GeminiDB

Service Overview

Issue	07
Date	2025-09-04

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions

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....5

3 Highlights....6

4 Typical Application Scenarios....8

5 Security....10

5.1 Shared Responsibilities....10

5.2 Identity Authentication and Access Control....12

5.3 Data Protection....12

5.4 Audit and Logs....14

5.5 Resilience....15

5.6 Risk Monitoring....16

5.7 Fault Recovery....17

5.8 Certificates....18

6 Billing....20

7 Permissions....22

8 Regions and AZs....35

9 Related Services....37

1 What Is GeminiDB?

GeminiDB is a distributed, multi-model NoSQL database service with decoupled storage and compute. GeminiDB instances can be deployed, backed up, or restored quickly. They deliver high availability, top-tier security, seamless scalability, and exceptional performance to meet diverse needs. Alarms will be triggered whenever certain metric thresholds are reached by your instance.

GeminiDB is compatible with mainstream NoSQL databases (Redis, DynamoDB, Cassandra, HBase, InfluxDB, and MongoDB). GeminiDB delivers high I/O performance at low costs, suitable 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: Redis	GeminiDB provides high-concurrency and low-latency service access and exceptional scalability to effortlessly absorb large bursts of traffic. Common user scenarios include gaming, RTA-based advertising, recommendation systems, e-commerce, and education.	GeminiDB Redis API is a scalable key-value (KV) database service compatible with Redis. A GeminiDB Redis instance delivers superior performance and can process data volumes even when they exceed the available 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 DynamoDB-Compatible API	Key-value: DynamoDB	GeminiDB DynamoDB-Compatible API is a key-value database service and is suitable for scenarios such as high-concurrency web applications, IoT, e-commerce, and retail. GeminiDB DynamoDB-Compatible API delivers consistent low-latency performance along with robust data storage and advanced query capabilities.	GeminiDB DynamoDB-Compatible API communicates with DynamoDB over HTTPS. You can run CQL queries and access GeminiDB DynamoDB-Compatible API through the SDK or CLI. GeminiDB DynamoDB-Compatible API is completely compatible with Amazon DynamoDB, so you can smoothly migrate data from DynamoDB to GeminiDB without refactoring.

API	Compatible API	Scenario	Description
GeminiDB Cassandra API	Wide-column: Cassandra	GeminiDB Cassandra instances can store terabytes of data, handle millions of QPS, and ensure strong consistency. It is ideal for high-capacity storage needs across industries that rely on large clusters, such as industrial manufacturing, meteorology, and Internet.	GeminiDB Cassandra API is a cloud-native in-house NoSQL database service with decoupled storage and compute. It is compatible with the Cassandra ecosystem and supports CQL, which gives you SQL-like syntax. It is secure, reliable, scalable, and easy to manage and can provide high I/O performance.
GeminiDB HBase API	Wide-column: HBase	GeminiDB HBase API excels in scenarios demanding extensive data storage, real-time queries, exceptional reliability, and seamless scalability. It is particularly suited for environments leveraging Hadoop and other advanced big data processing technologies.	GeminiDB HBase API is a distributed NoSQL database service compatible with the HBase ecosystem. It features high performance, high reliability, and robust scalability.

API	Compatible API	Scenario	Description
GeminiDB Influx API	Time series: 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 service, which uses an in-house architecture with decoupled storage and compute. GeminiDB Influx API supports highly-concurrent I/O requests, compressed storage, SQL-like queries, multi-dimensional aggregation computing, and GUI-based data analysis. It provides high write performance, robust scalability, high compression rate, and high query performance.
GeminiDB Mongo API	Document-oriented: 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 volumes of IoT/IoV data, making it an excellent choice for sectors like Internet, IoT, gaming, and finance.	GeminiDB Mongo API is a cloud-native NoSQL database service using Huawei's proprietary architecture with decoupled storage and compute. This scalable reliable database provides enterprise-grade performance and can be managed on the console.

