GeminiDB

Service Overview

Issue	01
Date	2024-09-30

HUAWEI

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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.

Contents

1 What Is GeminiDB?..... 1
2 System Architecture..... 4
3 Highlights..... 5
4 Typical Application Scenarios..... 7
5 Security..... 9
5.1 Shared Responsibilities..... 9
5.2 Identity Authentication and Access Control..... 10
5.3 Data Protection..... 11
5.4 Audit and Logs..... 12
5.5 Resilience..... 13
5.6 Risk Monitoring..... 14
5.7 Fault Recovery..... 14
5.8 Certificates..... 15
6 Billing..... 17
7 Permissions..... 19
8 Regions and AZs..... 32
9 Related Services..... 34

1 What Is GeminiDB?

GeminiDB is a distributed, multi-model NoSQL database service with decoupled storage and compute. It is highly available, reliable, secure, and scalable and delivers excellent performance. It can be deployed quickly and provide capabilities like backup, restoration, monitoring, and alarm reporting.

GeminiDB is compatible with mainstream NoSQL APIs, such as Cassandra, DynamoDB, MongoDB, InfluxDB, and Redis and provides high read/write performance at low costs. It is specifically suited for IoT, meteorology, Internet, and gaming sectors.

How Do I Select an API?

Different APIs provide different functions. You can select one of them based on your service requirements and scenarios.

Table 1-1 Scenario description

API	Compatible API	Scenario	Description
GeminiDB Redis API	Key-value API: Redis	GeminiDB provides high-concurrency and low-latency service access, and ultimate elastic scaling capabilities to cope with service peaks. Common user scenarios include gaming, RTA-based advertising, recommendation systems, e-commerce, and education.	GeminiDB Redis API is a scalable Key-Value (KV) database that supports the Redis protocol and offers a larger storage capacity than memory. It is stable, cost-effective, and has low latency. Unlike other databases, it does not require standby nodes and has an impressive data compression ratio of 4:1. Additionally, it includes enterprise-level features such as hash field expiration, Bloom filter, FastLoad, and memory acceleration.
GeminiDB Influx API	Time series API: InfluxDB	GeminiDB Influx API is widely used to monitor resources, workloads, IoT devices, and industrial production processes, evaluate production quality, and trace faults.	GeminiDB Influx API is a cloud-native, InfluxDB-compatible NoSQL time series database, which uses Huawei-developed architecture with decoupled storage and compute. GeminiDB Influx API can read and write a large volume of time series data concurrently and store it with compression algorithms. The database service then allows you to query the data using SQL-like statements and supports multi-dimension aggregate computing and visual analysis on it. GeminiDB Influx API also provides high write performance, scalability, a good compaction rate, and excellent query performance.

API	Compatible API	Scenario	Description
GeminiDB Cassandra API	Wide-column API: Cassandra and DynamoDB	GeminiDB Cassandra API can store terabytes of data and can handle millions of queries per second. It provides strong consistency, making it well suited to massive storage scenarios, such as manufacturing, meteorology, and the Internet.	GeminiDB Cassandra API is a cloud-native NoSQL database that uses a decoupled storage and compute architecture developed by Huawei. It is compatible with the Cassandra ecosystem and supports Cassandra Query Language (CQL), which gives you SQL-like syntax. It is secure, reliable, scalable, and easy to manage and can provide high read/write performance.
GeminiDB Mongo API	Document-oriented API: MongoDB	GeminiDB Mongo API gives you 3 times the read and write performance of MongoDB and can handle millions of queries per second. It can also store a huge number of documents, images, and social video/audios, and massive IoT/IoV data, making it an excellent choice for sectors like the Internet, IoT, gaming, and finance.	GeminiDB Mongo API is a cloud-native, MongoDB-compatible NoSQL database service, which uses Huawei-developed architecture with decoupled storage and compute. This flexible, scalable, and reliable database is designed for enterprise-grade performance and can be managed on a visual platform.

2 System Architecture

Overview

GeminiDB is a distributed database service with decoupled storage and compute. One compute cluster may consist of multiple homogeneous nodes, and data is stored in a distributed, shared storage pool. You can scale compute and storage resources separately without having to migrate any data.

