

Global Accelerator

User Guide

Issue 01
Date 2025-02-27



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions



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

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

Notice

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

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

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Global Accelerators.....	1
1.1 Global Accelerator Overview.....	1
1.2 Creating a Global Accelerator.....	1
1.3 Viewing a Global Accelerator.....	5
1.4 Modifying a Global Accelerator.....	6
1.5 Deleting a Global Accelerator.....	6
1.6 Managing Global Accelerator Tags.....	7
2 Listeners.....	10
2.1 Listener Overview.....	10
2.2 Adding a Listener.....	11
2.3 Access Control.....	14
2.4 Viewing a Listener.....	15
2.5 Modifying a Listener.....	16
2.6 Deleting a Listener.....	17
2.7 Managing Listener Tags.....	18
3 Endpoint Groups.....	20
3.1 Endpoint Group Overview.....	20
3.2 Adding an Endpoint Group.....	20
3.3 Viewing an Endpoint Group.....	22
3.4 Modifying an Endpoint Group.....	22
3.5 Deleting an Endpoint Group.....	23
4 Endpoints.....	25
4.1 Endpoint Overview.....	25
4.2 Adding an Endpoint.....	25
4.3 Viewing an Endpoint.....	26
4.4 Modifying an Endpoint.....	26
4.5 Removing an Endpoint.....	26
5 Health Checks.....	28
5.1 Health Check Overview.....	28
5.2 Configuring a Health Check.....	30
5.3 Viewing Health Check Settings.....	31
5.4 Modifying Health Check Settings.....	32

5.5 Disabling Health Check.....	33
6 IP Address Groups.....	34
6.1 IP Address Group Overview.....	34
6.2 Creating an IP Address Group.....	34
6.3 Viewing an IP Address Group.....	35
6.4 Modifying an IP Address Group	35
6.5 Deleting an IP Address Group.....	37
7 Cross-Border Permits.....	38
7.1 Cross-Border Permits Application.....	38
7.2 Application Progress Enquiry.....	40
8 Cloud Eye Monitoring.....	41
8.1 Overview.....	41
8.2 Supported Metrics.....	41
8.3 Setting an Alarm Rule.....	44
8.4 Viewing Metrics.....	45
9 Using CTS to Collect Global Accelerator Key Operations.....	46
9.1 Key Operations Recorded by CTS.....	46
9.2 Viewing Traces.....	47
10 Permissions Management.....	49
10.1 Creating a User and Granting Permissions.....	49
10.2 Custom Policy.....	50
11 Appendix.....	52
11.1 Configuring the TOA Module.....	52

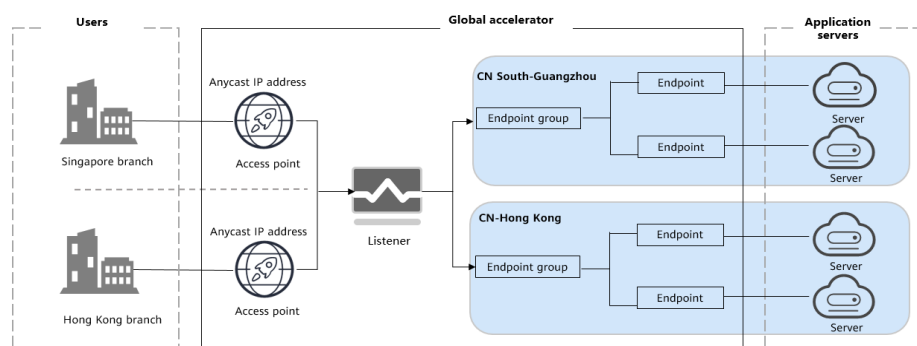
1 Global Accelerators

1.1 Global Accelerator Overview

A global accelerator directs user requests to endpoints through the Huawei backbone network.

You can create a global accelerator and select where you would like to use the global accelerator. The system will assign an anycast IP address for access from the nearest access point. When a client sends a request, the request will first go to the nearest access point, then to the Huawei backbone network, and finally to the optimal endpoints.

Figure 1-1 How Global Accelerator works



1.2 Creating a Global Accelerator

Scenario

To use Global Accelerator for faster application access, you first need to create a global accelerator.

Procedure

1. Log in to the [Global Accelerator console](#).

2. Click **Buy Global Accelerator**.
3. Specify the parameters listed in [Table 1-1](#).

Figure 1-2 Creating a global accelerator

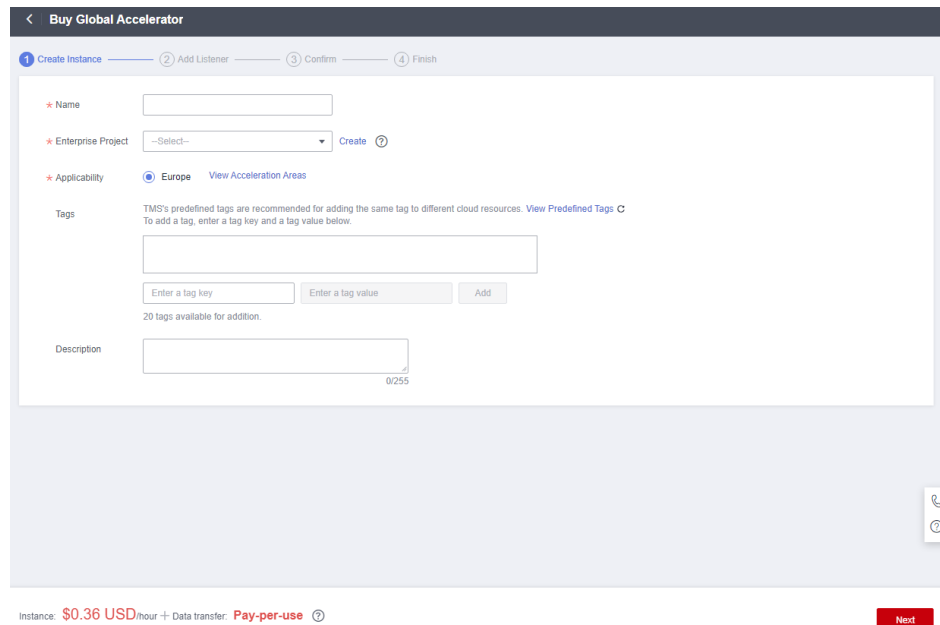


Table 1-1 Creating a global accelerator

Parameter	Description
Name	Name of the global accelerator you want to create. Only letters, digits, and hyphens are allowed. You can enter up to 64 characters.
Enterprise Project	An enterprise project you would like to use to centrally manage your Global Accelerator resources. You can use an existing enterprise project or create one.
Applicability	Where the global accelerator will be used. Default value: Europe .
Tags	The identifier of a global accelerator. Each tag consists of a key and a value. You can add 20 tags for a global accelerator.
Description	Supplementary information about the global accelerator. You can enter up to 255 characters.

4. Click **Next**.
5. Add one or more listeners to the global accelerator. For details about the parameters, see [Table 1-2](#).

Figure 1-3 Adding a listener

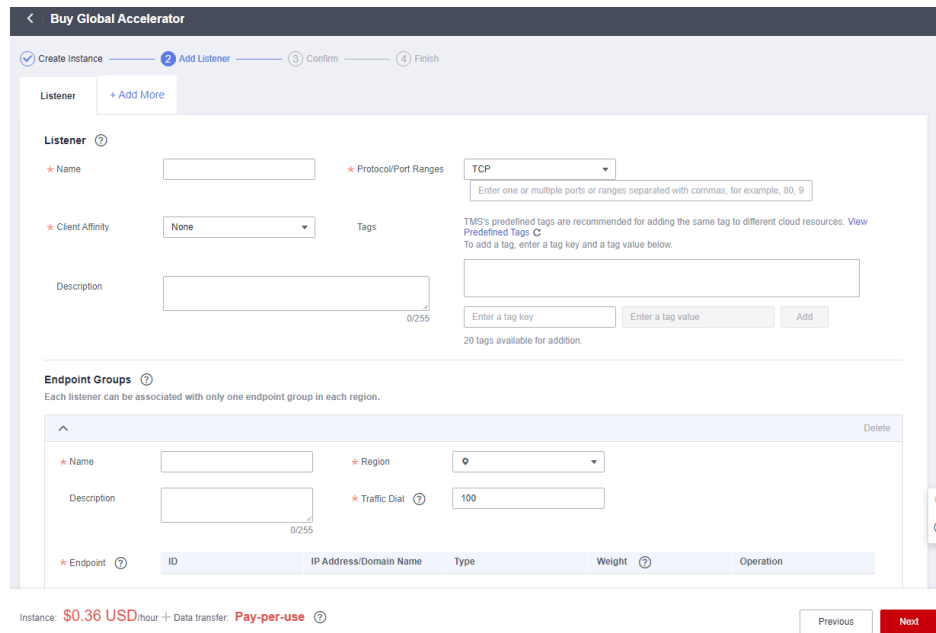


Table 1-2 Parameters required for adding a listener

Item	Parameter	Description
Basic Configuration	Name	Listener name. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
	Protocol	The protocol used by the listener to receive requests from clients. The protocol can be TCP or UDP.
	Port	The ports or port ranges used by the listener to receive requests from clients. The port number ranges from 1 to 65535. You can enter one or more ports or port ranges separated by commas (,). Example: 1-10,11-50,51,52-200