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

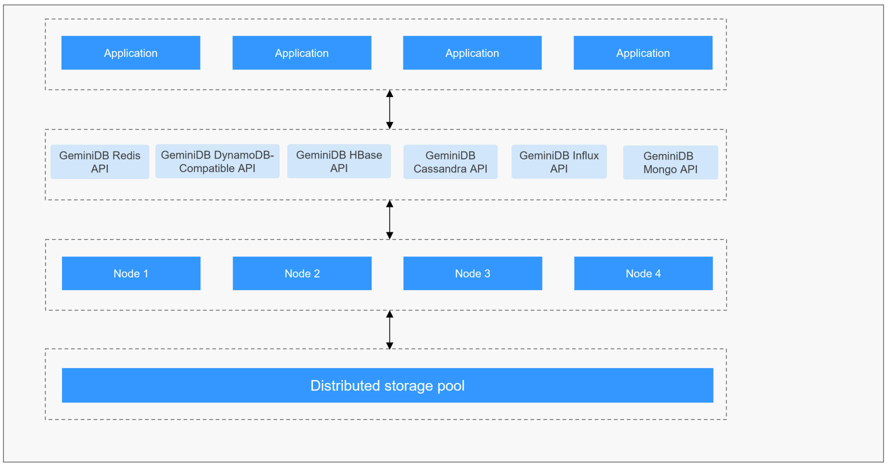

For details, see **Architecture Overview**.

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 improve the DR capability and save storage. 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 VPCs and 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 console and upload it for authentication when connecting 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 a server 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**: Within the first 6 hours after a game launches, 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**: You can 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 round-the-clock writes of massive volumes of data.

Internet

E-commerce and entertainment websites that include product catalogs, recommendations, and personalization engines use GeminiDB when they need to read and write data faster and demand 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 real-time recommendations.

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 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.

Unlike traditional on-premises data centers, cloud computing separates operators from users. This approach not only enhances flexibility and control for users but also greatly reduces their operational workload. For this reason, cloud security cannot be fully ensured by one party. Cloud security requires joint efforts of Huawei Cloud and you, as shown in **Figure 5-1**.

- **Huawei Cloud**: Huawei Cloud is responsible for infrastructure security, including security and compliance, regardless of cloud service categories. The infrastructure consists of physical data centers, which house compute, storage, and network resources, virtualization platforms, and cloud services Huawei Cloud provides for you. In PaaS and SaaS scenarios, Huawei Cloud is responsible for security settings, vulnerability remediation, security controls, and detecting any intrusions into the network where your services or Huawei Cloud components are deployed.
- **Customers**: As our customer, your ownership of and control over your data assets will not be transferred under any cloud service category. Without your explicit authorization, Huawei Cloud will not use or monetize your data, but you are responsible for protecting your data and managing identities and access. This includes ensuring the legal compliance of your data on the cloud, using secure credentials (such as strong passwords and multi-factor authentication), and properly managing those credentials, as well as monitoring and managing content security, looking out for abnormal account behavior, and responding to it, when discovered, in a timely manner.