Figure 2-1 System architecture

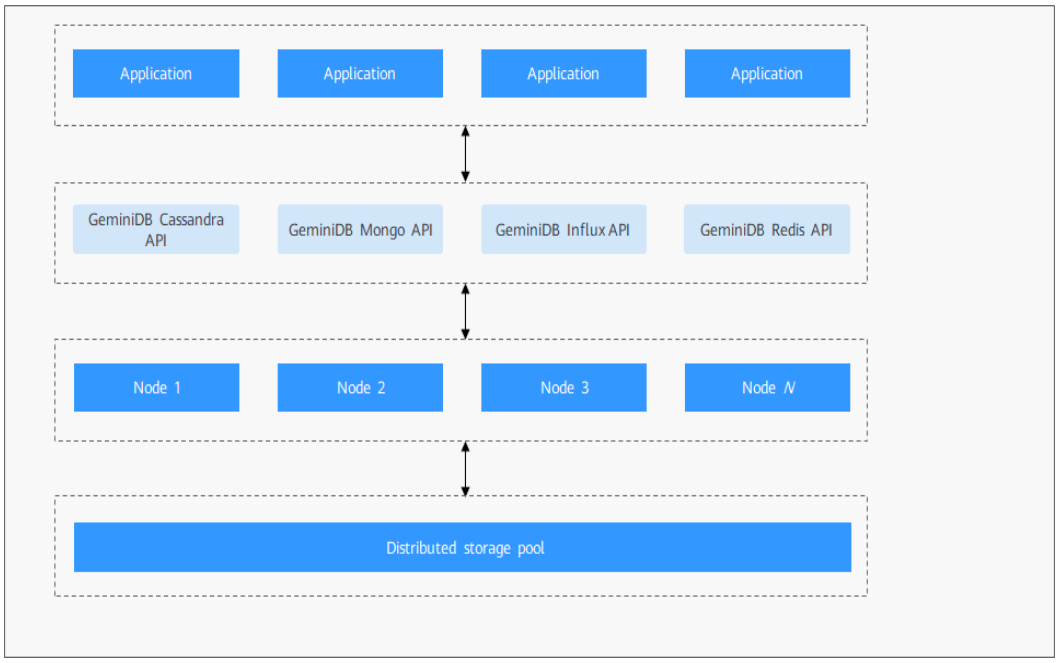

3 Highlights

High Reliability

Data backup

There are automated and manual backups. An automated backup is a full backup of an instance and is created by the system automatically while a manual backup is a full backup that you create yourself. Both of these backups can be used to restore instance data.

Backups are stored in Object Storage Service (OBS) buckets to provide for future disaster recovery and to save storage space. When you create an instance, automated backup is enabled by default. After the creation is complete, an automated full backup is created instantly and can be retained for 7 days by default. You can set a retention cycle or modify the backup policy. In addition, you can initiate a manual backup whenever you want to. Manual backups are saved until you manually delete them.

High Security

Network isolation

GeminiDB uses Virtual Private Clouds (VPCs) and network security groups to isolate instances. VPCs allow you to define which IP addresses are allowed to access a given instance. Running an instance in a VPC improves security. To further secure the instance, you can configure subnets and security groups to control access to it.

Access control

You can configure VPC security groups with inbound and outbound rules to control traffic to and from your instance.

Encryption

GeminiDB uses Secure Sockets Layer (SSL) to encrypt transmitted data. You can download the root CA certificate from the GeminiDB console and upload it for authentication on the database server when you connect to a database.

Security

GeminiDB uses a multi-layer security system. The system consists of VPCs, subnets, security groups, Anti-DDoS, and SSL, which collectively defend against a wide range of attacks to keep your data secure.

- VPCs isolate tenants and control access.
- SSL connections ensure data security and integrity.
- Security group rules control traffic to and from specific IP addresses and ports, protecting connections between GeminiDB and other services.

Performance monitoring

GeminiDB monitors instance performance and can take over more than 60% of tedious O&M activities. It can monitor instance data like the CPU usage, IOPS, and network throughput, allowing you to check instance status at any time.

Great Convenience

Ready to use out of the box

You can create an instance on the console and access the instance using an ECS over a private network to reduce response time and avoid the cost of using a public network.

Compatibility

GeminiDB is compatible with the Cassandra, MongoDB, InfluxDB, and Redis protocols.

Superior Scalability

GeminiDB, as a distributed database service with decoupled storage and compute, allows you to add compute nodes in minutes and scale up storage in just seconds.

4 Typical Application Scenarios

Gaming

GeminiDB allows you to keep track of gaming data like game items or points earned. Adding compute nodes is so easy, making it an excellent choice for high-concurrency scenarios often involved in online gaming.

Advantages:

- **Flexibility**: In the first 6 hours of when a game goes online, game databases have to scale out multiple times. GeminiDB Mongo API is a great choice for you because it lets you add nodes fast to ensure performance as new players come online.
- **Fast data recovery**: Table-level restoration helps you restore gaming data in specific tables to any point in time.
- **Stability**: Storage scaling does not affect your gaming experience.

IoT

GeminiDB is compatible with Cassandra APIs. It can deliver high write performance and is designed for write-intensive scenarios, specifically applied in sectors like manufacturing, logistics, health care, real estate, energy production, and agriculture. It can process the data sent by different types of sensors for further analysis.

Advantages:

- **High write performance**: GeminiDB provides higher write performance than other NoSQL services.
- **High scalability**: Compute nodes can be added in minutes and storage can be scaled up in seconds to handle traffic surges or large write loads at any time.

Internet

E-commerce and entertainment websites that include product catalogs, recommendations, and personalization engines use GeminiDB when they require fast reading and writing of data and high scalability. GeminiDB stores visitor activities, making it easy for analysis tools to access data fast and to generate recommendations.

Advantages:

- **High write performance**: GeminiDB provides higher write performance than other NoSQL services.
- **Big data analysis**: GeminiDB can work with big data tools, such as Spark, to provide recommendations to your customers in real time.

Finance

With Spark's big data analysis capabilities, GeminiDB helps companies in the finance sector build risk control systems and mitigate fraud.

Advantages:

Big data analysis: GeminiDB can work with tools like Spark to help you detect and prevent fraud in real time.

