GeminiDB HBase

User Guide

Issue 01

Date 2025-07-30





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

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

Trademarks and Permissions

HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd. All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, quarantees 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 Service Overview	1
1.1 What Is GeminiDB HBase API?	1
1.2 Compatible API and Version	3
1.3 Instance Specifications	6
1.4 Database Rules	8
1.4.1 Basic Design	8
1.5 Constraints	8
2 Billing	12
2.1 Billing Overview	12
2.2 Billing Modes	13
2.2.1 Overview	13
2.2.2 Yearly/Monthly Billing	14
2.2.3 Pay-per-Use Billing	18
2.3 Billing Items	22
2.4 Billing Examples	24
2.5 Changing Billing Modes	27
2.5.1 Overview	27
2.5.2 Changing a Pay-per-Use Instance to Yearly/Monthly	27
2.5.3 Changing a Yearly/Monthly Instance to Pay-per-Use	29
2.6 Renewal	30
2.6.1 Overview	30
2.6.2 Manually Renewing an Instance	32
2.6.3 Auto-Renewal	35
2.7 Bills	37
2.8 Arrears	41
2.9 Billing Termination	43
2.10 Cost Management	45
2.10.1 Costs	45
2.10.2 Cost Allocation	45
2.10.3 Cost Analysis	
2.10.4 Cost Optimization	
2.11 Billing FAQs	
2.11.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing?	47

2.11.2 Can I Switch Between Yearly/Monthly and Pay-per-use Billing Modes?	47
2.11.3 How Do I Renew a Single or Multiple Yearly/Monthly Instances?	48
2.11.4 How Do I Unsubscribe from a Yearly/Monthly Instance?	49
3 Getting Started with GeminiDB HBase API	52
3.1 Getting to Know GeminiDB HBase API	
3.2 Buying and Connecting to a GeminiDB HBase Instance	52
3.3 How Can I Delete Rows Based On Prefixes?	62
4 Working with GeminiDB HBase API	64
4.1 Using IAM to Grant Access to GeminiDB HBase API	64
4.1.1 Creating a User and Granting Permissions to Use GeminiDB HBase API	
4.1.2 Custom Policies for GeminiDB HBase API	65
4.2 Buying a GeminiDB HBase Instance	67
4.3 Instance Connection and Management	74
4.3.1 Connecting to a GeminiDB HBase Instance	74
4.3.2 Connecting to a GeminiDB HBase Instance over a Private Network	76
4.3.3 Connecting to a GeminiDB HBase Instance Using Program CodeCode	77
4.3.3.1 Connecting to a GeminiDB HBase Instance Using Java	77
4.3.4 How Can I Connect to a GeminiDB HBase Instance over TLS (SSL)?	78
4.3.5 Connection Information Management	80
4.3.5.1 Setting Security Group Rules for a GeminiDB HBase Instance	80
4.3.5.2 Viewing the IP Address and Port of a GeminiDB HBase Instance	83
4.3.5.3 Changing the Security Group of a GeminiDB HBase Instance	84
4.4 Data Migration	85
4.4.1 Solution Overview	85
4.5 Instance Lifecycle Management	85
4.5.1 Restarting a GeminiDB HBase instance	86
4.5.2 Exporting Instance Information	87
4.5.3 Deleting a Pay-per-Use Instance	87
4.5.4 Recycling an Instance	88
4.6 Instance Modifications	89
4.6.1 Upgrading a Minor Version	89
4.6.2 Changing an Instance Name	91
4.6.3 Resetting the Administrator Password	92
4.6.4 Changing vCPUs and Memory	93
4.6.5 Setting a Maintenance Window	96
4.6.6 Adding and Deleting Instance Nodes	98
4.6.6.1 Overview	98
4.6.6.2 Manually Adding Instance Nodes	
4.6.6.3 Automatically Adding Instance Nodes	102
4.6.6.4 Manually Deleting Instance Nodes	104
4.6.7 Scaling Storage Space	105
4.6.7.1 Overview	105

4.6.7.2 Manually Scaling Up Storage Space	107
4.6.7.3 Automatically Scaling Up Storage Space	108
4.6.7.4 Manually Scaling Down Storage Space	112
4.7 Data Backup	113
4.7.1 Overview	113
4.7.2 Managing Automated Backups	116
4.7.3 Managing Manual Backups	125
4.7.4 Managing Cross-Region Backups	127
4.7.5 Managing Table-level Backups	131
4.8 Data Restoration	135
4.8.1 Restoration Methods	135
4.8.2 Restoring a Backup to a New Instance	135
4.8.3 Restoring a Backup to a Specified Point in Time	137
4.9 Logs and Audit	138
4.10 Viewing Metrics and Configuring Alarms	138
4.10.1 Supported Metrics	138
4.10.2 Configuring Alarm Rules	145
4.10.3 Viewing Metrics	151
4.10.4 Event Monitoring	152
4.10.4.1 Introduction	152
4.10.4.2 Viewing Event Monitoring Data	152
4.10.4.3 Creating an Alarm Rule for Event Monitoring	153
4.10.4.4 Events Supported by Event Monitoring	155
4.11 Enterprise Project	168
4.11.1 Overview	168
4.11.2 Quota Management	168
4.12 Tag Management	170
4.13 User Resource Quotas	173
5 Best Practices	175
5.1 How Do I Set Pre-partition Keys When Creating a Table on a GeminiDB HBase Instance?	
6 Performance White Paper	
6.1 Performance Test Methods	
6.2 Performance Test Data	179
7 FAQs	181
7.1 Product Consulting	181
7.1.1 What Are the Precautions for Using GeminiDB HBase API?	181
7.1.2 What Is GeminiDB HBase Instance Availability?	181
7.2 Billing	181
7.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing Modes?	182
7.2.2 Can I Switch Between Yearly/Monthly and Pay-per-use Billing Modes?	182
7.3 Database Connection	182

7.3.1 How Can I Create and Connect to an ECS?	182
7.3.2 Can I Change the VPC After Buying a GeminiDB HBase Instance?	182
7.4 Backup and Restoration	182
7.4.1 How Long Can GeminiDB HBase Instance Backups Be Stored?	183
7.5 Regions and AZs	183
7.5.1 What Is AZ and How Can I Select an AZ?	183
7.5.2 Can Different AZs Communicate with Each Other?	183
7.5.3 Can I Change the Region After Buying a GeminiDB HBase Instance?	183
7.6 Instance Freezing, Release, Deletion, and Unsubscription	184

Service Overview

1.1 What Is GeminiDB HBase API?

GeminiDB HBase API is based on GeminiDB Cassandra API. They share some cluster components, but GeminiDB HBase API introduces its own advanced features. You can use the open-source HBase Java SDK or HBase Shell to access GeminiDB HBase API. Apache HBase Driver can be directly connected using the right protocol, so you can smoothly migrate data to GeminiDB HBase instances without refactoring.

GeminiDB HBase API strictly complies with the HBase syntax and data model. Therefore, Apache HBase applications can be easily migrated to GeminiDB HBase instances. GeminiDB HBase API also provides multiple automated management and O&M functions, such as cluster scaling in minutes, automated backup, fault detection, and multi-AZ fault tolerance. GeminiDB HBase API frees you from complex O&M and parameter tuning of open-source HBase clusters.

Architecture

Database programs on each GeminiDB HBase instance node provide standard HBase functions. A shared storage pool provides file services. You only need to connect to an open port of HBase on any node to use its functions.




Figure 1-1 GeminiDB HBase API architecture

Highlights

GeminiDB HBase API is a cloud-native NoSQL database service compatible with open-source HBase. With robust security and reliability, GeminiDB HBase API offers ultra-high performance and addresses pain points of open-source HBase.

- Security and reliability
 - A multi-layer security system, including VPCs, subnets, security groups, SSL, and fine-grained permission control, ensures database security and user privacy.
 - Cross-region active-active DR is supported. You can deploy an instance across three AZs and quickly back up or restore data.
 - The distributed architecture ensures fault tolerance for *N-1* nodes.
- Future-proof ecosystem
 - GeminiDB HBase API can directly connect to Apache HBase Driver. For details about the compatibility list, see Compatible API and Version.
 - Data can be seamlessly migrated from open-source HBase and HBaselike databases. GeminiDB HBase API can interact with peripheral components.
- Enhanced capabilities
 - Data recovery capabilities such as second-level flashback and Point-In-Time Recovery (PITR) ensure high data reliability.
 - Data can be quickly deleted by searching for prefixes.
 - Data shard failover takes just a few seconds. GeminiDB HBase API quarantees a shorter MTTR than open-source HBase.
- Superb performance
 - GeminiDB HBase API provides more than twice higher performance than open-source HBase. You can check the performance white paper.
- No pain points of open-source software
 - This out-of-the-box service eliminates challenges like complex parameter tuning in open-source software.
 - Storage can be scaled in seconds without interrupting services.
 - Compute nodes can be added in minutes. A jitter may last only a few seconds.

NOTICE

GeminiDB HBase API is based on GeminiDB Cassandra API.

Typical Scenarios

Internet

GeminiDB HBase API offers superior read/write performance, high availability, dynamic scalability, and high fault tolerance. It can handle concurrent requests at low latency, ideal for Internet websites with large data volumes,

for example, product catalogs, recommendations, personalized engines, and transaction records.

Advantages

Large-scale clusters

A single cluster can contain up to 100 nodes, which is ideal for Internet applications with large data volumes and write workloads.

High availability and scalability

A faulty node does not affect the entire cluster. Compute nodes and storage can be quickly added without interrupting services.

High-concurrency writes

Powerful write performance helps you handle a huge number of concurrent e-commerce transactions.

Industrial data collection

GeminiDB HBase API is compatible with the HBase ecosystem and integrates data collected from multiple terminal sources. It imports data and stores collected metrics in real time. GeminiDB HBase API works with upstream and downstream big data components to implement aggregation analysis and real-time statistics.

Advantages:

Large-scale clusters

GeminiDB HBase API is suitable for industrial manufacturing scenarios where a large number of metrics need to be collected and stored.

High availability and performance

Data can be written to databases around the clock.

Fast backup and restoration

Storage snapshots speed up backup and recovery.

Scale-out in minutes

Nodes can be added in minutes to effortlessly handle surges in jobs and projects.

1.2 Compatible API and Version

This section describes the compatibility of GeminiDB HBase API.

Currently, GeminiDB HBase API does not support APIs related to cluster O&M and management. You can perform operations on an instance on the console, for example, restarting an instance.

