

Global Accelerator

Service Overview

Issue 01
Date 2025-01-08



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 What Is Global Accelerator?	1
2 Why Huawei Cloud Global Accelerator	3
3 When to Use Global Accelerator	4
4 Acceleration Area	6
5 How Global Accelerator Works	7
6 Constraints	8
7 Billing	10
8 Permissions	12
9 Using Global Accelerator with Other Services	16
10 Concepts	18
11 Region and AZ	20

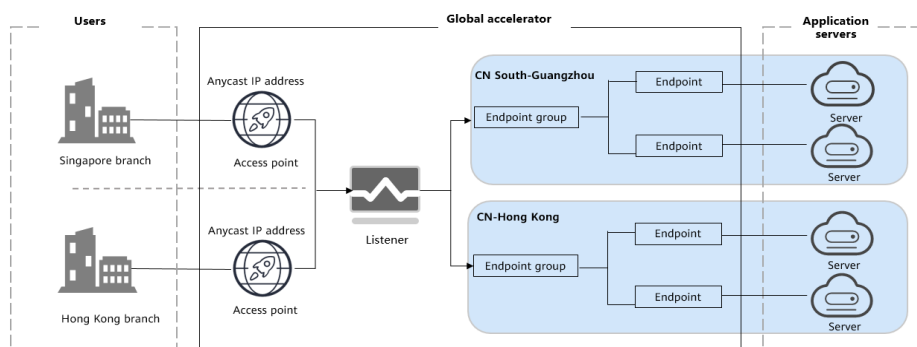
1 What Is Global Accelerator?

Global Accelerator allows users around the world to access cloud applications faster through anycast IP addresses and highly reliable, low-latency, and secure networking services.

As shown in [Figure 1-1](#), a multinational enterprise has branches all over the world. The Singapore branch has deployed an application on two servers in the CN South-Guangzhou region, and the Hong Kong branch has deployed an application on two servers in the CN-Hong Kong region.

With Global Accelerator, each branch can access their application faster from the nearest access point.

Figure 1-1 How Global Accelerator works



Components

- **Global accelerator:** an instance that routes traffic over the Huawei backbone network to accelerate your access to applications. You can create a global accelerator and select the acceleration area 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.
- **Acceleration area:** a location where the global accelerator is used in Europe.
- **Listener:** checks requests from clients based on the protocol and port you have configured and routes the requests to the optimal endpoint in the

associated endpoint group. Each listener can have one endpoint group associated in each region.

- **Endpoint group:** a collection of endpoints that receive traffic from the associated listener. Each endpoint group can include one or more endpoints in a given region.
- **Endpoint:** a resource that traffic is directed to. You can add an EIP as an endpoint. You can assign a weight to each endpoint. Global Accelerator will distribute traffic to the endpoints based on their weights.

2 Why Huawei Cloud Global Accelerator

Compliant: Global Accelerator complies with all relevant laws and regulations worldwide, making it easier for you to meet compliance requirements and stay focused on business innovation.

Reliable: The Huawei backbone network consists of lines from multiple carriers and provides multiple egresses, so no individual region of connection can be a single point of failure.

Fast: Global Accelerator routes client requests over the Huawei backbone network to endpoints. This optimizes the path to your application to keep packet loss, jitter, and latency consistently low.

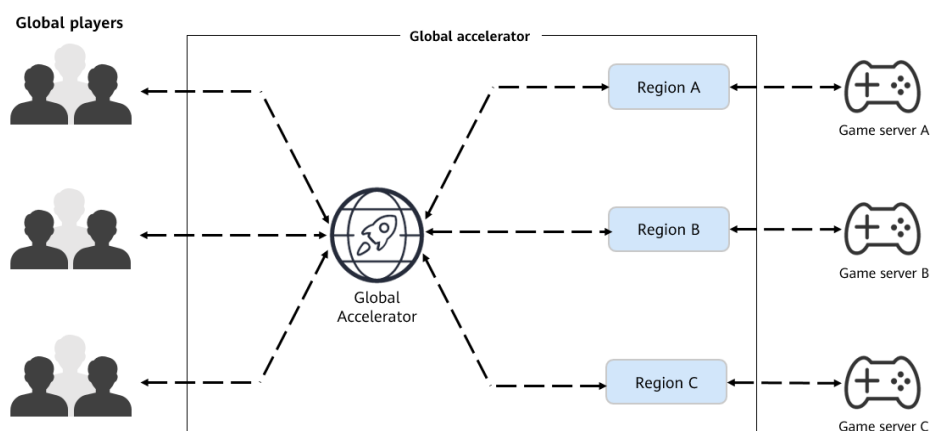
Simple: You can configure global accelerators within minutes and flexibly adjust the traffic sent to each endpoint group and endpoint for better load balancing or active/standby deployment.