5 Security

5.1 Shared Responsibilities

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

The shared responsibility model for Huawei Cloud and the tenants who use Huawei Cloud services is illustrated in **Figure 5-1**. Responsibilities are as follows:

- **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 more widely speaking, the security 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 OSs 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 behind and measures used to ensure 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 5-1 Huawei Cloud shared security responsibility model

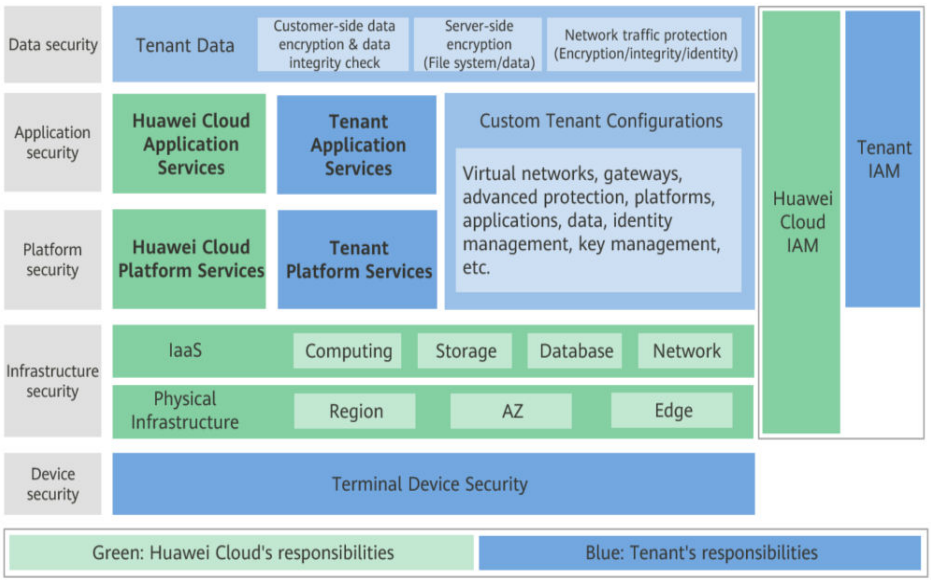

5.2 Identity Authentication and Access Control

Identity Authentication

When you access GeminiDB, the system authenticates your identity using two authentication modes: password and IAM.

- **Password authentication**

  When you want to manage your instances on the GeminiDB console, you need to log in to the console first using your account name and password.

- **IAM authentication**

  You can use **Identity and Access Management (IAM)** to provide fine-grained control over GeminiDB permissions. IAM provides identity authentication, permissions management, and access control, helping you efficiently manage access to your Huawei Cloud resources. You can create IAM users and use them to manage GeminiDB resources. When you log in using an IAM user, password authentication is required. For details, see **Step 2: Create IAM Users and Log In**.

Access Control

- **Permissions control**

  If you need to assign different permissions to different employees in your enterprise to access your instance resources, IAM is a good choice. For details, see **Permissions**.

- **VPCs and subnets**

  A VPC is a logically isolated, configurable, and manageable virtual network. It helps improve the security of cloud resources and simplifies network deployment. You can define security groups, virtual private networks (VPNs), IP address ranges, and bandwidth for a VPC. This makes it easy for you to manage and configure private networks and improves network security.

A subnet provides dedicated network resources that are logically isolated from other networks for security.

For details, see **Creating a VPC**.

- **Security groups**

  A security group is a logical group. It provides access control policies for ECSs and GeminiDB instances that have the same security protection requirements and are mutually trusted. To ensure database security and reliability, you need to configure security group rules to allow specified IP addresses and ports to access your GeminiDB instances.

  For details, see **Configuring Security Group Rules**.

5.3 Data Protection

GeminiDB provides a series of features to ensure data security and reliability.

Table 5-1 Features for data security

Feature	Description	Reference
Secure Sockets Layer (SSL)	GeminiDB Redis, GeminiDB Mongo, GeminiDB Cassandra, and GeminiDB Influx instances support both SSL and non-SSL connections. SSL is recommended to secure data transmission.	• GeminiDB Redis: **Configuring an SSL Connection** • GeminiDB Influx: **Enabling and Disabling the SSL Connection** • GeminiDB Cassandra: **Configuring an SSL Connection** • GeminiDB Mongo: **Configuring an SSL Connection**
Cross-AZ deployment	GeminiDB can deploy nodes of an instance evenly across three AZs to provide cross-AZ DR. These AZs are physically isolated and can be interconnected over private networks.	• GeminiDB Redis: **Buying an Instance** • GeminiDB Influx: **Buying a Cluster Instance** • GeminiDB Cassandra: **Buying an Instance** • GeminiDB Mongo: **Buying a Replica Set Instance**

Feature	Description	Reference
Load balancing	GeminiDB Redis API supports load balancing. Data access requests can be evenly routed to different nodes in a cluster to avoid hotspots and maximize overall throughput of the cluster.	GeminiDB Redis: **Connecting to an Instance Using a Load Balancer Address**
Intra-region DR	Primary GeminiDB Cassandra instances support HA. If a primary instance fails to be connected due to a natural disaster, services can be switched over to its DR instance.	GeminiDB Cassandra: **Creating a DR Instance**
Cross-region Dual-active DR	GeminiDB Cassandra API supports dual-active DR and bidirectional synchronization between two instances at different regions. Once one instance becomes faulty, the other takes over read/write traffic to ensure service continuity.	GeminiDB Cassandra: **Cross-region Dual-active DR**
Deletion protection	GeminiDB can move unsubscribed yearly/monthly instances and deleted pay-per-use instances to Recycle Bin. This function enables you to restore instances deleted up to 7 days ago.	• GeminiDB Redis: **Recycling an Instance** • GeminiDB Influx: **Recycling an Instance** • GeminiDB Cassandra: **Recycling an Instance** • GeminiDB Mongo: **Recycling an Instance**

5.4 Audit and Logs

Audit

Cloud Trace Service (CTS) records operations on 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 you enable CTS and configure a tracker, CTS can record management and data traces of GeminiDB for auditing.

