Global Accelerator

Service Overview

Issue	01
Date	2024-04-30

HUAWEI TECHNOLOGIES CO., LTD.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process.* For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

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
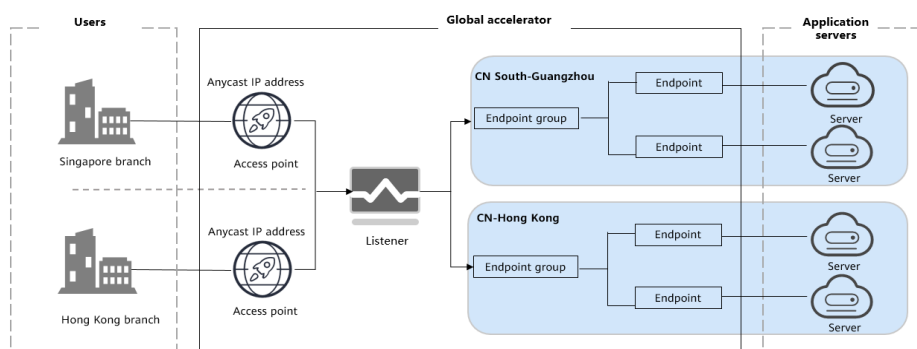
Components

- **Global accelerator**: an instance that routes traffic over the Huawei Cloud 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 Huawei Cloud backbone network, and finally to the optimal endpoints.
- **Acceleration area:** a location where a global accelerator is used in or outside the Chinese mainland. You can configure the location where your users are located as an acceleration area to improve application performance for your users.

- **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. Only EIPs can be added as endpoints. 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 laws and regulations worldwide, allowing you to focus on business innovation.

Reliable: The Huawei Cloud backbone network consists of lines from multiple carriers and provides multiple egresses. This prevents a single point of failure in a single region or on a single line and ensure high service continuity.

Fast: Anycast IP addresses are provided for access from the nearest access point. When a client sends a request, the request will first go to the nearest access point and then to Huawei Cloud's high-speed private network. 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 traffic dial of endpoint groups or endpoint weights 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 Cloud backbone network. There are fewer hops on the public network, so in-game latency, jitter, and packet loss are kept as low as possible.
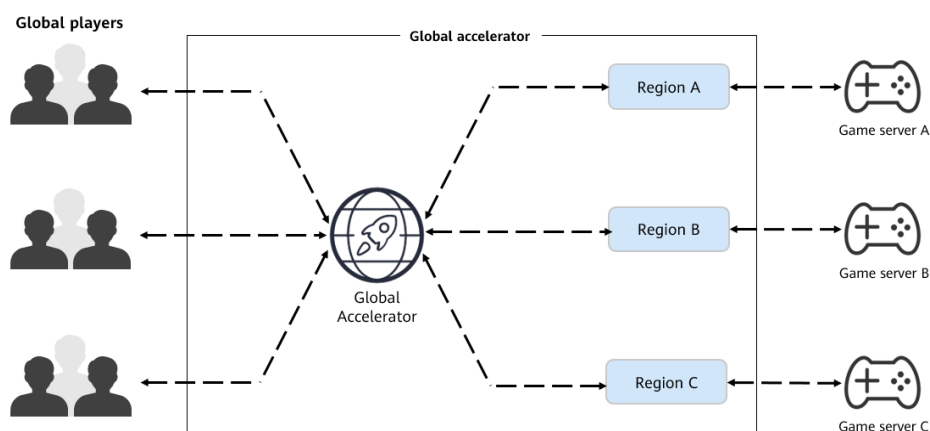
Multinational Enterprises

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 Huawei Cloud's high-speed backbone network to keep the latency and jitter consistently low.
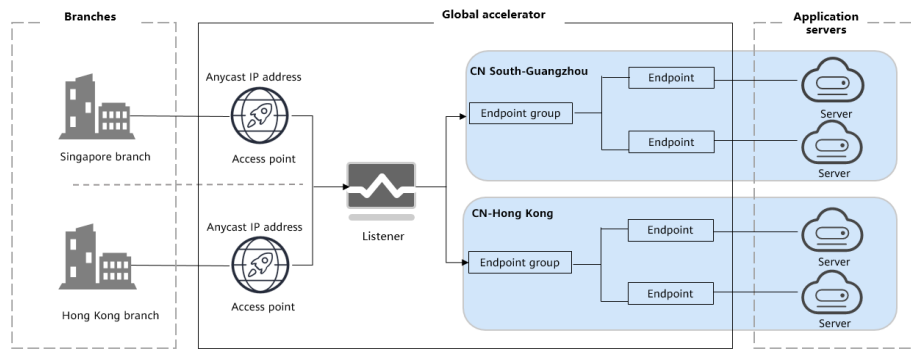
4 Acceleration Area

An acceleration area is where your clients or end users use global accelerators to access your application faster. You can deploy a global accelerator in the Chinese mainland or outside the Chinese mainland. For example, if you deploy your application in the Chinese mainland and want users outside the Chinese mainland to access your application faster, you can select **Outside the Chinese mainland** for **Applicability** when creating a global accelerator.