3 When to Use Global Accelerator

Gaming Acceleration

Pain points: Online gaming faces problems such as high latency and packet loss, and frequent disconnections, severely affecting gaming experience. Industries, such as online education and livestreaming, face the same problems.

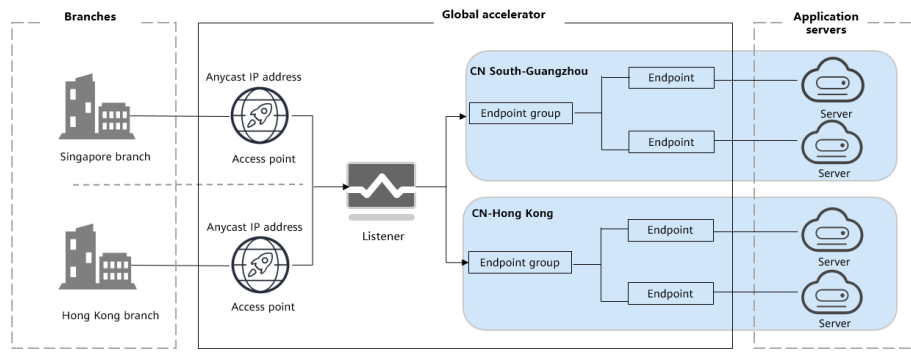
Solution: Global Accelerator can improve global player's gaming experience by routing requests from the nearest access point to optimal game servers through the Huawei backbone network. There are fewer hops on the public network, so in-game latency, jitter, and packet loss are kept as low as possible.



Multinational Enterprises

Pain points: Multinational enterprises generally store data at their headquarters. When employees in the branches around the world access the data stored at the headquarters, the latency is high and the access is slow, which drags down the productivity and increases the OPEX.

Solution: Global Accelerator can connect enterprise servers around the world through the Huawei backbone network to keep the latency and jitter consistently low.



4 Acceleration Area

An acceleration area is where your clients or end users use global accelerators to access your application faster. Global Accelerator deploys access points in Europe as detailed in [Table 4-1](#).

Table 4-1 Access points

Applicability	Geographic Region	Access Point
Europe	Europe	Paris (France), Frankfurt (Germany), and Dublin (Ireland)

5 How Global Accelerator Works

The Huawei backbone network used by Global Accelerator is like an expressway network that connects countries around the world. An access point is like an expressway entrance. An acceleration area is a country or geographic region where an access point is deployed. A region is like an exit where you leave an expressway. You can access an application using an EIP over the Internet, which is like driving on a common road. Alternatively, you can access the application over the Huawei backbone network from the nearest access point. This is where Global Accelerator is used. It provides an anycast IP address that allows you to access the application faster.

Suppose you have an application deployed in Shanghai and you use Global Accelerator to accelerate access from users around the globe. Users in Singapore can access the application faster over the Huawei backbone network from the Singapore access point. However, there is no access point in Malaysia. Users in Malaysia can access your application over the Huawei backbone network from the Singapore access point.

Global Accelerator can also accelerate access to an IP address that is not used on Huawei Cloud. You only need to select the region nearest to this IP address. If your application is deployed in your on-premises data center in Wuxi, you can select the Shanghai region nearest to Wuxi. In this way, users in Singapore can access your application from the Singapore access point. Access requests are then transmitted to Shanghai over the Huawei backbone network and finally to Wuxi data center over the public network.

6 Constraints

Table 6-1 lists default resource quotas of Global Accelerator.