Table 1-1 Compatible version

Compatible API	Instance Type	Version
Open-source HBase	Cluster	2.6.1

Table 1-2

API Function Category	АРІ	Function	Supported
Data APIs	Get	Get: queries a single row.	Yes
		Filter: filters server data.	SingleColumnValu eFilter and PageFilter are supported. Other filters are in OBT.
		Consistency: ensures strong data consistency.	Yes
		Versions: queries multiple versions.	Yes
		Batch: reads rows in batches.	Yes
	Put	Put: inserts a single row.	Yes
		Condition: inserts a condition.	Yes
		TTL: automatically deletes expired data.	Yes
		Batch: writes rows in batches.	Yes
	Versions: multiversion write Delete Delete: single row deletion Delete: deletes data by specifying ColumnFamily or Qualifier.	Yes	
			Yes
		data by specifying ColumnFamily or	Yes
		Versions: multi- version deletion	Yes
			No

API Function Category	API	Function	Supported
		Prefix Delete: deletes data based on prefixes.	Yes
	Scan	Scan: performs a full scan.	Yes
		Scan: scans data in specified rows.	Yes
		TimeRange: filters data in a time range.	Yes
			SingleColumnValu eFilter and PageFilter are supported. Other filters are in OBT.
		Versions: queries multiple versions.	Yes
		Reversed: queries data in reverse order.	Yes
	Increment	Performs a common Increment operation.	No
	Append	Performs a common Append operation.	No
	Bulk Load	Imports data from files in batches.	No
Metadata operations	Table	Create: creates a table.	Yes
		SplitKey: specifies a key during table creation.	Yes
		Region: automatically splits regions.	Yes
		Disable/Delete: disables or deletes a table.	Yes

API Function Category	API	Function	Supported
		Compress: specifies a compression algorithm (data compression is available).	No
		Alter Table: adds ColumnFamily.	Yes
		Alter Table: reduces ColumnFamily.	No
	Cluster management API	You can ignore this	API.
	Namespace	Creates or deletes a namespace.	Yes
Others	Coprocessor	Supports user- defined plugins.	No
	Secondary index	Supports diverse column queries.	No

1.3 Instance Specifications

Each instance type comes with various specifications based on memory configurations.

This section describes GeminiDB HBase instance specifications. The instance specifications depend on the selected CPU model.



The specifications listed in the following table are the same as those of GeminiDB Cassandra instances. For details, see the specifications displayed on the console.

Table 1-3 GeminiDB HBase cluster instance specifications

CPU Type	Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)
x86	geminidb.cassandra.la rge.4	2	8	10	96,000

CPU Type	Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)
	geminidb.cassandra.xl arge.4	4	16	10	96,000
	geminidb.cassandra.xl arge.8	4	32	10	96,000
	geminidb.cassandra.2x large.4	8	32	10	96,000
	geminidb.cassandra.2x large.8	8	64	10	96,000
	geminidb.cassandra.4x large.4	16	64	10	96,000
	geminidb.cassandra.4x large.8	16	128	10	96,000
	geminidb.cassandra.6x large.8	24	192	10	144,000
	geminidb.cassandra.8x large.4	32	128	10	192,000
	geminidb.cassandra.8x large.8	32	256	10	192,000

Table 1-4 Specifications of a GeminiDB HBase instance with cloud native storage

Data Node Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)
geminidb.cassandra- geminifs.large.4	2	8	10	64000
geminidb.cassandra- geminifs.xlarge.4	4	16	10	64000
geminidb.cassandra- geminifs.2xlarge.4	8	32	10	64000
geminidb.cassandra- geminifs.4xlarge.4	16	64	10	64000
geminidb.cassandra- geminifs.8xlarge.4	32	128	10	64000

1.4 Database Rules

1.4.1 Basic Design

Design Rules

Rule 1: Do not store big data such as images and files in databases.

Rule 2: The maximum size of the key and value in a single row cannot exceed 64 KB, and the average size of rows cannot exceed 10 KB.

Rule 3: A data deletion policy must be specified for a table to prevent data from growing infinitely.

Rule 4: Design partition keys to evenly distribute workloads and avoid data skew. For details about how to set partition keys, see **How Do I Set Pre-partition Keys When Creating a Table on a GeminiDB HBase Instance?**

1.5 Constraints

The following tables list the constraints designed to ensure stability and security of GeminiDB HBase instances.

Specifications

Table 1-5 Specifications

Resource Type	Specifications	Description
CPU and memory	GeminiDB HBase cluster instances are supported.	 For details about specifications of different instance types, see Instance Specifications.
		You can change the specifications to meet your service requirements by following Changing vCPUs and Memory.
Storage	The storage space depends on the selected instance specifications.	Storage can be scaled up or down. For details, see Overview .

Quotas

Table 1-6 Quotas

Resource Type	Constraint	Description
Tag	A maximum of 20 tags can be added for each instance.	For more information, see Tag Management .
Free backup storage	GeminiDB HBase instances provide free backup storage.	For more information, see Backup Storage .
Retention period	The default value is 7 days. The value ranges from 1 to 3660 days.	For more information, see Configuring an Automated Backup Policy.

Naming Rules

Table 1-7 Naming rules

Item	Description	
Instance name	 Can contain 4 to 64 characters. Must start with a letter. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed. 	
Backup name	 Can contain 4 to 64 characters. Must start with a letter. Only letters (case sensitive), digits, hyphens (-), and underscores (_) are allowed. 	
Parameter template name	 Can contain 1 to 64 characters. Only letters (case sensitive), digits, hyphens (-), underscores (_), and periods (.) are allowed. 	

Security

Table 1-8 Security

Item	Description	
Password of database administrator rwuser	 Can contain 8 to 32 characters. Can contain at least two types of the following characters: uppercase letters, lowercase letters, digits, and special characters ~!@#%^*=+? For more information, see Resetting the Administrator Password. Keep your password secure. The system cannot retrieve it if it is lost. 	
Database port	Database port number. When you create a GeminiDB HBase instance, the port cannot be specified. By default, three ports are occupied: 2181, 16000, and 16020. The port cannot be changed after the instance is created. Use the default port 2181 for connection.	
VPC	After a GeminiDB HBase instance is created, the VPC where the instance is deployed cannot be changed.	
Security group	A security group controls access between GeminiDB HBase API and other services. Ensure that the security group you selected allows your client to access the instance.	

Instance Operations

Table 1-9 Instance operations

Function	Constraint	
Database access	A GeminiDB HBase instance can access an ECS only when they are in the same VPC and subnet.	
	The security group must allow access from the associated ECS. By default, a GeminiDB HBase instance cannot be accessed through an ECS in a different security group. You need to add an inbound rule to the security group.	
	The default port of a GeminiDB HBase instance is 2181, which is also used during the connection.	
	The database port cannot be set during instance creation and cannot be changed after the instance is created.	

Function	Constraint	
Instance deployment	The servers where instances are deployed are not directly visible to you. You can only access the instances through IP addresses and database ports.	
Restarting a GeminiDB HBase instance	GeminiDB HBase instances cannot be rebooted through commands. They must be rebooted on the console.	
	 Restarting an instance will interrupt services, so off- peak hours are the best time. Ensure that your application can be reconnected. 	
Viewing GeminiDB HBase instance backups	GeminiDB HBase instance backups are stored in OBS buckets and are invisible to you.	
Changing the CPU or memory of a GeminiDB HBase instance	Second-level intermittent disconnection occurs once when the specifications are changed on a single node. Therefore, the entire instance is intermittently disconnected several times. Ensure that the client can be reconnected. You are advised to change the specifications during off-peak hours.	
	 For a node whose specifications are being changed, its computing tasks are handed over to other nodes. Change specifications of nodes during off-peak hours to prevent instance overload. 	
Data restoration	To prevent data loss, you are advised to back up key data before data restoration.	
Storage	If the instance storage is full, data cannot be written to databases. You are advised to periodically check the storage.	
	GeminiDB HBase instance storage can be automatically scaled up in case of a sudden surge in data volumes. Enable autoscaling by following Automatically Scaling Up Storage Space.	
Recycle bin	• You can move unsubscribed yearly/monthly instances and deleted pay-per-use instances to the recycle bin. You can restore an instance that was deleted up to 7 days ago from the recycle bin.	
	• The recycle bin is enabled by default and cannot be disabled. Instances in the recycle bin can be retained for 7 days by default. This function is free of charge.	
	Currently, you can put a maximum of 100 instances into the recycle bin.	
	If you delete an instance running out of storage, it will not be moved to the recycle bin.	

2 Billing

2.1 Billing Overview

In this document, you will learn about how GeminiDB HBase instances are billed, how you can renew subscriptions and manage costs, and what happens if your account goes into arrears.

◯ NOTE

GeminiDB HBase and GeminiDB Cassandra instances have the same billing modes.

Billing modes

There are two billing modes: yearly/monthly and pay-per-use.

- Yearly/Monthly: You are billed in advance based on the duration you plan to use the instance. Ensure your top-up account has enough funds or a valid payment method pre-configured.
- Pay-per-use: You are billed for your immediate needs based on how long you have actually used a product.

For details about the two billing modes, see Overview.

You can also change the billing mode later if it no longer meets your needs. For details, see **Overview**.

• Billing items

You will be billed for instance specifications, storage, backup storage, and EIP bandwidth. For details about the billing factors and formulas of each billing item, see **Billing Items**.

For more information about the billing samples and the billing for each item, see **Billing Examples**.

Renewal

After a yearly/monthly instance expires, it cannot run properly. To continue using the instance, renew it within the specified period. Otherwise, compute and storage resources will be automatically released and data may be lost.

You can renew the instance manually or automatically. For details, see **Overview**.

Bills

You can choose **Billing & Costs** > **Bills** to check the instance transactions and bills. For details, see **Bills**.

Arrears

If the available account balance is less than amount to be settled and there is no other payment method configured, your account will go into arrears. To continue using your cloud services, top up your account in a timely manner. For details, see **Arrears**.

• Billing termination

If your instance is no longer used, you can unsubscribe from or delete it to avoid unexpected bills. For details, see **Billing Termination**.

• Cost management

Runtime costs are primarily composed of resource and O&M costs. You can allocate, analyze, and optimize instance costs to save more money. For details, see **Cost Management**.

2.2 Billing Modes

2.2.1 Overview

There are two billing modes: yearly/monthly and pay-per-use.

- Yearly/Monthly is a prepaid billing mode. You pay in advance for a subscription term, and in exchange, you get a discounted rate. The longer the subscription term, the bigger the discount. Yearly/Monthly billing is a good option for long-term, stable services.
- Pay-per-use is a postpaid billing mode. You pay as you go and just pay for what you use. The instance usage is calculated by the second but billed every hour. Pay-per-use billing is a good option for scenarios where there are sudden traffic bursts, such as e-commerce promotions.

Table 2-1 lists their differences.

Table 2-1 Billing modes

Billing Mode	Yearly/Monthly	Pay-per-use
Payment	Prepaid Settled based on the subscription period specified in your order.	Postpaid Settled based on how long you have actually used your instance.
Billing Cycle	Billed based on the subscription period specified in your order.	Pricing is listed on a per-hour basis, and bills are calculated down to the second.
Billing Item	Instance specifications (vCPUs and memory), storage, backup storage, and EIP	Instance specifications (vCPUs and memory), storage, backup storage, and EIP

Billing Mode Change	A yearly/monthly instance can be changed to pay-per-use. The change is applied only after the yearly/monthly instance expires. For details, see Changing a Yearly/Monthly Instance to Pay-per-Use.	A pay-per-use instance can be changed to yearly/monthly. For details, see Changing a Pay-per-Use Instance to Yearly/Monthly.
Specificati on Change	Supported	Supported
Scenario	Maintains a long-term commitment and offers significant discounts compared to pay-per-use billing.	Aligns with your scaling needs of compute resources. You can create or delete your instance anytime.

