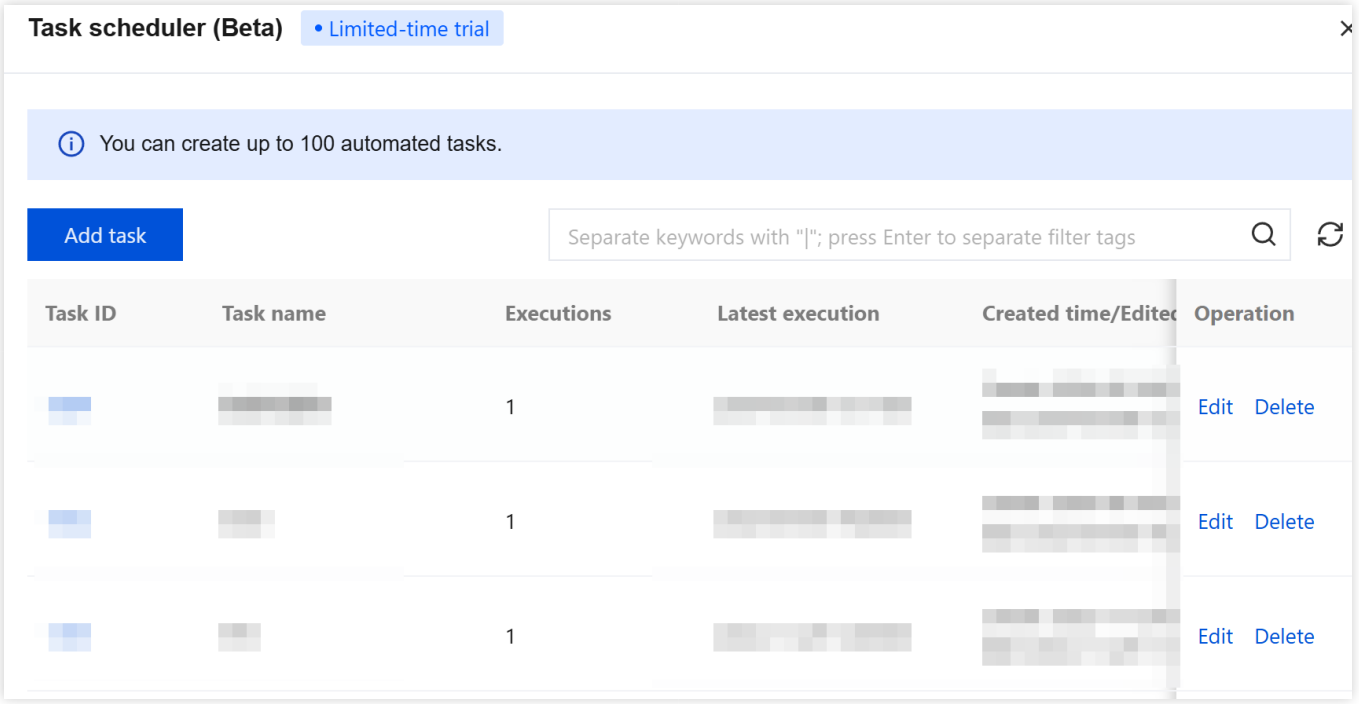
Adding an Automatic Task

1. On the Task Scheduler page, click Add Task.
2. In the Create Automatic Task pop-up window, enter the automatic task name.



3. Click Click to Select on the right of the task object. You can select the task object to associate with the automatic task via Category Selection or Custom Selection, then click OK.

4. Click Add Action to add the task actions to be executed. You can configure temporary tasks or periodic tasks along with the execution time. Multiple actions will be executed sequentially in the specified order. Click Duplicate to add more actions, up to 10 in total.
5. Click Submit to view the configured automatic tasks and their status in the task list.



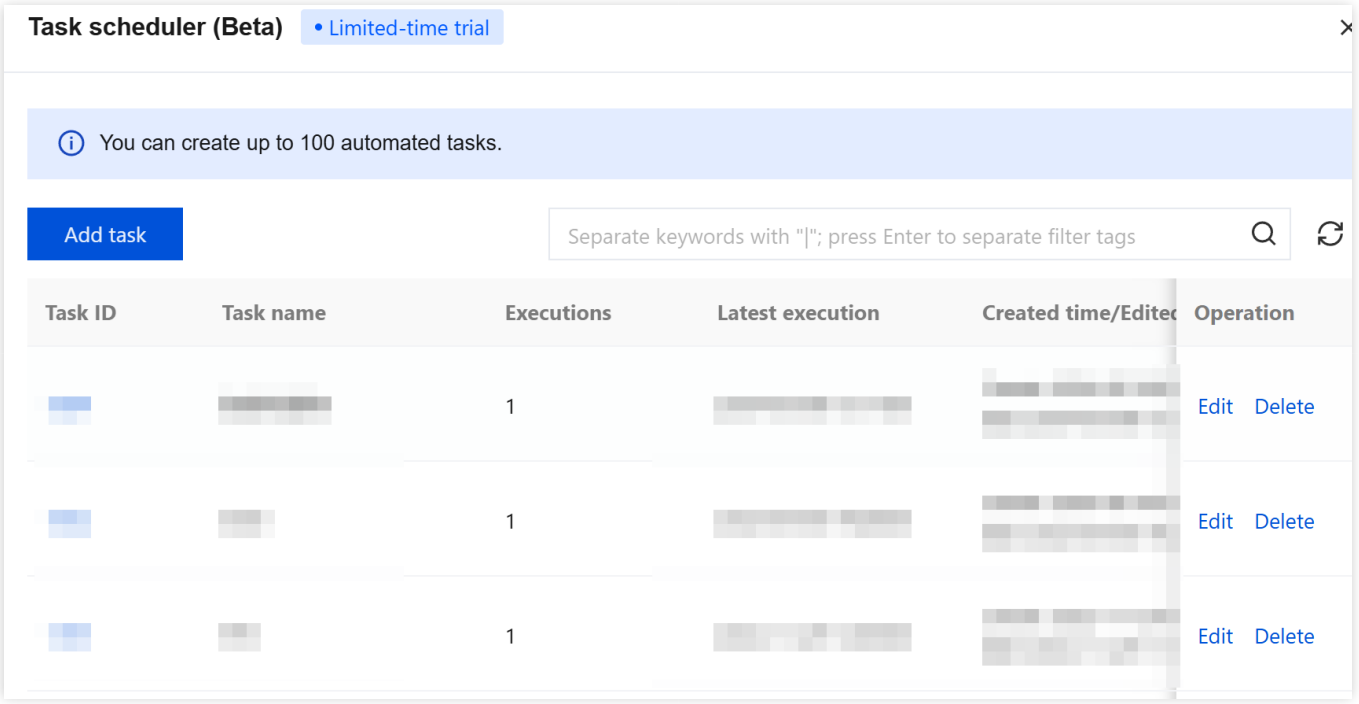
Quickly Adding an Automatic Task

You can quickly add tasks from the access control rule list. The following uses NAT firewall rules as an example.

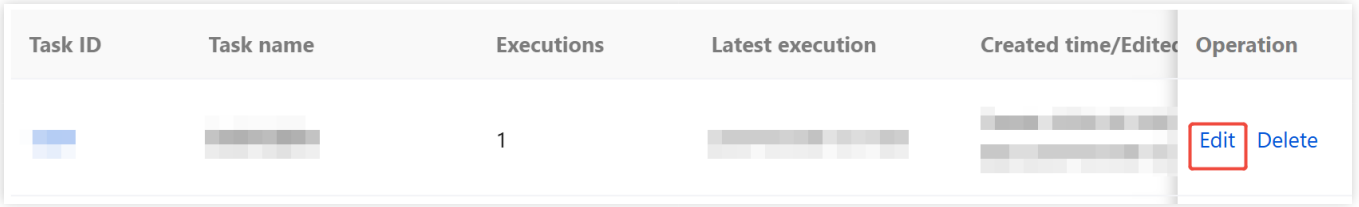
1. Log in to the CFW console. In the left sidebar, click Access Control > NAT Boundary Rule.
2. On the NAT Firewall Rules page, click Batch Operations.
3. Check the rules for which you want to execute automatic tasks, then click Add Automatic Task.
4. In the new automatic task window, fill in the task name, click Add Action, configure other parameters, and click Submit.

Editing an Automatic Task

1. On the Common Tools page, click Go to Configuration under Task Scheduler.

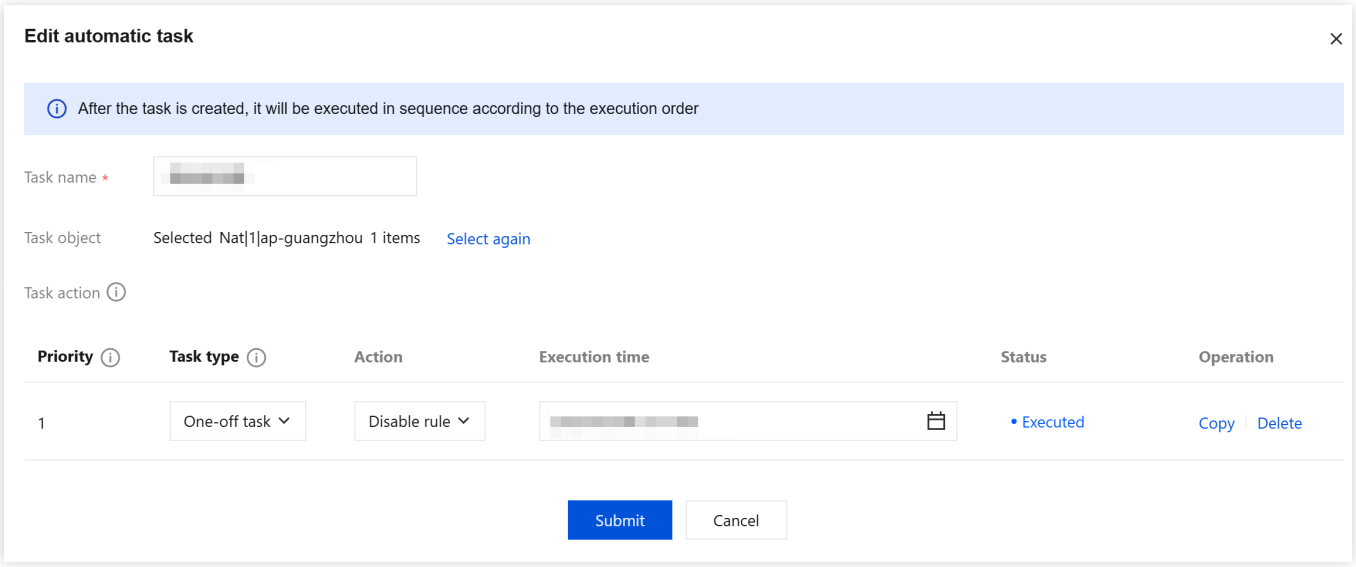


2. On the Task Scheduler page, locate the task you want to edit and click Edit.



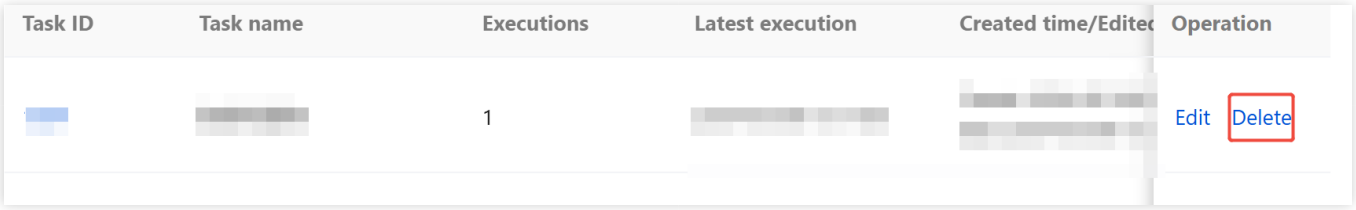
3. In the Edit Automatic Task pop-up window, you can edit the task name, task type, task actions, and execution time, as well as duplicate or delete tasks. Click Submit to save.