Figure 5-1 Huawei Cloud shared security responsibility model

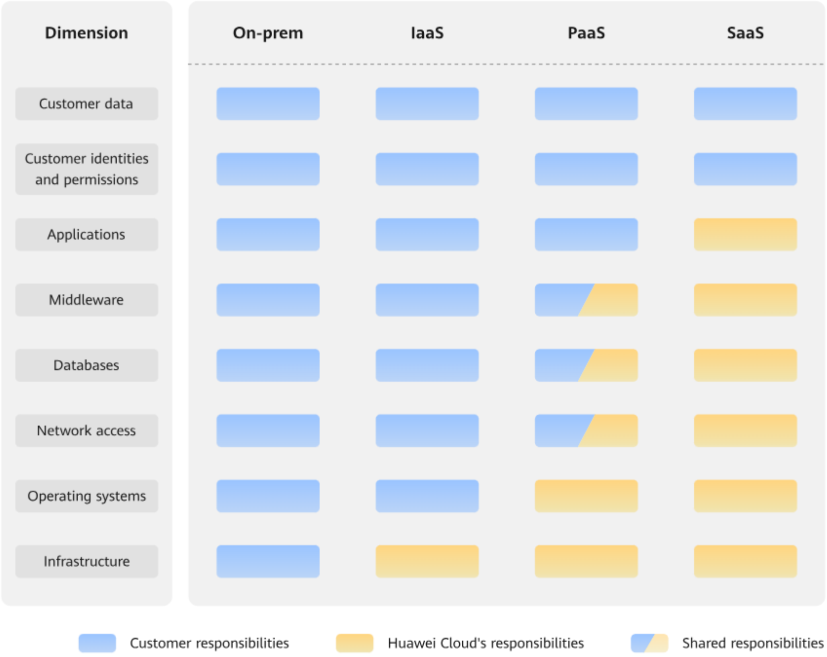

Dimension	On-prem	IaaS	PaaS	SaaS
Customer data	🟦	🟦	🟦	🟦
Customer identities and permissions	🟦	🟦	🟦	🟦
Applications	🟦	🟦	🟦	🟨
Middleware	🟦	🟦	◩	🟨
Databases	🟦	🟦	◩	🟨
Network access	🟦	🟦	◩	🟨
Operating systems	🟦	🟦	🟨	🟨
Infrastructure	🟦	🟨	🟨	🟨

🟦 Customer responsibilities 🟨 Huawei Cloud's responsibilities ◩ Shared responsibilities

Cloud security responsibilities are determined by control, visibility, and availability. When you migrate services to the cloud, assets, such as devices, hardware, software, media, VMs, OSs, and data, are controlled by both you and Huawei Cloud. This means that your responsibilities depend on the cloud services you select. As shown in **Figure 5-1**, customers can select different cloud service types (such as IaaS, PaaS, and SaaS services) based on their service requirements. As control over components varies across different cloud service categories, the responsibilities are shared differently.

- In on-premises scenarios, customers have full control over assets such as hardware, software, and data, so tenants are responsible for the security of all components.
- In IaaS scenarios, customers have control over all components except the underlying infrastructure. So, customers are responsible for securing these components. This includes ensuring the legal compliance of the applications, maintaining development and design security, and managing vulnerability remediation, configuration security, and security controls for related components such as middleware, databases, and operating systems.
- In PaaS scenarios, customers are responsible for the applications they deploy, as well as the security settings and policies of the middleware, database, and network access under their control.
- In SaaS scenarios, customers have control over their content, accounts, and permissions. They need to protect their content, and properly configure and protect their accounts and permissions in compliance with laws and regulations.

5.2 Identity Authentication and Access Control

Identity Authentication

When you access GeminiDB, the system authenticates your identity in either of the following ways.

- **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, GeminiDB Influx, GeminiDB DynamoDB-Compatible, and GeminiDB HBase instances support both SSL and non-SSL connections. SSL is recommended to secure data transmission.	- GeminiDB Redis instance: **Configuring an SSL Connection** - GeminiDB Influx instance: **Configuring an SSL Connection** - GeminiDB Cassandra instance: **Configuring an SSL Connection** - GeminiDB DynamoDB-Compatible instance: **Configuring an SSL Connection** - GeminiDB HBase instance: **Configuring an SSL Connection** - GeminiDB Mongo instance: **Configuring an SSL Connection**
Cross-AZ deployment	GeminiDB allows you to deploy nodes of an instance evenly across three AZs to provide cross-AZ DR. These AZs are physically isolated using hash functions and can be connected over private networks.	- GeminiDB Redis instance: **Buying an Instance** - GeminiDB Influx instance: **Buying an Instance** - GeminiDB Cassandra instance: **Buying an Instance** - GeminiDB DynamoDB-Compatible instance: **Buying an Instance** - GeminiDB HBase instance: **Buying an Instance** - GeminiDB Mongo instance: **Buying an Instance**
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 instance: **Connecting to an Instance Using a Load Balancer Address**