2.2.2 Yearly/Monthly Billing

Yearly/Monthly is ideal for users who are confident in their long-term needs and want to secure a lower price. You will enjoy a discount on a yearly/monthly instance. This section describes billing rules for yearly/monthly GeminiDB HBase instances.

Application Scenarios

If you want to ensure resource stability over a certain period of time, yearly/monthly billing is a good choice for the following types of workloads:

- Long-term workloads with stable resource requirements, such as official websites, online malls, and blogs.
- Long-term projects, such as scientific research projects and large-scale events.
- Workloads with predictable traffic bursts, for example, e-commerce promotions or festivals.
- Workloads with high data security requirements.

Billing Items

You are billed for the following items on a yearly/monthly basis.


Table 2-2 Items billed on a yearly/monthly basis

Billing Item	Description
Instance specificatio ns	vCPUs and memory

Billing Item	Description
Storage	If the actual storage usage exceeds your purchased storage, you will be billed for additional storage on a pay-per-use basis.
Backup storage	GeminiDB HBase API provides free backup storage equal to the amount of storage you purchased.
	After the free backup storage is used up, additional usage will incur bills based on the backup storage pricing details. Pricing is listed on a per-hour basis, and bills are calculated down to the second. If the instance has been used less than one hour, you will be billed based on the actual duration.

Assume that you want to buy a three-node GeminiDB HBase instance with 2 vCPUs, 8 GB of memory, and 100 GB of storage. At the bottom of the page for buying an instance, price details (excluding the backup storage fee) will be displayed.

Figure 2-1 Example price



Price \$433.06 USD ③

You are billed for:

- Selected specifications
- Storage

□ NOTE

The backup storage fee is not included. After the free backup storage is used up, additional usage will incur bills based on the backup storage pricing details. For details, see **Product Pricing Details**.



Billing Cycle

The billing cycle of a yearly/monthly GeminiDB HBase instance is determined by the duration (UTC+8) you commit to the subscription. The billing starts from when you activated or renewed the subscription and ends at 23:59:59 of the expiry date.

For example, if you bought a GeminiDB HBase instance for one month at 15:50:04 on March 8, 2023, the billing cycle is from 15:50:04 on March 8, 2023 to 23:59:59 on April 8, 2023.

Billing Examples

Assume that you bought a three-node GeminiDB HBase instance with 2 vCPUs, 8 GB of memory, 100 GB of storage, 110 GB (100 GB for free) of backup storage for one month at 15:50:04 on March 8, 2023 and renewed the subscription for another month before the instance expired.

- 15:50:04, March 8, 2023 to 23:59:59, April 8, 2023
- 23:59:59, April 8, 2023 to 23:59:59, May 8, 2023
 - From 23:59:59 on April 8, 2023 to 23:59:59 on May 1, 2023, 50 GB of free backup storage was used.
 - From 23:59:59 on May 1, 2023 to 23:59:59 on May 8, 2023, another 10 GB of backup storage was used, which was billed for 168 hours.

You will be billed for each billing cycle. Resources of the GeminiDB HBase instance are billed separately.

Table 2-3 Formulas for billing yearly/monthly resources


Resource	Formula	Unit Price
Instance specifications (including vCPUs and memory)	Unit price of the instance specifications x Required duration x Number of nodes	See the estimated price of a cluster instance in GeminiDB Price Calculator.
Storage space	Storage space unit price x Required duration x Storage space (GB)	See the estimated price of a cluster instance with specified storage in GeminiDB Price Calculator .
Backup storage	Backup storage unit price x Required duration x (Backup storage – Storage) (GB) NOTE The billed duration refers to	See the estimated price of an instance with specified backup storage in GeminiDB Price Calculator .
	the length of time the billed backup storage was used for.	

Figure 2-2 shows how the price is calculated.

NOTICE

The prices in the following figure are for reference only. For the actual prices, see **GeminiDB Price Calculator**.

Figure 2-2 Pricing example of a yearly/monthly GeminiDB HBase instance



Impact on Billing After Specification Changes

If the specifications of a yearly/monthly instance no longer meet your needs, you can change the specifications on the console. The system will re-calculate the price based on the following rules:

- If you upgrade instance specifications, you need to pay the difference in price.
- If you downgrade instance specifications, Huawei Cloud will refund you the difference.

Downgrading instance specifications will affect instance performance. You are not advised to do so. Assume that you bought a yearly/monthly three-node GeminiDB HBase instance with 2 vCPUs and 8 GB of memory on April 8, 2023 and upgraded its specifications to 4 vCPUs and 16 GB of memory on the same day. The old specifications cost USD587.06/month and the new specifications USD981.62/month. The calculation formula is as follows:

Price difference for the specification upgrade = Price for the new specifications × Remaining period - Price for the original specifications × Remaining period


The remaining period in the formula is the number of remaining days in each calendar month divided by the maximum number of days in the calendar month. In this example, the remaining period is 0.6581 (12/30 + 8/31). The upgrade fee is USD259.66 (981.62 x $0.6581 - 587.06 \times 0.6581$).

For more information, see **Pricing of a Changed Specification**.

Impact of Expiration

Figure 2-3 shows the statuses of a yearly/monthly GeminiDB HBase instance. After an instance is purchased, it enters the valid period and runs normally during this period. If the instance is not renewed after it expires, before being deleted, it first enters a grace period and then a retention period.

Figure 2-3 Lifecycle of a yearly/monthly GeminiDB HBase instance



Expiration reminder

From the 7th day before a yearly/monthly GeminiDB HBase instance expires, the system will send an expiration reminder to the creator of the Huawei Cloud account by email, SMS, and in-app message.

Impact

If your yearly/monthly instance is not renewed after it expires, its status changes to **Expired** and it enters a grace period. During the grace period, you can access the instance but cannot:

- Change instance specifications.
- Change the billing mode from yearly/monthly to pay-per-use.
- Unsubscribe from the instance.

If the yearly/monthly instance is not renewed after the grace period ends, its status turns to **Frozen** and it enters a retention period. You cannot perform any operations on the instance while it is in the retention period.

If the yearly/monthly instance is not renewed by the time the retention period ends, it will be released and data cannot be restored.

◯ NOTE

• For details about the renewal, see Overview.

2.2.3 Pay-per-Use Billing

Pay-per-use billing means you pay nothing up front and are not tied into any contract or commitment. This section describes billing rules of pay-per-use GeminiDB HBase instances.

Application Scenarios

Pay-per-use billing is good for short-term, bursty, or unpredictable workloads that cannot tolerate any interruptions, such as applications for e-commerce flash sales, temporary testing, and scientific computing.

Billing Items

You are billed for the following items on a pay-per-use basis.

Table 2-4 Items billed on a pay-per-use basis

Billing Item	Description
Instance specificatio ns	vCPUs and memory
Storage	Instance storage space, which is billed hourly on a pay-per-use basis.
Backup storage	GeminiDB HBase API provides free backup storage equal to the amount of storage you purchased.
	After the free backup storage is used up, additional usage will incur bills based on the backup storage pricing details. These fees are settled by the hour. If it has been used less than one hour, you will be billed based on the actual duration.

Assume that you want to buy a three-node GeminiDB HBase instance with 2 vCPUs, 8 GB of memory, and 500 GB of storage. At the bottom of the page for buying an instance, price details (excluding the backup storage fee) will be displayed.

Figure 2-4 Example price

Price \$1.17 USD/hour ②

The price includes:

- Instance specifications (including vCPUs and memory)
- Selected storage space

■ NOTE

The backup space fee is not included. For details about the backup price, see **Product Pricing Details**.





Billing Cycle

Pricing is listed on a per-hour basis (GMT+8), and bills are calculated down to the second. The billing starts when the instance is created and ends when it is deleted.

■ NOTE

It takes a certain time to create an instance. The billing starts from the time when the instance is successfully created. You can view the two time points on the **Basic Information** page. You can view the time when the instance is created beside the **Created** field.

For example, if you buy a pay-per-use GeminiDB HBase instance at 8:45:30 and deleted it at 8:55:30, you are billed for the 600 seconds from 8:45:30 to 8:55:30. The billing items include compute resources (vCPUs and nodes), storage, and backup storage.

Billing Examples

Assume that you bought a pay-per-use 3-node instance with 2 vCPUs, 8 GB of memory, 100 GB of storage, and 110 GB of backup storage (100 GB for free) at 09:59:30 on April 18, 2023 and deleted the instance at 10:45:46 on the same day. The billing items include compute resources (vCPUs and nodes) and storage.

- Usage of 30 seconds from 9:59:30 to 10:00:00
- Usage of 2,746 seconds from 10:00:00 to 10:45:46
 - The free backup storage is used from 10:00:00 to 10:45:00.
 - 10 GB of backup storage is billed for 46 seconds from 10:45:00 to 10:45:46.

You will be billed for each billing cycle. Resources of the GeminiDB HBase instance are billed separately. The price per hour is shown in the pricing details. To calculate the price per second, divide it by 3,600.

Table 2-5 Billing formulas

Resource	Formula	Unit Price
Compute resources (including vCPUs and nodes)	Unit price of instance specifications x Required duration	See the estimated price of a cluster instance in GeminiDB Price Calculator.
Storage	Storage unit price x Required duration	See the estimated price of a cluster instance with specified storage in GeminiDB Price Calculator .
Backup storage	Backup storage unit price x Required duration x (Backup storage – Storage) (GB) NOTE The billed duration refers to the length of time the billed backup storage was used for.	See the estimated price of an instance with specified backup storage in GeminiDB Price Calculator.


Figure 2-2 shows how the price is calculated.

NOTICE

The prices in the following figure are for reference only. For the actual prices, see **GeminiDB Price Calculator**.

For pay-per-use billing, decimal numerals on the price calculator are rounded off and are accurate to two decimal places. If the rounded price falls below USD0.01, it will be shown as USD0.01.

Figure 2-5 Billing example of a pay-per-use GeminiDB HBase instance



Impact on Billing After Specification Changes

If you change the specifications of a pay-per-use instance, the original order will become invalid and a new order will be placed. You will be billed based on the new specifications.

If you change instance specifications within a given hour, multiple records will be generated. Different records record the billing for different specifications.


For example, if you buy a pay-per-use instance with 2 vCPUs and 8 GB of memory at 9:00:00 and upgrade it to 4 vCPUs and 16 GB of memory at 9:30:00, two billing records are generated between 9:00:00 and 10:00:00:

- 2 vCPUs and 8 GB of memory from 9:00:00 to 9:30:00
- 4 vCPUs and 16 GB of memory from 9:30:00 to 10:00:00

Impact of Arrears

Figure 2-6 shows the statuses of a pay-per-use GeminiDB HBase instance throughout its lifecycle. After you buy a pay-per-use GeminiDB HBase instance, it is valid within the billing cycle. If your account falls into arrears because of automatic fee deduction, it will transition into a grace period followed by a retention period.