> Note:
>
> If the task object changes (such as being edited or deleted), the automatic task will no longer be executed.

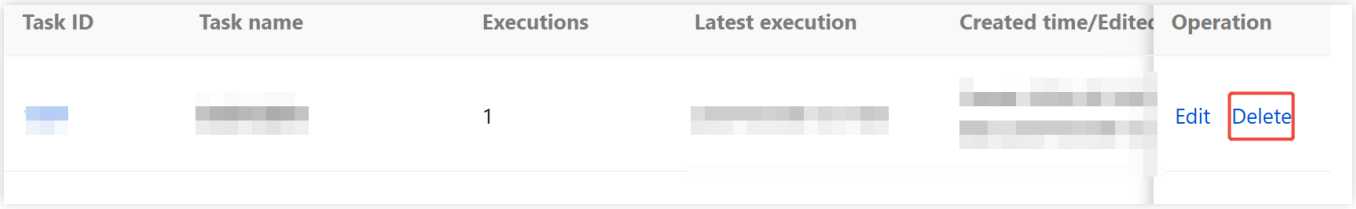


Deleting a Task

1. On the Common Tools page, click Go to Configuration under Task Scheduler.



2. On the Task Scheduler page, locate the task you want to delete and click Delete.



3. In the Confirm Deletion pop-up window, click OK to delete the task.

Querying Execution Records

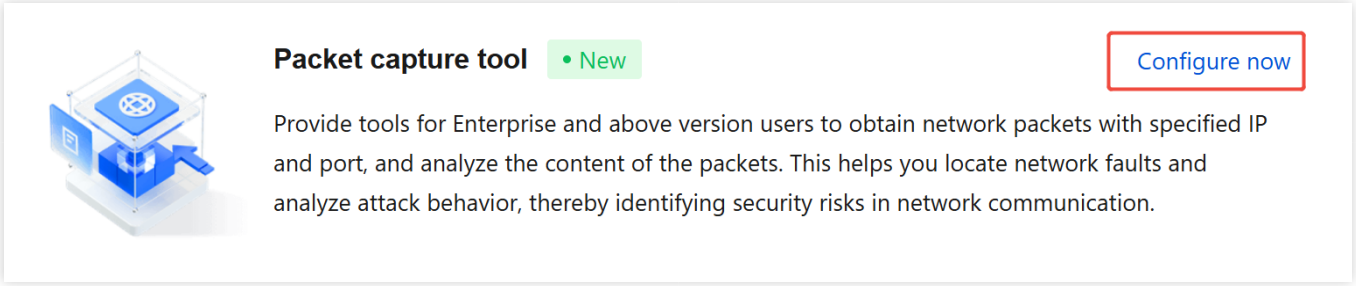
In the operation logs of the task-associated object, you can view operation records where the operation account is "Automatic Task".

Packet Capture Tool

The packet capture tool allows users to capture network packets for specified IP addresses and ports, and analyze their contents. It helps you locate network failures and analyze attack behaviors to identify security risks in network communications.

Feature Entry

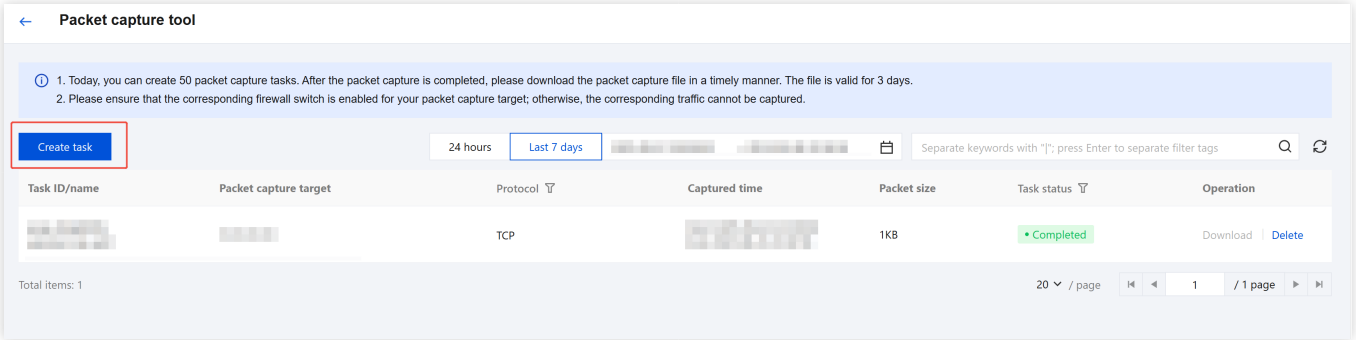
1. Log in to the CFW console. In the left sidebar, click Common Tools.
2. On the Common Tools page, click Go to Configuration under Network Packet Capture Tool.



3. On the Network Packet Capture Tool page, users can capture packets for traffic from the NAT firewall and inter-VPC firewall and download them for analysis.

Creating a Packet Capture Task

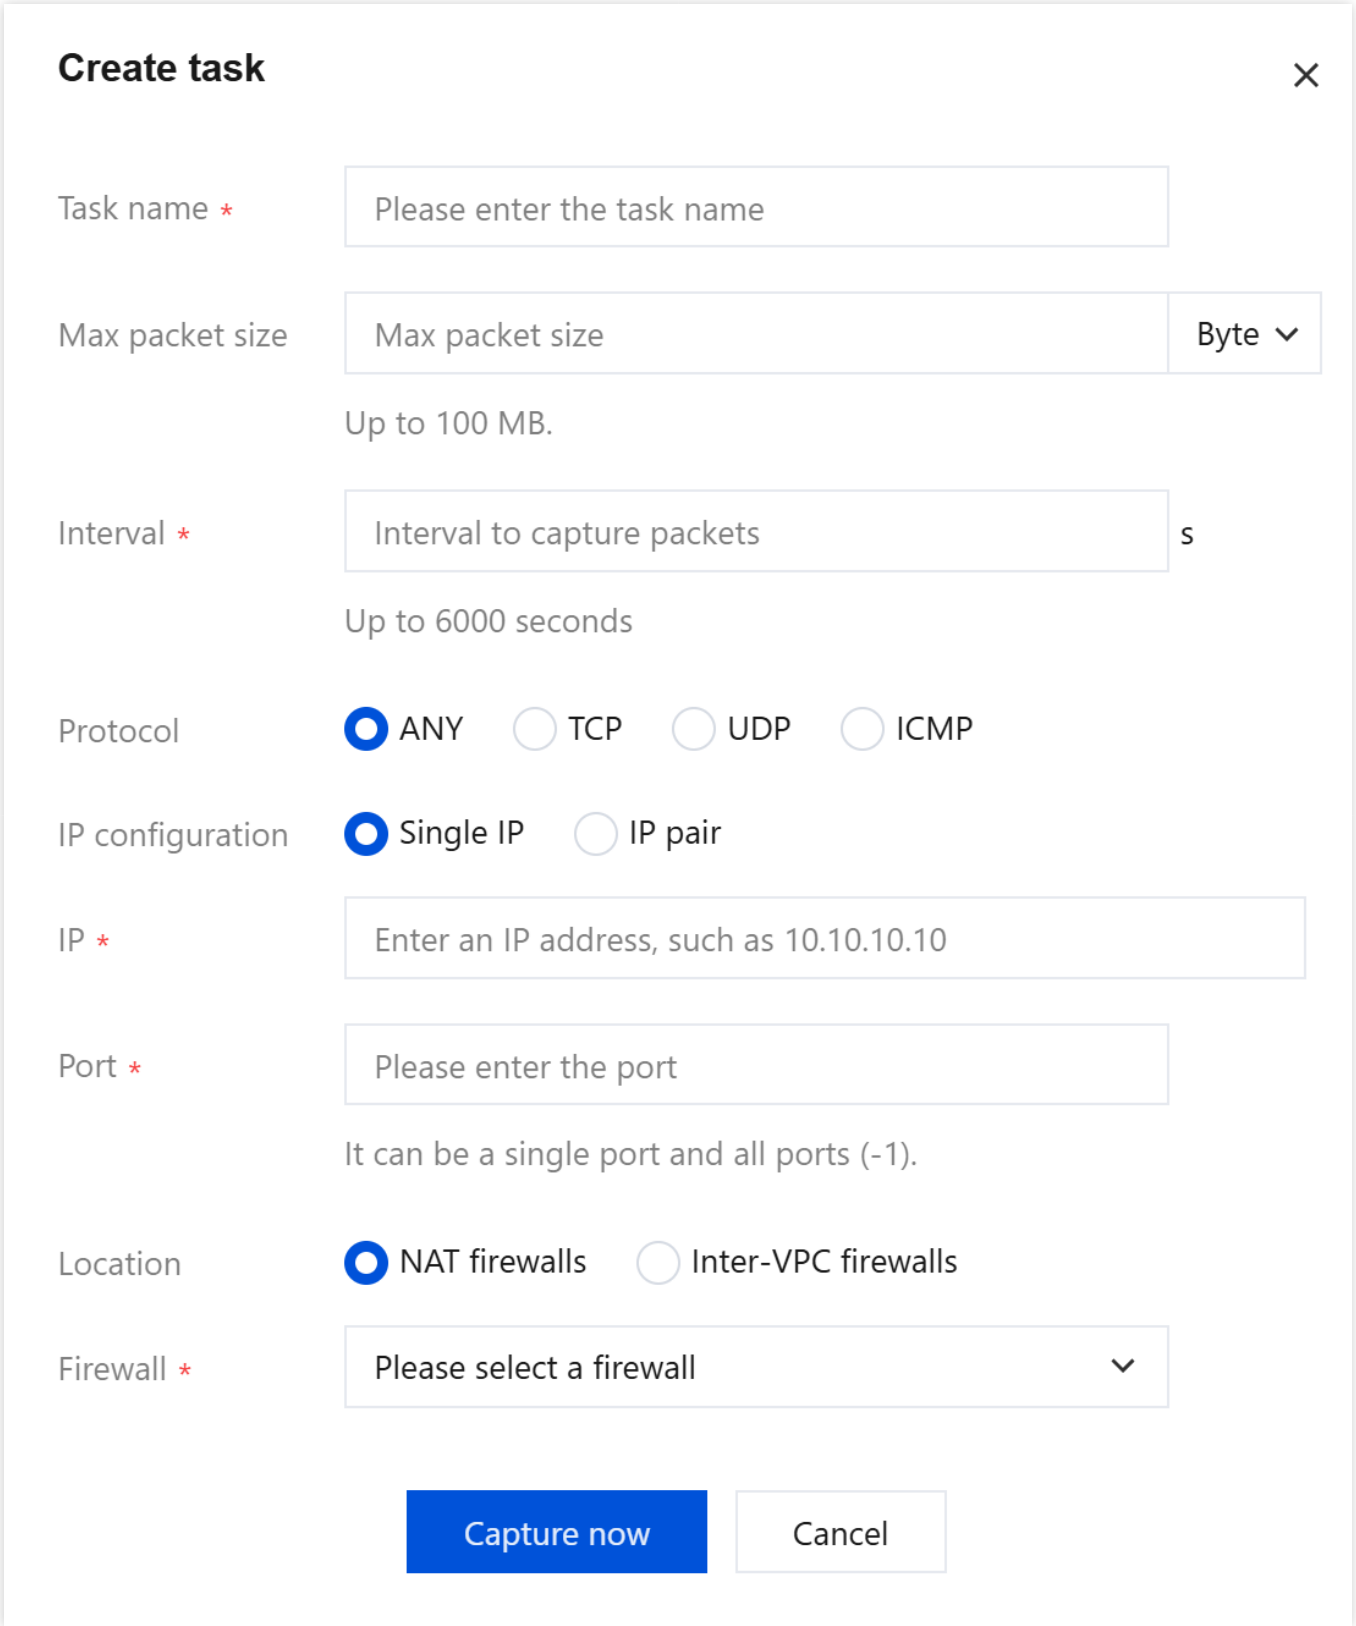
1. On the Network Packet Capture Tool page, click Create Packet Capture Task.