Feature	Description	Reference
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 instance: **Creating a DR Instance**
Cross-region active-active DR	GeminiDB Cassandra API supports active-active DR and bidirectional synchronization between two instances in different regions. Once an instance becomes faulty, the other takes over its workloads to ensure service continuity.	GeminiDB Cassandra instance: **Cross-region Dual-active DR**
Deletion protection	You can move unsubscribed yearly/monthly and deleted pay-per-use GeminiDB instances to the recycle bin. Instances in the recycle bin can be retained for up to 7 days.	● GeminiDB Redis instance: **Recycling an Instance** ● GeminiDB Influx instance: **Recycling an Instance** ● GeminiDB Cassandra instance: **Recycling an Instance** ● GeminiDB DynamoDB-Compatible instance: **Recycling an Instance** ● GeminiDB HBase instance: **Recycling an Instance** ● GeminiDB Mongo instance: **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 API management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.
- For details about GeminiDB Influx API management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.
- For details about GeminiDB Cassandra API management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**
- For details about GeminiDB DynamoDB-Compatible API management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.
- For details about GeminiDB Mongo API management and data traces that can be tracked by CTS, see **Key Operations Supported by CTS**.

Logs

- GeminiDB Redis

  You can view slow query logs of databases to optimize slow SQL statements. Any query that takes longer than an execution time threshold (in milliseconds) will be logged.

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

  You can view slow query logs of databases to optimize slow SQL statements. Any query that takes longer than an execution time threshold (in milliseconds) will be logged.

  For details about slow query logs, see **Slow Query Logs**.
- GeminiDB DynamoDB-Compatible API

  You can view slow query logs of databases to optimize slow SQL statements. Any query that takes longer than an execution time threshold (in milliseconds) will be logged.

  For details about slow query logs, see **Slow Query Logs**.
- GeminiDB Mongo
  - You can view slow query logs of databases to optimize slow SQL statements. Any query that takes longer than an execution time threshold (in milliseconds) will be logged.

    For details about slow query logs, see **Slow Query Logs**.
  - You can view warning- and error-level logs generated while databases are running to analyze system issues.

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

5.5 Resilience

- GeminiDB Redis API uses the DFV pool with three-copy storage, so data can be persistently stored in real time. It offers solutions like multi-AZ deployment, failover in seconds, load balancing, and node scale-in, to guarantee instance reliability and availability.
- GeminiDB Influx API uses the DFV pool with three-copy storage and high write performance. It offers solutions like multi-AZ deployment and autoscaling to guarantee instance reliability and availability.

- GeminiDB Cassandra API uses the DFV pool with three-copy storage and supports 24/7 online writes. It offers solutions like intra-region DR, cross-region dual-active DR, multi-AZ deployment, *N*-1 fault tolerance, and autoscaling to guarantee instance reliability and availability.
- GeminiDB DynamoDB-Compatible API uses the DFV pool with three-copy storage. Data can be seamlessly migrated without refactoring. GeminiDB DynamoDB-Compatible API also offers solutions like multi-AZ deployment, *N*-1 fault tolerance, and autoscaling to guarantee instance reliability and availability.
- GeminiDB HBase API uses the DFV pool with three-copy storage. Failover can be completed within seconds. GeminiDB HBase API also offers solutions like multi-AZ deployment, *N*-1 fault tolerance, and autoscaling to guarantee instance reliability and availability.
- GeminiDB Mongo API uses the DFV pool with three-copy storage and supports 24/7 real-time writes. It offers solutions like multi-AZ deployment, *N*-1 fault tolerance, and autoscaling to guarantee instance reliability and availability.

Figure 5-2 Reliability architecture

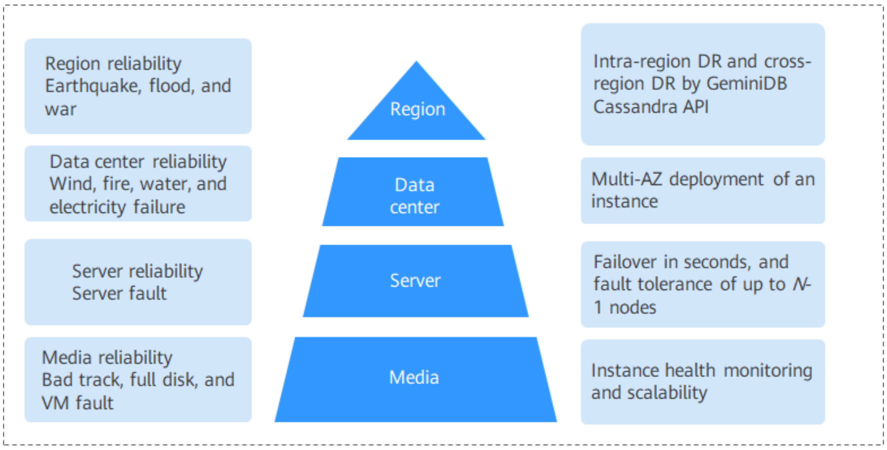
Region reliability
Earthquake, flood, and war
Region
Intra-region DR and cross-region DR by GeminiDB Cassandra API
Data center reliability
Wind, fire, water, and electricity failure
Data center
Multi-AZ deployment of an instance
Server reliability
Server fault
Server
Failover in seconds, and fault tolerance of up to N-1 nodes
Media reliability
Bad track, full disk, and VM fault
Media
Instance health monitoring and scalability