For details about how to enable and configure CTS, see **Enabling CTS**.

- For details about GeminiDB Redis management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.

- For details about GeminiDB Influx management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.
- For details about GeminiDB Cassandra management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**
- For details about GeminiDB Mongo management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.

Logs

- GeminiDB Redis

  Allows you to view slow query logs of databases. The unit of the execution time is ms. Allows you to identify the SQL statements that take a long time to execute and tune them based on slow query logs.

  For details about slow query logs, see **Slow Query Logs**.
- GeminiDB Cassandra

  Allows you to view slow query logs of databases. The unit of the execution time is ms. Allows you to identify the SQL statements that take a long time to execute and tune them based on slow query logs.

  For details about slow query logs, see **Slow Query Logs**.
- GeminiDB Mongo
  - Allows you to view slow query logs of databases. The unit of the execution time is ms. Allows you to identify the SQL statements that take a long time to execute and tune them based on slow query logs.

    For details about slow query logs, see **Slow Query Logs**.
  - Allows you to view error logs of databases, including warning logs and error logs generated during database running, helping you analyze system issues.

    For details about error logs, see **Error Logs**.

5.5 Resilience

- GeminiDB Redis API uses the DFV storage pool that provides three copies of data, so data can be persistently stored in real time. GeminiDB Redis API also provides solutions like multi-AZ deployment, failover in seconds, load balancing, and node reduction, to guarantee reliability and availability of your instances.
- GeminiDB Influx API uses the DFV storage pool that provides three copies of data and high write performance. This database product offers you solutions like multi-AZ deployment and elastic scaling to guarantee reliability and availability of your instances.
- GeminiDB Cassandra API uses the DFV storage pool that provides three copies of data and supports 24/7 online writes. GeminiDB Cassandra API also offers you solutions like intra-region DR, cross-region dual-active DR, multi-AZ deployment, *N*-1 fault tolerance, and elastic scaling to guarantee reliability and availability of your instances.
- GeminiDB Mongo API uses the DFV storage pool that provides three copies of data and supports 24/7 online writes. GeminiDB Mongo API also offers you

solutions like multi-AZ deployment, N-1 fault tolerance, and elastic scaling to guarantee reliability and availability of your instances.

Figure 5-2 Reliability architecture

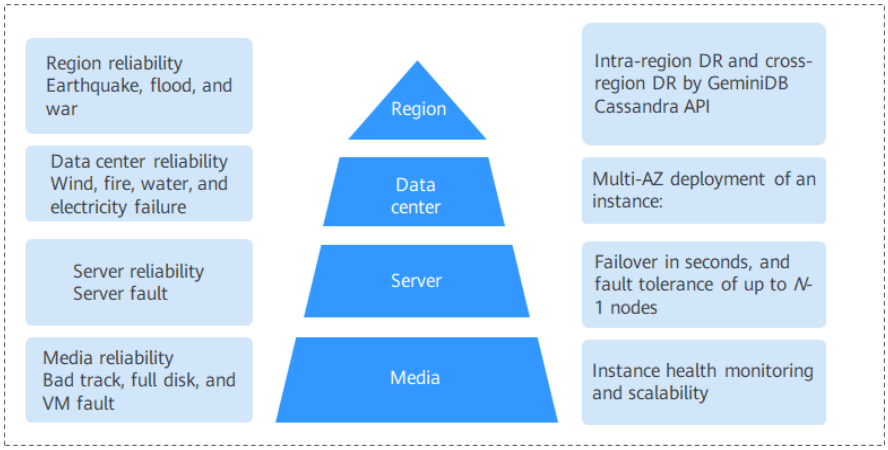

5.6 Risk Monitoring

Metrics

GeminiDB works with Cloud Eye to monitor instances in your account in real time, reporting alarms and sending notifications based on your settings. You can get details about running metrics and storage usage of your instances in real time.

- For details about GeminiDB Redis metrics and how to create alarm rules, see **GeminiDB Redis Metrics**.
- For details about GeminiDB Influx metrics and how to create alarm rules, see **GeminiDB Influx Metrics**.
- For details about GeminiDB Cassandra metrics and how to create alarm rules, see **GeminiDB Cassandra Metrics**.
- For details about GeminiDB Mongo metrics and how to create alarm rules, see **GeminiDB Mongo Metrics**.

Protection for Critical Operations

With this function enabled, the system authenticates your identity when you perform critical operations like deleting an instance, to further secure your data and configurations. For more information, see **Critical Operation Protection**.

5.7 Fault Recovery

GeminiDB can automatically create backups for your instances during the backup time window that you specified. The backups are stored based on the preset retention period (1 to 35 days).

- GeminiDB Redis provides methods for restoring instance data. For details, see **Restoring Data to a New Instance**.