2. In the New Packet Capture Task pop-up window, configure the relevant parameters and click Start Packet Capture.

> Note:
>
> Please confirm that the corresponding firewall switch for your capture object is enabled; otherwise,

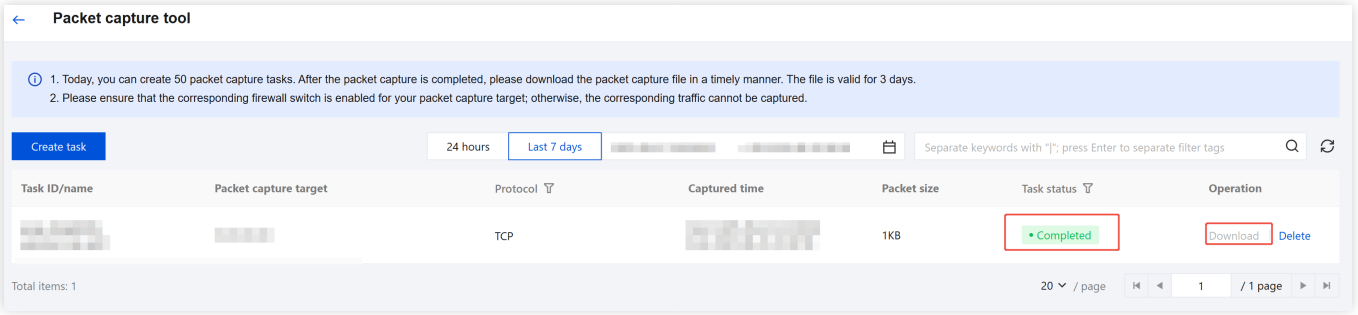
> traffic cannot be captured.



3. After the packet capture task is completed, click Download to download the packet capture results.

> Note:

> After the packet capture is completed, please download the packet capture file promptly; it is valid for 3 days.



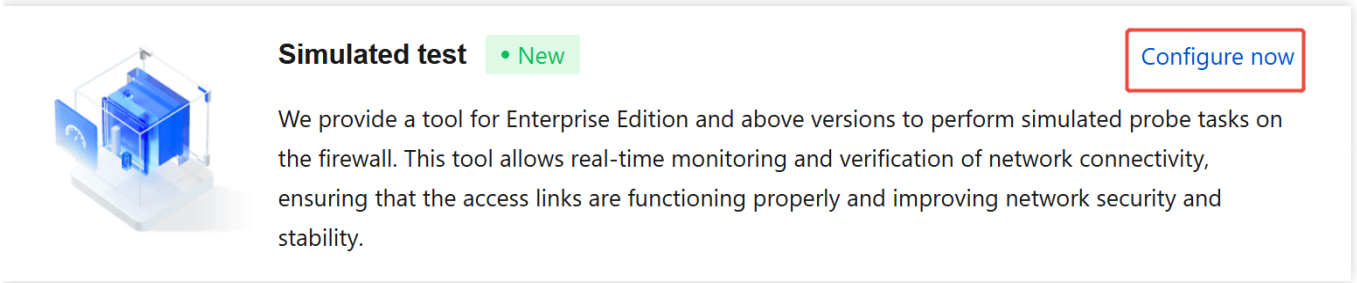
4. You can use network packet capture software such as Wireshark to open the packet capture results for further analysis.

Simulation Test

Choose Common Tools > Simulation Test to make related settings. The simulation test tool can monitor and verify network connectivity in real time, and detect access path anomalies in a timely manner, improving the Ops and availability of network linkages.

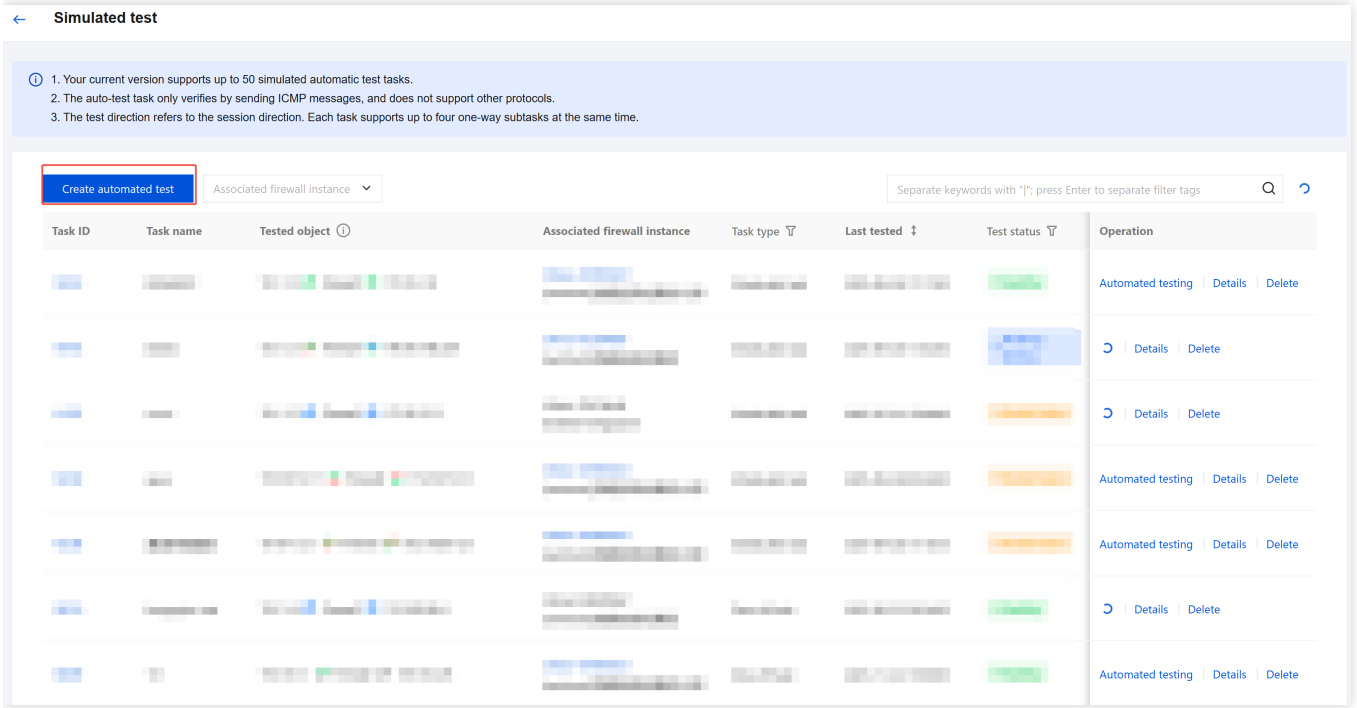
Entry to Simulation Test

Choose Common Tools > Simulation Test.



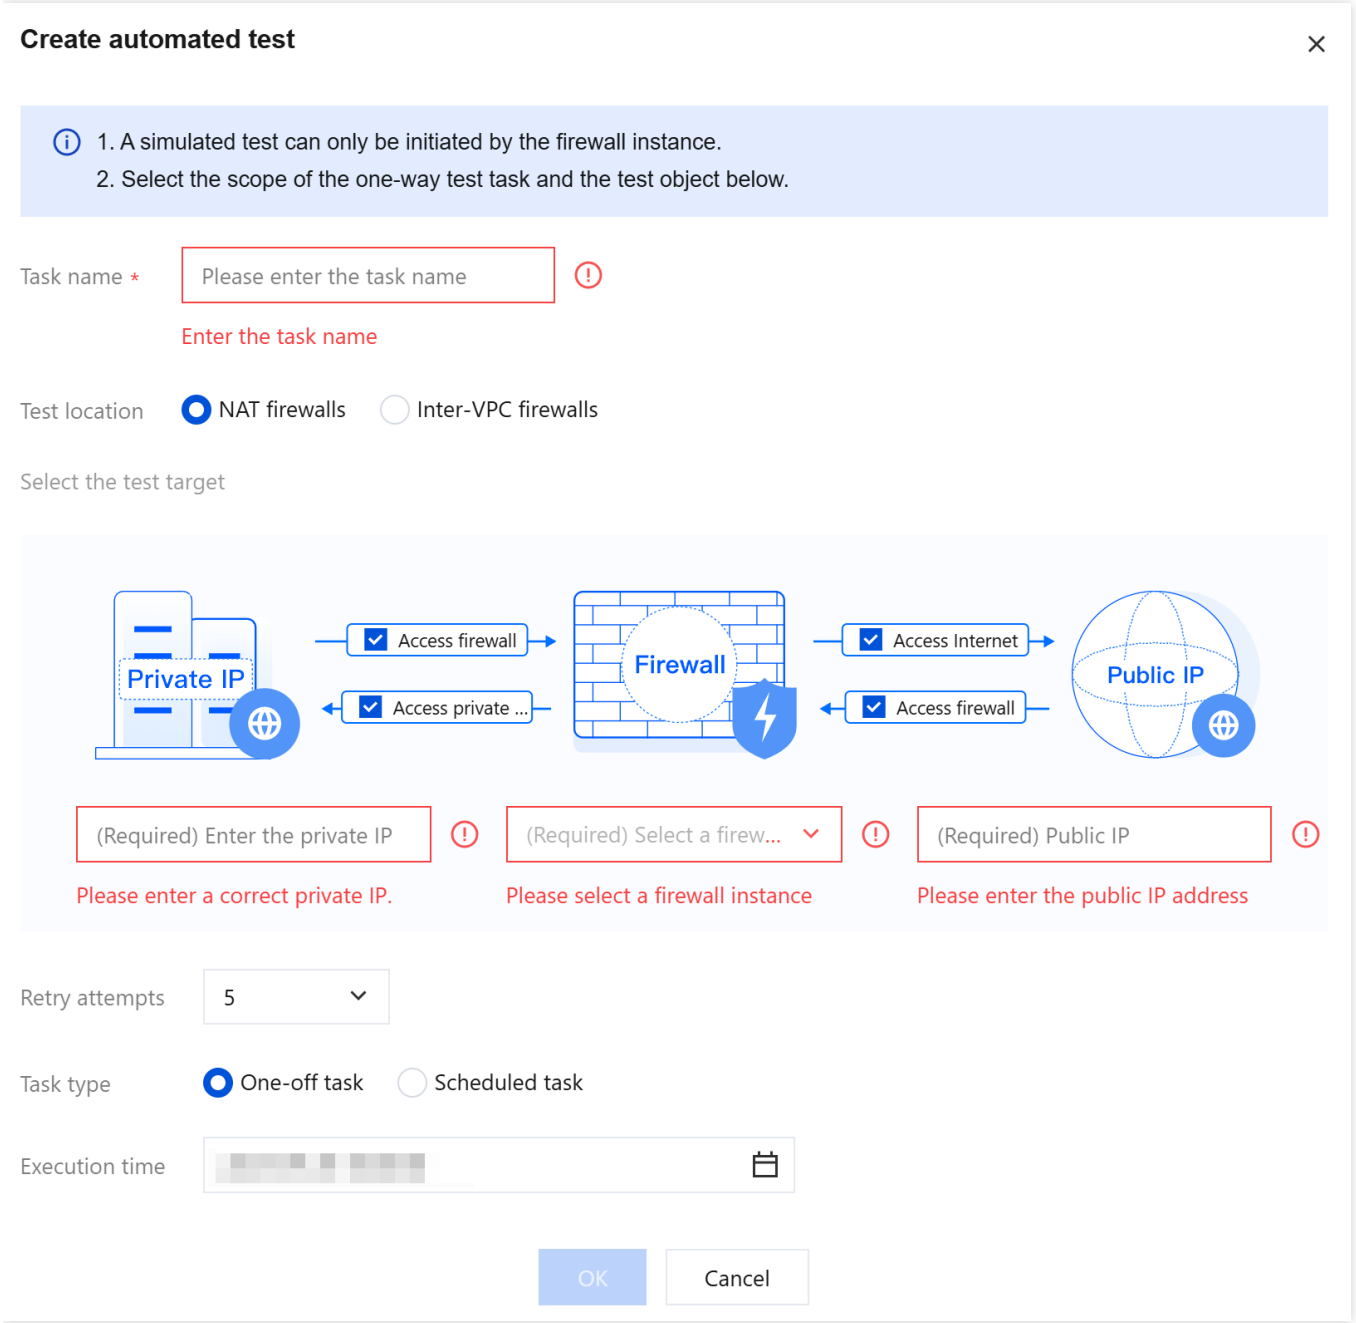
Creating a Simulation Test Task

Click Create Simulation Test Task.