Item	Parameter	Description
	Client Affinity	<ul style="list-style-type: none"> If you select None, the listener routes requests evenly among the endpoints in the endpoint group. If you select Source IP address, the source IP address of each request is calculated using the consistent hashing algorithm to obtain a unique hash key, and all the endpoints are numbered and mapped to the hash keys. Requests from the same IP address are forwarded to the same endpoint for processing. <p>TCP and UDP listeners support only Source IP address.</p>
	Tags	The identifier of a listener. Each tag consists of a key and a value. You can add up to 20 tags to a listener.
	Description	Supplementary information about the listener. You can enter up to 255 characters.
Endpoint Groups	Name	Name of the endpoint group. Each listener can be associated with only one endpoint group in a given region. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
	Region	Region where the endpoint group is used.
	Description	Supplementary information about the endpoint group. You can enter up to 255 characters.
	Traffic Dial	<p>The percentage of traffic directed to the endpoint group.</p> <p>If you increase the traffic dial, more requests will be distributed to this endpoint group.</p> <p>If you set the traffic dial to 0, no requests will be distributed to this endpoint group.</p> <p>The weight ranges from 0 to 100.</p> <p>NOTE If a listener has multiple endpoint groups, traffic will be first distributed to the endpoint group with the lowest latency and then to other endpoint groups based on the traffic dial value you set.</p>

Item	Parameter	Description
	Endpoint	An endpoint serves as a single point of contact for clients, and Global Accelerator distributes incoming traffic across healthy endpoints.
Health Check	Health Check	Whether to enable health check. If you disable health check, requests may be forwarded to unhealthy endpoints.
	Protocol	The health check protocol can be TCP. Default value: TCP .
	Port	The port used for health check. The port number ranges from 1 to 65535.
	Advanced Settings	
	Interval (s)	The maximum time between two consecutive health checks, in seconds. The interval ranges from 1 to 60 .
	Timeout (s)	The maximum time required for waiting for a response to a health check request, in seconds. The timeout ranges from 1 to 60 .
	Maximum Retries	The maximum number of health check retries allowed. The value ranges from 1 to 10 .

6. Click **Save**.
7. Click **Next** and confirm the configuration.
8. Click **Submit**.
9. If message "Accelerator xxx created successfully" is displayed, click **Finish**.

1.3 Viewing a Global Accelerator

Scenario

You can view the basic information about a global accelerator, including its name/ID, status, IP address, listener (frontend protocol/port), billing mode, tag, description, enterprise project, and acceleration area.

You can also view the following information about the global accelerator:

- Listeners added to the global accelerator, including the listener name, status, protocol, port, and client affinity. For details, see [Viewing a Listener](#).
- Endpoint groups associated with each listener, including the protocol, health check settings, and endpoints. For details, see [Viewing an Endpoint Group](#).

- Endpoints in each endpoint group, including the status, type, IP address, health check result, and weight. For details, see [Viewing an Endpoint](#).
- Monitoring information, including monitoring dimensions and metrics. For details, see [Supported Metrics](#).
- Tag information, including tag keys and values. For details, see [Managing Global Accelerator Tags](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
The basic information about the global accelerator is displayed. To view more information, go to [4](#).
3. Click the name of the global accelerator to go to the details page and view more information.

1.4 Modifying a Global Accelerator



Scenario

You can modify the name and description of a global accelerator.

You can also modify the following information about the global accelerator:

- Listeners added to the global accelerator, including the listener name, port, client affinity, and description. For details, see [Modifying a Listener](#).
- Endpoint groups associated with each listener, including the name, traffic dial, health check settings, and description. For details, see [Modifying an Endpoint Group](#).
- Endpoints in each endpoint group, including their weights. For details, see [Modifying an Endpoint](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click  next to the name or description.
5. Modify the name or description and click .

1.5 Deleting a Global Accelerator

Scenario

If you no longer need your global accelerator and its associated listeners, endpoint groups, endpoints, and health checks, you can delete them in a few clicks.

You can delete a global accelerator if you no longer need it.

 **CAUTION**

If you delete the accelerators, their associated resources, such as listeners, endpoint groups, health checks, and endpoints, will also be deleted and cannot be restored.

Procedure

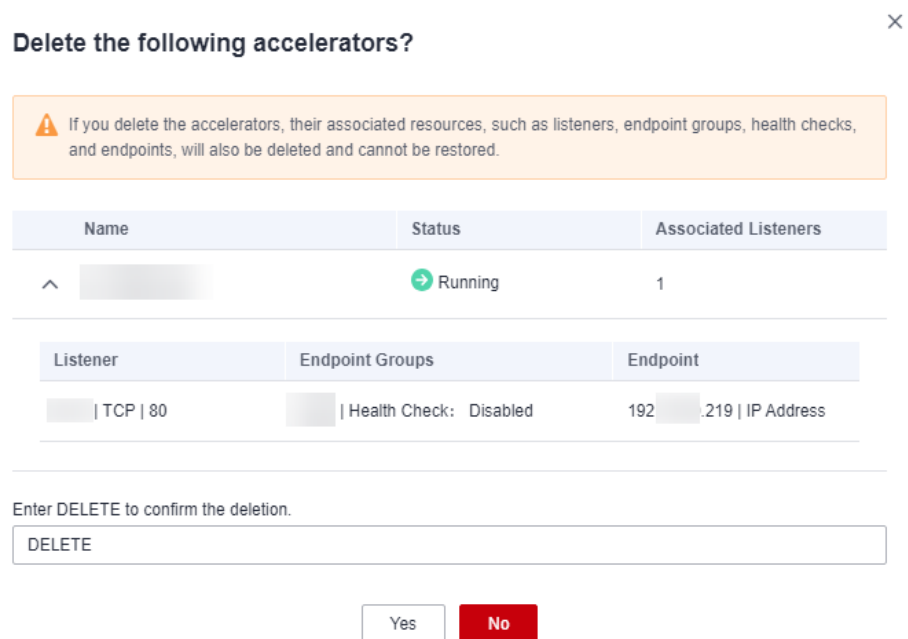
1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click **Delete** in the **Operation** column.

 **NOTE**

You can select more than one accelerator at a time and click **Delete** above the list to delete them all.

4. In the displayed dialog box, confirm the information.
Enter **DELETE** in the confirmation box as prompted.

Figure 1-4 Deleting a global accelerator



5. Click **Yes**.
6. In the displayed dialog box, click **Yes**.

1.6 Managing Global Accelerator Tags

Scenarios

After a global accelerator is created, you can view its tags or add, edit or delete a tag.

A tag is the identifier of a global accelerator and consists of a key and a value. You can add 20 tags for a global accelerator.

NOTE

If a predefined tag has been created in TMS, you can select the corresponding tag key and value.

For details about predefined tags, see [Predefined Tag Overview](#).

Adding a Tag

Add a tag to an existing global accelerator.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. On the displayed page, select the **Tags** tab.
5. Click **Add Tag**.
6. In the displayed dialog box, enter a key and a value, and click **Add**.

[Table 1-3](#) describes the tag key and value requirements.

Table 1-3 Tag key and value requirements

Parameter	Requirements
Tag key	<ul style="list-style-type: none">• Cannot be left blank.• Can contain letters, digits, spaces, underscores (_), periods (.), colons (:), equal signs (=), plus signs (+), hyphens (-), and at signs (@).• Cannot start or end with a space.• Cannot start with <code>_sys_</code>.• Can contain a maximum of 128 characters.
Tag value	<ul style="list-style-type: none">• Can be left blank.• Can contain letters, digits, spaces, underscores (_), periods (.), colons (:), slashes (/), equal signs (=), hyphens (-), and at signs (@).• Cannot start or end with a space.• Can contain a maximum of 255 characters.

7. Click **OK**.

Editing a Tag

Modify the value of a tag added to a global accelerator.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.

4. On the displayed page, select the **Tags** tab.
5. In the tag list, locate the tag you want to modify and click **Edit** in the **Operation** column.
6. Enter a new value.
7. Click **OK**.

Deleting a Tag

Delete a tag added to a global accelerator.

 **CAUTION**

Deleted tags cannot be recovered.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. On the displayed page, select the **Tags** tab.
5. In the tag list, locate the tag you want to delete and click **Delete** in the **Operation** column.
6. Click **Yes**.

2 Listeners

2.1 Listener Overview

Each global accelerator has at least one listener for listening to requests and distributing the requests to endpoints based on the client affinity and weight you set.