- GeminiDB Influx provides methods for restoring instance data. For details, see **Restoring Data to a New Instance**.
- GeminiDB Cassandra provides methods for restoring instance data. For details, see **Restoring Data to a New Instance** and **Restoring a Backup to a Point in Time**.
- GeminiDB Mongo provides methods for restoring instance data. For details, see **Restoring Data to a New or Existing Instance**.

Intra-region DR

Primary GeminiDB Cassandra instances support HA. If a primary instance fails to be connected due to a natural disaster, services can be switched over to its DR instance.

Cross-region Dual-active DR

GeminiDB Cassandra supports dual-active DR and bidirectional synchronization between two instances at different regions. Once one instance becomes faulty, the other takes over read/write traffic to ensure service continuity.

Multiple-AZ Deployment

An AZ is a physical region where resources have their own independent power supply and networks. AZs are physically isolated but interconnected over a private network. Nodes of one GeminiDB instance can be evenly deployed across three AZs to provide DR support.

5.8 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), System and Organization Controls (SOC), and Payment Card Industry (PCI) compliance standards. These certifications are available for **download**.

Figure 5-3 Downloading compliance certificates

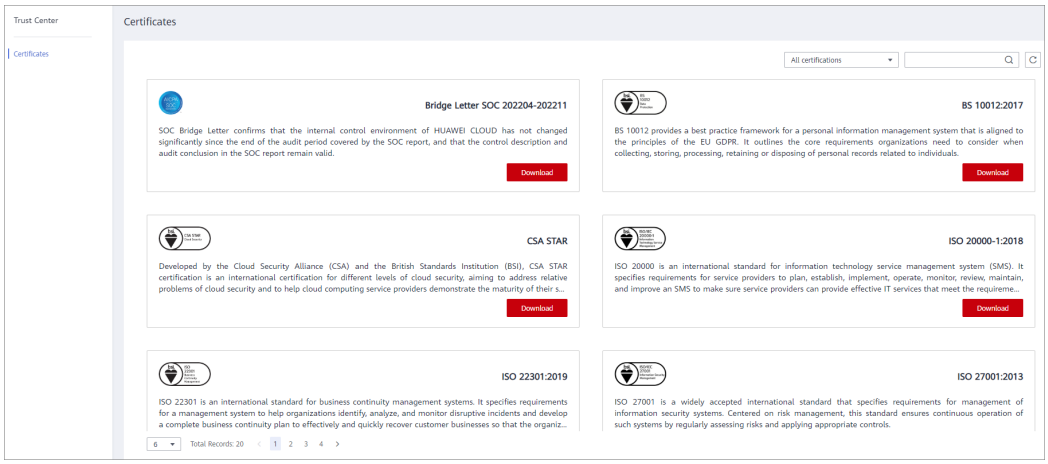

Resource Center

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

Figure 5-4 Resource center

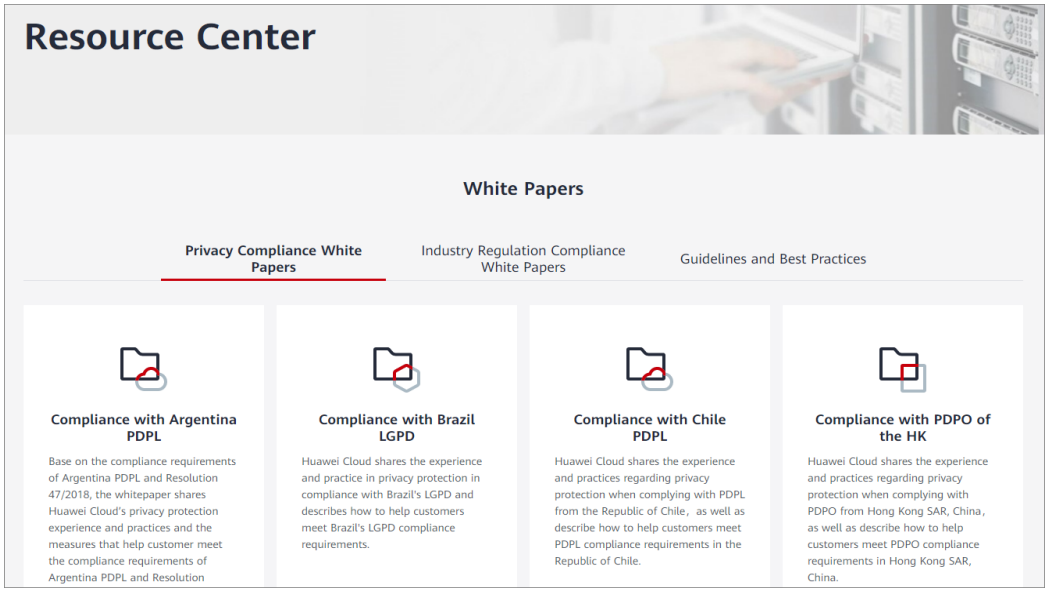

6 Billing

GeminiDB allows you to pay only for what you use. There are no minimum fee requirements.

Billing Items

You are billed based on specifications and storage space of your instance. The total cost of an instance is the price of the instance specifications plus the cost of storage.

Table 6-1 Billing item description

Item	Description
DB instance	• The instance specifications that you select. • The yearly/monthly or pay-per-use (hourly) billing mode.
Database storage	The amount of storage that you select.
(Optional) Backup storage	The backup data of GeminiDB instances is stored on OBS. After you purchase an instance, GeminiDB will provide additional backup storage of the same size as the storage space you purchased. For example, if you purchase an instance with 100 GB of storage space, you will obtain additional 100 GB of backup space free of charge. If your backup data exceeds 100 GB, the backup storage consumed in excess will be charged according to OBS billing rules.
Public network traffic	GeminiDB is public accessible, and you are charged for the generated public network traffic, but not for private network traffic.