Figure 2-6 Lifecycle of a pay-per-use GeminiDB HBase instance



Arrears reminder

The system will bill you for pay-per-use resources after each billing cycle ends. If your account goes into arrears, the system will send an email, SMS message, or in-app message to the one who created the Huawei Cloud account.

Impact

When your account falls into arrears due to automatic fee deduction for the payper-use GeminiDB HBase instance, your account will be marked as in arrears, and the instance will enter the grace period. After you top up your account, Huawei Cloud will bill you for expenditures generated during the grace period. You can view the charges on the **Billing Center** > **Overview** page.

If you do not pay the arrears within the grace period, your instance enters the retention period and its status changes to **Frozen**. You cannot perform any operations on the instance in the retention period.

If you do not pay the arrears within the retention period, your instance will be released, and data will be lost.

- During the retention period, you cannot access or use your instance but the data stored in it can be retained. The retention period for Huawei Cloud International website is 15 days.
- During the grace period, you can access and use only some resources of your instance. The grace period for Huawei Cloud International website is 15 days.
- For details about top-up, see Topping Up an Account.

2.3 Billing Items

Billing

You will be billed for instance specifications, storage, and backup storage. The following table lists the billing items.

□ NOTE

Billing items marked with an asterisk (*) are mandatory.

Table 2-6 Billing items of a GeminiDB HBase instance

Billing Item	Description	Billing Mode	Formula
* Specific ations	Billed by instance specifications, including vCPUs and memory. Compute and storage capabilities vary by the number of vCPUs and memory size.	Yearly/ Monthly and pay- per-use	Unit price x Required duration See the estimated price of a cluster instance in GeminiDB Price Calculator.
* Storage space	Billed based on unified standards.	Yearly/ Monthly and pay- per-use	Unit price x Storage space x Required duration See the estimated price of a cluster instance with specified storage in GeminiDB Price Calculator.
Backup storage	Billed based on unified standards.	Pay-per- use	Unit price x Billed backup storage x Required duration See the estimated price of an instance with specified backup storage in GeminiDB Price Calculator. NOTE The billed duration refers to the length of time the billed backup storage was used for.
(Option al) Cross- region backup	Billed based on unified standards.	Pay-per- use	Unit price x Storage x Required duration Unit price of storage: CNY0.0009/GB/hour
	Billed based on unified standards.	Billed by storage	Unit price of cross- region backup traffic x Storage Unit price of cross- region backup traffic: CNY0.5/GB

Billing Examples

Assume that you bought a three-node GeminiDB HBase instance with 2 vCPUs, 8 GB of memory, 100 GB of storage, 110 GB (100 GB for free) of backup storage for

one month at 15:50:04 on March 8, 2023 and renewed the subscription for another month before the instance expired.

15:50:04 on March 8, 2023 to 23:59:59 on April 8, 2023

23:59:59 on April 8, 2023 to 23:59:59 on May 8, 2023

From 23:59:59 on April 8, 2023 to 23:59:59 on May 1, 2023, 50 GB of free backup storage was used.


From 23:59:59 on May 1, 2023 to 23:59:59 on May 8, 2023, another 10 GB of backup storage was used, which was billed for 168 hours.

Figure 2-7 shows how the total price is calculated.

NOTICE

Prices in the figure are only for reference. For details, see **Product Pricing Details**.

Figure 2-7 Pricing example of a yearly/monthly GeminiDB HBase instance



For more billing examples of a pay-per-use instance, see **Billing Examples**.

2.4 Billing Examples

Billing Scenario

A user bought a pay-per-use GeminiDB HBase instance with the following configurations at 15:30:00 on March 18, 2023:

- 2 vCPUs and 8 GB of memory
- Nodes: 3
- EIP bandwidth: 6 Mbit/s

After a period of time, the user found that the current instance specifications could not meet service requirements and expanded the specifications to 4 vCPUs and 16 of GB at 9:00:00 on March 20, 2023. The user changed the billing mode to

yearly/monthly at 10:30:00 on March 20, 2023 and paid the instance in advance for one month. So how much will the user be billed for this instance in March and April?

Billing Analysis

This GeminiDB HBase instance is billed in two phases:


- 15:30:00 on March 18, 2023 to 10:30:00 on March 20, 2023: pay-per-use
 - 15:30:00 on March 18, 2023 to 9:00:00 on March 20, 2023
 - Instance specifications: 2 vCPUs and 8 GB of memory
 - Nodes: 3
 - Storage: 100 GB
 - Backup storage: 100 GB
 - EIP bandwidth: 6 Mbit/s
 - 9:00:00 on March 20, 2023 to 10:30:00 on March 20, 2023
 - Instance specifications: 4 vCPUs and 16 GB of memory
 - Nodes: 3
 - Storage: 200 GB
 - Backup storage: 210 GB (pay-per-use from 10:00:00 to 10:30:00 on March 20, 2023)
 - EIP bandwidth: 6 Mbit/s
- 10:30:00 on Mar 20, 2023 to 23:59:59 on Apr 20, 2023: yearly/monthly
 - Instance specifications: 4 vCPUs and 16 GB of memory
 - Nodes: 3
 - Storage: 200 GB
 - Backup space: 300 GB (pay-per-use from 23:59:59 on April 10, 2023 to 23:59:59 on April 20, 2023)
 - EIP bandwidth: 6 Mbit/s
 - Billed duration: one month

NOTICE


Unit prices in this example are used for reference only. The prices shown here are only estimates. The unit price may be different from the fee calculated in the actual scenario. For details, see the data released on the Huawei Cloud official website.

Pay-per-use

From 15:30:00 on March 18, 2023 to 09:00:00 on March 20, 2023, an instance with 2 vCPUs and 8 GB of memory was used for 41.5 hours, so the price is calculated as follows:




From 9:00:00 to 10:30:00 on March 20, 2023, an instance with 4 vCPUs and 16 GB of memory was used for 1.5 hours, so the price is calculated as follows:



Yearly/Monthly

From March 20, 2023, 10:30:00 to April 20, 2023, 23:59:59, an instance purchased using yearly/monthly billing was used for one month, so the price would be calculated as follows.



From March to April, the total price of this instance is USD950.91 (40.51 + 2.82 + 907.58).

2.5 Changing Billing Modes

2.5.1 Overview

After buying a GeminiDB HBase instance, you can change its billing mode if it no longer meets your needs. **Table 2-7** lists changeable billing items of a GeminiDB HBase instance.

Table 2-7 Changeable billing items of a GeminiDB HBase instance

Billing Item	Change Description	Reference
Instance specification s (vCPUs and nodes)	 If you change the billing mode of a GeminiDB HBase instance, the billing mode of compute resources (vCPUs and nodes) will also be changed. Change a pay-per-use instance to yearly/monthly to enjoy discounts. Change a yearly/monthly instance to pay-per-use to use it more flexibly. NOTE Such a change is only applied after the yearly/monthly subscription expires. 	 Changing a Pay-per- Use Instance to Yearly/Monthly Changing a Yearly/ Monthly Instance to Pay-per-Use

2.5.2 Changing a Pay-per-Use Instance to Yearly/Monthly

To commit to a long subscription and save costs, you can change a pay-per-use GeminiDB HBase instance to yearly/monthly. Doing so will create an order. After you pay for the order, yearly/monthly billing will be applied immediately.

Suppose you bought a pay-per-use GeminiDB HBase instance at 15:29:16 on April 18, 2023 and changed it to yearly/monthly at 16:30:30 on the same day. After you

paid for the order, yearly/monthly billing was applied immediately. On the **Billing Center** > **Billing** page, three line items were generated.

- Pay-per-use expenditures for 15:29:16 to 16:00:00 on April 18, 2023
- Pay-per-use expenditures for 16:00:00 to 16:30:30 on April 18, 2023
- A yearly/monthly expenditure generated at 16:30:30 on April 18, 2023

Prerequisites

- The billing mode of the instance is pay-per-use.
- The instance status is Available.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and click **Change to Yearly/ Monthly** in the **Operation** column.

Figure 2-8 Change to Yearly/Monthly



Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Change to Yearly/Monthly** in the **Billing Mode** field.

Figure 2-9 Change to Yearly/Monthly



□ NOTE

The billing mode of multiple instances can be changed in batches. Perform the following steps:

- 1. Select the instances whose billing mode you want to change.
- 2. Click Change to Yearly/Monthly above the instance list.
- **Step 4** On the displayed page, specify a subscription duration in month. The minimum duration is one month.

If you do not need to modify your settings, click Pay to go to the payment page.

- **Step 5** Select a payment method and click **Confirm**.
- **Step 6** View the results on the **Instances** page.

In the upper right corner of the instance list, click G to refresh the list. The instance status will become **Available** after the change is successful. The billing mode changes to **Yearly/Monthly**.

----End

2.5.3 Changing a Yearly/Monthly Instance to Pay-per-Use

You can change a yearly/monthly GeminiDB HBase instance to pay-per-use to reclaim a portion of your subscription costs while adapting to your usage needs.

Suppose you bought a yearly/monthly GeminiDB HBase instance at 15:29:16 on April 18, 2023 and changed it to pay-per-use at 16:30:00 on May 18, 2023. On the **Billing Center > Billing** page, bill information is generated as follows:

- Yearly/Monthly expenditures for 15:29:16 on April 18 to 23:59:59 on May 18, 2023
- Pay-per-use expenditures for 23:59:59 on May 18, 2023 to the end time of pay-per-use billing. A bill was generated every hour.

The pay-per-use billing mode will take effect only after the yearly/monthly subscription has expired. Auto-renewal will not be in effect.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3 On the Instances page, locate the target instance and choose More > Change to Pay-per-Use in the Operation column.

Figure 2-10 Change to Pay-per-Use



The billing mode of multiple pay-per-use instances can be changed in batches. Perform the following steps:

- 1. Select the instances whose billing mode you want to change.
- 2. Click More > Change to Pay-per-Use in the Operation column
- **Step 4** On the displayed page, confirm the instance information and click **Change to Pay- per-Use**. The billing mode will change to pay-per-use after the instance expires. After the billing mode is changed, auto-renewal will be disabled.
- **Step 5** After you submit the change, check whether a message is displayed in the **Billing Mode** column, indicating that the billing mode will be changed to pay-per-use after the subscription expires.
- **Step 6** To cancel the change, choose **Billing** > **Renewal** to enter the Billing Center. On the **Renewals** page, locate the instance and click **More** > **Cancel Change to Payper-Use**.
- **Step 7** In the displayed dialog box, click **Yes**.

----End

2.6 Renewal

2.6.1 Overview

When to Renew Subscriptions

Expired yearly/monthly instances cannot run properly. To continue using it, renew your subscription within a specified period. Otherwise, resources such as vCPUs and memory will be automatically released, and data will be lost and cannot be restored.

Only yearly/monthly instance subscriptions can be renewed. If you use pay-per-use instances, just ensure that your account has a valid payment method configured or a top-up account with a sufficient balance.

If you renew the instance before it expires, resources will be retained and you can continue using the instance. For details about statuses of the expired instance, see **Impact of Expiration**.

How to Renew Subscriptions


You can renew a yearly/monthly instance manually or automatically.