Protocols Supported by Listeners

Table 2-1 Protocols supported by listeners

OSI Layer	Protocol	Description	Scenarios
Layer 4	TCP	<ul style="list-style-type: none">• Source IP address-based sticky sessions• Fast data transfer	<ul style="list-style-type: none">• File transfer, email sending and receiving, remote login, and other scenarios that require high reliability and high data accuracy• Web applications that need to be robust and require high performance to process a large number of concurrent requests
Layer 4	UDP	<ul style="list-style-type: none">• Relatively low reliability• Fast data transfer	Video chats, gaming, real-time financial quotations, and other scenarios that require quick response

Listening Ports

Table 2-2 Listening ports

Protocol	Port Range	Description
TCP	1-65535	Port 22 is used by the system and is not recommended. Multiple ports or port ranges are separated by commas (,).
UDP	1-65535	Port 4789 is used by the system and is not recommended. Multiple ports or port ranges are separated by commas (,).

2.2 Adding a Listener

Scenario

You can add a listener to a global accelerator so that the listener can listen to requests and distribute them to the associated endpoints based on the client affinity and weight you set.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. Click **Add Listener**.
6. Specify the parameters listed in [Table 2-3](#).

Table 2-3 Parameters required for adding a listener

Item	Parameter	Description
Basic Configuration	Name	Listener name. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
	Protocol	The protocol used by the listener to receive requests from clients. The protocol can be TCP or UDP.

Item	Parameter	Description
	Port	The ports or port ranges used by the listener to receive requests from clients. The port number ranges from 1 to 65535. You can enter one or more ports or port ranges separated by commas (.). Example: 1-10,11-50,51,52-200
	Client Affinity	<ul style="list-style-type: none"> If you select None, the listener routes requests evenly among the endpoints in the endpoint group. If you select Source IP address, the source IP address of each request is calculated using the consistent hashing algorithm to obtain a unique hash key, and all the endpoints are numbered and mapped to the hash keys. Requests from the same IP address are forwarded to the same endpoint for processing. <p>TCP and UDP listeners support only Source IP address.</p>
	Tags	The identifier of a listener. Each tag consists of a key and a value. You can add up to 20 tags to a listener.
	Description	Supplementary information about the listener. You can enter up to 255 characters.
Endpoint Groups	Name	Name of the endpoint group. Each listener can be associated with only one endpoint group in a given region. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
	Region	Region where the endpoint group is used.
	Description	Supplementary information about the endpoint group. You can enter up to 255 characters.

Item	Parameter	Description
	Traffic Dial	<p>The percentage of traffic directed to the endpoint group.</p> <p>If you increase the traffic dial, more requests will be distributed to this endpoint group.</p> <p>If you set the traffic dial to 0, no requests will be distributed to this endpoint group.</p> <p>The weight ranges from 0 to 100.</p> <p>NOTE If a listener has multiple endpoint groups, traffic will be first distributed to the endpoint group with the lowest latency and then to other endpoint groups based on the traffic dial value you set.</p>
	Endpoint	An endpoint serves as a single point of contact for clients, and Global Accelerator distributes incoming traffic across healthy endpoints.
Health Check	Health Check	<p>Whether to enable health check.</p> <p>If you disable health check, requests may be forwarded to unhealthy endpoints.</p>
	Protocol	The health check protocol can be TCP. Default value: TCP .
	Port	The port used for health check. The port number ranges from 1 to 65535.
	Advanced Settings	
	Interval (s)	The maximum time between two consecutive health checks, in seconds. The interval ranges from 1 to 60 .
	Timeout (s)	The maximum time required for waiting for a response to a health check request, in seconds. The timeout ranges from 1 to 60 .
	Maximum Retries	The maximum number of health check retries allowed. The value ranges from 1 to 10 .

7. Click **OK**.

2.3 Access Control

Scenario

Access control allows you to whitelist certain IP addresses to allow them to access a listener or blacklist certain IP addresses to deny them to access a listener.

You can modify or disable the access control option as needed.

CAUTION

Once a whitelist is added, only IP addresses in the whitelist can access the listener. After a blacklist is added, IP addresses in the blacklist cannot access the listener.

Constraints

- Access control does not restrict the ping command. You can still ping endpoints from the blacklisted IP addresses.
If a whitelist is configured for a listener but IP addresses that are not in the whitelist can access the endpoints associated with the listener, one possible reason is that a persistent connection is established between the client and the endpoints. To deny such IP addresses from accessing the listener, the persistent connection needs to be disconnected.
- You can add up to 20 CIDR blocks at a time and 200 CIDR blocks in total to an IP address group. Each CIDR block must be unique.
- An IP address group can be configured for an access control policy of up to 10 listeners.
- Access control policies take effect only for new connections, but not for connections that have been established.

Prerequisites

If you want to use a whitelist or blacklist for access control, you must select an IP address group. If do not have an IP address group, create one by referring to [Creating an IP Address Group](#).

The IP address group must be in the **Running** state.

Configuring Access Control

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. Click the name of the target listener.
6. On the **Basic Information** page, click **Configure** on the right of **Access Control**.

7. Configure the parameters. For details, see [Table 2-4](#).

Table 2-4 Parameters for configuring access control

Parameter	Description
Access Control	If you have set Access Control to Whitelist or Blacklist , you can enable or disable access control. <ul style="list-style-type: none">• Only after you enable access control, the whitelist or blacklist takes effect.• If you disable access control, the whitelist or blacklist does not take effect.
Access Control	Specifies how access to the listener is controlled. Three options are available: <ul style="list-style-type: none">• Whitelist: Only IP addresses in the IP address group can access the listener.• Blacklist: IP addresses in the IP address group are not allowed to access the listener.
IP Address Group	CIDR blocks that are added to the whitelist or blacklist for access control.

8. Click **OK**.

2.4 Viewing a Listener

Scenario

You can view basic information about a listener, including the name, status, protocol/port, client affinity, and access control.

You can also view the following information about the listener:

- Endpoint groups associated with the listener, including the protocol, health check settings, and endpoints. For details, see [Viewing an Endpoint Group](#).
- Endpoints in each endpoint group, including the status, type, IP address, health check result, and weight. For details, see [Viewing an Endpoint](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.

3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. Click the name of the listener and view its details.

2.5 Modifying a Listener

Scenario

You can modify the name, port, client affinity, access control, and description of a listener.

You can also modify the following information about the listener:

- Endpoint groups associated with each listener, including the name, traffic dial, health check, and description. For details, see [Modifying an Endpoint Group](#).
- Endpoints in each endpoint group, including their weights. For details, see [Modifying an Endpoint](#).

Procedure


1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. Locate the listener you want to modify and click .
6. Modify the parameters. For details, see [Table 2-5](#).

Table 2-5 Listener parameters that you can modify

Parameter	Description
Name	Listener name. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
Port	The ports or port ranges used by the listener to receive requests from clients. The port number ranges from 1 to 65535. You can enter one or more ports or port ranges separated by commas (.). Example: 1-10,11-50,51,52-200

Parameter	Description
Client Affinity	<p>If you select None, the listener routes requests evenly among the endpoints in the endpoint group.</p> <p>If you select Source IP address, the source IP address of each request is calculated using the consistent hashing algorithm to obtain a unique hash key, and all the endpoints are numbered and mapped to the hash keys. Requests from the same IP address are forwarded to the same endpoint for processing.</p> <p>TCP and UDP listeners support only Source IP address.</p>
Access Control	<p>You can add IP addresses to a whitelist or blacklist to control access to a listener.</p> <p>A whitelist allows specified IP addresses to access the listener, while a blacklist denies access from specified IP addresses.</p> <p>For details, see Access Control.</p>
Description	<p>Supplementary information about the listener.</p> <p>You can enter up to 255 characters.</p>

7. Click **OK**.

2.6 Deleting a Listener

Scenario

You can delete a listener if you no longer need it.


NOTE

If the listener has an endpoint group with endpoints, remove the endpoints and delete the endpoint group before you delete the listener.

- For details about how to remove an endpoint, see [Removing an Endpoint](#).
- For details about how to delete an endpoint group, see [Deleting an Endpoint Group](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.

5. Locate the listener you want to delete and click .
6. In the displayed dialog box, click **Yes**.

2.7 Managing Listener Tags

Scenarios

After a listener is created, you can view its tags or add, edit or delete a tag.

A tag is the identifier of a listener and consists of a key and a value. You can add up to 20 tags to a listener.

NOTE

If a predefined tag has been created in TMS, you can select the corresponding tag key and value.

For details about predefined tags, see [Predefined Tag Overview](#).

Adding a Tag

Add a tag to an existing listener.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click the **Listeners** tab.
5. In the right pane of the listener page, click the **Tags** tab.
6. Click **Add Tag**.
7. In the displayed dialog box, enter a key and a value, and click **Add**.

[Table 2-6](#) describes the tag key and value requirements.