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 API metrics and how to create alarm rules, see **Supported Metrics**.
- For details about GeminiDB Influx API metrics and how to create alarm rules, see **Supported Metrics**.
- For details about GeminiDB Cassandra API metrics and how to create alarm rules, see **Supported Metrics**.
- For details about GeminiDB DynamoDB-Compatible API metrics and how to create alarm rules, see .**Supported Metrics**.

- For details about GeminiDB HBase API metrics and how to create alarm rules, see **Supported Metrics**.
- For details about GeminiDB Mongo API metrics and how to create alarm rules, see **Supported Metrics**.

Critical Operation Protection

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).

- For details about how to restore GeminiDB Redis instance data, see **Restoring Data to a New Instance**.
- For details about how to restore GeminiDB Influx instance data, see **Restoring Data to a New Instance**.
- For details about how to restore GeminiDB Cassandra instance data, see **Restoring Data to a New Instance** and **Restoring a Backup to a Specified Point in Time**.
- For details about how to restore GeminiDB DynamoDB-Compatible instance data, see **Restoring Data to a New Instance** and **Restoring a Backup to a Specified Point in Time**.
- For details about how to restore GeminiDB HBase instance data, see **Restoring Data to a New Instance** and **Restoring a Backup to a Specified Point in Time**.
- For details about how to restore GeminiDB Mongo instance data, see **Restoring Data to a New or an 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 Active-Active DR

GeminiDB Cassandra supports active-active DR and bidirectional synchronization between two instances in different regions. Once an instance becomes faulty, the other takes over its workloads 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