Table 2-8 Renewing a yearly/monthly instance

Method	Description
Manually Renewing an Instance	You can renew a yearly/monthly instance anytime on the console before it is automatically deleted.
Auto-Renewal	You can enable auto-renewal to automatically renew the instance before it expires. This prevents resources from being deleted in case you forget to renew a subscription.

You can renew a yearly/monthly instance in different phases of its lifecycle. For details, see **Figure 2-11**.

Figure 2-11 Selecting a renewal method based on the instance's current phase



- An instance is in the **Provisioned** state after it is provisioned.
- When an instance subscription expires, the status will change from **Provisioned** to **Expired**.
- If an expired instance is not renewed, it enters a grace period. If it is not renewed by the time the grace period expires, the instance will be frozen and enter a retention period.
- If you do not renew the subscription before the retention period expires, your resources will be automatically deleted.

□ NOTE

- During the retention period, you cannot access or use your instance but the data stored in it can be retained. The retention period for Huawei Cloud International website is 15 days.
- During the grace period, you can access and use only some resources of your instance. The grace period for Huawei Cloud International website is 15 days.

You can enable auto-renewal any time before an instance expires. By default, the system will make the first attempt to charge your account for the renewal at 03:00, seven days before the expiry date. If this attempt fails, it will make another attempt at 03:00 every day until the subscription is renewed or expired. You can change the auto-payment date for renewal as required.

2.6.2 Manually Renewing an Instance

You can renew a yearly/monthly instance anytime on the console before it is automatically deleted.

Renewing an Instance on the Console


- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and click **Renew** in the **Operation** column.

Figure 2-12 Renewing an instance



Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 2-13 Renewing an instance



Ⅲ NOTE

You can also renew multiple instances all at once:

- 1. Select the yearly/monthly instances to be renewed.
- 2. Click **Renew** above the instance list.

Step 4 On the displayed page, renew the instances.

----End

Renewing an Instance in Billing Center


- Step 1 Log in to the Huawei Cloud console.
- **Step 2** Hover over **Billing & Costs** in the upper part of the console and choose **Renewal** from the drop-down list.

The **Renewals** page is displayed.

Step 3 Select the search criteria.

On the Manual Renewals, Auto Renewals, Pay-per-Use After Expiration, and Renewals Canceled pages, you can view the instances to be renewed.


Figure 2-14 Renewal management



You can move all resources to be manually renewed to the **Manual Renewals** tab. For details, see **Restoring to Manual Renewal**.


- **Step 4** Manually renew resources.
 - Individual renewal: Locate an instance that you want to renew and click **Renew** in the **Operation** column.

Figure 2-15 Individual renewal




• Batch renewal: Select multiple instances that you want to renew and click **Batch Renew** in the upper left corner.

Figure 2-16 Batch renewal



Step 5 Select a renewal duration and optionally select Renew on the standard renewal date. For details, see Setting the Same Renewal Day for Yearly/Monthly Resources. Confirm the price and click Pay.

Figure 2-17 Confirming renewal



Step 6 Select a payment method and make your payment. Once the order is paid for, the renewal is complete.

----End

Setting the Same Renewal Day for Yearly/Monthly Resources

If the instances have different expiry dates, you can set the same renewal day, for example, the first day of each month, to make it easier to manage renewals.

In Figure 2-18, a user sets the same renewal day for two resources that will expire at different dates.

2. Select resources for operations. 3. Renew to the renewal date. 1. Configure a renewal date. Procedure For example, the renewal date is the 1 day of each month. Resource A Additional renewal Renewal for 1 month for 14 days Expiration: April 17 Rules Additional renewal for 24 Resource B Renewal for 1 month days Expiration:

Figure 2-18 Setting the same renewal day for resources with different expiry dates

For more details, see **Setting a Renewal Date**.

2.6.3 Auto-Renewal

Auto-renewal can prevent instances from being automatically deleted if you forget to manually renew them. The auto-renewal rules are as follows:

- The first auto-renewal date is based on when an instance expires and the billing cycle.
- The auto-renewal period of an instance depends on the subscription term.
 - Monthly subscriptions renew each month.
 - Yearly subscriptions renew each year.
- You can enable auto-renewal anytime before an instance expires. By default, the system will make the first attempt to renew your account at 03:00 seven days before the expiry date. If this attempt fails, it will make another attempt at 03:00 every day until the subscription is renewed or expired.
- After auto-renewal is enabled, you can still renew the instance manually if you want to. After a manual renewal is complete, auto-renewal is still valid, and the renewal fee will be deducted from your account seven days before the new expiry date.
- By default, the renewal fee is deducted from your account seven days before the new expiry date. You can change this auto-renewal payment date as required.

For more information about auto-renewal rules, see Auto-Renewal Rules.

Prerequisites

Your yearly/monthly instance is not expired.

Enabling Auto-Renewal When Buying an Instance

You can enable auto-renewal on the page for buying an instance. For details, see **Buying an Instance**.


Figure 2-19 Enabling auto-renewal



Enabling Auto-Renewal on the Renewals Page


- Step 1 Log in to the Huawei Cloud console.
- **Step 2** Hover over **Billing & Costs** in the upper part of the console and choose **Renewal** from the drop-down list.
- **Step 3** Select the search criteria.
 - On the **Auto Renewals** page, you can view the resources that auto-renewal has been enabled for.
 - You can enable auto-renewal for resources on the **Manual Renewals**, **Payper-Use After Expiration**, and **Renewals Canceled** pages.

Figure 2-20 Renewal management



- **Step 4** Enable auto-renewal for yearly/monthly resources.
 - Enabling auto-renewal for a single instance: Locate the instance that you
 want to enable auto-renewal for and choose More > Enable Auto-Renew in
 the Operation column.

Figure 2-21 Enabling auto-renewal for an instance



 Enabling auto-renewal for multiple instances at a time: Select the instances that you want to enable auto-renewal for and click Enable Auto-Renew above the list.

ewals (9) Auto Renewals (0) Pay-per-Use After Expiration (0) Renewals Canceled (0) Batch Renew Enable Auto-Renew Change to Pay-per-Use After Expiration Cancel Renewal Batch Export Export Renewal Prices C Instance NameID Product Type/Specifications Region Provisioned/Expires Status Validity Period Operation · -Apr 26, 2023 10:38:08 GMT+08:00 May 26, 2023 23:59:59 GMT+08:00 3 days until deletion

Delete after retention period

Cancel Renewal | More ▼ Frozen · 🗆 3 days until deletion

Delete after retention period Renew | More ▼ ® Frozen · 🗆 4 days until deletion
Delete after retention period Renew | More ▼ May 16, 2023 15:29:36 GMT+08:00 Jun 16, 2023 23:59:59 GMT+08:00 6 days until deletion

Delete after retention period

Renew | More ▼ v 🗆 May 18, 2023 18:19:32 GMT+08:00 Jun 18, 2023 23:59:59 GMT+08:00 May 18, 2023 17:06:19 GMT+08:00 Jun 18, 2023 23:59:59 GMT+08:00 **~** 🗆 6 days until deletion

Delete after retention period

Renew | More ▼ **~** 🗆 Jun 01, 2023 22:51:24 GMT+08:00 Jul 01, 2023 23:59:59 GMT+08:00 19 days until deletion

Delete after retention period

Cancel Renewal | More ▼ Jun 02, 2023 11:34:42 GMT+08:00 Jul 02, 2023 23:59:59 GMT+08:00 v 🛮 6 hours 48 minutes until del...


Delete after retention period Cancel Renewal | More ▼ Jul 18, 2023 11:34:01 GMT+08:00 Aug 18, 2023 23:59:59 GMT+08:00 25 days until expiration
Delete after retention period Renew | More ▼ Jul 24, 2023 15:54:35 GMT+08:00 Jul 24, 2024 23:59:59 GMT+08:00 366 days until expiration
Delete after retention period

Renew | More ▼

Figure 2-22 Enabling auto-renewal for multiple instances

Step 5 Select a renewal period, specify the auto-renewal times, and click **Pay**.

Figure 2-23 Enabling auto-renewal



----End

2.7 Bills

You can view the resource usage and bills for different billing cycles on the **Bills** page in the Billing Center.

Bill Generation

A bill is generated after a yearly/monthly instance is paid.

The usage of pay-per-use instances is reported to the billing system at a fixed interval. Pay-per-use instances can be settled by the hour, day, or month based on the usage type. For details, see **Bill Run for Pay-per-Use Resources**. Pricing is listed on a per-hour basis.

The fee deduction time of a pay-per-use instance may be later than the settlement period. For example, if an instance is deleted at 08:30, the fees generated from 08:00 to 09:00 are usually deducted at about 10:00. In Billing Center, choose **Billing** > **Transactions and Detailed Bills** > **Transaction Bills**.

Expenditure Time in the bill indicates the time when the pay-per-use resource is used.

Viewing Bills of a Specific Resource

[Method 1: Use the instance ID to search for a bill.]

- Step 1 Log in to the Huawei Cloud console. Choose Databases > GeminiDB.
- **Step 2** On the **Instances** page, locate the instance whose bill you want to view and click its name.
- **Step 3** Click the icon shown in the figure below to copy the instance ID.

Figure 2-24 Copying the instance ID




Step 4 On the top menu bar, choose **Billing & Costs** > **Bills**.

The **Bills** page is displayed.

Step 5 Choose Transactions and Detailed Bills > Bill Details. On the displayed page, select Resource ID as the filter, enter the obtained instance ID, and click Q.

Figure 2-25 Searching for a bill



By default, the bill details are displayed by usage and billing cycle. You can choose other display option. For details, see **Bills**.

----End

[Method 2: Use the resource name to search for a bill.]

- Step 1 Log in to the Huawei Cloud console. Choose Databases > GeminiDB.
- **Step 2** On the **Instances** page, locate the instance whose bill you want to view and click its name.
- **Step 3** On the **Basic Information** > **Instance Information** page, obtain the instance name.

Figure 2-26 Obtaining an instance name




Step 4 On the top menu bar, choose **Billing & Costs** > **Bills**.

The Bills page is displayed.

Step 5 Choose Transactions and Detailed Bills > Bill Details. On the displayed page, select Resource Name as the filter, enter the obtained instance ID, and click Q.

Figure 2-27 Searching for a bill



By default, the bill details are displayed by usage and billing cycle. You can choose other display option. For details, see **Bills**.

----End

Scenario Example: Checking the Consistency of the Actual Usage and Billed Usage

Assume that you bought a pay-per-use GeminiDB HBase instance at 10:09:06 on April 8, 2023 and deleted it at 12:09:06 on April 8, 2023.

GeminiDB HBase instance transaction records
 Pricing is listed on a per-hour basis, and bills are calculated down to the second. You can check the transaction records against the actual usage.
 Resources are billed separately in the transaction records.