Configure the location to be tested, the access source and destination of the test request, the direction of the traffic request, and the position of entering and exiting a firewall. You can also set the task type to a temporary task or a periodic task.

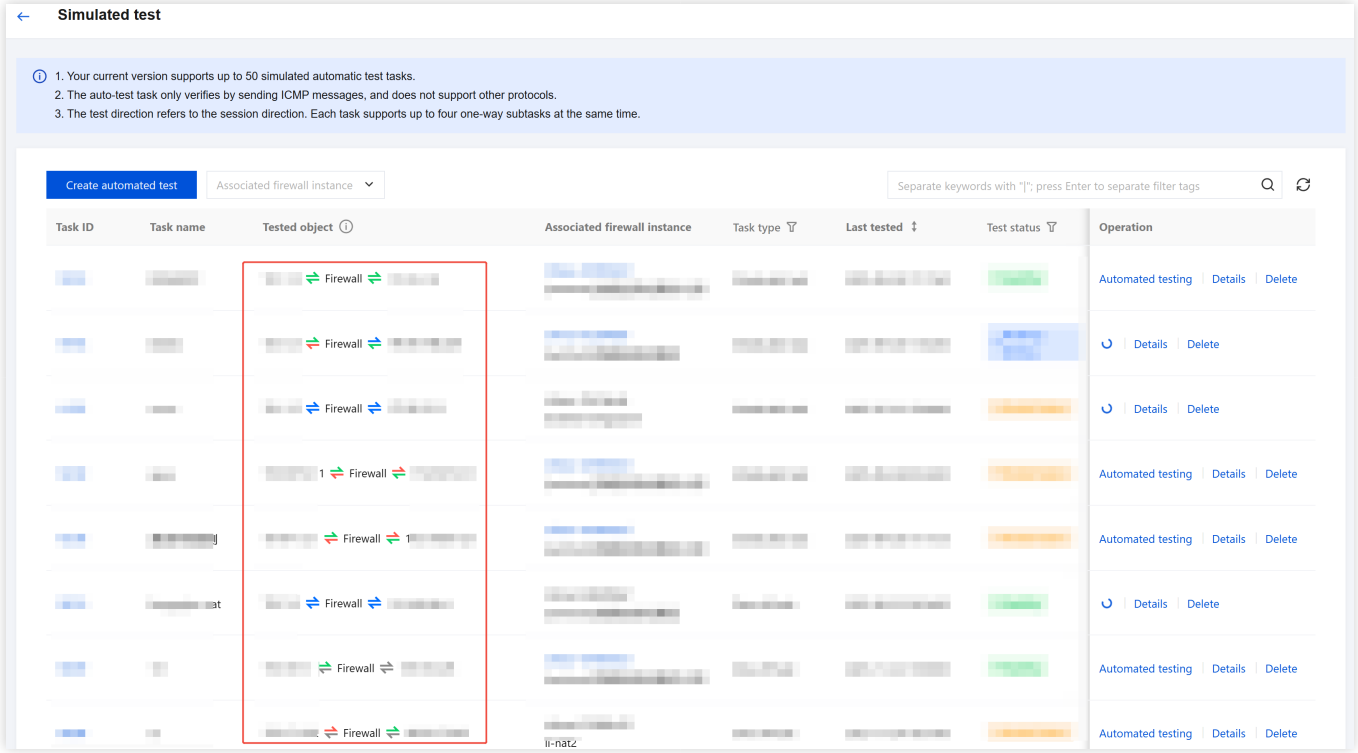
Note that a simulation test can only be initiated in one direction via a firewall instance. The principle is that the firewall monitors the health status of the engine access path by simulating packet sending and receiving, which helps verify whether the traffic is normal when it passes through the firewall engine sequentially.



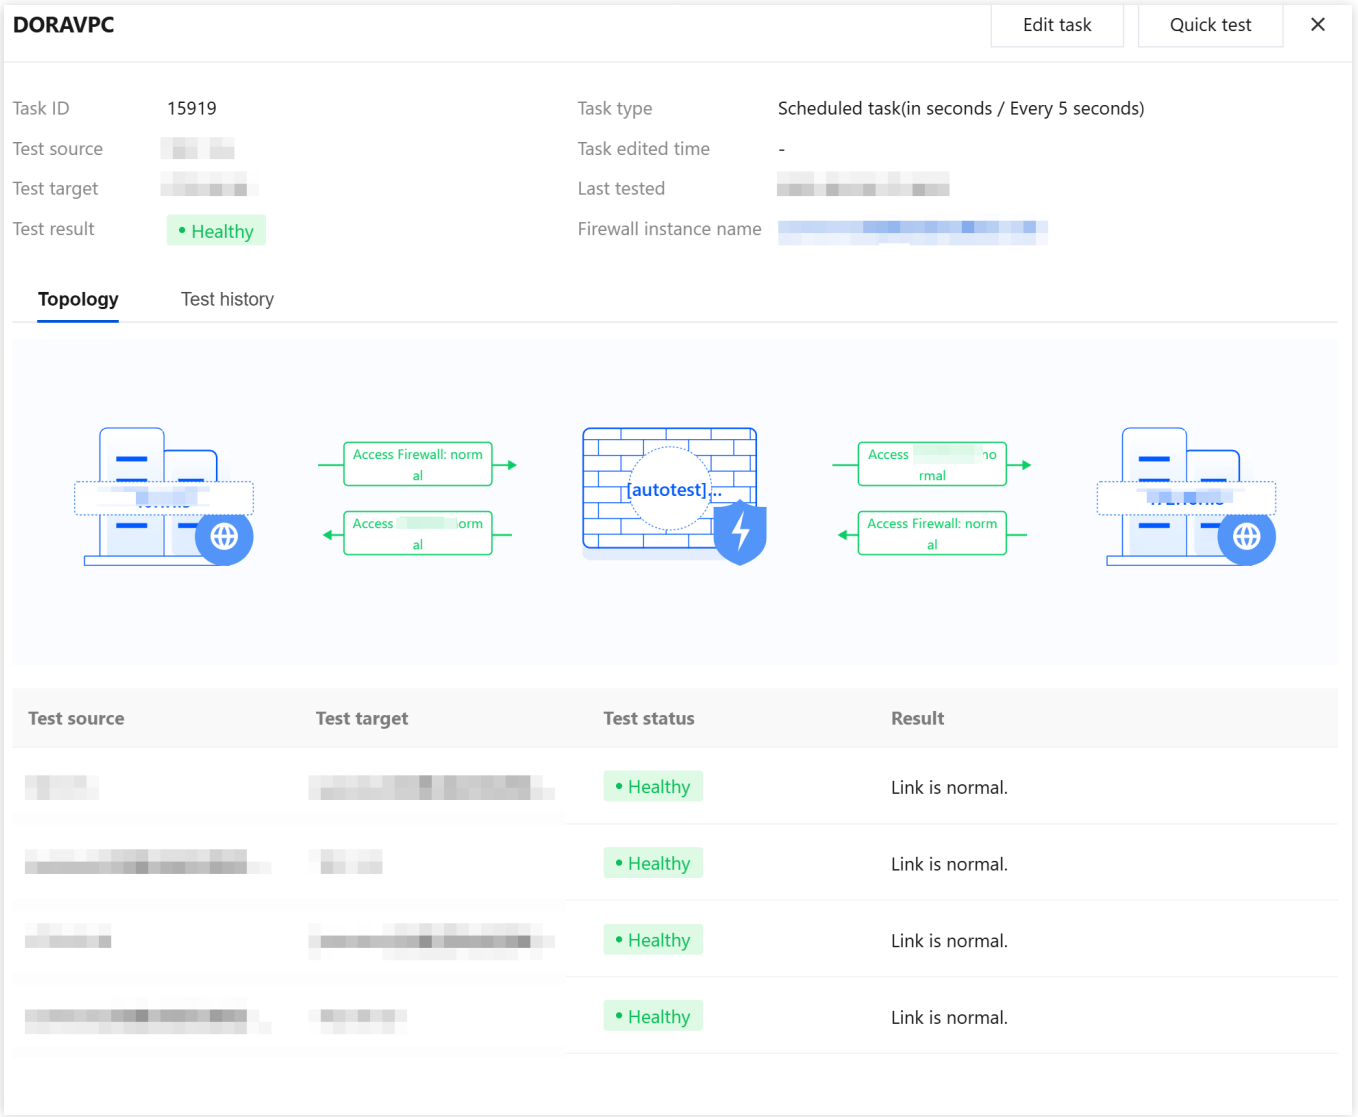
Viewing Test Tasks

1. View the test status of each task in the list. The Test Object column displays the current status of each test

subtask in real time. The Test Status column displays the current health status of the entire linkage.



2. Click a task ID or the Details button on the right side to view the details of a task.



When a test task is abnormal, system logs are recorded and exception information is pushed according to "system log push configurations".

FAQs for Beginners

Is the Default Rule of CFW "Allow" or "Block" If No Rule Is Configured?

CFW allows all traffic by default. When the CFW toggle is enabled, CFW will start recording traffic logs and generating intrusion prevention alarms. However, no traffic will be blocked at this time since no rules are configured.

How Long Does It Take for an Access Control Rule to Take Effect After It Is Configured?

When CFW rules are configured, it takes about 10 seconds to 1 minute for the rules to take effect.

What Protocols Are Supported by Access Control Rules?

- NAT Firewall supports TCP, UDP, ICMP, HTTP, HTTPS, SMTP, SMTPS, and FTP protocols.
- The inter-VPC firewall supports TCP, UDP, and ICMP protocols.

What Policy Modes Are Supported by Access Control Rules?

CFW supports three policies: Allow, Observe, and Block.

- In Allow mode: Allow the hit traffic and record the rule hit count and traffic logs, but not record access controls.
- In Observe mode: Allow the hit traffic and record the rule hit count, access control logs, and traffic logs.
- In Block mode: Block the access data and record the rule hit count and access control logs, but not

record traffic logs.

Under What Circumstances Does Traffic That Parses Domain Names Not Pass Through CFW?

CVM may use a cloud platform's self-built DNS service, and the generated DNS messages will not pass through the Internet boundary, resulting in missing alarms and logs for that part of the domain name access. If you want a CVM to be able to use the domain name detection and alarm features of CFW normally, you can manually change the resolution address in the /etc/resolv.conf file to 8.8.8.8.

Why Cannot Some Inter-VPC Firewall Switches Be Enabled?