Table 6-1 Global Accelerator resource quotas

Resource	Default Quota	Description
Global accelerators that can be created by each account	20	To increase the quota, submit a service ticket .
Listeners that can be added to each global accelerator	10	To increase the quota, submit a service ticket .
Endpoints allowed in each endpoint group	10	To increase the quota, submit a service ticket .
Connections that can be handled by each global accelerator	20,000	A global accelerator can handle 20,000 connections in an access point. To increase the quota, submit a service ticket .
Peak bandwidth	100 Mbit/s	The maximum bandwidth for a server in an access point is 100 Mbit/s. To increase the quota, submit a service ticket .
IP address groups that can be created by each account	50	N/A
CIDR blocks that can be added to an IP address group	200	N/A

Resource	Default Quota	Description
Listeners that can be associated with an IP address group	10	N/A

7 Billing

Billing Mode

Only pay-per-use is available. The bill for each month is generated at the beginning of the next calendar month.

Billing Items

Table 7-1 Global Accelerator pricing details

Billing Item	Description	Price
Global accelerators	You are charged based on how long each global accelerator is retained in your account. The smallest billing unit is one hour. Partial hours are counted as full hours. Global accelerator price = Unit price x Required duration	\$0.356 USD/hour

Billing Item	Description	Price
Data transfer	<p>You are charged for either the inbound or outbound traffic, whichever direction has more traffic.</p> <p>Data transfer price = Unit price x Traffic used</p>	<p>The actual price is subject to what is displayed on the Global Accelerator console.</p> <p>NOTE</p> <ul style="list-style-type: none"> • If an EIP is added as an endpoint, the traffic will not be repeatedly billed. • See Acceleration Area for available acceleration areas. • The regions where an endpoint group can be deployed are those you can select on the Global Accelerator console.

Billing Examples

Suppose you have an application deployed in Guangzhou, if you want end users in Türkiye to be able to access your application faster, you need a global accelerator.

If end users in Türkiye access your application, inbound traffic to your application is 1 GB and outbound traffic from your application is 20 GB, you are only charged for the 20 GB of outbound traffic but not the 1 GB of inbound traffic.

The total price for using this global accelerator for an hour can be calculated using the following formula: Global accelerator price + Data transfer price = \$0.356 USD/hour x 1 hour + 1.098 x 20 GB = \$22.316 USD

Changes Between Billing Modes

The billing mode cannot be changed.

Renewal

For details, see [Renewal Management](#).

Expiration and Overdue Payment

For details, see [Service Suspension and Resource Release](#) and [Payment and Repayment](#).

8 Permissions

If you need to assign different permissions to employees in your enterprise to access your Global Accelerator resources, IAM is a good choice for fine-grained permissions management. IAM provides identity authentication, permissions management, and access control, helping you secure access to your cloud resources.

With IAM, you can use your HUAWEI ID to create IAM users, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use Global Accelerator resources but should not delete them or perform any other high-risk operations. In this scenario, you can create IAM users for the software developers and grant them only the required permissions.

Skip this section if your HUAWEI ID does not require individual IAM users for permissions management.

IAM can be used free of charge. You pay only for the resources in your account. For more information about IAM, see the [What Is IAM?](#)

Global Accelerator Permissions

By default, new IAM users do not have permissions assigned. You need to add a user to one or more groups, and attach permissions policies or roles to these groups. Users inherit permissions from the groups that they are added to, and can perform specified operations on cloud services.

Global Accelerator is a global service for access from any region. You can assign IAM permissions to users in the global service project. In this way, users do not need to switch regions when they access IAM.

You can grant permissions by using roles or policies.

- **Roles:** A coarse-grained authorization strategy provided by IAM to assign permissions based on users' job responsibilities. Only a limited number of service-level roles are available for authorization. When using roles to grant permissions, you may need to also assign other dependency roles. However, roles are not an ideal choice for fine-grained authorization and secure access control.
- **Policies:** A type of fine-grained authorization mechanism that defines permissions required to perform operations on specific cloud resources under

certain conditions. This mechanism allows for more flexible policy-based authorization for secure access control. For example, you can grant users only the permissions for managing Global Accelerator resources.

Table 8-1 lists the system-defined roles or policies supported by Global Accelerator.