Table 2-9 GeminiDB HBase instance transaction records

Produc	GeminiDB HBase API
t Type	

Resour ce Type	Storage
Billing Mode	Pay-per-use
Expend iture Time	For the period of time from 10:09:06 to 12:09:06 on April 08, 2023, 6 transaction records would be generated for the resource usage in the following periods: • 10:09:06 – 11:00:00 • 11:00:00 – 12:00:00 • 12:00:00 – 12:09:06
List Price	List price on the official website = Usage x Unit price x Capacity The instance was used for 3,054 seconds in the first period. You can check its unit price in GeminiDB Price Calculator . The list price in the first period is USD0.02375333 = $(3054/3600) \times 0.0007 \times 40$. You can also calculate the list price in the other periods.
Discou nted Amoun t	You can enjoy discounts on cloud services, such as business, partner-authorized, and promotional discounts. The discounts are calculated based on the list price.
Truncat ed Amoun t	The billing of Huawei Cloud is calculated to the 8th decimal place (USD). However, the amount due is truncated to the 2nd decimal place. The third and later decimal places are referred to as the truncated amounts. For example, in the first billing cycle, the truncated amount is USD0.00375333.
Amoun t Due	Amount due = List price – Discount amount – Truncated amount Take the first period as an example. If the discount amount is 0, the amount due is \$0.02 USD (0.02375333 – 0 – 0.00375333).

GeminiDB HBase instance bill details

Bill details can be displayed in multiple ways. By default, the bill details are displayed by usage and billing cycle. You can check the information listed in **Table 2-10** against the actual usage.

Table 2-10 GeminiDB HBase instance bill details

Produc t Type	GeminiDB HBase API
Resour ce Type	Storage
Billing Mode	Pay-per-use

Resour	Name and ID	
ce	Example: nosql-b388 and	
Name/I D	21e8811a64bf4de88bc2e2556da17983in12	
Specific ations	Storage	
Usage Type	Duration	
Unit Price	When pay-per-use billing is used, the unit price is only provided if the amount is equal to the usage multiplied by the unit price. No unit price is provided in other pricing modes, for example, tiered pricing.	
	You can search for the unit price of a pay-per-use GeminiDB HBase instance in GeminiDB Price Calculator .	
Unit	USD/GB/Hour in GeminiDB Price Calculator	
Usage	Depends on the unit of the unit price, which is USD/GB/hour. Storage usage is priced by the hour. In this example, the total duration is 2 hours.	
Usage Unit	Hour	
List Price	List price on the official website = Usage x Unit price x Capacity The instance has been used for 2 hours. Its unit price is displayed in GeminiDB Price Calculator . The list price is USD0.056 (2 \times 0.0007 \times 40).	
Discou nted Amoun t	You can enjoy discounts on cloud services, such as business, partner-authorized, and promotional discounts. The discounts are calculated based on the list price.	
Amoun t Due	Amount that should be paid for used cloud services after discounts are applied.	

2.8 Arrears

If the available account balance is less than the amount to be settled, the account will be in arrears. To continue using your instances, you need to top up your account in a timely manner.

Arrears Reason

If you do not have yearly/monthly instances, your account falls into arrears any time your configured payment method is unable to pay for the used resources on the pay-per-use basis.

Impact of Arrears

Yearly/Monthly

You have prepaid for yearly/monthly GeminiDB HBase instances. If your account is in arrears, the existing yearly/monthly instances can still be used. However, you cannot perform operations that involve fees, such as buying another GeminiDB HBase instance, upgrading specifications, and renewing the existing instances.


Pay-per-use

When your account falls into arrears due to automatic fee deduction for the pay-per-use GeminiDB HBase instance, your account will be marked as in arrears, and the pay-per-use instances enter the grace period. After you top up your account, Huawei Cloud will bill you for expenditures generated during the grace period. You can view the charges on the **Billing Center** > **Overview** page.

If your account is still in arrears after the grace period ends, the resources enter the retention period and their status turns to **Frozen**. You cannot perform any operations on these resources.

After the retention period ends, the compute resources (vCPUs and memory) and EIPs will be released and cannot be restored.

Figure 2-28 Lifecycle of a pay-per-use instance



◯ NOTE

The grace period and retention period are both 15 days.

Avoiding and Handling Arrears

Top up your account in time. For details, see **Topping Up an Account**.

If you no longer need GeminiDB HBase instances, delete them to avoid additional fees.

To help make sure your account never falls into arrears, you can configure the **Balance Alert** on the **Overview** page of the Billing Center. Then, anytime an expenditure quota drops to below the threshold you specify, Huawei Cloud automatically notifies you by SMS or email.

2.9 Billing Termination

Yearly/Monthly Resources

You are billed upfront for yearly/monthly resources, such as yearly/monthly GeminiDB HBase instances. Your access to these instances will automatically cease upon expiration.

- You can unsubscribe from a yearly/monthly resource before it expires.
 Depending on whether coupons or discounts were used, Huawei Cloud may issue you a refund. For details about unsubscription rules, see
 Unsubscriptions.
- If you have enabled auto-renewal but no longer wish to automatically renew the subscription, disable it before the auto-renewal date (7 days before the expiration date by default) to avoid unexpected expenditures.

Pay-per-Use Resources

If you no longer need the pay-per-use instances and want to stop being billed for them, simply delete them.

Searching for Resources from Bills and Stopping Billing

To ensure that all related resources are deleted, you can search the billing records by resource ID, and then delete the resources you identify in this way. The procedure is as follows:

[Method 1: Use the resource ID in the bill to search for the resource.]

Step 1 Log in to the Huawei Cloud console. On the top menu bar, choose Billing & Costs > Bills.

The **Bills** page is displayed.

Step 2 Choose **Transactions and Detailed Bills** > **Bill Details**, and click the icon shown in the following figure to copy the resource ID.

Figure 2-29 Copying the resource ID



- Step 3 Log in to the Huawei Cloud console. Choose Databases > GeminiDB.
- **Step 4** Select the region where the resource is located. Select **Instance ID**, enter the resource ID copied from **Step 2**, and click Q to search for the resource.

Figure 2-30 Searching for resources



Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

◯ NOTE

You are billed one hour after the resource usage is calculated, so a bill may still be generated after the pay-per-use resource is deleted. For example, if you delete an instance (which is billed on an hourly basis) at 08:30, the expenditures for that hour from 08:00 to 09:00 are usually not billed until about 10:00.

----End


[Method 2: Use the resource name in the bill to search for the resource.]

Step 1 Log in to the Huawei Cloud console. On the top menu bar, choose Billing & Costs > Bills.

The **Bills** page is displayed.

Step 2 Choose **Transactions and Detailed Bills** > **Bill Details**, and click the icon shown in the following figure to copy the resource name.

Figure 2-31 Copying the resource name



- Step 3 Log in to the Huawei Cloud console. Choose Databases > GeminiDB.
- **Step 4** Enter the instance name copied from **Step 2** in the search box and click ${}^{\mathbf{Q}}$.

Figure 2-32 Searching for resources



Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

You are billed one hour after the resource usage is calculated, so a bill may still be generated after the pay-per-use resource is deleted. For example, if you delete an instance (which is billed on an hourly basis) at 08:30, the expenditures for that hour from 08:00 to 09:00 are usually not billed until about 10:00.


----End

2.10 Cost Management

2.10.1 Costs

The costs of a GeminiDB HBase instance include:

- Resource costs: costs of compute and storage resources. For details, see Billing Modes.
- O&M costs: labor costs generated when a GeminiDB HBase instance is used





2.10.2 Cost Allocation

A good cost accountability system is a prerequisite for cost management. It ensures that departments, business teams, and owners are accountable for their respective cloud costs. An enterprise can allocate cloud costs to different teams or projects so as to have a clear picture of their respective costs.

Huawei Cloud **Cost Center** provides various tools for you to group costs in different ways. You can experiment with these tools and find a way that works best for you.

By linked account

The enterprise master account can manage costs by grouping the costs of its member accounts by linked account. For details, see **Viewing Costs by Linked Account**.

• By enterprise project

Before allocating costs, enable Enterprise Project Management Service (EPS) and plan your enterprise projects based on your organizational structure or service needs. When purchasing cloud resources, select an enterprise project so that the costs of resources will be allocated to the selected enterprise project. For details, see **Viewing Costs by Enterprise Project**.


Figure 2-33 Selecting an enterprise project



By cost tag

You use tags to sort your Huawei Cloud resources in a variety of different ways, for example, by purpose, owner, or environment. The following is the process of managing costs by predefined tags (recommended).

Figure 2-34 Adding a tag



For details, see Viewing Costs by Cost Tag.

By cost category

You can use cost categories provided by **Cost Center** to split shared costs. Shared costs are the costs of resources (compute, network, storage, or resource packages) shared across multiple departments or the costs that cannot be directly split by cost tag or enterprise project. These costs are not directly attributable to a singular owner, and they cannot be categorized into a singular cost type. In this case, you can define cost splitting rules to fairly allocate these costs among teams or business units. For details, see **Viewing Cost By Cost Category**.

2.10.3 Cost Analysis

To precisely control and optimize your costs, you need a clear understanding of what parts of your enterprise incurred different costs. **Cost Center** visualizes your original costs and amortized costs using various dimensions and display filters for cost analysis so that you can analyze the trends and drivers of your service usage and costs from a variety of perspectives or within different defined scopes.

You can also use cost anomaly detection provided by **Cost Center** to detect unexpected expenses in a timely manner. In this way, costs can be monitored, analyzed, and traced.

For details, see Performing Cost Analysis to Explore Costs and Usage and Enabling Cost Anomaly Detection to Identify Anomalies.

2.10.4 Cost Optimization

You can identify resources with high costs based on the analysis results in the cost center, determine the causes of high costs, and take optimization measures accordingly.

Resource optimization

- On Cloud Eye, view metrics of your GeminiDB HBase instance, such as CPUs, memory, and disk usage. If the current configurations are too high, you can use lower specifications instead.
- Monitor idle GeminiDB HBase instances and delete them in a timely manner.

Billing mode selection

Different types of services have different requirements on resource usage periods, so the most economical billing mode for one resource may not be the best option for another resource.

- For mature services that tend to be stable for the long term, select yearly/ monthly billing.
- For short-term, unpredictable services that experience traffic bursts and cannot afford to be interrupted, select pay-per-use billing.
- Monitor the lifecycle of instances and renew yearly/monthly resources that are about to expire in a timely manner.

2.11 Billing FAQs

2.11.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing?

Yearly/Monthly is a prepaid mode. 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 is a postpaid mode. You are billed for your immediate needs based on how long you have actually used a product, so you can start or stop an instance at any time. 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.

2.11.2 Can I Switch Between Yearly/Monthly and Pay-per-use Billing Modes?

Yes.

- For details about how to change a yearly/monthly GeminiDB HBase instance to pay-per-use, see Changing a Yearly/Monthly Instance to Pay-per-UseChanging a Yearly/Monthly Instance to Pay-per-Use.
- For details about how to change a pay-per-use GeminiDB HBase instance to yearly/monthly, see **Changing a Pay-per-Use Instance to Yearly/Monthly**.

2.11.3 How Do I Renew a Single or Multiple Yearly/Monthly Instances?

This section describes how to renew your yearly/monthly GeminiDB HBase instance.

Usage Notes

Pay-per-use instances do not support this function.