Table 2-6 Tag key and value requirements

Parameter	Requirements
Tag key	<ul style="list-style-type: none">• Cannot be left blank.• Can contain letters, digits, spaces, underscores (_), periods (.), colons (:), equal signs (=), plus signs (+), hyphens (-), and at signs (@).• Cannot start or end with a space.• Cannot start with _sys_.• Can contain a maximum of 128 characters.

Parameter	Requirements
Tag value	<ul style="list-style-type: none">• Can be left blank.• Can contain letters, digits, spaces, underscores (_), periods (.), colons (:), slashes (/), equal signs (=), hyphens (-), and at signs (@).• Cannot start or end with a space.• Can contain a maximum of 255 characters.

8. Click **OK**.

Editing a Tag

Modify the value of a tag added to a listener.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click the **Listeners** tab.
5. In the right pane of the listener page, click the **Tags** tab.
6. In the tag list, locate the tag you want to modify and click **Edit** in the **Operation** column.
7. Enter a new value.
8. Click **OK**.

Deleting a Tag

Delete a tag added to a listener.

⚠ CAUTION

Deleted tags cannot be recovered.

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click the **Listeners** tab.
5. In the right pane of the listener page, click the **Tags** tab.
1. In the tag list, locate the tag you want to delete and click **Delete** in the **Operation** column.
2. Click **Yes**.

3 Endpoint Groups

3.1 Endpoint Group Overview

An endpoint group includes one or more endpoints in a given region. You can set a weight for each endpoint group, and Global Accelerator will route requests based on the weights you specified.

You need to associate an endpoint group with each listener, which will route traffic to the endpoints in the associated endpoint group.

For details about the endpoint types supported by Global Accelerator, see [Endpoint Overview](#).

3.2 Adding an Endpoint Group

Scenario

A listener needs to have at least one endpoint group where it can route traffic to.

Prerequisites

A global accelerator has been created, and a listener has been added to it.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. On the right of the page, click **Endpoint Groups**.
6. Click **Add Endpoint Group**.
7. Specify the parameters listed in [Table 3-1](#).

Table 3-1 Parameters required for adding an endpoint group

Parameter	Description
Name	Name of the endpoint group. Each listener can be associated with only one endpoint group in a given region. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
Region	Region where the endpoint group is used.
Description	Supplementary information about the endpoint group. You can enter up to 255 characters.
Traffic Dial	The percentage of traffic directed to the endpoint group. If you increase the traffic dial, more requests will be distributed to this endpoint group. If you set the traffic dial to 0, no requests will be distributed to this endpoint group. The weight ranges from 0 to 100. NOTE If a listener has multiple endpoint groups, traffic will be first distributed to the endpoint group with the lowest latency and then to other endpoint groups based on the traffic dial value you set.
Health Check	
Health Check	Whether to enable health check. If you disable health check, requests may be forwarded to unhealthy endpoints.
Protocol	The health check protocol can be TCP. Default value: TCP .
Port	The port used for health check. The port number ranges from 1 to 65535.
Interval (s)	The maximum time between two consecutive health checks, in seconds. The interval ranges from 1 to 60 .
Timeout (s)	The maximum time required for waiting for a response to a health check request, in seconds. The timeout ranges from 1 to 60 .
Maximum Retries	The maximum number of health check retries allowed. The value ranges from 1 to 10 .

8. Click **OK**.

3.3 Viewing an Endpoint Group

Scenario

You can view basic information about an endpoint group, including the protocol, health check settings, and endpoints.

To view details about the endpoints in the endpoint group, including their status, type, IP address, health check result, and weight, see [Viewing an Endpoint](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. Click the name of the endpoint group to view its basic information.

3.4 Modifying an Endpoint Group

Scenario

You can modify the name, traffic dial, and description of an endpoint group.

You can also modify the following information about the endpoint group:

- Health check settings of the endpoint group, including whether the health check is enabled, protocol/port, interval, timeout, and maximum number of retries. For details about how to modify the health check settings, see [Modifying Health Check Settings](#).
- Endpoints in each endpoint group, including their weight. For details, see [Modifying an Endpoint](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Listeners**.
5. On the right of the page, click **Endpoint Groups**.
6. Locate the endpoint group you want to modify and click **Modify** in the **Operation** column.
7. Modify the parameters. For details, see [Modifying an Endpoint Group](#).

Table 3-2 Endpoint group parameters that you can modify

Parameter	Description
Name	Name of the endpoint group. Each listener can be associated with only one endpoint group in a given region. You can enter up to 64 characters. Only letters, digits, and hyphens are allowed.
Description	Supplementary information about the endpoint group. You can enter up to 255 characters.
Traffic Dial	The percentage of traffic directed to the endpoint group. If you increase the traffic dial, more requests will be distributed to this endpoint group. If you set the traffic dial to 0, no requests will be distributed to this endpoint group. The weight ranges from 0 to 100. NOTE If a listener has multiple endpoint groups, traffic will be first distributed to the endpoint group with the lowest latency and then to other endpoint groups based on the traffic dial value you set.

8. Click **OK**.

3.5 Deleting an Endpoint Group

Scenario

You can delete an endpoint group if you no longer need it.

After the endpoint group is deleted, the global accelerator will no longer forwards requests to it.


NOTE

If the endpoint group has endpoints or health check configured, you need to remove the endpoints and delete the health check before deleting the endpoint group.

- For details about how to remove an endpoint, see [Removing an Endpoint](#).
- For details about how to delete a health check, see [Deleting Health Check Settings](#).

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.

5. Locate the endpoint group you want to delete and click .
6. In the displayed dialog box, click **Yes**.

4 Endpoints

4.1 Endpoint Overview

An endpoint is a destination to which requests are routed, and up to 10 endpoints can be added to each endpoint group.

You can add an EIP as an endpoint.

If there are multiple endpoints in an endpoint group, you can set a weight for each endpoint to specify the proportion of requests to distribute to each endpoint. The global accelerator adds up the weights of all endpoints in the endpoint group and routes requests to each endpoint based on the ratio of its weight to the total weights.

4.2 Adding an Endpoint

Scenario

You can add endpoints to an endpoint group as needed to receive requests from the associated listener.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. On the right of the page, click **Add Endpoint**.
6. Select an endpoint, set a weight, and click **OK**.

4.3 Viewing an Endpoint

Scenario

You can view details about an endpoint, including the status, type, IP address, health check result, and weight.

Procedure


1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. On the right of the page, view the endpoint information.

4.4 Modifying an Endpoint

Scenario

You can modify the weight assigned to an endpoint.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. On the right of the page, locate the endpoint and click  in the **Weight** column.
6. Set a new weight for the endpoint.

 **CAUTION**

If the weight is set to 0, the global accelerator will stop distributing traffic to the endpoint.

7. Click **OK**.

4.5 Removing an Endpoint

Scenario

You can remove an endpoint if you no longer need it.

After the endpoint is removed, the global accelerator will no longer forwards requests to it.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. On the right of the page, locate the endpoint and click **Remove** in the **Operation** column.

NOTE

You can select more than one endpoint at a time and click **Remove** above the list to delete them all.

6. In the displayed dialog box, click **Yes**.

5 Health Checks

5.1 Health Check Overview

Global Accelerator provides health check to monitor the health of endpoints to help improve service reliability and availability.

After you enable health check, the global accelerator periodically sends requests to endpoints to check their status. If any endpoints become unavailable, the global accelerator stops sending requests to these endpoints. After the endpoints recover from failure, the global accelerator continue to route requests to them.

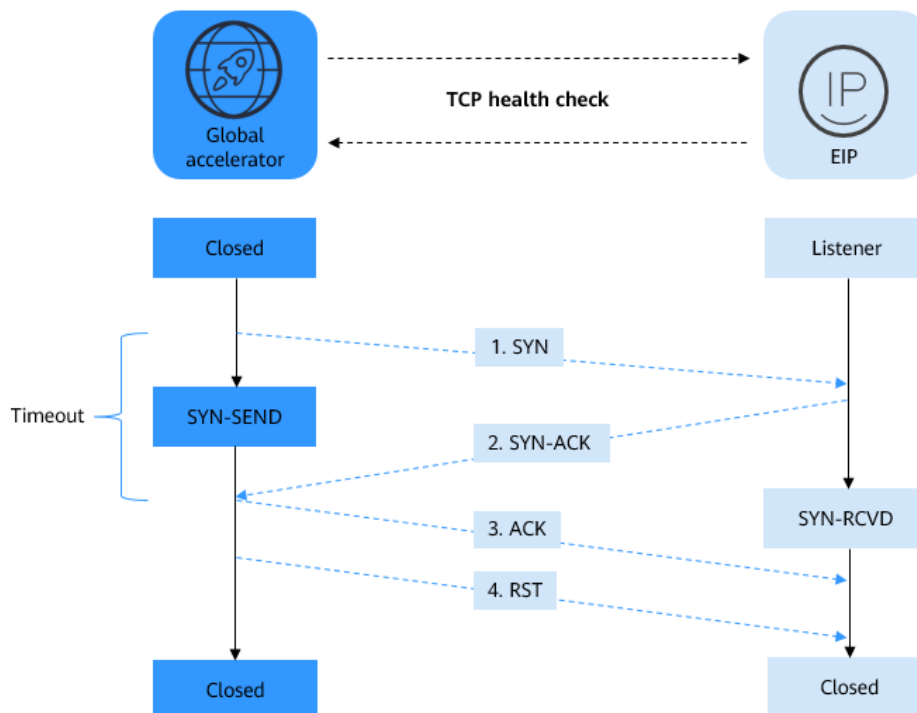
TCP can be used for health checks.