Table 8-1 System-defined roles and policies supported by Global Accelerator

Role/Policy Name	Description	Type	Dependency
GA FullAccess	Permissions: all permissions for Global Accelerator Scope: Global-level service	System-defined policy	-
GA ReadOnlyAccess	Permissions: read-only permissions for Global Accelerator Scope: Global-level service	System-defined policy	-

Table 8-2 lists the common operations supported by each system-defined permission for Global Accelerator.

Table 8-2 Common operations supported by system-defined permissions

Operation	GA FullAccess	GA ReadOnlyAccess
Creating a global accelerator	√	×
Viewing a global accelerator	√	√
Modifying a global accelerator	√	×
Deleting a global accelerator	√	×
Adding a listener	√	×
Viewing a listener	√	√
Modifying a listener	√	×
Deleting a listener	√	×
Adding an endpoint group	√	×

Operation	GA FullAccess	GA ReadOnlyAccess
Viewing an endpoint group	√	√
Modifying an endpoint group	√	×
Deleting an endpoint group	√	×
Adding an endpoint	√	×
Viewing an endpoint	√	√
Modifying an endpoint	√	×
Removing an endpoint	√	×
Configuring a health check	√	×
Viewing health check settings	√	√
Modifying health check settings	√	×
Disabling a health check	√	×
Deleting a health check	√	×
Creating an IP address group	√	×
Querying IP address groups	√	√
Querying the details of an IP address group	√	√
Modifying an IP address group	√	×
Deleting an IP address group	√	×
Adding CIDR blocks to an IP address group	√	×
Removing CIDR blocks from an IP address group	√	×
Associating an IP address group with a listener	√	×
Disassociating an IP address group from a listener	√	×

Operation	GA FullAccess	GA ReadOnlyAccess
Adding tags to a resource	√	×
Querying tags of a specific resource	√	√
Deleting tags from a resource	√	×

References

- [What Is IAM?](#)
- [Creating a User and Granting Permissions](#)

9 Using Global Accelerator with Other Services

Figure 9-1 Using Global Accelerator with other services




Table 9-1 Interaction between Global Accelerator and other services

Interaction	Cloud Service	Reference
An EIP can be added as an endpoint to receive requests from an associated listener.	Elastic IP (EIP)	Buying an EIP
If your application uses a domain name to provide services, you can configure DNS records to translate your domain name to an anycast IP address provided by Global Accelerator to accelerate access to your application.	Domain Name Service (DNS)	DNS
On the Cloud Eye console, you can view the bandwidth usage and the total number of connections established by a global accelerator.	Cloud Eye	Cloud Eye

10 Concepts

Global Accelerator

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 to each area where the global accelerator will be used. The clients can access the cloud from the nearest access points, and their traffic will then be routed to the optimal endpoints through the Huawei backbone network.

Anycast IP Address

Global Accelerator assigns an anycast IP address for each acceleration area. Clients can access the backbone network from the nearest access point through this anycast IP address.

Listener

A listener checks requests from clients based on the protocol and port you have configured and routes the requests to the optimal endpoint in the associated endpoint group. Each listener can have one endpoint group associated in each region.

Endpoint

An endpoint is where requests are routed to. 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.

Endpoint Group

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 weight you specified.

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

Health Check

Global Accelerator provides health check to ensure service reliability and availability. TCP and UDP can be used for health checks.

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 starts routing requests to them again.

11 Region and AZ

Concept

A region and availability zone (AZ) identify the location of a data center. You can create resources in a specific region and AZ.

- Regions are divided based on geographical location and network latency. Public services, such as Elastic Cloud Server (ECS), Elastic Volume Service (EVS), Object Storage Service (OBS), Virtual Private Cloud (VPC), Elastic IP (EIP), and Image Management Service (IMS), are shared within the same region. Regions are classified into universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides specific services for specific tenants.
- An AZ contains one or more physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electricity facilities. Within an AZ, computing, network, storage, and other resources are logically divided into multiple clusters.

Selecting a Region

If your target users are in Europe, select the **EU-Dublin** region.

Selecting an AZ

When deploying resources, consider your applications' requirements on disaster recovery (DR) and network latency.

- For high DR capability, deploy resources in different AZs within the same region.
- For lower network latency, deploy resources in the same AZ.