Global Accelerator deploys corresponding access points in different acceleration areas. **Table 4-1** shows the access points.

Table 4-1 Access points

Geographic Region	Access Point
Chinese mainland	Guangzhou, Beijing, and Shanghai
Asia Pacific	Hong Kong (China), Singapore, Bangkok (Thailand), Jakarta (Indonesia), Manila (Philippines), Tokyo (Japan), and Ho Chi Minh (Vietnam)
Latin America	Bogota (Colombia) and Mexico City (Mexico)
Europe	Istanbul (Türkiye)
Africa	Johannesburg (South Africa)
Europe	Paris (France), Frankfurt (Germany)

5 Constraints

Table 5-1 lists default resource quotas of Global Accelerator.

Table 5-1 Global Accelerator resource quotas

Resource	Default Quota	How to Increase Quotas
Global accelerators that can be created by each account	20	**Submit a service ticket**.
Listeners that can be added to each global accelerator	10	**Submit a service ticket**.
Endpoints allowed in each endpoint group	10	**Submit a service ticket**.
Connections that can be handled by each global accelerator	20,000	**Submit a service ticket**.
Peak bandwidth	100 Mbit/s	**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
Listeners that can be associated with an IP address group	10	N/A

6 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 6-1 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
Data transfer	You are charged for either the inbound or outbound traffic, whichever direction has more traffic. Data transfer price = Unit price x Traffic used	For details, see "Specifications Price" in **Global Accelerator Pricing Details**. **NOTE** ● 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 Hong Kong and the Philippines to be able to access your application faster, you need a global accelerator.

- If end users in Hong Kong 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.

- Similarly, if end users in the Philippines access your application, inbound traffic to your application is 1 GB and outbound traffic from your application is 5 GB, you are only charged for the 5 GB of outbound 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 + 1.098 x 5 GB = $27.806 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**.

7 Security

7.1 Shared Responsibilities

Huawei guarantees that its commitment to cyber security will never be outweighed by the consideration of commercial interests. To cope with emerging cloud security challenges and pervasive cloud security threats and attacks, Huawei Cloud builds a comprehensive cloud service security assurance system for different regions and industries based on Huawei's unique software and hardware advantages, laws, regulations, industry standards, and security ecosystem.

Figure 7-1 illustrates the responsibilities shared by Huawei Cloud and users.

- **Huawei Cloud**: Ensure the security of cloud services and provide secure clouds. Huawei Cloud's security responsibilities include ensuring the security of our IaaS, PaaS, and SaaS services, as well as the physical environments of the Huawei Cloud data centers where our IaaS, PaaS, and SaaS services operate. Huawei Cloud is responsible for not only the security functions and performance of our infrastructure, cloud services, and technologies, but also for the overall cloud O&M security and, in the broader sense, the security and compliance of our infrastructure and services.

- **Tenant**: Use the cloud securely. Tenants of Huawei Cloud are responsible for the secure and effective management of the tenant-customized configurations of cloud services including IaaS, PaaS, and SaaS. This includes but is not limited to virtual networks, the OS of virtual machine hosts and guests, virtual firewalls, API Gateway, advanced security services, all types of cloud services, tenant data, identity accounts, and key management.

Huawei Cloud Security White Paper elaborates on the ideas and measures for building Huawei Cloud security, including cloud security strategies, the shared responsibility model, compliance and privacy, security organizations and personnel, infrastructure security, tenant service and security, engineering security, O&M security, and ecosystem security.