TCP Health Check

TCP health check is performed on the network layer through three-way handshakes.

[Figure 5-1](#) shows the TCP health check process.

Figure 5-1 TCP health check



The TCP health check process is as follows:

1. The global accelerator sends a TCP SYN packet to the endpoint.
2. The endpoint returns an SYN-ACK packet.
 - If the global accelerator does not receive the SYN-ACK packet within the timeout duration, it declares that the endpoint is unhealthy and sends an RST packet to the endpoint to terminate the TCP connection.
 - If the global accelerator receives the SYN-ACK packet from the endpoint within the timeout duration, it declares that the endpoint is healthy and sends an ACK packet and an RST packet to the endpoint to terminate the TCP connection.

Health Check Time Window

Health check helps ensure service availability. To avoid frequent health checks on endpoints, you can disable health check after several consecutive health checks that declare endpoints healthy or unhealthy.

The time required for declaring endpoints healthy or unhealthy is determined by the following factors:

- **Interval:** how often health checks are performed.
- **Timeout:** how long the global accelerator waits for the response from the endpoint.
- **Maximum Retries:** the maximum number of consecutive health checks after which an endpoint is declared healthy.

Endpoints can be declared unhealthy after three consecutive health checks that detect the endpoints are unhealthy, regardless of the value set for **Maximum Retries**.

The following is a formula for you to calculate the time:

- Time required for declaring endpoints healthy = Timeout x Maximum retries + Interval x (Maximum retries - 1)
- Time required for declaring endpoints unhealthy = Timeout x 3 + Interval x (3 - 1)

For example, if the interval is set to 4s and the timeout is set to 2s, the time required for declaring endpoints unhealthy is $2 \times 3 + 4 \times (3 - 1) = 14s$

5.2 Configuring a Health Check

Scenario

You can configure a health check to detect unhealthy endpoints and ensure service reliability and availability.

Constraints

- For UDP listeners, ensure that the security group rules configured for each endpoint allow ICMP traffic over the health check port.
- Ensure that the security group rules configured for each endpoint allow traffic from given CIDR blocks over the health check port.

For details, see [Table 5-1](#).

Table 5-1 CIDR blocks that need to be allowed

Endpoint Type	Endpoint Group Region	CIDR Blocks to Be Allowed
EIP Custom EIP	N/A	122.9.234.0/23 116.196.216.0/23 124.71.242.0/23
ECS ELB IP address Custom domain name	EU-Dublin	101.46.32.0/25 101.46.33.0/25

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.

4. Click **Endpoint Groups**.
5. Click the name of the endpoint group.
6. In the basic information area, click **Configure** next to **Health Check**.
7. Specify the parameters listed in [Table 5-2](#).

Table 5-2 Parameters required for configuring a health check

Parameter	Description
Health Check	Whether to enable health check. If you disable health check, requests may be forwarded to unhealthy endpoints.
Protocol	The health check protocol can be TCP. Default value: TCP .
Port	The port used for health check. The port number ranges from 1 to 65535.
Advanced Settings	
Interval (s)	The maximum time between two consecutive health checks, in seconds. The interval ranges from 1 to 60 .
Timeout (s)	The maximum time required for waiting for a response to a health check request, in seconds. The timeout ranges from 1 to 60 .
Maximum Retries	The maximum number of health check retries allowed. The value ranges from 1 to 10 .

8. Click **OK**.

5.3 Viewing Health Check Settings

Scenario

You can view the health check settings and the health check result of each endpoint.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. Click the name of the endpoint group whose health check settings you want to view.

6. In the basic information area, click **Configure** next to **Health Check** to view the health check configuration.

In the endpoint list, view the health check result of each endpoint.

5.4 Modifying Health Check Settings

Scenario

You can modify the health check configuration of an endpoint group.

Procedure

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. Click the name of the endpoint group whose health check configuration you want to modify.
6. In the basic information area, click **Configure** next to **Health Check**.
7. Configure the parameters. For details, see [Table 5-3](#).

Table 5-3 Health check parameters that you can modify

Parameter	Description
Health Check	Whether to enable health check. If you disable health check, requests may be forwarded to unhealthy endpoints.
Protocol	The health check protocol can be TCP. Default value: TCP .
Port	The port used for health check. The port number ranges from 1 to 65535.
Advanced Settings	
Interval (s)	The maximum time between two consecutive health checks, in seconds. The interval ranges from 1 to 60 .
Timeout (s)	The maximum time required for waiting for a response to a health check request, in seconds. The timeout ranges from 1 to 60 .
Maximum Retries	The maximum number of health check retries allowed. The value ranges from 1 to 10 .

8. Click **OK**.

5.5 Disabling Health Check

Scenario

You can disable health check or delete health check settings for an endpoint group.

After health check is disabled or health check settings are deleted, the health of the endpoints in the endpoint group will not be checked and the endpoints are always considered healthy. The global accelerator will still forward requests to the endpoints even if they are unhealthy. As a result, services will become unavailable.

To ensure service continuity, do not disable the health check or delete health check settings.

Disabling Health Check

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. Click the name of the endpoint group whose health check you want to disable.
6. In the basic information area, click **Configure** next to **Health Check**.
7. Disable health check.
8. Click **OK**.

NOTE

Disabling health check does not change the health check settings. You can enable health check with the same settings when you need the health check later.

Deleting Health Check Settings

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Endpoint Groups**.
5. Click the name of the endpoint group whose health check you want to delete.
6. In the basic information area, click **Delete** next to **Health Check**.
7. Click **Yes**.

NOTE

After the health check is deleted, the global accelerator will stop checking the health of the endpoints. To check their health again, you need to configure another health check.

6 IP Address Groups

6.1 IP Address Group Overview

An IP address group is a collection of IP addresses. You can use IP address groups to manage IP addresses with the same security requirements or whose security requirements change frequently.

You can configure a [whitelist or blacklist](#) to allow or deny accesses from IP addresses in an IP address group to listeners.

You can add IPv4 or IPv6 CIDR blocks to an IP address group and associate the IP address group with a maximum of 10 listeners.

6.2 Creating an IP Address Group

Scenario

If you want to use a [whitelist or blacklist](#) to control access to listeners, you must create an IP address group first.

Constraints

- Each user can create up to 50 IP address groups.
- An IP address group can have up to 200 unique CIDR blocks.
- Up to 20 unique CIDR blocks can be added to an IP address group at a time.
- An IP address group can be associated with up to 10 listeners.

Procedure

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, click **Create IP Address Group**.
3. Configure the parameters. For details, see [Table 6-1](#).

Table 6-1 Parameters required for creating an IP address group

Parameter	Description
Name	The name of the IP address group.
CIDR Blocks	CIDR blocks that are added to the whitelist or blacklist for access control. <ul style="list-style-type: none">• Each IP address or CIDR block must be on a separate line and end with a carriage return.• Each line can include a description with a vertical bar () separating it from the IP address or CIDR block, for example, 192.168.1.0/24 GROUP01. The description can contain up to 255 characters long and cannot contain angle brackets (<>).• You can add a maximum of 20 CIDR blocks at a time and 200 CIDR blocks in total.
Description	Provides supplementary information about the IP address group.

4. Click **OK**.

6.3 Viewing an IP Address Group

Scenario

You can view the details of an IP address group on the Global Accelerator console, including its name/ID, status, description, creation time, CIDR blocks, associated listeners, listener protocol/port, and access control policies.

Procedure

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, search for the target IP address group by name or ID.
3. Click the name of the IP address group to view the following information:
 - Basic information: name/ID, status, description, and creation time
 - CIDR blocks: IP addresses and description
 - Associated listeners: name, listener protocol/port, and access control policies

6.4 Modifying an IP Address Group

Scenario

You can modify the name and description of an IP address group, add a new CIDR block, or delete an existing CIDR block.

 **NOTE**

Once deleted, the CIDR block will no longer be restricted by the access control policy.

Constraints

Before **adding a CIDR block to** or **deleting a CIDR block from** an IP address group, ensure that the IP address group is in the **Running** state. If the IP address group has been associated with a listener, the listener must also be in the **Running** state.