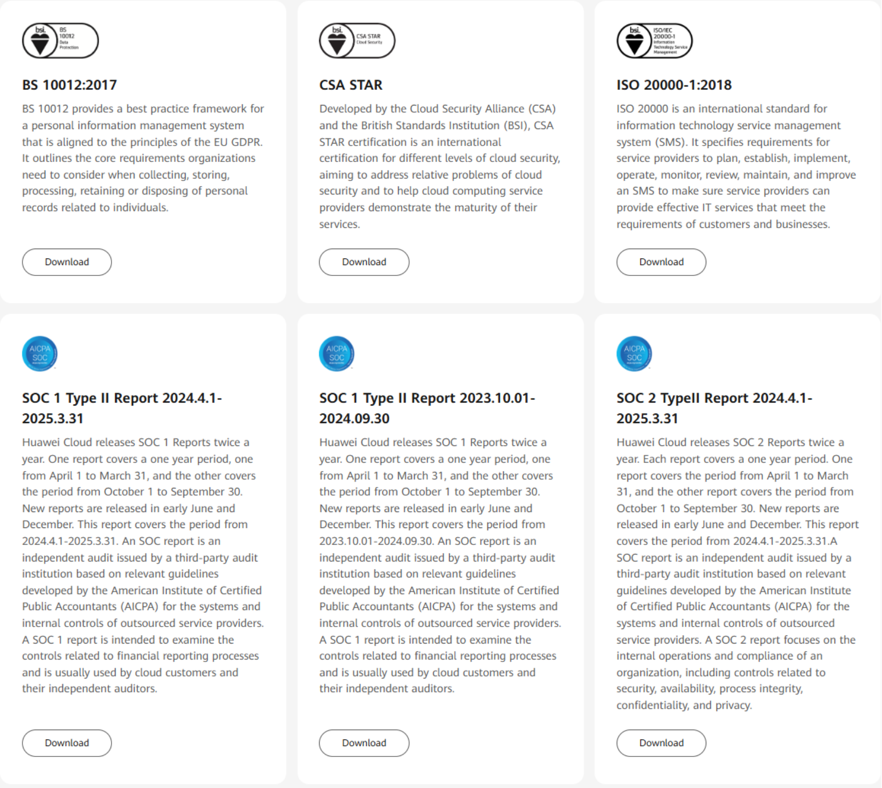

BS 10012:2017

BS 10012 provides a best practice framework for a personal information management system that is aligned to the principles of the EU GDPR. It outlines the core requirements organizations need to consider when collecting, storing, processing, retaining or disposing of personal records related to individuals.

CSA STAR

Developed by the Cloud Security Alliance (CSA) and the British Standards Institution (BSI), CSA STAR certification is an international certification for different levels of cloud security, aiming to address relative problems of cloud security and to help cloud computing service providers demonstrate the maturity of their services.

ISO 20000-1:2018

ISO 20000 is an international standard for information technology service management system (SMS). It specifies requirements for service providers to plan, establish, implement, operate, monitor, review, maintain, and improve an SMS to make sure service providers can provide effective IT services that meet the requirements of customers and businesses.

SOC 1 Type II Report 2024.4.1-2025.3.31

Huawei Cloud releases SOC 1 Reports twice a year. One report covers a one year period, one from April 1 to March 31, and the other covers the period from October 1 to September 30. New reports are released in early June and December. This report covers the period from 2024.4.1-2025.3.31. An SOC report is an independent audit issued by a third-party audit institution based on relevant guidelines developed by the American Institute of Certified Public Accountants (AICPA) for the systems and internal controls of outsourced service providers. A SOC 1 report is intended to examine the controls related to financial reporting processes and is usually used by cloud customers and their independent auditors.

SOC 1 Type II Report 2023.10.01-2024.09.30

Huawei Cloud releases SOC 1 Reports twice a year. One report covers a one year period, one from April 1 to March 31, and the other covers the period from October 1 to September 30. New reports are released in early June and December. This report covers the period from 2023.10.01-2024.09.30. An SOC report is an independent audit issued by a third-party audit institution based on relevant guidelines developed by the American Institute of Certified Public Accountants (AICPA) for the systems and internal controls of outsourced service providers. A SOC 1 report is intended to examine the controls related to financial reporting processes and is usually used by cloud customers and their independent auditors.

SOC 2 TypeII Report 2024.4.1-2025.3.31

Huawei Cloud releases SOC 2 Reports twice a year. Each report covers a one year period. One report covers the period from April 1 to March 31, and the other report covers the period from October 1 to September 30. New reports are released in early June and December. This report covers the period from 2024.4.1-2025.3.31.A SOC report is an independent audit issued by a third-party audit institution based on relevant guidelines developed by the American Institute of Certified Public Accountants (AICPA) for the systems and internal controls of outsourced service providers. A SOC 2 report focuses on the internal operations and compliance of an organization, including controls related to security, availability, process integrity, confidentiality, and privacy.