Figure 7-1 Huawei Cloud shared security responsibility model
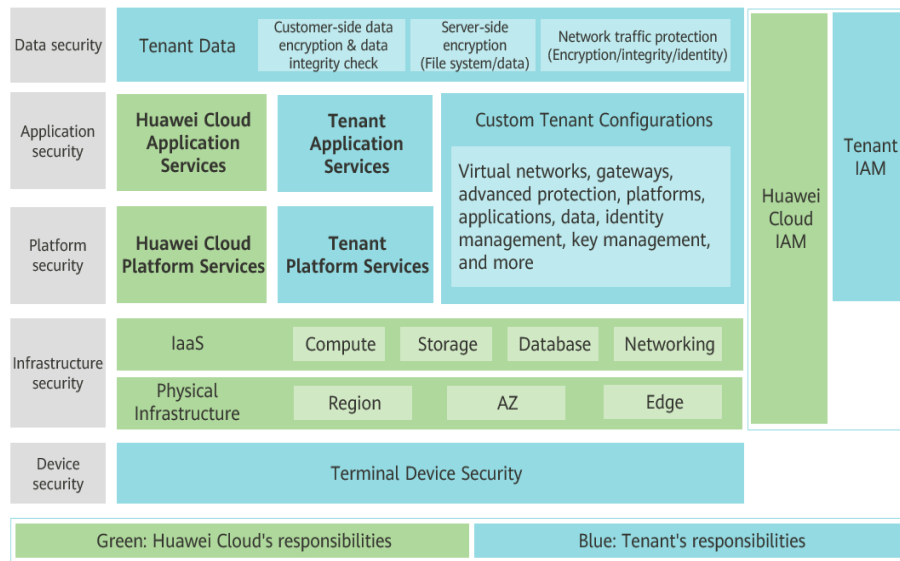
7.2 Identity and Access Management

You can use Identity and Access Management (IAM) to control access to your Global Accelerator resources. IAM permissions define which actions on your cloud resources are allowed or denied to control access to your resources. After creating an IAM user, the administrator needs to add it to a user group and grant the permissions required by Global Accelerator to the user group. And then, all users in this group automatically inherit the granted permissions.

For details, see **Permissions**.

7.3 Auditing and Logging

Cloud Trace Service (CTS) records operations on the cloud resources in your account. You can use the logs generated by CTS to perform security analysis, track resource changes, audit compliance, and locate faults.

After CTS is enabled, traces can be generated for operations performed on service resources.

- For details about how to enable and configure CTS, see **Enabling CTS**.
- For details about key operations, see **Key Operations Recorded by CTS**.
- For details about traces, see **Viewing Traces**.

7.4 Service Resilience

Global Accelerator provides health check to ensure service reliability and availability.

You can enable health check for the endpoint group. After you enable health check, the global accelerator periodically sends requests to endpoints to check their status. If any endpoints become unhealthy, the global accelerator stops

sending requests to these endpoints. After the endpoints recover from failure, the global accelerator starts routing requests to them again.

For details about the health check protocols and how to configure a health check, see **Health Check**.

7.5 Monitoring Security Risks

You can use Cloud Eye to monitor the Global Accelerator status and resource usage. You can also configure Cloud Eye to alert you of any potential issues in Global Accelerator in real time.

For details about supported metrics and setting alarm rules, see **Monitoring**.

7.6 Certificates

Compliance Certificates

Huawei Cloud services and platforms have obtained various security and compliance certifications from authoritative organizations, such as International Organization for Standardization (ISO). You can **download** them from the console.

Figure 7-2 Downloading compliance certificates
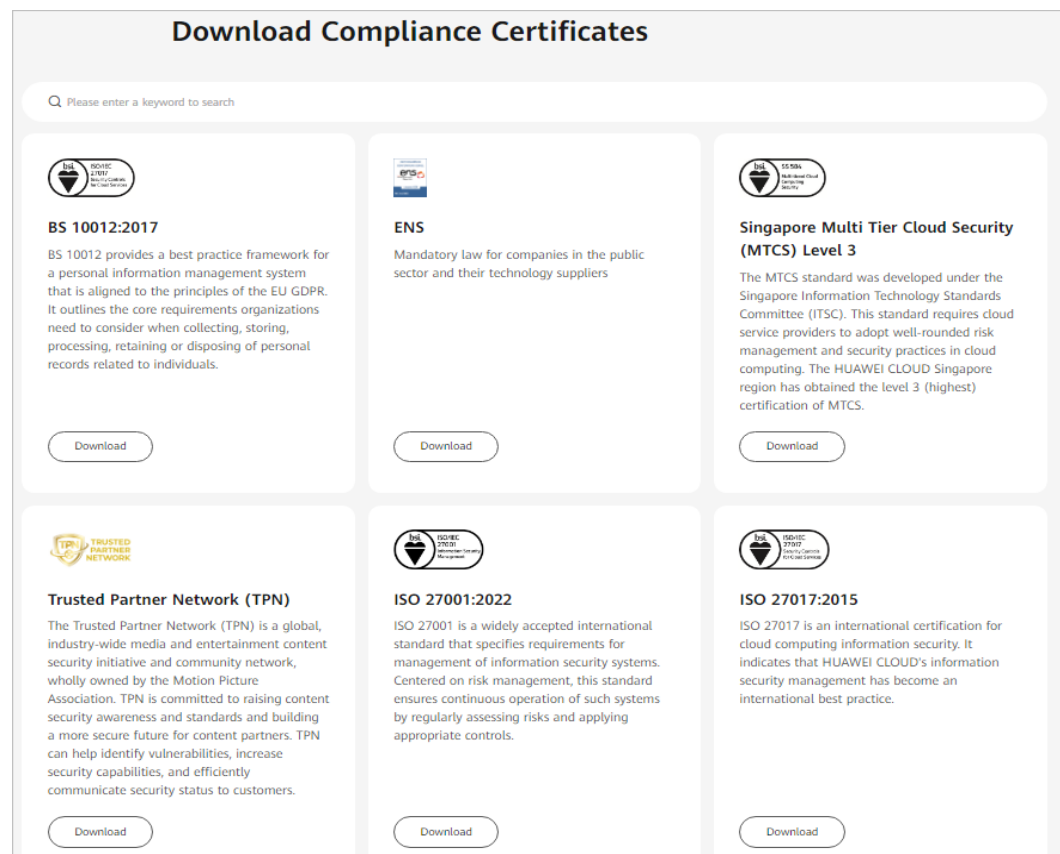
Resource Center