Due to routing and IP range conflicts or other reasons, CFW restricts the toggles that cause conflicts. You can eliminate the conflicts according to the toggle error prompts before trying to enable Inter-VPC Firewall.

Ops Guide
Overview

The Ops of CFW has three UIs.

- Tenant side: You can use the tenant account and password to log in to the tenant side, and then go to CFW.
- Operation side: You can use the operation account and password to log in to the tenant side, and then go to CFW and Monitoring System.

CFW Monitoring Metrics and Exception Alarm Handling

Presentation of Alarm Events

[Cloud Platform Operation Side] - [Monitoring System] - [Cloud Product Metrics] - [CFW]

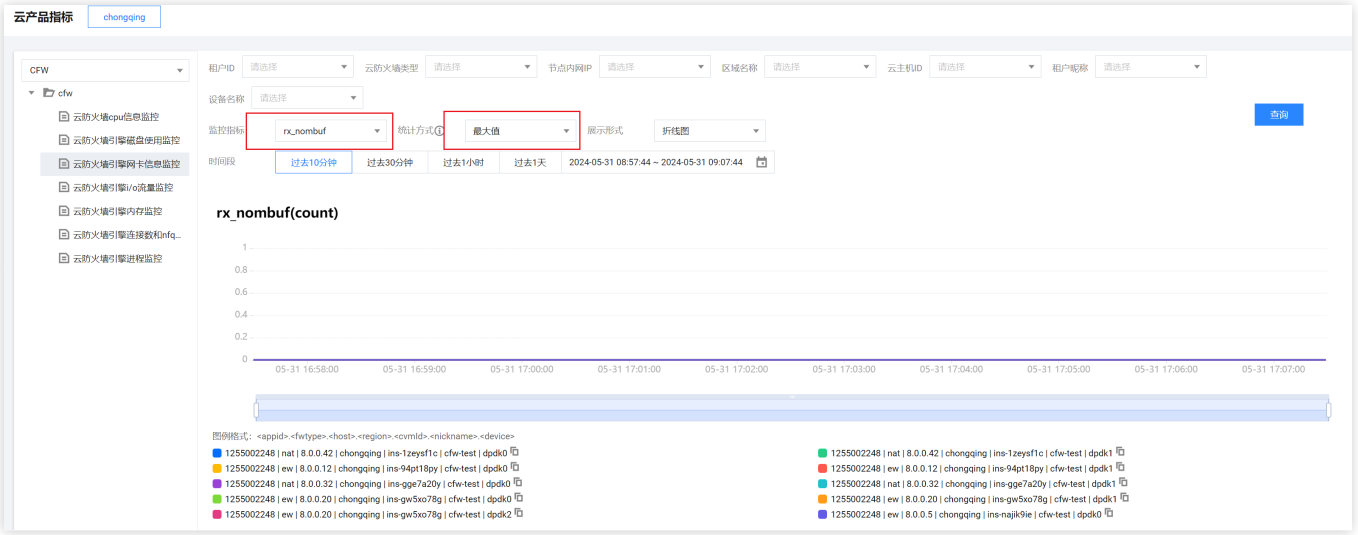
[Monitoring System] Alarm Events and Metrics:

No Mbuf Is Available When Receiving Messages [Critical]

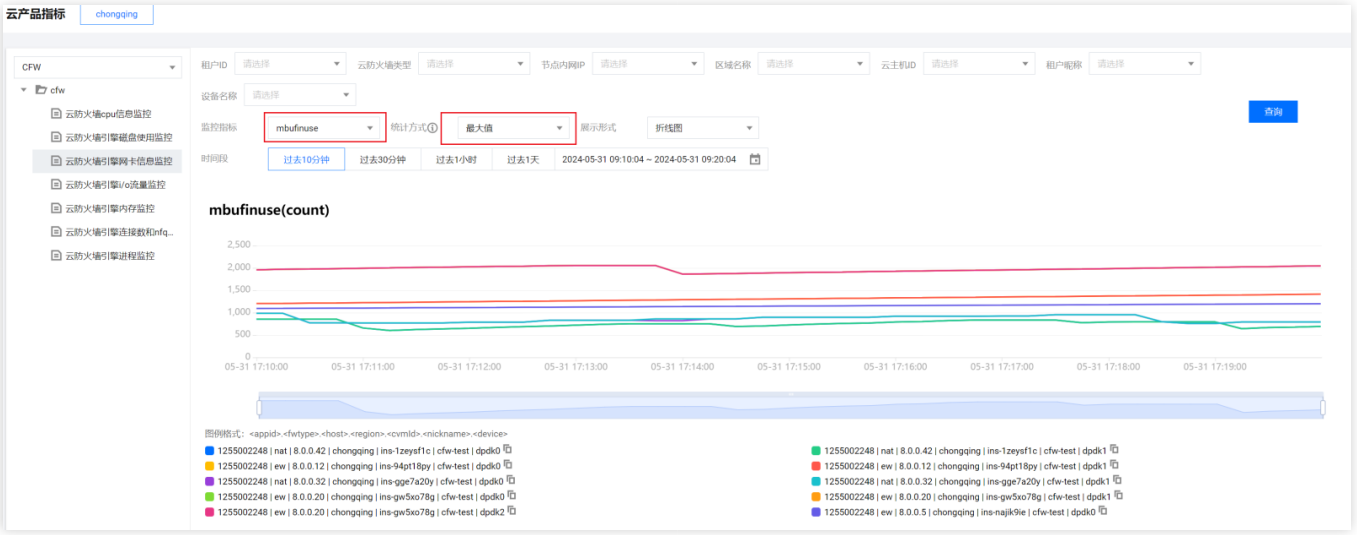
Threshold:
View metrics: View monitoring folding line graphs and monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  No Mbuf is available when receiving messages tcs_type = "cfw_event_err_info", event_name = "NO_MBUF"

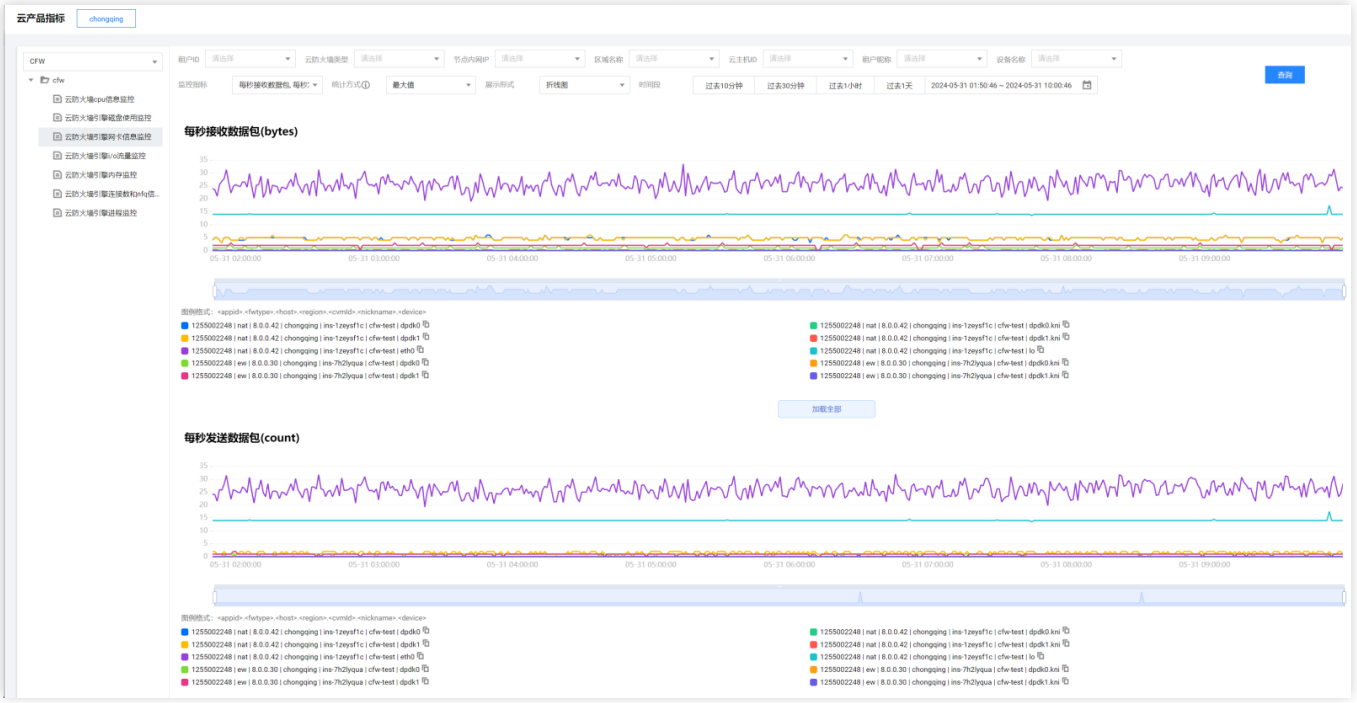
- Monitoring folding line graph: [CFW engine NIC information - rx_nombuf] When rx_nombuf is not 0, it indicates that an exception exists.



- Monitoring folding line graph: [CFW engine NIC information - mbufinuse]



- Monitoring folding line graph: [CFW engine NIC information - sending or receiving data packets per second]



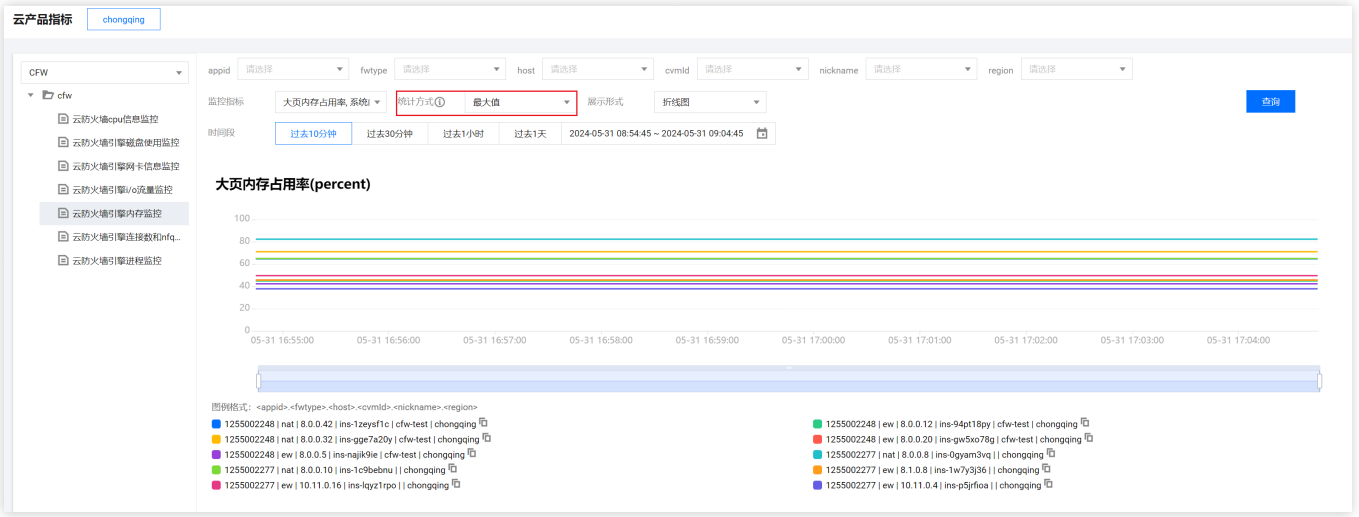
Insufficient Huge Page Memory Leads to an Automatic Primary-Secondary Switch [Critical]

Threshold: The primary engine memory usage exceeds 98%.
View metrics: View monitoring folding line graphs and monitoring events.