Billing Mode

GeminiDB provides two billing modes: pay-per-use and yearly/monthly. You can change the billing mode from pay-per-use to yearly/monthly or vice versa.

- Yearly/Monthly: This billing option offers a larger discount than pay-per-use and is recommended for users who can predict long-term resource usage.
- Pay-per-use (hourly): You pay only for the resources you actually use. Pricing is listed on a per-hour basis, but your actual bill is calculated down to the second.

Modifying Instance Configurations

- You can change specifications of an instance based on service requirements. After the change, you will be billed based on the new specifications.
- You can scale up storage space based on service requirements. After the scaling, you will be billed based on the new storage space. Storage space can only be scaled up. The minimum increment is 1 GB.

7 Permissions

If you need to assign different permissions to employees in your enterprise to access your cloud 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 resources.

With IAM, you can use your Huawei Cloud account 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 resources but cannot delete them or perform any high-risk operations. To achieve this result, you can create IAM users for the developers and grant them only the permissions for using resources.

You can skip this section if you do not need fine-grained permissions management.

IAM is a free service. You only pay for the resources in your account.

For more information about IAM, see **IAM Service Overview**.

GeminiDB Permissions

By default, new IAM users do not have any permissions assigned. To assign permissions to these new users, you need to add them to one or more groups, and attach permissions policies or roles to these groups.

GeminiDB is a project-level service deployed in specific physical regions. When assigning GeminiDB permissions to a user group, you need to specify region-specific projects where the permissions will take effect. If you select **All projects**, the permissions will be granted for all region-specific projects. To access GeminiDB, you need to switch to the region where you are authorized.

You can grant users permissions using roles and policies.

- Roles: A type of coarse-grained authorization mechanism that provides only a limited number of service-level roles. When using roles to grant permissions, you also need to assign other dependent roles. Roles are not ideal for fine-grained authorization and secure access control.
- Policies: A 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 and more secure access control. For example, you can grant ECS users only the permissions for managing a certain type of ECSs. For API actions supported by GeminiDB, see **Permissions Policies and Supported Actions**.

Table 7-1 lists all the system-defined roles and policies supported by GeminiDB.

Table 7-1 GeminiDB policies or roles

Policy Name/ System Role	Description	Type	Dependency
GeminiDB FullAccess	All permissions for GeminiDB	System-defined policy	To create a yearly/monthly instance, you need to configure the following CBC actions: • bss:balance:view • bss:balance:update • bss:order:view • bss:order:pay • bss:order:update • bss:renewal:view • bss:renewal:update To unsubscribe from a yearly/monthly instance, you need to configure the following CBC action: • bss:unsubscribe:update To use storage autoscaling, configure the following actions for IAM users: • Creating a custom policy: – iam:agencies:listAgencies

Policy Name/ System Role	Description	Type	Dependency
			– iam:agencies:createAgency – iam:permissions:listRolesForAgencyOnProject – iam:permissions:grantRoleToGroupOnProject – iam:roles:listRoles – iam:roles:createRole • Adding system role **Security Administrator**: 1. Select a user group to which the user belongs. 2. Click **Authorize** in the **Operation** column. 3. Add the **Security Administrator** role. GeminiDB FullAccess already contains the iam:agencies:listAgencies, iam:roles:listRo

Policy Name/ System Role	Description	Type	Dependency
			les, and iam:agencies:pass permissions. GeminiDB is a region-level service, and IAM is a global service. If you want to grant GeminiDB FullAccess to a project, grant BSS ServiceAgency ReadPolicy (global service) to it as well. Granting GeminiDB FullAccess to all projects eliminates the need for additional configuration when using IAM actions. BSS ServiceAgency CreatePolicy contains the following permissions: • iam:agencies:createAgency • iam:permissions:grantRoleToAgency
GeminiDB ReadOnlyAccess	Read-only permissions for GeminiDB	System-defined policy	None

Table 7-2 lists the common operations supported by each system-defined policy or role of GeminiDB. Select the policies or roles as required.

Table 7-2 Common operations supported by each system-defined policy or role

Operation	GeminiDB FullAccess	GeminiDB ReadOnlyAccess
Creating an instance	√	x
Querying instances	√	√
Querying details of an instance	√	√
Querying tasks	√	√
Deleting an instance	√	x
Restarting an instance	√	x
Resetting a password	√	x
Changing a security group	√	x
Changing a database port	√	x
Binding or unbinding an EIP	√	x
Scaling up storage space	√	x
Changing specifications	√	x
Adding nodes	√	x
Deleting nodes	√	x
Modifying a backup policy	√	x
Renaming an instance	√	x
Creating a manual backup	√	x
Querying backups	√	√
Restoring data to a new instance	√	x
Deleting a backup	√	x
Creating a parameter template	√	x
Querying parameter templates	√	√

Operation	GeminiDB FullAccess	GeminiDB ReadOnlyAccess
Modifying a parameter template	√	x
Deleting a parameter template	√	x
Querying enterprise project quotas	√	√
Modifying enterprise project quotas	√	x
Enabling or disabling SSL	√	x
Stopping a backup	√	x

Table 7-3 lists common GeminiDB operations and corresponding actions. You can refer to this table to customize permission policies.