Huawei Cloud also provides the following resources to help users meet compliance requirements. For details, see **Resource Center**.

Figure 7-3 Resource center

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.
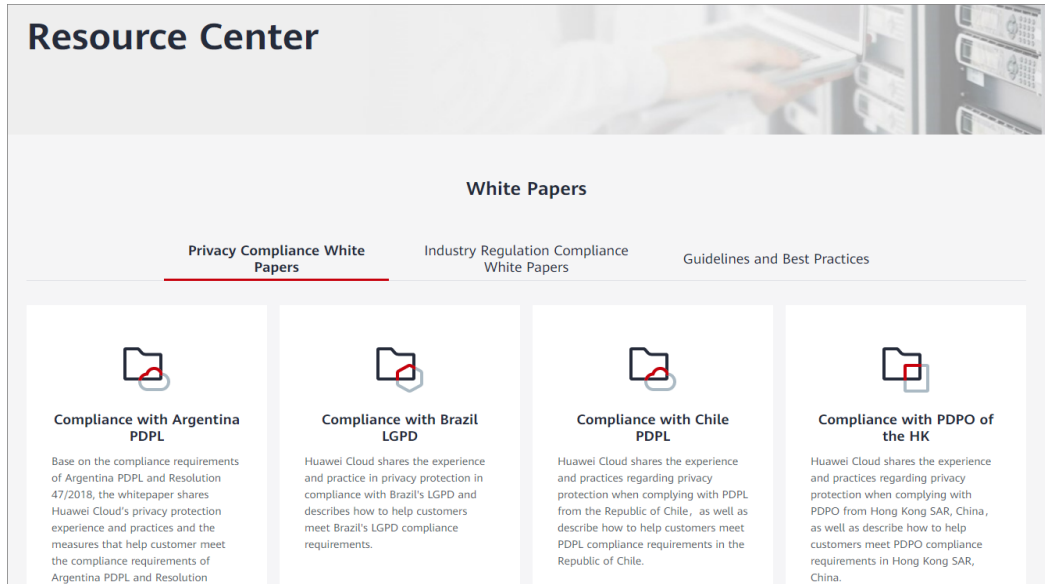
With IAM, you can use your HUAWEI ID to create IAM users, and assign permissions to the users to control their access to specific resources. For example, some software developers in your enterprise need to use Global Accelerator resources but should not delete them or perform any other high-risk operations. 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
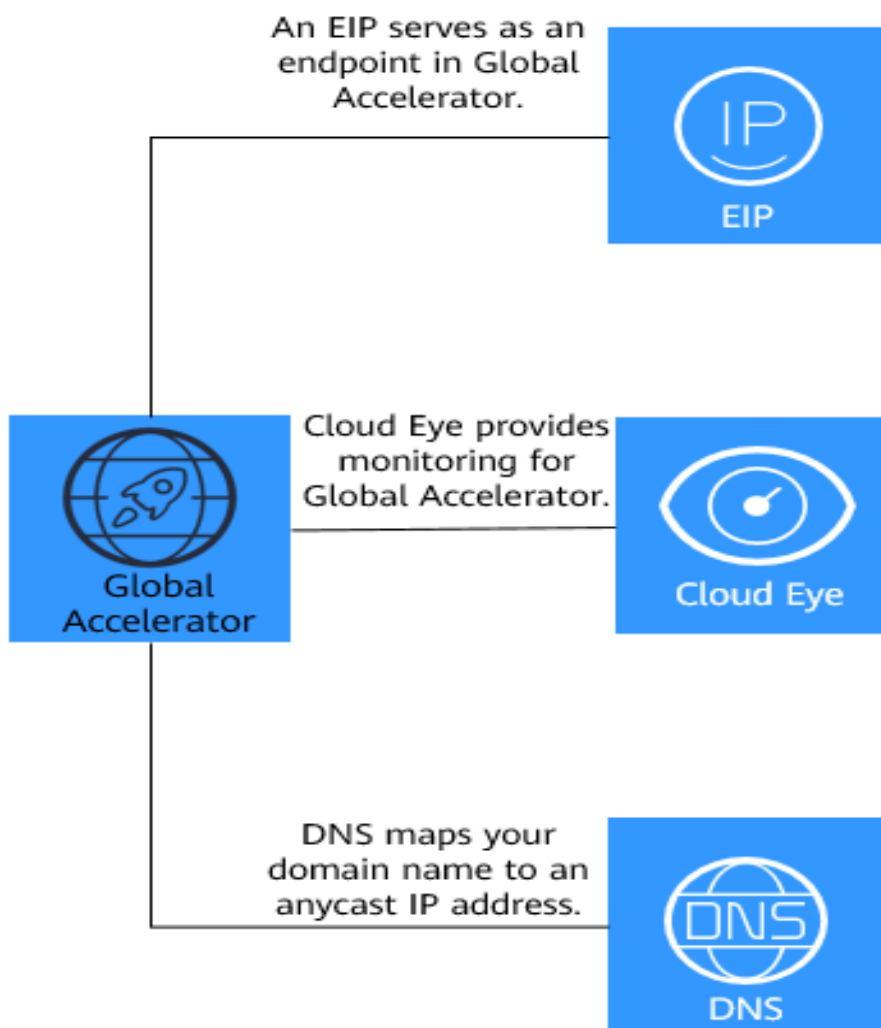
Table 9-1 Interaction between Global Accelerator and other services