Modifying the Basic Information

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, search for the target IP address group by name or ID.
3. Locate the target IP address group and click **Modify** in the **Operation** column.
4. Configure related parameters as prompted.

Table 6-2 IP address group parameters that you can modify

Parameter	Description
Name	The name of the IP address group.
Description	Provides supplementary information about the IP address group.

5. Click **OK**.

Adding a CIDR Block

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, search for the target IP address group by name or ID.
3. Click the name of the target IP address group.
The CIDR blocks are shown on the displayed page.
4. Click **Add CIDR Block**.
5. Add CIDR blocks as prompted.
6. Click **OK**.

Deleting a CIDR Block

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, search for the target IP address group by name or ID.
3. Click the name of the target IP address group.
The CIDR blocks are shown on the displayed page.
4. Locate the target CIDR block, click **Delete** in the **Operation** column.
5. Confirm the information about the CIDR block and click **OK**.

 **NOTE**

Once deleted, the CIDR block will no longer be restricted by the access control policy.

6.5 Deleting an IP Address Group

Scenario

If an IP address group is no longer used for access control, you can delete it on the management console.

Constraints

Before deleting an IP address group, you need to disassociate the IP address group from its listener and ensure that the IP address group is in the **Running** state.

For details about how to disassociate the IP address group from its listener, see [Configuring Access Control](#).

Procedure

1. Go to the [IP Address Groups](#) page.
2. On the displayed page, search for the target IP address group by name or ID.
3. Locate the target IP address group and click **Delete** in the **Operation** column.
4. Confirm the information and click **Yes**.

7 Cross-Border Permits

7.1 Cross-Border Permits Application

Scenario

In accordance with the laws and administrative regulations of the Ministry of Industry and Information Technology (MIIT) of the People's Republic of China, only China Mobile, China Telecom, and China Unicom are allowed for cross-border network communications, and a cross-border permit is required if you carry out business activities outside the Chinese mainland.

To comply with laws and regulations on cross-border network communications, you need to apply for a cross-border permit in the following scenarios:

The acceleration areas are in Europe, but the endpoints are running inside the Chinese mainland.

Procedure

1. Log in to the [Cross-border Permits](#) page.
2. Click **Request a Cross-border Permit**.
The **Cross-Border Service Application System** page is displayed.
3. On the application page, set related parameters and upload related materials.

Table 7-1 Online cross-border permit application

Parameter	Description
Applicant Name	The applicant name, which must be the same as the company name in the <i>Letter of Commitment to Information Security</i> .

Parameter	Description
Huawei Cloud UID	The account ID to log in to the management console. You can take the following steps to obtain your account ID. <ol style="list-style-type: none"> 1. Log in to the management console. 2. Move your cursor over the username in the upper right corner and select My Credentials from the drop-down list. 3. On the API Credentials page, view the Account ID.
Bandwidth(M)	The bandwidth size, which must be the same as the bandwidth in the <i>Letter of Commitment to Information Security</i> . The information is for reference only and does not affect the actual service bandwidth.
Start Date	For reference only.
Termination Date	For reference only.
Customer Type	The customer type. Select a type as required.
Country of the Customer	Country where the applicant is located.
Contact Name	-
Contact Number	-
Type of ID	-
ID Number	-
Scope of Business	Briefly describe the main business.
Number of Employees	For reference only.
Branch Location Country	Country where the applicant branch is located. Set this parameter as required.

Table 7-2 Required materials

Material	Signature	Seal	Description
A scanned copy of your company's business license	-	√	See the template Huawei Cloud provides for the position of the seal.

Material	Signature	Seal	Description
A scanned copy of <i>Huawei Cloud Cross-Border Circuit Service Agreement</i>	√	√	<ul style="list-style-type: none">• Sign the material on the signature block.• Stamp the seal over the signature.
A scanned copy of <i>China Unicom Letter of Commitment to Information Security of the Cross-Border Circuit Service</i>	√	√	<ul style="list-style-type: none">• Sign the material on the signature block.• Stamp the seal over the signature.• Specify the bandwidth you estimated and your company name.

4. Click **Submit**.

7.2 Application Progress Enquiry

Scenario

You can query the progress after you submit the application for cross-border permit.

Procedure

1. Log in to the [Cross-border Permits](#) page.
2. Click **Request a Cross-border Permit**.
The **Cross-Border Service Application System** page is displayed.
3. On the application page, click **Application Progress Enquiry** in the upper right corner.
4. On the **Self-inquiry System** page, enter the **Huawei Cloud UID** and **Contact Number**, and click **Query**.

8 Cloud Eye Monitoring

8.1 Overview

Monitoring is key to ensuring the performance, reliability, and availability of Global Accelerator. You can use Cloud Eye to monitor the Global Accelerator status and resource usage on a single pane of glass. You can also configure Cloud Eye to alert you of any potential issues in Global Accelerator in real time.

You can learn more about Global Accelerator monitoring by exploring the following topics:

- [Monitoring Metrics](#)
- [Setting an Alarm Rule](#)
- [Viewing Metrics](#)

8.2 Supported Metrics

Description

This topic describes the Global Accelerator namespace and the metrics to report to Cloud Eye. You can use the management console or call APIs provided by Cloud Eye to query the metrics and alarms for Global Accelerator.

Namespace

SYS.GA

Metrics

Table 8-1 Global Accelerator metrics

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval
m1_cps	Top Concurrent Connections	Largest number of connections processed by the monitored object Unit: Count	≥ 0	<ul style="list-style-type: none"> Global accelerator Listener Access region Destination cloud region Access point Destination geographic region 	60s
m2_act_conn	Active Connections	Number of active connections processed by the monitored object Unit: Count	≥ 0		60s
m3_inact_conn	Inactive Connections	Number of inactive connections processed by the monitored object Unit: Count	≥ 0		60s
m4_ncps	New Connections	Number of new connections processed by the monitored object per second Unit: Count/s	≥ 0		60s
m5_in_pps	Incoming Packets	Number of incoming data packets to the monitored object per second Unit: Count/s	≥ 0		60s

ID	Name	Description	Value Range	Monitored Object	Monitoring Interval
m6_out_pps	Outgoing Packets	Number of outgoing data packets from the monitored object per second Unit: Count/s	≥ 0		60s
m7_in_Bps	Inbound Rate	Incoming traffic per second to the monitored object Unit: byte/s	≥ 0		60s
m8_out_Bps	Outbound Rate	Outgoing traffic per second from the monitored object Unit: byte/s	≥ 0		60s
m9_in_bandwidth	Inbound Bandwidth	Bandwidth used for accessing the monitored object from the Internet Unit: bit/s	≥ 0		60s
ma_out_bandwidth	Outbound Bandwidth	Bandwidth used by the monitored object to access the Internet Unit: bit/s	≥ 0		60s
mb_in_Bytes	Inbound Traffic	Network traffic going into the cloud platform Unit: byte	≥ 0		60s
mc_out_Bytes	Outbound Traffic	Network traffic going out of the cloud platform Unit: byte	≥ 0		60s

Dimensions


Key	Value
ga_accelerator_id	ID of a global accelerator
ga_listener_id	ID of the listener added to a global accelerator
ga_source_pop	The access point of a global accelerator
ga_source_area	The access region where a global accelerator is used for fast access
ga_destination_region	The destination cloud region where a global accelerator is used for fast access
ga_destination_area	The destination geographic region where a global accelerator is used for fast access
ga_listener_region	The listener and destination cloud region
ga_pop_listener	The listener and access point
ga_pop_region	The access point and destination cloud region
ga_pop_listener_region	The listener, access point, and destination cloud region
ga_source_destination_area	The access region and destination geographic region
ga_outbound_region	The egress region where a global accelerator is used for fast access

8.3 Setting an Alarm Rule

Scenario

You can set alarm rules with customized metrics and notification policies to keep track of your Global Accelerator in real time.

Procedure

1. Log in to the management console.
2. Click  on the upper left corner to display **Service List** and choose **Management & Governance > Cloud Eye**.
3. In the navigation pane, choose **Alarm Management > Alarm Rules**.
4. On the **Alarm Rules** page, click **Create Alarm Rule** to create one. You can also modify an existing alarm rule.

5. After configuring the parameters, click **Create**.
After the alarm rule is created, the system automatically notifies you if an alarm is triggered for your global accelerator.

 **NOTE**


For more information about Global Accelerator alarm rules, see [Cloud Eye User Guide](#).

8.4 Viewing Metrics

You can view the monitoring metrics for global accelerators on either the Global Accelerator console or the Cloud Eye console.

You can view data from the last 1, 3, 12, or 24 hours or the last 7 days.

On the Cloud Eye Console

1. Log in to the management console.
2. Click  on the upper left corner to display **Service List** and choose **Management & Governance > Cloud Eye**.
3. In the navigation pane on the left, choose **Cloud Service Monitoring > Global Accelerator**.
4. Locate the row containing the target global accelerator, click **View Metric** in the **Operation** column.
You can view data of the last 1, 3, 12, or 24 hours, or last 7 days. You can also specify a time period.