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

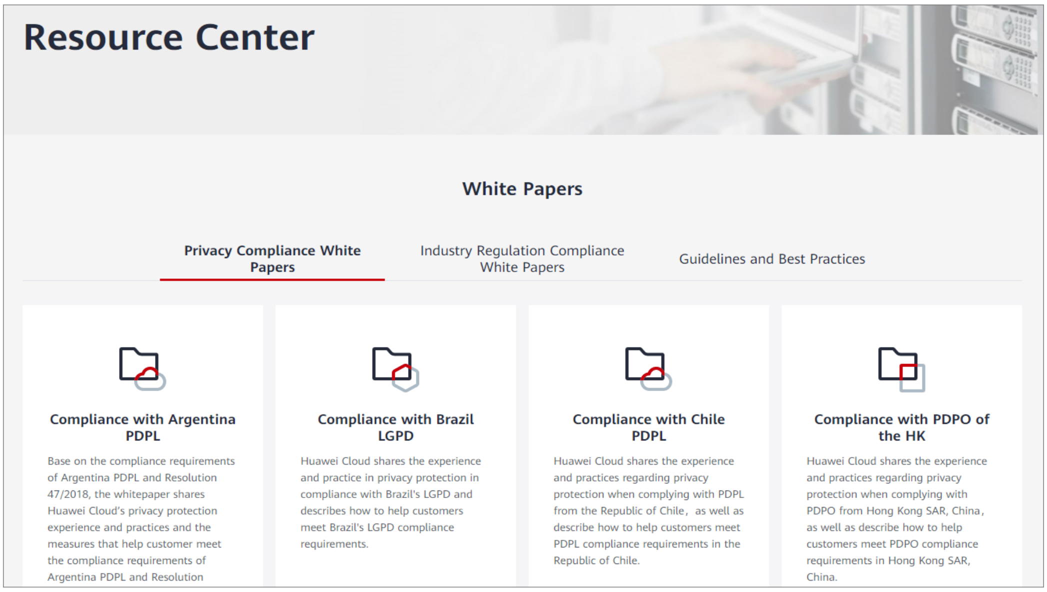

Resource Center

White Papers

Privacy Compliance White Papers
Industry Regulation Compliance White Papers
Guidelines and Best Practices

Compliance with Argentina PDPL

Base on the compliance requirements of Argentina PDPL and Resolution 47/2018, the whitepaper shares Huawei Cloud's privacy protection experience and practices and the measures that help customer meet the compliance requirements of Argentina PDPL and Resolution

Compliance with Brazil LGPD

Huawei Cloud shares the experience and practice in privacy protection in compliance with Brazil's LGPD and describes how to help customers meet Brazil's LGPD compliance requirements.

Compliance with Chile PDPL

Huawei Cloud shares the experience and practices regarding privacy protection when complying with PDPL from the Republic of Chile，as well as describe how to help customers meet PDPL compliance requirements in the Republic of Chile.

Compliance with PDPO of the HK

Huawei Cloud shares the experience and practices regarding privacy protection when complying with PDPO from Hong Kong SAR, China，as well as describe how to help customers meet PDPO compliance requirements in Hong Kong SAR, China.

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 buy an instance with 100 GB of storage, you will get additional 100 GB of backup storage at no extra cost. If your backup data exceeds 100 GB, the backup storage consumed in excess will be charged according to OBS billing rules.
Public network traffic	You will be billed for the bandwidth you use to access GeminiDB instances over a public network. You are not billed for traffic generated when accessing GeminiDB instances within your own private network.

Billing Mode

GeminiDB provides two billing modes: pay-per-use and yearly/monthly. You can switch between the two billing modes.

- Yearly/Monthly: If your future usage is predictable, this billing mode is generally less expensive than pay-per-use. Longer subscriptions offer larger discounts.
- Pay-per-use (hourly): You are only billed for how long you have actually used your instance. This mode can be a good option when future requirements are unpredictable. Pay-per-use instances are priced by the hour, but if an instance is used for less than one hour, you will be billed based on the actual duration.

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 by at least 1 GB each time. After that, you will be billed for the new storage. Storage cannot be scaled down.

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 GeminiDB, but they are not allowed to delete the resources or perform any high-risk operations. To achieve this, you can create IAM users for the developers and grant them only the permissions for using GeminiDB 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 system permissions of GeminiDB.

Table 7-1 GeminiDB system permissions

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 system permissions

Operation	GeminiDB FullAccess	GeminiDB ReadOnlyAccess
Creating an instance	√	x
Querying the instance list	√	√
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 the backup list	√	√
Restoring data to a new instance	√	x
Deleting a backup	√	x
Creating a parameter template	√	x
Querying the parameter template list	√	√

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 restore data to an encrypted instance, you need to configure the KMS administrator permission for the selected project.
Querying the instance list	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	You need to delete the data node IP address.
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 created EIPs. ● 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	You can delete 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 the backup list	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	To restore data to an encrypted instance, you need to configure the KMS administrator permission for the selected 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 the parameter template list	nosql:param:list	Supported: ● IAM projects ● Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Changing parameter values in a parameter template	nosql:param:modify	Supported: - IAM projects - Enterprise projects	-
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	-
Viewing the 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 audit logs	nosql:instance:switchAuditLog	Supported: - IAM projects - Enterprise projects	-