Renewing a Single Yearly/Monthly Instance


- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Renew** in the **Operation** column.

Figure 2-35 Renewal



Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 2-36 Renewal



Step 4 On the displayed page, renew the instance.

----End

Renewing Instances In Batches

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, select the instances that you want to renew and click **Renew** above the instance list.

Figure 2-37 Renewing instances in batches



Step 4 In the displayed dialog box, click **Yes**.

----End

2.11.4 How Do I Unsubscribe from a Yearly/Monthly Instance?

If you do not need a yearly/monthly instance any longer, unsubscribe from it.

Usage Notes

- The unsubscription action cannot be undone. To retain data, create a manual backup before unsubscription. For details, see Creating a Manual Backup.
- After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved. Ensure that the manual backup is complete before submitting the unsubscription request.

Unsubscribing from a Single Yearly/Monthly Instance

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Unsubscribe** in the **Operation** column.

Figure 2-38 Unsubscribe



Step 4 In the displayed dialog box, click **Yes**.

Step 5 On the displayed page, confirm the order to be unsubscribed and select a reason. Then, click **Confirm**.

For details, see Unsubscription Rules.

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

NOTICE

- 1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
- 2. Ensure that the manual backup is complete before submitting the unsubscription request.
- **Step 7** View the unsubscription result. After you unsubscribe from the instance order, the instance is no longer displayed in the instance list on the **Instances** page.

----End

Unsubscribing from Yearly/Monthly Instances In Batches

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** Choose **Instances** in the navigation pane on the left, select the instances you want to unsubscribe from and click **Unsubscribe** above the instance list.

Figure 2-39 Unsubscribe



- **Step 4** In the displayed dialog box, click **Yes**.
- **Step 5** On the displayed page, confirm the order to be unsubscribed and select a reason. Then, click **Confirm**.

For details, see **Unsubscription Rules**.

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

NOTICE

- 1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
- 2. Ensure that the manual backup is complete before submitting the unsubscription request.

Step 7 View the unsubscription result. After you unsubscribe from the instance order, the instance is no longer displayed in the instance list on the **Instances** page.

----End

3 Getting Started with GeminiDB HBase API

3.1 Getting to Know GeminiDB HBase API

This section instructs you to create and connect to a GeminiDB HBase instance.

Connection Methods

You are advised to use Java code to connect to a GeminiDB HBase instance.

Table 3-1 Connection methods

Method	Scenario	Description
Connecti ng to a GeminiD B HBase Instance Using Java	An example of connecting to a GeminiDB HBase instance using Java	The default port is 2181.

More Connection Operations

See Connecting to a GeminiDB HBase Instance.

3.2 Buying and Connecting to a GeminiDB HBase Instance

This section describes how to buy and connect to a GeminiDB HBase instance on the GeminiDB console.

Viewing the Instance IP Address

- Step 2: Buying an Instance
- Connecting to an Instance Using Java

For details about other connection methods, see **Connecting to a GeminiDB HBase Instance**

Prerequisites

- A GeminiDB HBase instance has been created and is running properly. For details about how to create a GeminiDB HBase instance, see Getting to Know GeminiDB HBase API.
- For details about how to create an ECS, see Purchasing an ECS in Getting Started with Elastic Cloud Server.
- JDK has been installed on the ECS.
- Download the **HBase client**. Click a directory of the latest version 2.6.X and download **hbase-2.6.X-client-bin.tar.gz**. For example, if the latest version is 2.6.1, click that directory and download **hbase-2.6.1-client-bin.tar.gz**. HBase 1.X is not recommended due to compatibility issues.

Viewing the Instance IP Address

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** On the **Instances** page, click the target instance.

Method 1:

In the **Node Information** area on the **Basic Information** page, view the private IP address of each GeminiDB HBase instance node.

□ NOTE

Currently, GeminiDB HBase instances cannot be bound to EIPs.

Figure 3-1 Viewing the IP address



In the **Network Information** area, view the GeminiDB HBase instance port. The default port for connecting to the client is 2181.


Figure 3-2 Viewing the port



Method 2

In the navigation pane, choose **Connections** to view the private IP address and EIP of the instance.

Figure 3-3 Viewing the IP address and port



----End

Step 2: Buying an Instance

For details, see **Buying a GeminiDB HBase Instance**.

- 1. Log in to the Huawei Cloud console.
- 2. In the service list, choose **Databases** > **GeminiDB**.
- 3. On the **Instances** page, click **Buy DB Instance**.
- 4. On the displayed page, select a billing mode, configure instance parameters, and click **Next**.

Figure 3-4 Billing mode and basic information (classic storage)






Figure 3-5 Billing mode and basic information (cloud native storage)

Table 3-2 Billing mode

Parameter	Description
Billing Mode	Select Yearly/Monthly or Pay-per-use. • Yearly/Monthly
	 Specify Required Duration. The system deducts fees from your account based on the service price.
	 If you do not need such an instance any longer after it expires, change the billing mode to pay-per- use. For details, see Changing a Yearly/Monthly Instance to Pay-per-Use.
	NOTE
	 Yearly/Monthly instances cannot be deleted directly. If such an instance is no longer required, unsubscribe from it. For details, see How Do I Unsubscribe from a Yearly/Monthly Instance?.
	 Yearly/Monthly instances with cloud native storage are now in OBT. To use such an instance, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.
	Pay-per-use
	 You do not need to set Validity Period because you are billed based on how long you have actually used your instance.
	 To use an instance for a long time, change its billing mode to yearly/monthly to reduce costs. For details, see Changing a Pay-per-Use Instance to Yearly/Monthly.

Table 3-3 Basic information

Parameter	Description
Region	Region where a tenant is located NOTE To reduce network latency, select a region nearest from which you will access the instance. Instances deployed in different regions cannot communicate with each other over a private network. After you buy an instance, you cannot change its region.
DB Instance Name	 The instance name: Can be the same as an existing instance name. Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).
Compatible API	HBase GeminiDB is compatible with mainstream NoSQL databases, including Redis, DynamoDB, Cassandra, HBase, MongoDB, and InfluxDB. You can select GeminiDB APIs by following How Do I Select an API?
Storage Type	 Classic: classic architecture with decoupled storage and compute Cloud native: new, more flexible, new-gen version with support for more AZs NOTE The way you use instances with classic or cloud native storage is similar. Cloud native storage supports more AZs. If both classic and cloud native are supported, you can select any of them. Cloud native storage is now in OBT. To use it, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.
DB Instance Type	Cluster One cluster consists of at least three nodes. A cluster is easy to scale out to meet increasing data growth needs. A cluster is recommended when dealing with stringent availability demands, substantial data volumes, and the need for seamless scalability.
AZ	Availability zone where the instance is created. An AZ is a part of a region with its own independent power supplies and networks. AZs are physically isolated but can communicate with each other over a private network.

Figure 3-6 Specifications and storage




Table 3-4 Specifications and storage

Parameter	Description
Instance Specification s	Higher CPU specifications provide better performance. Select specifications as needed.
Nodes	The number of nodes ranges from 3 to 80.
Storage Space	Instance storage space. The range depends on the instance specifications.
	To scale up storage, you need to add at least 1 GB each time. The value must be an integer.

Parameter	Description	
Disk	You can determine whether to encrypt disks.	
Encryption	Disable: Data is not encrypted.	
	• Enable: Your data will be encrypted on disks and stored in ciphertext after you create an instance. When you download encrypted objects, the ciphertext will be decrypted into plain text and then sent to you. Disk encryption can improve data security and may have slight impacts on database writes and reads.	
	 Key Name: Select an existing key or create one. 	
	 To use a shared key, ensure that you have created an agency. For details, see Creating an Agency (by a Delegating Party). Select another account from the drop-down list to share the key of the current account. VPC owners can share the keys with one or multiple accounts through Resource Access Manager (RAM). For details, see Creating a Resource Share. 	
	- Enter a key ID. The key must be in the current region.	
	NOTE	
	 This function is now in CBT. To use it, choose Service Tickets Create Service Ticket in the upper right corner of the console and contact the customer service. 	
	 After an instance is created, the disk encryption status and key cannot be changed. 	
	 The key cannot be disabled, deleted, or frozen when used, or the database becomes unavailable. 	
	 For details about how to create a key, see "Creating a CMK" in <i>Data Encryption Workshop User Guide</i>. 	

Figure 3-7 Network configuration



Table 3-5 Network configuration

Parameter	Description
VPC	Virtual private network where your instances are located. A VPC isolates networks for different services. You can select an existing VPC or create a VPC.
	If no VPC is available, the system creates one for you.
	For details, see "Creating a VPC" in <i>Virtual Private Cloud User Guide</i> .
	You can also use VPCs and subnets shared by other accounts.
	VPC owners can share the subnets in a VPC with one or multiple accounts through Resource Access Manager (RAM), which ensures cost efficiency of network resources.
	For more information about VPC subnet sharing, see VPC Sharing in <i>Virtual Private Cloud User Guide</i> .
	NOTE
	After the instance is created, its VPC cannot be changed.
	 If you want to connect to an instance using an ECS over a private network, ensure that the instance and the ECS are in the same VPC. If they are not, create a VPC peering connection between them.
Subnet	A subnet where your instance is created. The subnet provides dedicated and isolated networks, improving network security.
	NOTE An IPv6 subnet cannot be associated with your instance. Select an IPv4 subnet.
Security Group	A security group controls access between instances and other services. Ensure that the security group you selected allows your client to access the instance.
	If no security group is available, the system creates one for you.

Figure 3-8 Database configuration




Table 3-6 Database configuration

Parameter	Description	
Administrator	Username of the administrator account. The default value is rwuser .	
Administrator Password	 Password of the administrator account. The password: Can include 8 to 32 characters. Must contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!#%^*= +? Cannot contain @ or / For security reasons, set a strong password. The system will verify the password strength. Keep your password secure. The system cannot retrieve it if it is lost. 	
Confirm Password	Enter the administrator password again.	
Enterprise Project	This parameter is provided for enterprise users. An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is default . Select an enterprise project from the drop-down list. For more information about enterprise projects, see <i>Enterprise Management User Guide</i> .	
Parameter Template	A parameter template contains the configuration values of a database engine. The parameters can be applied to one or more instances of the same type. After creating an instance, you can modify the parameter template. You can modify the instance parameters as required after the instance is created.	

Retain the default values for other parameters.

- On the displayed page, confirm instance details. To modify the configurations, click **Previous**.
- 6. If no modification is required, read and agree to the service agreement and click **Submit**.
- 7. Click **Back to Instance Management** to go to the instance list.
- 8. On the **Instances** page, view and manage the created instance.
 - It takes about 5 to 9 minutes to create an instance. During the process, the instance status is **Creating**.
 - After the instance is created, its status becomes **Available**.