- Monitoring events: [CFW core workload event monitoring]

Insufficient huge page memory leads to an automatic primary-secondary switch tcs_type = "cfw_event_err_info", event_name = "HUGE_PAGES_NOT_ENOUGH_SWITCH"

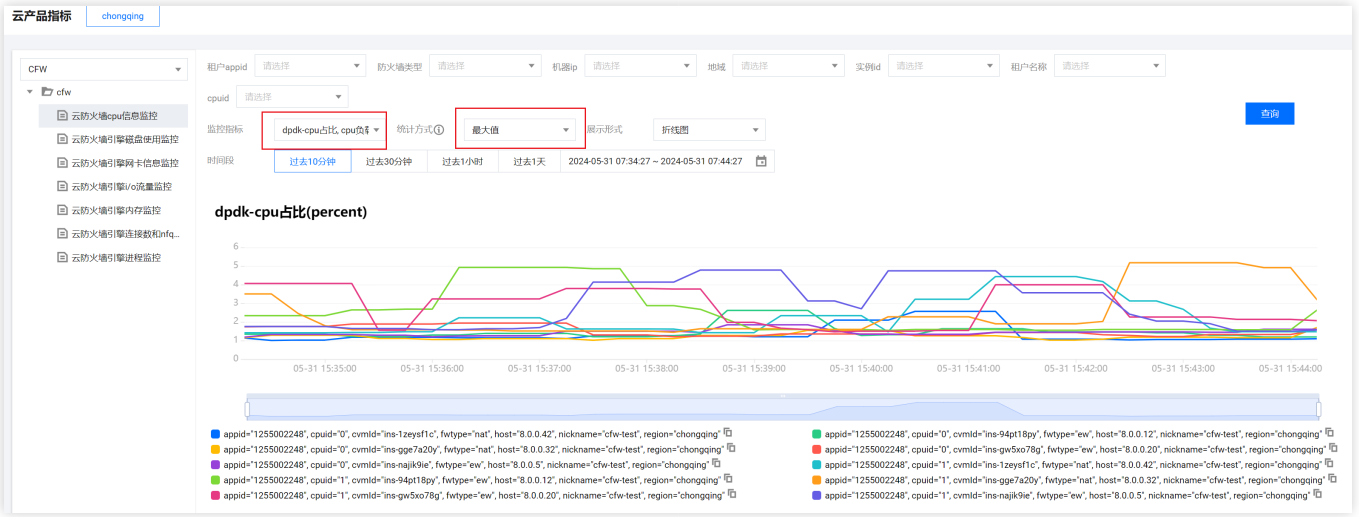
- Monitoring folding line graph: [CFW engine memory monitoring]



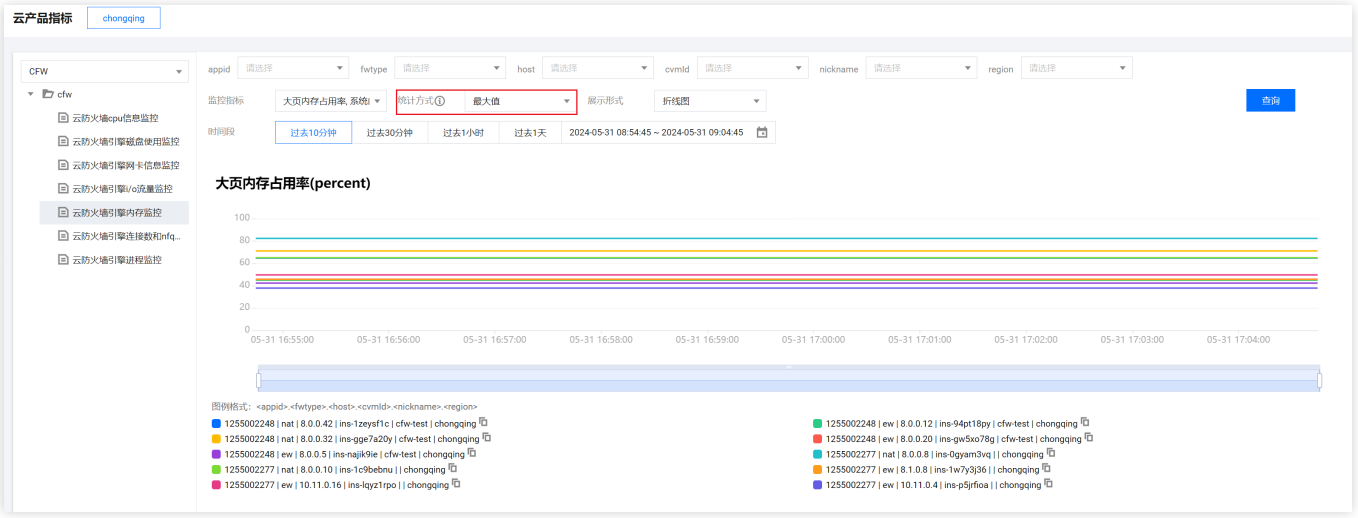
DPDK CPU Usage Exceeds 95% [Critical]

Threshold: A continuous full load of 95% has existed for 3 minutes.
View metrics: View monitoring folding line graphs and monitoring events.

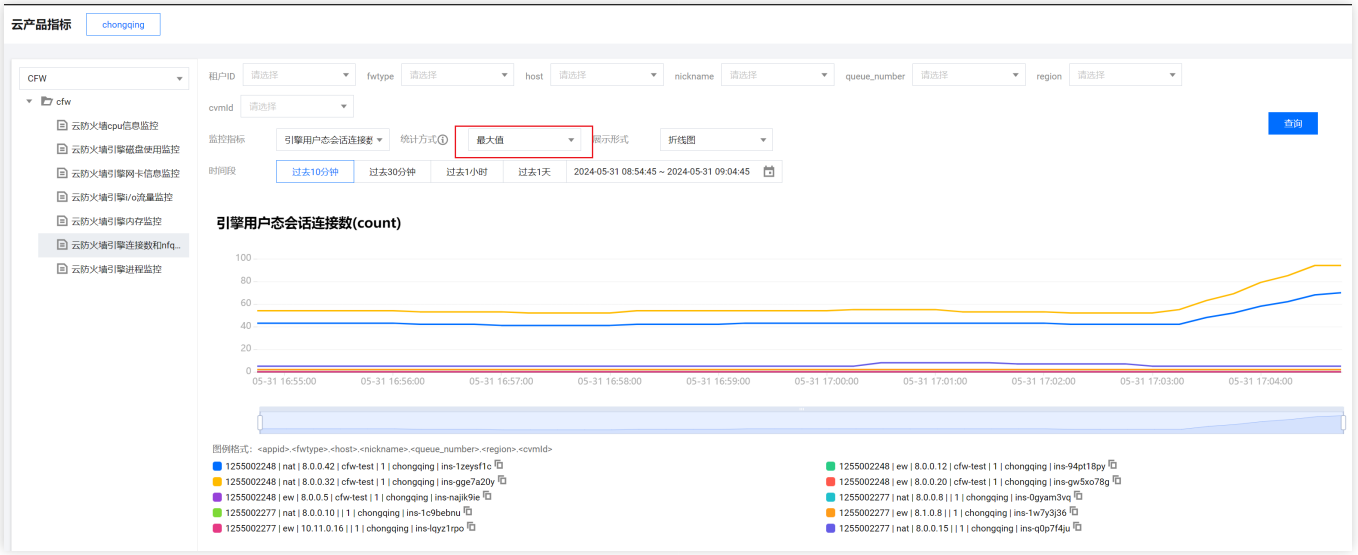
- Monitoring events: [CFW core workload event monitoring]
  DPDK CPU usage exceeds 95% tcs_type = "cfw_event_err_info", event_name = "DPDK CPU. HIGH"
- Monitoring folding line graph: [CFW CPU information monitoring]
  The CPU metric is reported once every 5 seconds, The minimum scale of Monitoring System is 15 seconds. Therefore, you need to select the statistical method as the maximum value.



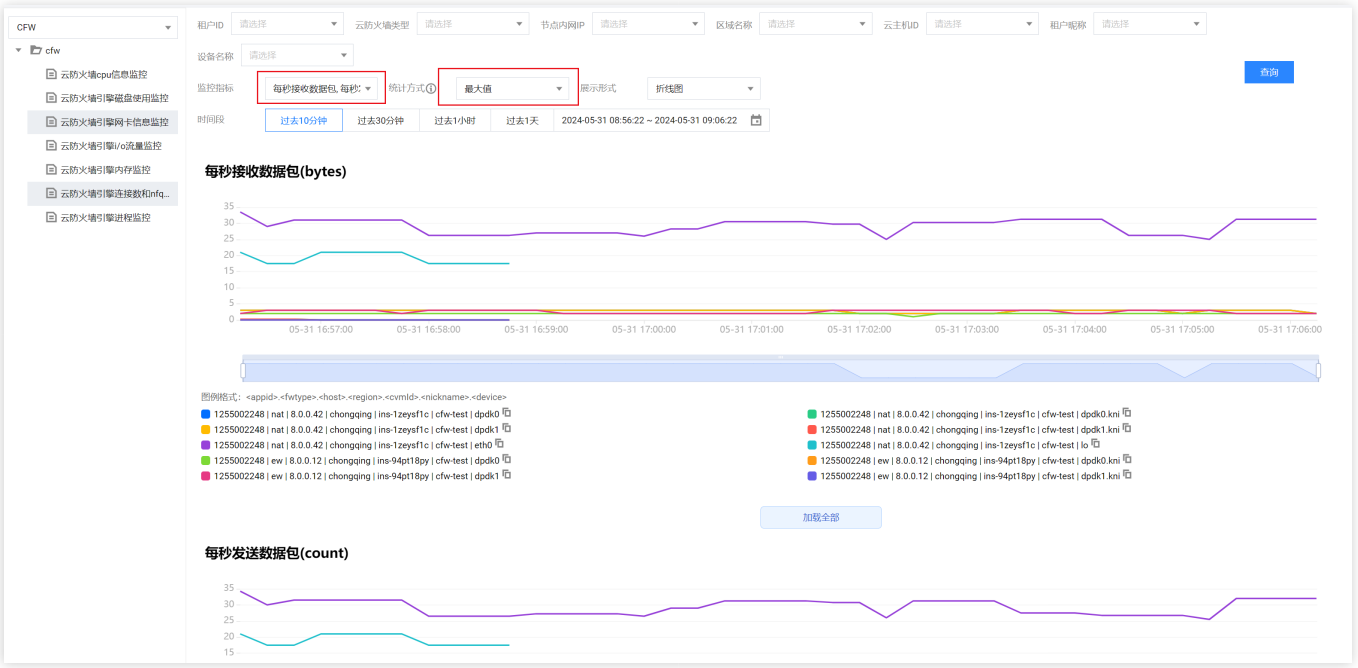
- Monitoring folding line graph: [CFW engine memory monitoring - huge page memory usage]



- Monitoring folding line graph: [number of CFW engine connections - number of engine user-mode session connections]



- Monitoring folding line graph: [CFW engine NIC information monitoring - sending or receiving data packets per second]



Primary-Secondary Engine Switch Occurs [Failed to Take Effect] [Critical]

Multiple alarms are generated in a short time, and the alarm that a primary-secondary engine switch occurs [successfully effective] is not received.
View metrics: View monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  A primary-secondary engine switch occurs [failed to take effect] tcs_type = "cfw_event_err_info", event_name = "HAVIP_INVALID_CHECK"
  A primary-secondary engine switch occurs [successfully effective] tcs_type = "cfw_event_err_info", event_name = "HAVIP_VALID_CHECK"

Firewall HAVIP Is Inconsistent with the Effective CVM [Critical]

View metrics: View monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  The firewall HAVIP is inconsistent with the effective CVM barad_err_status = 23