On the Global Accelerator Console

1. Log in to the [Global Accelerator console](#).
2. Search for the global accelerator by name or ID.
3. Click the name of the global accelerator to go to the details page.
4. Click **Monitoring**.
5. On the **Monitoring** tab page, choose **Period** and **Time Range** as you need to view metrics of the global accelerator.

9 Using CTS to Collect Global Accelerator Key Operations

9.1 Key Operations Recorded by CTS

Scenario

With Cloud Trace Service (CTS), you can record operations associated with Global Accelerator for later query, audit, and backtracking.

Constraints

You have enabled CTS.

Key Operations Recorded by CTS

Table 9-1 Global Accelerator operations recorded by CTS

Operation	Resource Type	Trace
Creating a global accelerator	accelerator	createAccelerator
Updating a global accelerator	accelerator	updateAccelerator
Deleting a global accelerator	accelerator	deleteAccelerator
Adding a listener	listener	createListener
Modifying a listener	listener	updateListener
Deleting a listener	listener	deleteListener
Adding an endpoint group	endpointGroup	createEndpointGroup

Operation	Resource Type	Trace
Modifying an endpoint group	endpointGroup	updateEndpointGroup
Deleting an endpoint group	endpointGroup	deleteEndpointGroup
Configuring a health check	healthCheck	createHealthCheck
Updating a health check	healthCheck	updateHealthCheck
Deleting health check settings	healthCheck	deleteHealthCheck
Adding an endpoint	endpoint	createEndpoint
Updating an endpoint	endpoint	updateEndpoint
Removing an endpoint	endpoint	deleteEndpoint



9.2 Viewing Traces

Scenarios

After CTS is enabled, CTS starts recording operations on cloud resources. The CTS management console stores the last seven days of operation records.

This section describes how to query or export the last seven days of operation records on the management console.

Procedure

1. Log in to the management console.
2. Click  in the upper left corner and select the desired region and project.
3. In the upper left corner of the page, click  to go to the service list. Under **Management & Deployment**, click **Cloud Trace Service**.
4. In the navigation pane on the left, choose **Trace List**.
5. Specify filters as needed. The following filters are available:
 - **Trace Type**: Set it to **Management** or **Data**.
 - **Trace Source**, **Resource Type**, and **Search By**
Select filters from the drop-down list.
If you select **Trace name** for **Search By**, select a trace name.
If you select **Resource ID** for **Search By**, select or enter a resource ID.

If you select **Resource name** for **Search By**, select or enter a resource name.

- **Operator:** Select a specific operator (a user other than an account).
 - **Trace Status:** Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
 - **Search time range:** In the upper right corner, choose **Last 1 hour**, **Last 1 day**, or **Last 1 week**, or specify a custom time range.
6. Click arrow on the left of the required trace to expand its details.
 7. Locate the required trace and click **View Trace** in the **Operation** column.
A dialog box is displayed, showing the trace content.

10 Permissions Management

10.1 Creating a User and Granting Permissions

Use [IAM](#) to implement fine-grained permissions control for your Global Accelerator resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing Global Accelerator resources.
- Grant only the permissions required for users to perform a specific task.
- Delegate another HUAWEI ID or cloud service to perform professional and efficient O&M on your Global Accelerator resources.

Skip this part if your HUAWEI ID does not require individual IAM users.

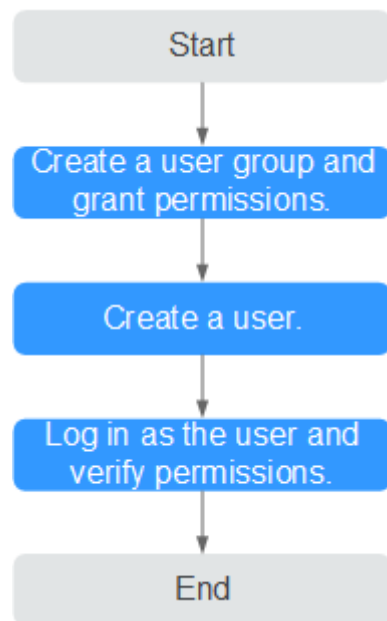
[Figure 10-1](#) shows the process for granting permissions.

Prerequisites

You have learned about Global Accelerator permissions that can be assigned to the user group and select permissions based on actual requirements. For details about the system permissions of Global Accelerator, see [Permissions](#). For the system policies of other services, see [System Permissions](#).


Process Flow

Figure 10-1 Process for granting permissions



1. On the IAM console, **create a user group and grant it permissions** (GA **FullAccess** as an example).
2. **Create an IAM user and add it to the user group.**
3. **Log in as the IAM user** and verify permissions.

In the authorized region, perform the following operations:

- Click  on the upper left corner to display **Service List** and choose **Networking > Global Accelerator**. Click **Buy Global Accelerator** in the upper right corner. If the global accelerator is created, the **GA FullAccess** policy is in effect.
- Choose another service in **Service List**. If a message appears indicating that you do not have permissions to access the service, the **GA FullAccess** policy is in effect.

10.2 Custom Policy

Custom policies can be created to supplement the system-defined policies of Global Accelerator.

You can create custom policies in either of the following ways:

- Visual editor: Select cloud services, actions, resources, and request conditions. You do not need to have knowledge of the policy syntax.
- JSON: Create a policy in the JSON format from scratch or based on an existing policy template.

For details, see **Creating a Custom Policy**. The following are examples of custom policies created for Global Accelerator.

Example Custom Policies

- Example 1: Allowing users to update a global accelerator

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ga:accelerator:update"
      ]
    }
  ]
}
```

- Example 2: Denying users to delete a global accelerator

A deny policy must be used in conjunction with other policies to take effect. If the permissions granted to an IAM user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

If you grant the system policy **GA FullAccess** to a user but do not want the user to have the permission to delete global accelerators, you can create a custom policy that denies the deletion of global accelerators. Then you can grant the **GA FullAccess** and deny policies to the user, so that the user can perform all operations on global accelerators except deleting them.

The following is an example deny policy:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Deny",
      "Action": [
        "ga:accelerator:delete"
      ]
    }
  ]
}
```

- Example 3: Defining actions for multiple services in a policy

A custom policy can contain the actions of multiple services that are of the global or project-level type.