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 Cloud 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 Cloud 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. Currently, only EIPs can be added as endpoints.

You can assign a weight to each endpoint. The global accelerator will distribute traffic to the endpoints based on their weights.

Endpoint Group

An endpoint group includes one or more endpoints in a given region. The global accelerator routes traffic to the endpoints in an endpoint group based on the load balancing algorithm.

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. Only TCP 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.
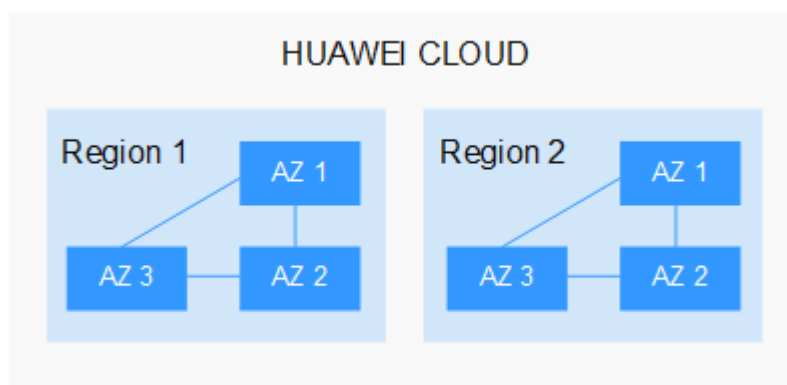
11 Region and AZ

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. AZs within a region are interconnected using high-speed optical fibers, to support cross-AZ high-availability systems.

Figure 11-1 shows the relationship between regions and AZs.

Figure 11-1 Regions and AZs

Huawei Cloud provides services in many regions around the world. You can select a region and an AZ based on requirements. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  It is recommended that you select the closest region for lower network latency and quick access.

  - If your target users are in Asia Pacific (excluding the Chinese mainland), select the **CN-Hong Kong**, **AP-Bangkok**, or **AP-Singapore** region.
  - If your target users are in Africa, select the **AF-Johannesburg** region.
  - If your target users are in Latin America, select the **LA-Santiago** region.

    ☐ NOTE

    The **LA-Santiago** region is located in Chile.

- Resource price

  Resource prices may vary in different regions. For details, see **Product Pricing Details**.

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.

Regions and Endpoints

Before you use an API to call resources, specify its region and endpoint. For more details, see **Regions and Endpoints**.