Operation	Actions	Authorization Scope	Remarks
Downloading audit logs	nosql:instance:downloadAuditLog	Supported: ● IAM projects ● Enterprise projects	-
Deleting audit logs	nosql:instance:deleteAuditLog	Supported: ● IAM projects ● Enterprise projects	-
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	-
Patching a database	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	-

Operation	Actions	Authorization Scope	Remarks
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, compute, 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

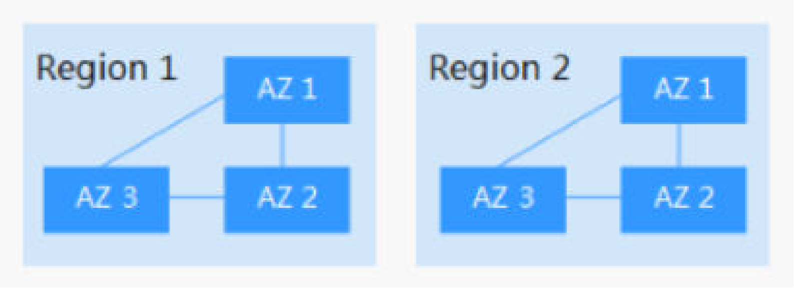

Huawei Cloud provides services in many regions around the world. You can select a region and AZ as needed. For more information, see **Huawei Cloud Global Regions**.

Selecting a Region

When selecting a region, consider the following factors:

- Location

  Select a region closest to your target users for low network latency and quick access. 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.

- To achieve robust DR, deploy resources in different AZs in the same region.
- To reduce 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

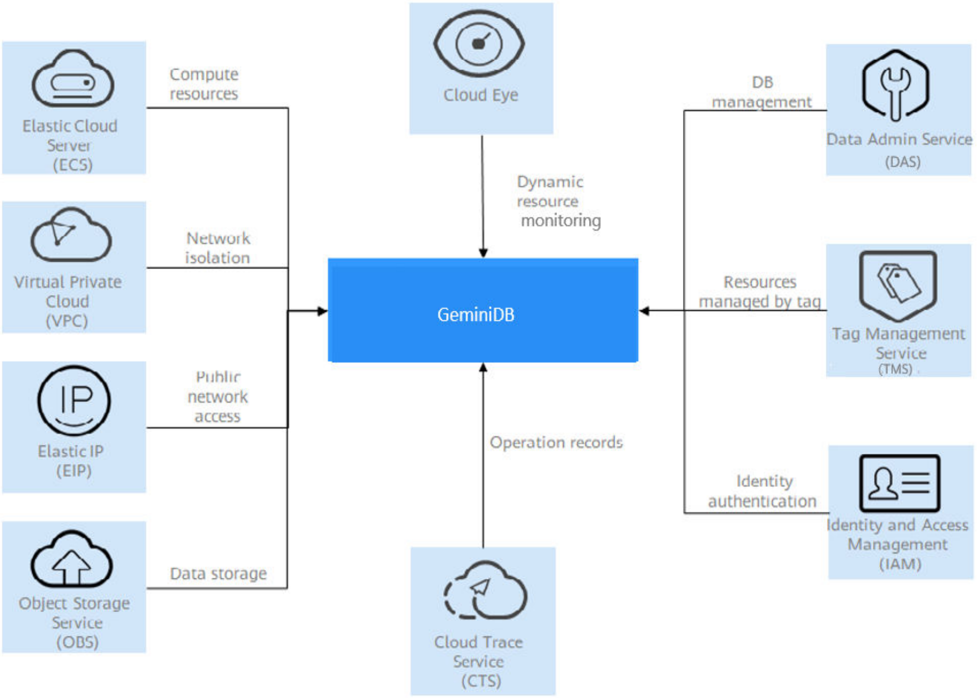

Table 9-1 Relationship between GeminiDB and other services

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

Service	Relationship Between GeminiDB and Other Services
VPC	GeminiDB uses VPCs and 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.
EIP	The Elastic IP (EIP) service provides independent public IP addresses and bandwidth for Internet access.
OBS	Backups are stored in OBS buckets to improve the DR capability and save storage.
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.