Table 7-3 Common operations and supported actions

Operation	Actions	Authorization Scope	Remarks
Accessing the instance creation page	• vpc:vpcs:list • vpc:subnets:get • vpc:securityGroups:get	Supported: • IAM projects • Enterprise projects	The VPC, subnet, and security group are displayed on the instance creation page.
Creating an instance	• nosql:instance:create • vpc:vpcs:list • vpc:vpcs:get • vpc:subnets:get • vpc:securityGroups:get • vpc:ports:get	Supported: • IAM projects • Enterprise projects	If the default VPC, subnet, and security group are used, the **vpc:*:create** permission must be configured. To create an encrypted instance, you need to configure the **KMS Administrator** permission for the project.
Querying instances	nosql:instance:list	Supported: • IAM projects • Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Querying details of an instance	nosql:instance:list	Supported: • IAM projects • Enterprise projects	If the VPC, subnet, and security group need to be displayed on the instance details page, add the **vpc:*:get** and **vpc:*:list** actions.
Querying tasks	nosql:task:list	Supported: • IAM projects • Enterprise projects	-
Deleting an instance	nosql:instance:delete	Supported: • IAM projects • Enterprise projects	When deleting an instance, you need to delete the IP address on the data side.
Restarting an instance	nosql:instance:restart	Supported: • IAM projects • Enterprise projects	-
Resetting a password	nosql:instance:modifyPasswd	Supported: • IAM projects • Enterprise projects	-
Changing a security group	nosql:instance:modifySecurityGroup	Supported: • IAM projects • Enterprise projects	-
Changing a database port	nosql:instance:modifyPort	Supported: • IAM projects • Enterprise projects	-
Binding an EIP	nosql:instance:bindPublicIp	Supported: • IAM projects	When binding an EIP, you need to query the created EIP. • Enterprise projects are not supported. • Fine-grained authorization is not supported. For details, see **Floating IP Address**.

Operation	Actions	Authorization Scope	Remarks
Unbinding an EIP	nosql:instance:unbindPublicIp	Supported: • IAM projects	• Enterprise projects are not supported. • Fine-grained authorization is not supported. For details, see **Floating IP Address**.
Scaling up storage space	nosql:instance:modifyStorageSize	Supported: • IAM projects • Enterprise projects	-
Changing specifications	nosql:instance:modifySpecification	Supported: • IAM projects • Enterprise projects	-
Adding nodes	• nosql:instance:extendNode • vpc:vpcs:list • vpc:vpcs:get • vpc:subnets:get • vpc:securityGroups:get • vpc:ports:get	Supported: • IAM projects • Enterprise projects	-
Deleting nodes	nosql:instance:reduceNode	Supported: • IAM projects • Enterprise projects	Deleting nodes from a cluster
Modifying a backup policy	nosql:instance:modifyBackupPolicy	Supported: • IAM projects • Enterprise projects	-
Renaming an instance	nosql:instance:rename	Supported: • IAM projects • Enterprise projects	-
Creating a manual backup	nosql:backup:create	Supported: • IAM projects • Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Querying backups	nosql:backup:list	Supported: • IAM projects • Enterprise projects	-
Downloading a backup file	nosql:backup:download	Supported: • IAM projects • Enterprise projects	-
Restoring data to a new instance	• nosql:backup:restoreToNewInstance • vpc:vpcs:list • vpc:vpcs:get • vpc:subnets:get • vpc:securityGroups:get • vpc:ports:get	Supported: • IAM projects • Enterprise projects	The KMS Administrator permission needs to be configured for the encrypted instance in the project.
Restoring data to an existing instance	nosql:backup:restoreToExistInstance	Supported: • IAM projects • Enterprise projects	-
Deleting a backup	nosql:backup:delete	Supported: • IAM projects • Enterprise projects	-
Creating a parameter template	nosql:param:create	Supported: • IAM projects • Enterprise projects	-
Querying parameter templates	nosql:param:list	Supported: • IAM projects • Enterprise projects	-
Changing parameter values in a parameter template	nosql:param:modify	Supported: • IAM projects • Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Changing parameter settings of an instance node	nosql:instance:modifyParameter	Supported: • IAM projects • Enterprise projects	-
Deleting a parameter template	nosql:param:delete	Supported: • IAM projects • Enterprise projects	-
Performing an operation on tags	• nosql:instance:tag • tms:resourceTags:list	Supported: • IAM projects • Enterprise projects	-
Tag list	• nosql:tag:list • tms:resourceTags:list	Supported: • IAM projects • Enterprise projects	-
Querying enterprise project quotas	nosql:quota:list	Supported: • IAM projects • Enterprise projects	-
Modifying enterprise project quotas	nosql:quota:modify	Supported: • IAM projects • Enterprise projects	-
Enabling or disabling local log auditing	nosql:instance:switchAuditLog	Supported: • IAM projects • Enterprise projects	-
Downloading audit logs	nosql:instance:downloadAuditLog	Supported: • IAM projects • Enterprise projects	-
Deleting audit logs	nosql:instance:deleteAuditLog	Supported: • IAM projects • Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Enabling or disabling the display of slow query logs in plaintext	nosql:instance:modifySlowLogPlaintextSwitch	Supported: • IAM projects • Enterprise projects	-
Enabling or disabling SSL	nosql:instance:switchSSL	Supported: • IAM projects • Enterprise projects	-
Changing a private IP address	nosql:instance:modifyPrivateIp	Supported: • IAM projects • Enterprise projects	-
Switching between primary and standby instances	nosql:instance:switchover	Supported: • IAM projects • Enterprise projects	-
Upgrading the minor version of an instance	nosql:instance:upgradeDatabaseVersion	Supported: • IAM projects • Enterprise projects	-
Stopping a backup	nosql:backup:stop	Supported: • IAM projects • Enterprise projects	-
Querying a log group	lts:groups:get	Supported: • IAM projects • Enterprise projects	-
Querying a log stream	lts:topics:get	Supported: • IAM projects • Enterprise projects	-