The HAVIP Engine Node IP Address of a NAT Firewall Is Lost [Critical]

View metrics: View monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  The HAVIP engine node IP address of NAT Firewall is lost. tcs_type = "cfw_event_err_info", event_name = "NATHAVIP_EMPTY_ALARM"

IP Address of the Server to Which HAVIP Belongs May Not Exist. Please Check the Traffic Status [Critical]

View metrics: View monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  The IP address of the server to which HAVIP belongs may not exist. Please check the traffic status. tcs_type = "cfw_event_err_info", event_name = "NO_IP_CHECK"

Insufficient Huge Page Memory [Major]

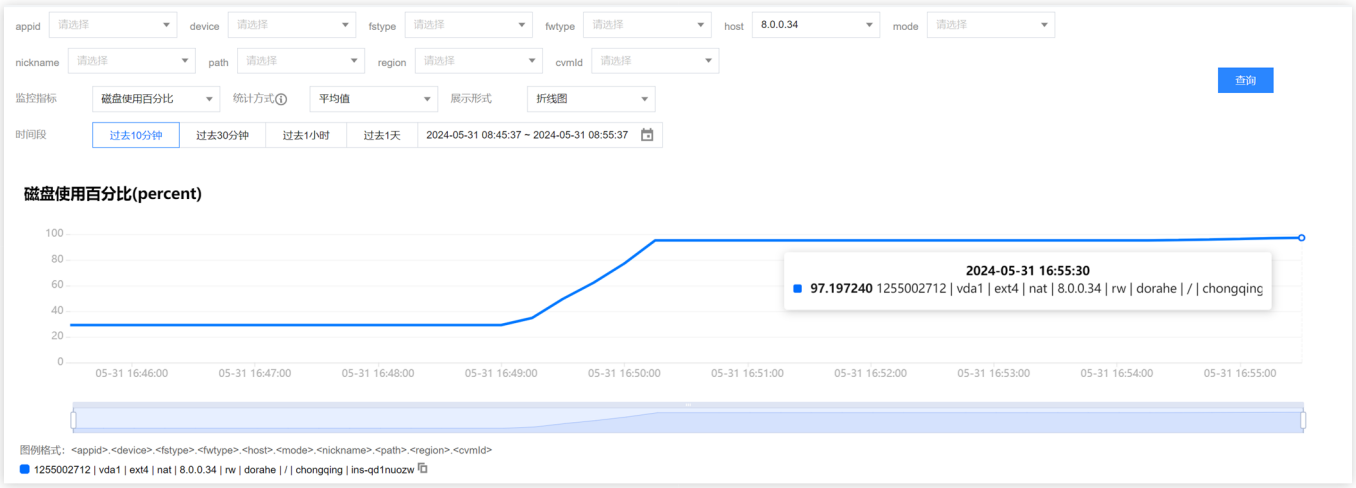
The memory usage exceeds 95%.
View metrics: View monitoring events.

- Monitoring events: [CFW core workload event monitoring]
  Insufficient huge page memory [severity: high] tcs_type = "cfw_event_err_info", event_name = "HUGE_PAGES_NOT_ENOUGH"

High Disk Usage [Major]

The disk usage exceeds 95%.
View metrics: View monitoring folding line graphs.

- Monitoring folding line graph: [CFW engine disk usage monitoring]



[Monitoring System] Alarm Event Handling

No Mbuf Is Available When Receiving Messages [Critical]

This alarm indicates that there is no extra memory for the NIC packet receiving queue. You can handle this alarm as follows:

1. Alarm acknowledgment: [CFW engine NIC information - rx_nombuf]. If there is a value, it means that it is valid.
2. Alarm handling: You need to contact Ops personnel to choose Tenant Side > CFW > Console or go to the cloud platform operation side to perform a primary-secondary engine switch, and then check whether the traffic has recovered.
3. Recovery confirmation: You can check the NIC traffic [CFW engine NIC information - sending or receiving data packets per second].
4. Exception handling: If the NIC traffic does not reach the desired level before the alarm, continues to drop to 0, or decreases sharply, you need to contact Ops personnel to choose Tenant Side > CFW > Console or go to the cloud platform operation side to trigger the engine bypass or disable the firewall toggle to restore your workload.

Insufficient Huge Page Memory Leads to an Automatic Primary-Secondary Switch [Critical]

When the engine's memory usage reaches 98%, the backend will automatically perform a primary-secondary switch. That is, the original primary engine becomes the current secondary engine. At this time, you need to log in to the current secondary engine to restart the engine, ensuring that huge page memory usage returns to a

healthy status, the secondary server remains healthy, and the engine works properly during the next switchover.
Alarm handling: You can log in to the engine's secondary server (see steps 2-2-1, 2-2-2, and 2-2-3) and then execute the instruction to restart the engine:
 kill -9 pidof fwdpdkengine``

DPDK CPU Usage Exceeds 95% [Critical]

When the usage proportion of a certain core of the engine reaches 95%, this alarm will be pushed.
Alarm acknowledgment: You can check the folding line graph: [CFW CPU information monitoring]. If the usage proportion has been higher than 90% for a long time, you need to find out the specific cause based on [CFW engine memory monitoring - huge page memory usage], [Number of CFW engine connections - number of engine user-mode session connections], and [CFW engine NIC information monitoring - sending or receiving data packets per second].
Here are some possible causes that have ever been troubleshooted:

- Huge page memory is exhausted.
- There is a sudden increase in traffic.
- Elephant flows enter the IPS.
  One or two alarms related to DPDK-CPU occur occasionally, which has a minor impact and can be considered a normal workload fluctuation. If the alarm persists for a long time (for example, 1 hour or 3 hours), it indicates that the workload traffic does not match the current instance. You can handle the alarm based on actual conditions:
- Perform vertical scaling on firewall instances.
- Disable some toggles to reduce firewall pressure.

Primary-Secondary Engine Switch Occurs [Failed to Take Effect] [Critical]

If you receive this alarm but do not receive the alarm that a primary-secondary engine switch occurs [successfully effective], you need to handle this alarm as follows:

1. View operation logs such as recent firewall toggles and ACL rules, and check whether this alarm is caused by some operations.
2. If you have not received a successful switchover alarm, check whether the workload NIC traffic decreases sharply or drops to 0 based on [CFW engine NIC information - sending or receiving data packets per second].
3. If the workload NIC traffic decreases sharply or drops to 0, you need to contact Ops personnel to perform a primary-secondary switch in the console or backend.
4. If the alarm persists, you can disable the firewall toggle or bypass the firewall.

Firewall HAVIP Is Inconsistent with the Effective CVM [Critical]

If you receive this alarm but do not receive the alarm that the firewall HAVIP is inconsistent with the effective CVM [recovery], you can handle this alarm as follows:

1. Check whether the workload NIC traffic drops to 0 [CFW engine NIC information - sending or receiving data packets per second].
2. If the workload NIC traffic drops to 0, you need to contact Ops personnel to trigger a primary-secondary switch in the console or backend.
3. If the alarm still persists, you can disable the firewall toggle or bypass the firewall.

The HAVIP Engine Node IP Address of a NAT Firewall Is Lost [Critical]

If you receive this alarm but do not receive the alarm that the HAVIP engine node IP address of NAT Firewall is lost, it indicates that the engine does not have a traffic routing NIC IP address. You can handle the alarm as follows:

1. Check whether the workload NIC traffic drops to 0 [CFW engine NIC information - sending or receiving data packets per second].
2. If the workload NIC traffic drops to 0, you need to contact Ops personnel to trigger a primary-secondary switch in the console or backend.
3. If the alarm still persists, you can disable the firewall toggle or bypass the firewall.

IP Address of the Server to Which HAVIP Belongs May Not Exist. Please Check the Traffic Status [Critical]

If you receive this alarm, it indicates that two engine NICs are abnormal and do not have traffic routing NIC IP addresses. You can handle the alarm as follows:

1. Check whether the workload NIC traffic drops to 0 [CFW engine NIC information - sending or receiving data packets per second].
2. If the workload NIC traffic drops to 0, you need to contact Ops personnel to trigger a primary-secondary switch in the console or backend.
3. If the alarm still persists, you can disable the firewall toggle or bypass the firewall.

Insufficient Huge Page Memory [Major]

You will receive this alarm when the engine's memory usage reaches 95%. At this point, an automatic primary-secondary switch will not occur. You can contact Ops personnel to view the Monitoring System metrics: view

[CFW engine memory monitoring], select the statistical method as the maximum value, and then check whether memory is continuously increasing or pulsating occasionally.
If memory is continuously increasing, it is recommended to contact Ops personnel to handle the alarm as follows in the console or backend:

1. View the curved line graphs related to huge page memory and CPU usage, and check whether the secondary server is healthy.
2. If the secondary server is healthy, you can contact Ops personnel to perform a primary-secondary engine switch in the console or backend.
3. After the switchover is successful, the original primary engine becomes the current secondary engine. At this point, you need to log in to the current secondary engine to restart the engine, ensuring that huge page memory usage returns to a healthy status, the secondary server remains healthy, and the engine works properly during the next switchover.
4. Log in to the secondary server of the engine (see steps 2-2-1, 2-2-2, and 2-2-3) and then execute the following instruction to restart the engine: kill -9 pidof fwdpdkengine``.
   If memory is pulsating occasionally, you can ignore this alarm and need to continuously check whether there are other fatal alarm metrics.

High Disk Usage [Major]

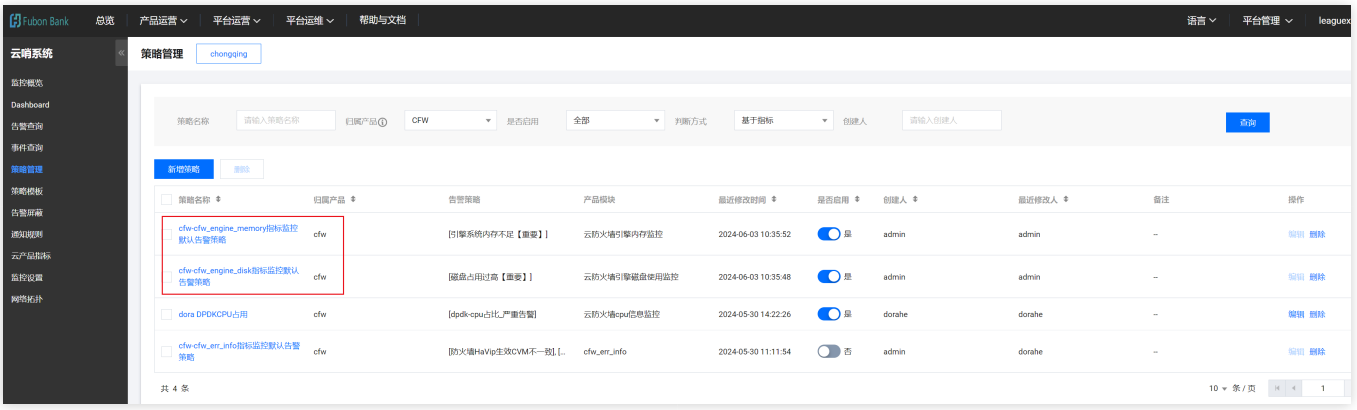
You will receive this alarm when disk usage exceeds 95%. Full disk usage will affect the delivery of firewall rules and routing policies and the normal operation of the engine. You need to log in to the engine to find and clean large files in time.

[Monitoring System] Alarm Configuration Guidelines

Confirming Policy Management

Console path: Choose Monitoring System > Policy Management. After CFW sends packages, the following two types of policies are provided by default.