The following is an example policy containing actions of multiple services:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ga:listener:create",
        "ga:healthcheck:create",
        "ga:endpointgroup:create",
        "ga:endpoint:create"
      ]
    },
    {
      "Effect": "Allow",
      "Action": [
        "eps:enterpriseProjects:enable",
        "eps:enterpriseProjects:update",
        "eps:enterpriseProjects:create"
      ]
    }
  ]
}
```

11 Appendix

11.1 Configuring the TOA Module

Scenario

Global Accelerator provides customized strategies for managing service access. Before these strategies can be customized, the clients' IP addresses contained in the requests are required. The TCP Option Address (TOA) kernel module is used to obtain the IP addresses of clients. It is installed on the server of the endpoint.

This section describes how you can compile the module in the OS if you use TCP to distribute IPv4 traffic.

The operations for Linux OSs with kernel version of 2.6.32 are different from those for Linux OSs with kernel version of 3.0 or later.

NOTE

- The TOA module cannot be used for UDP listeners.
- The TOA module can work properly in the following OSs, and the methods for installing other kernel versions are similar:
 - CentOS 6.8 (kernel version 2.6.32)
 - SUSE 11 SP3 (kernel version 3.0.76)
 - CentOS 7 or CentOS 7.2 (kernel version 3.10.0)
 - Ubuntu 16.04.3 (kernel version 4.4.0)
 - Ubuntu 18.04 (kernel version 4.15.0)
 - OpenSUSE 42.2 (kernel version 4.4.36)
 - Debian 8.2.0 (kernel version 3.16.0)

Constraints

- The development environment for compiling the module must be the same as that of the current kernel. For example, if the kernel version is kernel-3.10.0-693.11.1.el7, the kernel development package version must be kernel-devel-3.10.0-693.11.1.el7.
- The OS repositories are accessible to servers.

- Users other than **root** must have sudo permissions.

Procedure

- Linux kernel version 3.0 or later
1. Prepare the compilation environment.

NOTE

- During the installation, download the required module development package from the Internet if it cannot be found in the source.
- If the kernel development package (kernel-devel) cannot be obtained, contact the image provider.

The following are operations for compiling the module in different Linux OSs.

– CentOS

- i. Install the gcc compiler.

sudo yum install gcc

- ii. Install the make tool.

sudo yum install make

- iii. Install the module development package (the package header and module library must have the same version as the kernel).

sudo yum install kernel-devel-`uname -r`

NOTE

- Download the required module development package from the following address if it cannot be found in the source:
https://mirror.netcologne.de/oracle-linux-repos/ol7_latest/getPackage/
For example, run the following command to install 3.10.0-693.11.1.el7.x86_64:
rpm -ivh kernel-devel-3.10.0-693.11.1.el7.x86_64.rpm
- If the kernel development package (kernel-devel) cannot be obtained, contact the image provider.

– Ubuntu and Debian

- i. Install the gcc compiler.

sudo apt-get install gcc

- ii. Install the make tool.

sudo apt-get install make

- iii. Install the module development package (the package header and module library must have the same version as the kernel).

sudo apt-get install linux-headers-`uname -r`

– SUSE

- i. Install the gcc compiler.

sudo zypper install gcc

- ii. Install the make tool.

sudo zypper install make

- iii. Install the module development package (the package header and module library must have the same version as the kernel).

sudo zypper install kernel-default-devel

2. Compile the module.

- a. Enter the source code directory and compile the module.

cd src**make**

If no warning or error information is prompted, the compilation is successful. Verify that the **toa.ko** file has generated in the current directory.

 **NOTE**

If error message "config_retpoline=y but not supported by the compiler, Compiler update recommended" is displayed, the GCC version is too old. Upgrade the GCC to a later version.

3. Load the kernel module.

- a. Run the following command to load the kernel module:

sudo insmod toa.ko

- b. Check the module loading and view the kernel output information.

dmesg | grep TOA

If "TOA: toa loaded" is displayed in the command output, the module has been loaded.

 **NOTE**

After the CoreOS module is compiled in the container, copy it to the host system and then load it. The container for compiling the module shares the **/lib/modules** directory with the host system, so you can copy the module to this directory, allowing the host system to use it.

4. Set the script to enable the system to automatically load the module.

To make the module take effect when the system starts, add the command for loading the module to your startup script.

You can use either of the following methods to enable the module to automatically load:

- Add the command for the module to automatically load to the startup script as required.
- Perform the following operations to configure the startup script:
 - i. Create the **toa.modules** file in the **/etc/sysconfig/modules/** directory. This file contains the module loading script.

The following is an example of the content in the **toa.modules** file.

#!/bin/sh**/sbin/modinfo -F filename /root/toa/toa.ko > /dev/null 2>&1****if [\$? -eq 0]; then****/sbin/insmod /root/toa/toa.ko****fi**

/root/toa/toa.ko is the path of the module file. You need to replace it with their actual path.

- ii. Add execution permissions for the **toa.modules** startup script.

sudo chmod +x /etc/sysconfig/modules/toa.modules **NOTE**

If the kernel is upgraded, the current module will no longer match. Compile the module again.

5. Install the module on servers.

To load the module in the same OSs, copy the **toa.ko** file to VMs where the module is to be loaded and then perform the operations in [3](#).

After the module is loaded, the IP address of a client can be obtained.

 **NOTE**

The OS version of each server must be the same as that of the kernel.

6. Verify the module.

After the module is installed, the source IP address can be directly obtained. You can perform the following operations to verify:

Start SimpleHTTPServer on the server of the endpoint where Python is installed:

```
python -m SimpleHTTPServer port
```

The value of **port** must be a port used by the server, and the default value is **80**.

Access the anycast IP address provided by Global Accelerator. Access logs on the server are as follows:

```
192.168.0.90 -- [06/Aug/2020 14:24:21] "GET / HTTP/1.1" 200 --
```

 **NOTE**

192.168.0.90 is the source IP address and also the real IP address of the client that can be obtained by the backend server.

- **In the following operations, the Linux kernel version is 2.6.32.**

 **NOTE**

The TOA module supports OSs (CentOS 6.8 image) with a kernel of 2.6.32-xx.

1. Obtain the kernel source code package
Linux-2.6.32-220.23.1.el6.x86_64.rs.src.tar.gz containing the module from the following link:
http://kb.linuxvirtualserver.org/images/3/34/Linux-2.6.32-220.23.1.el6.x86_64.rs.src.tar.gz
2. Decompress the kernel source code package.
3. Modify compilation parameters.
 - a. Open the **linux-2.6.32-220.23.1.el6.x86_64.rs** directory.
 - b. Edit the **net/toa/toa.h** file.
Change the value of **#define TCPOPT_TOA200** to **#define TCPOPT_TOA254**.
 - c. On the Shell page, run the following commands:

```
sed -i 's/CONFIG_IPV6=m/CONFIG_IPV6=y/g' .config  
echo -e '\n# toa\nCONFIG_TOA=m' >> .config
```

The IPv6 module has been compiled into the kernel. TOA is compiled into a separate module and can be independently started and stopped.

d. Edit **Makefile**.

You can add a description after **EXTRAVERSION =**. This description will be displayed in **uname -r**, for example, **-toa**.

4. Compile the software package.

make -j *n*

 **NOTE**

n indicates the number of vCPUs. For example, if there are four vCPUs, *n* must be set to 4.

5. Install the module.

make modules_install

Figure 11-1 shows the command output.

Figure 11-1 Installing the module

```
INSTALL /lib/firmware/kaweth/trigger_code_fix.bin
INSTALL /lib/firmware/ti_3410.fw
INSTALL /lib/firmware/ti_5052.fw
INSTALL /lib/firmware/mts_cdma.fw
INSTALL /lib/firmware/mts_gsm.fw
INSTALL /lib/firmware/mts_edge.fw
INSTALL /lib/firmware/edgeport/boot.fw
INSTALL /lib/firmware/edgeport/boot2.fw
INSTALL /lib/firmware/edgeport/down.fw
INSTALL /lib/firmware/edgeport/down2.fw
INSTALL /lib/firmware/edgeport/down3.bin
INSTALL /lib/firmware/whiteheat_loader.fw
INSTALL /lib/firmware/whiteheat.fw
INSTALL /lib/firmware/keyspan_pda/keyspan_pda.fw
INSTALL /lib/firmware/keyspan_pda/xircom_pgs.fw
DEPMOD 2.6.32-toa
```

6. Install the kernel.

make install

Figure 11-2 shows the command output.

Figure 11-2 Installing the kernel

```
INSTALL /lib/firmware/keyspan_pda/xircom_pgs.fw
DEPMOD 2.6.32-toa
[root@SZX1000167219 linux-2.6.32-220.23.1.el6.x86_64.rs]# make install
sh /root/humin/linux-2.6.32-220.23.1.el6.x86_64.rs/arch/x86/boot/install.sh 2.6.32-toa arch/x86/boot/bzImage \
    System.map "/boot"
ERROR: modinfo: could not find module xen_procfs
ERROR: modinfo: could not find module ipv6
ERROR: modinfo: could not find module xen_scxifront
ERROR: modinfo: could not find module xen_hcall
ERROR: modinfo: could not find module xen_balloon
[root@SZX1000167219 linux-2.6.32-220.23.1.el6.x86_64.rs]#
```

7. Open the **/boot/grub/grub.conf** file and configure the kernel to start up when the system starts.

- a. Change the default startup kernel from the first kernel to the zeroth kernel. To do so, change the value of **default** to **0**.
- b. Add the **nohz** parameter (set it to **off**) to the end of the line containing the **vmlinuz-2.6.32-toa** kernel. If **nohz** is not disabled, the CPU0 utilization may be high and overload the kernel.

Figure 11-3 Configuration file

```
default=1
timeout=5
splashimage=(hd0,1)/boot/grub/splash.xpm.gz
hiddenmenu
title Red Hat Enterprise Linux Server (2.6.32-toa)
    root (hd0,1)
    kernel /boot/vmlinuz-2.6.32-toa ro root=UUID-
et nohz=ofi
    initrd /boot/initramfs-2.6.32-toa.img
```

- c. Save the modification and exit. Restart the OS.

During the restart, the system will load the **vmlinuz-2.6.32-toa** kernel.

8. After the restart, load the module.

modprobe toa

Add the **modprobe toa** command to both the startup script and the system scheduled monitoring script.

Figure 11-4 Adding the **modprobe toa** command

```
[root@SZX1000167219 ~]# modprobe toa
[root@SZX1000167219 ~]# lsmod |grep toa
toa                4203  0
[root@SZX1000167219 ~]#
```

After the module is loaded, query the kernel information.

Figure 11-5 Querying the kernel

```
[root@SZX1000167219 ~]# uname -a
Linux SZX1000167219 2.6.32-toa #1 SMP Sat Oct 15 11:50:05 CST 2016 x86_64 x86_64 x86_64 GNU/Linux
```

9. Verify the module.

After the module is installed, the source IP address can be directly obtained. You can perform the following operations to verify:

Start SimpleHTTPServer on the server of the endpoint where Python is installed:

```
python -m SimpleHTTPServer port
```

The value of **port** must be a port used by the server, and the default value is **80**.

Access the anycast IP address provided by Global Accelerator. Access logs on the server are as follows:

```
192.168.0.90 - - [06/Aug/2020 14:24:21] "GET / HTTP/1.1" 200 -
```

NOTE

192.168.0.90 is the source IP address and also the real IP address of the client that can be obtained by the backend server.