Helpful Links

- IAM Service Overview
- Create a User and Assign Permissions
- Supported Actions

8 Regions and AZs

Concepts

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 defined by their 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 as universal regions and dedicated regions. A universal region provides universal cloud services for common tenants. A dedicated region provides only services of the same type or provides services only for specific tenants.
- An AZ contains one or multiple physical data centers. Each AZ has independent cooling, fire extinguishing, moisture-proof, and electrical 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 so you can build cross-AZ high-availability systems.

Figure 8-1 shows the relationship between regions and AZs.

Figure 8-1 Regions and AZs

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

Selecting a Region

When selecting a region, consider the following factors:

- Location

  Select the region closest to you or your target users to reduce likelihood of latency issues. Chinese mainland regions provide basically the same infrastructure, BGP network quality, as well as operations and configurations on resources. If you or your target users are in the Chinese mainland, it is not necessary to consider network latency differences when selecting a region.

- Resource price

  Resource prices may vary by region. For details, see **Product Pricing Details**.

Selecting an AZ

When determining whether to deploy resources in the same AZ, consider your application's requirements on disaster recovery (DR) and network latency.

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

Selecting a Region and Endpoint

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

9 Related Services

Figure 9-1 shows the relationship between GeminiDB and other services.

Figure 9-1 Relationship between GeminiDB and other services

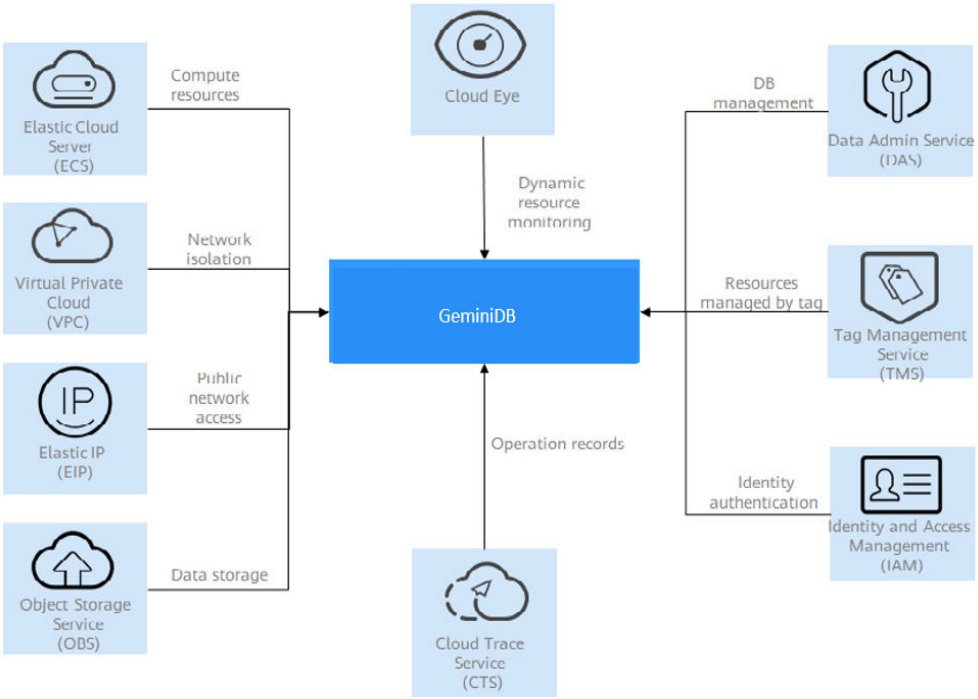

Table 9-1 Relationship between GeminiDB and other services

Service	Relationship with GeminiDB
ECS	Elastic Cloud Server (ECS) provides elastic computing resources for GeminiDB and a running environment for DB instances.

Service	Relationship with GeminiDB
VPC	GeminiDB uses Virtual Private Clouds (VPCs) and network security groups to keep instances isolated. VPCs allow you to define which IP addresses are allowed to access a given DB instance. Running an instance in a VPC improves security.
EIP	The Elastic IP (EIP) service provides independent public IP addresses and bandwidth for Internet access.
OBS	Backups are stored in Object Storage Service (OBS) to allow for disaster recovery and save space.
CTS	Cloud Trace Service (CTS) records operations related to GeminiDB for your later query, audit, and backtracking.
IAM	Identity and Access Management (IAM) provides permission management for GeminiDB.
TMS	Tag Management Service (TMS) enables you to use tags to manage resources on the management console. TMS works with other cloud services to manage global tags, and other cloud services manage their own tags.