- Metric-based policies:
  Policy name:
  Default alarm policy based on the cfw-cfw_engine_disk monitoring metric
  Default alarm policy based on the cfw-cfw_engine_memory monitoring metric



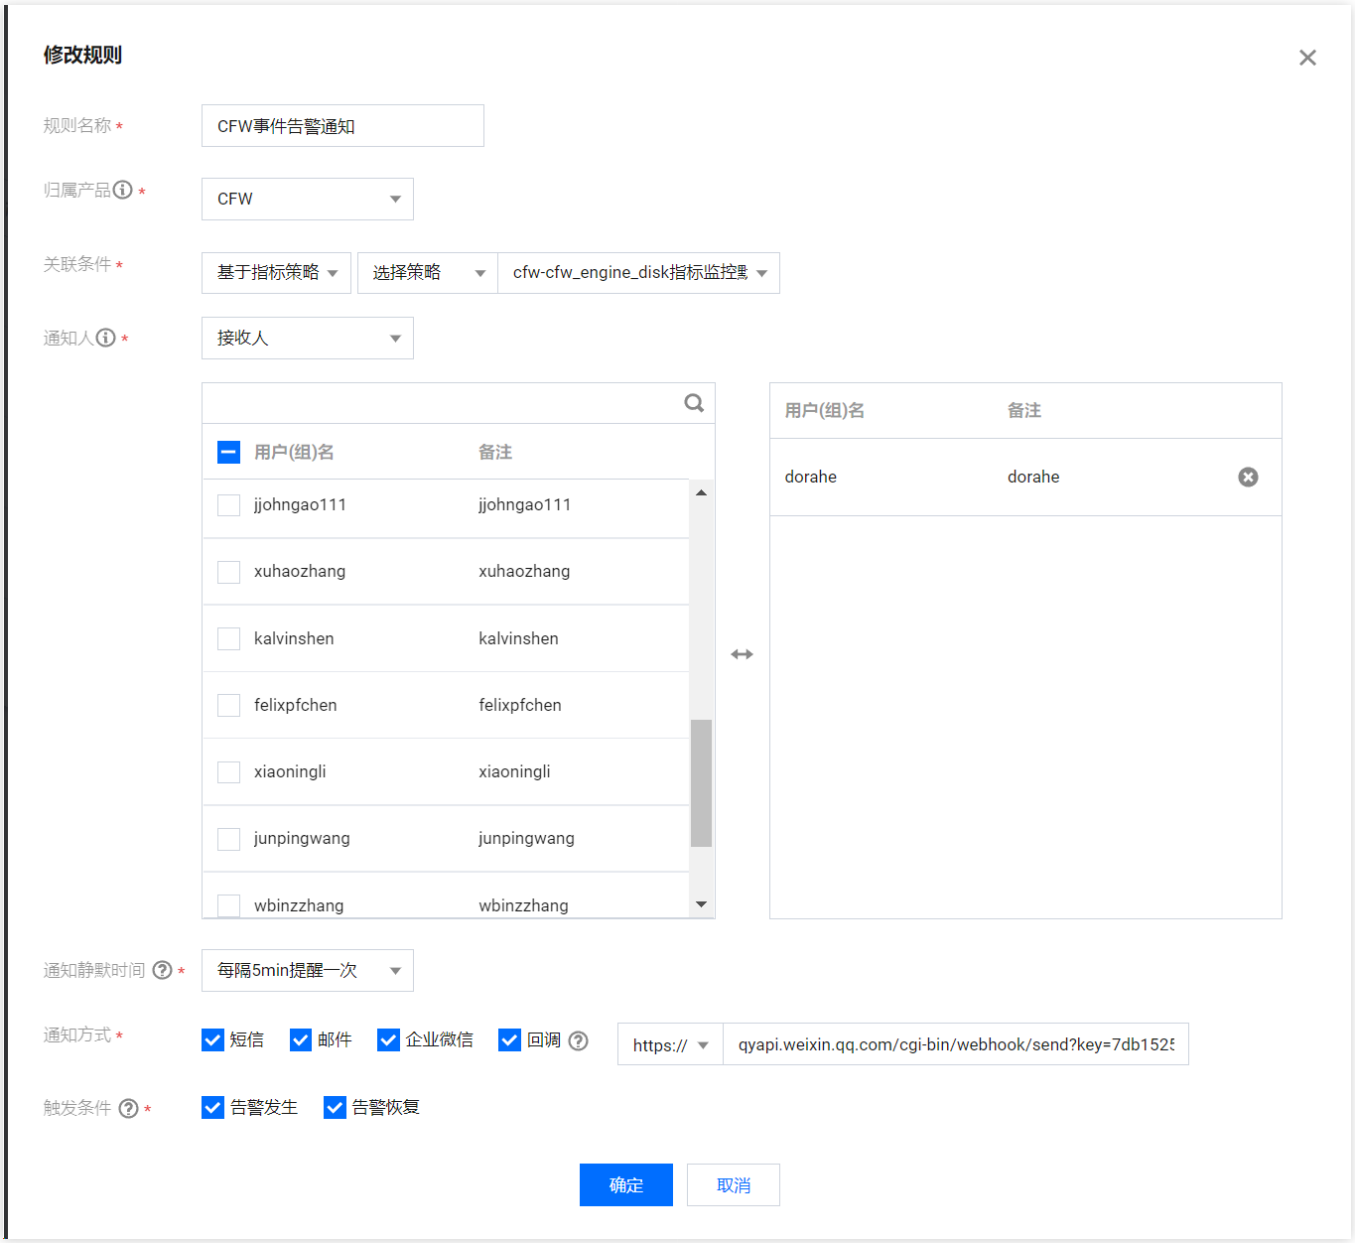
- Event-based policies:
  Policy name: default alarm policy based on the cfw-cfw_event_err_info event monitoring



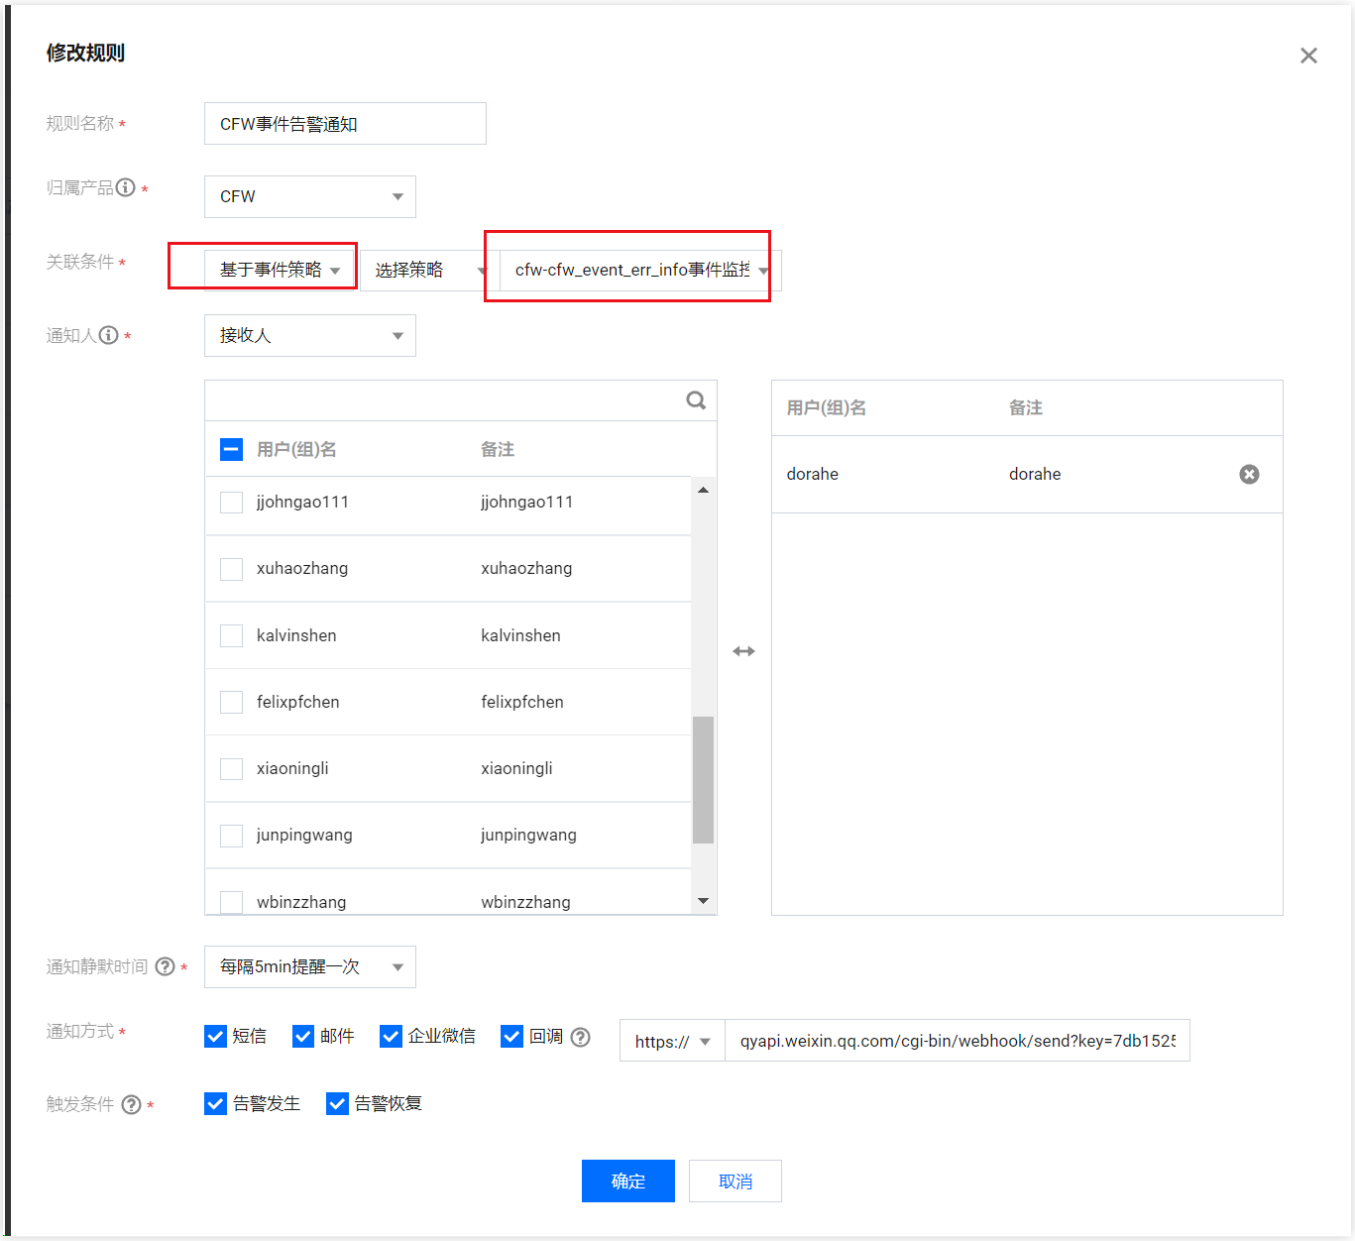
Configuring Notification Rules

Notification rules need to be configured onsite by customers after delivery.
Configured by: Ops personnel or onsite engineers
Console path: Choose Monitoring System > Notification Rules > Create

- Configuring metric-based notification rules:



- Configuring event-based notification rules:



Viewing Alarm Notifications

You can view the received historical alarms in the console, as shown in the following figure.


Fubon Bank
总览
产品运营
平台运营
平台运维
平台消息中心
站内信
消息订阅
通知渠道管理
短信渠道管理
消息统计
运营消息查询
邮件查询
短信查询
站内信查询
自定义渠道查询
企业微信查询
企业微信查询
内容详情
告警触发 | 产品: cfw | 等级: 严重 | 策略: dora DPDKCF
barad_dpdk_cpu_usage{tcs_product=~"cfw",tcs_type=