Connecting to an Instance Using Java

- Step 1 Obtain the private IP address and port of the GeminiDB HBase instance.
 For details about how to obtain the private IP address and port, see Viewing the Instance IP Address.
- **Step 2** Log in to the ECS. For details, see **Logging In to an ECS** in *Getting Started with Elastic Cloud Server*.
- **Step 3** Add the following Maven dependencies to the **pom.xml** file of your project. HBase 1.*X* is not recommended due to compatibility issues causing read and write failures. You are advised to use HBase 2.2.3 or later. SSL is supported only in HBase 2.6.0 or later.

```
<dependency>
    <groupId>org.apache.hbase</groupId>
    <artifactId>hbase-client</artifactId>
    <version>2.6.1</version>
</dependency>
```

Step 4 Edit the code for connecting to the GeminiDB HBase instance. Replace **your_hbase_instance_quorum** with the cluster IP address, **your_user_name** with the username (**rwuser** by default) you set when creating the cluster, and **your_password** with the password you set when creating the cluster. Before running the code, ensure that an HBase table has been created in the instance, and replace **your_table_name** with the table name.

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.*;
import org.apache.hadoop.hbase.util.Bytes;
import java.io.IOException;
public class HBaseExample {
  public static void main(String[] args) throws IOException {
     // Creates a configuration object and sets HBase connection parameters.
     Configuration config = HBaseConfiguration.create();
     config.set("hbase.zookeeper.quorum", "your_hbase_instance_quorum"); config.set("hbase.zookeeper.property.clientPort", "2181");
     // Enters a username and password.
     UserGroupInformation ugi = UserGroupInformation.createProxyUser("your_user_name",
UserGroupInformation.createRemoteUser("your_password"));
     // Establishes a connection to the HBase instance.
     Connection connection = ConnectionFactory.createConnection(config, User.create(ugi));
        // Obtains table objects.
        TableName tableName = TableName.valueOf("your_table_name");
        Table table = connection.getTable(tableName);
        // Inserts data.
        Put put = new Put(Bytes.toBytes("row_key"));
        put.addColumn(Bytes.toBytes("cf"), Bytes.toBytes("col"), Bytes.toBytes("value"));
        table.put(put);
        // Obtains a single row of data.
        Get get = new Get(Bytes.toBytes("row_key"));
        Result result = table.get(get);
        byte[] value = result.getValue(Bytes.toBytes("cf"), Bytes.toBytes("col"));
        System.out.println("Success: " + Bytes.toString(value));
     } finally {
```

```
// Closes the connection.
    connection.close();
}
}
```

Step 5 Run the sample code to check whether the result is normal. If it is, "Success: value" is displayed.

----End

3.3 How Can I Delete Rows Based On Prefixes?

You can delete rows with keys matching a given prefix on GeminiDB HBase instances. This function takes effect quickly. You do not need to scan data before deleting it anymore. Unlike Apache HBase that supports only single-row deletion, GeminiDB HBase allows you to delete rows based on prefixes, which is more convenient and faster. This capability provides significant advantages for tasks like range data clearance and historical data purging.



Incorrect use of this function may have a significant impact on data. Before using this function, ensure that you have fully understood the following usage notes.

Usage Notes (Mandatory)

A message is displayed, indicating the data has been deleted as expected, but it is not deleted immediately and is only marked for deletion. The marked data needs to be gradually deleted in stored procedures, and generated range tombstones will also be deleted. Therefore, to ensure that this does not affect database performance, you need to comply with the following conventions:

- Do not repeatedly delete and write a single range of data within a short period of time.
- Do not delete massive volumes of data within a short period of time.
- Do not scan data that has been deleted.
- Verify the range in advance to prevent accidental deletion of a large amount of data.

Typical Violations

- In a specific data range, rows with 1 billion keys matching a given prefix are deleted 50,000 times within a day. At the same time, a large amount of data in the same range is written.
- Rows are deleted based on a given prefix 100,000 or more times within a day.
- A large amount of data is deleted by mistake based on short prefixes (for example, **0** or **a**) without verification.

A CAUTION

Severe violations will increase the read latency, result in failed requests, and affect read and write performance. You need to check the service status in a timely manner. Before final deletion, verify the results in a test environment.

If the preceding issues occur for a large amount of data, stop using this function immediately and consult experts. In the upper right corner of the console, choose **Service Tickets > Create Service Ticket** and contact the customer service.

How to Use

You can add additional attributes to mark deletion requests as those based on prefixes. After the requests are marked, only **key** takes effect. Other parameters, such as the specified column and qualifier, will not take effect. Data that matches a prefix is deleted immediately.

Currently, this function can be used only through the Java HBase client. In the following Java code, all keys starting with **row1** will be deleted.

```
Delete delete = new Delete(Bytes.toBytes("row1"));
delete.setAttribute("PREFIXDELETE", "true".getBytes(StandardCharsets.UTF_8));
table.delete(delete);
```

FAQs

- 1. Q: If a request times out or fails, has my data been deleted?
 - A: GeminiDB HBase does not provide transactions and cannot ensure atomicity. In the case of a failed request, the target data may be completely or partially deleted. If the request is successful, all data is deleted. If the failure is caused by network disconnection or other reasons, try again.
- 2. Q: How can I perform a large number of deletions for historical data based on prefixes?
 - A: Specify the range of historical data to be deleted. Verification in a test environment is recommended to prevent unexpected data deletion. No more than 2,000 times per day are recommended for this function. In a short period of time, a few deletions based on prefixes can be applied to massive volumes of data, which can meet your requirements. Continuously check the read latency while data is deleted. If there is any exception, stop deleting data immediately.

4 Working with GeminiDB HBase API

4.1 Using IAM to Grant Access to GeminiDB HBase API

4.1.1 Creating a User and Granting Permissions to Use GeminiDB HBase API

This section describes how to use IAM to control fine-grained permissions for your GeminiDB resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing GeminiDB resources.
- Grant only the permissions required for users to perform a specific task.
- Entrust a Huawei Cloud account or cloud service to perform efficient O&M on your GeminiDB resources.

If you do not need to create an IAM user, skip this section.


The following describes the procedure for granting permissions (see Figure 4-1).

Prerequisites

You have learned about permissions that can be granted to the user group. For details, see **GeminiDB Permissions**. For system permissions of other services, see **Permission Policies**.

Process Flow

Figure 4-1 Process of granting GeminiDB permissions



1. Create a user group and assign permissions.

Create a user group on the IAM console and assign the **GeminiDB FullAccess** permission to the group.

□ NOTE

To use some interconnected services, you also need to configure permissions of such services.

For example, when connecting to your instance through DAS, you need to obtain the **GeminiDB FullAccess** and **DAS FullAccess** permissions.

2. Create an IAM user and add it to a user group.

Create a user on the IAM console and add the user to the group created in 1.

3. Log in and verify permissions.

Log in to the management console using the created user, and verify the user's permissions:

Choose **Service List** > **GeminiDB** and click **Buy DB Instance**. If you can buy an instance, the required permission policy has taken effect.

4.1.2 Custom Policies for GeminiDB HBase API

Custom policies can be created to supplement the system-defined policies of GeminiDB. For the actions supported for custom policies, see **Permissions Policies and Supported Actions**.

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

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.
- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**. The following describes examples of common GeminiDB custom policies.

Example Custom Policy

• Example 1: Allowing users to create GeminiDB instances

• Example 2: Refusing users to delete GeminiDB instances

A policy with only "Deny" permissions must be used in conjunction with other policies to take effect. If the policies assigned to a user contain both Allow and Deny actions, the Deny actions take precedence over the Allow actions.

The following method can be used if you need to assign permissions of the **GeminiDB FullAccess** policy to a user but you want to prevent the user from deleting GeminiDB instances. Create a custom policy for denying instance deletion, and attach both policies to the group to which the user belongs. Then, the user can perform all operations on GeminiDB instances except deleting GeminiDB instances. The following is an example of the deny policy:

• Example 3: Defining permissions for multiple services in a policy

A custom policy can contain the actions of multiple services that are of the global or project-level type. The following is an example policy containing actions of multiple services:

4.2 Buying a GeminiDB HBase Instance

This section describes how to buy a GeminiDB HBase instance.

GeminiDB Cassandra, DynamoDB-Compatible, and HBase instances of each tenant share a quota. Each tenant can create a maximum of 50 instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact customer service personnel.

Prerequisites

You have created a Huawei Cloud account.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** On the **Instances** page, click .
- **Step 3** On the displayed page, select a billing mode, set **Compatible API** to **HBase**, set other parameters, and click **Next**. If HBase is unavailable, contact the customer service.

Figure 4-2 Billing mode and basic information (classic storage)



Figure 4-3 Billing mode and basic information (cloud native storage)




Table 4-1 Billing parameters

Parameter	Description			
Billing Mode	Select Yearly/Monthly or Pay-per-use. • Yearly/Monthly			
	Specify Required Duration . The system deducts fees from your account based on the service price.			
	 If you do not need such an instance any longer after it expires, change the billing mode to pay-per-use. For details, see Changing a Yearly/Monthly Instance to Pay-per-Use. 			
	NOTE			
	Yearly/Monthly instances cannot be deleted directly. If such an instance is no longer required, unsubscribe from it. For details, see How Do I Unsubscribe from a Yearly/ Monthly Instance?.			
	Yearly/Monthly instances with cloud native storage are now in OBT. To use such an instance, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.			
	Pay-per-use			
	 If you select this billing mode, you are billed based on how much time the instance is in use. 			
	 To use an instance for a long time, change its billing mode to yearly/monthly to reduce costs. For details, see Changing a Pay-per-Use Instance to Yearly/ Monthly. 			

Table 4-2 Basic information

Parameter	Description
Region	Region where a tenant is located
	NOTE To reduce network latency, select a region nearest from which you will access the instance. Instances deployed in different regions cannot communicate with each other over a private network. After you buy an instance, you cannot change its region.
DB Instance	The instance name:
Name	Can be the same as an existing instance name.
	• Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).

Parameter	Description	
Compatible	HBase	
API	GeminiDB is compatible with mainstream NoSQL databases, including Redis, DynamoDB, Cassandra, HBase, MongoDB, and InfluxDB. You can select GeminiDB APIs by following How Do I Select an API?	
Storage Type	Classic: classic architecture with decoupled storage and compute	
	Cloud native: more flexible, new-gen version with support for more AZs	
	NOTE	
	The way you use instances with classic or cloud native storage is similar. Cloud native storage supports more AZs. If both classic and cloud native are supported, you can select any of them.	
	 Cloud native storage is now in OBT. To use it, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service. 	
DB Instance	Cluster	
Туре	One cluster consists of at least three nodes. A cluster is easy to scale out to meet increasing data growth needs. A cluster is recommended when dealing with stringent availability demands, substantial data volumes, and the need for seamless scalability.	
AZ	Availability zone where the instance is created. An AZ is a part of a region with its own independent power supplies and networks. AZs are physically isolated but can communicate with each other over a private network.	

Figure 4-4 Specifications and storage




Table 4-3 Specifications and storage

Parameter	Description		
Instance Specifications	Higher CPU specifications provide better performance. Select specifications as needed.		
Nodes	The number of nodes ranges from 3 to 80.		
Storage Space	Instance storage space. The range depends on the instance specifications.		
	To scale up storage, you need to add at least 1 GB each time. The value must be an integer.		
Disk Encryption	You can select to enable disk encryption based on service requirements.		
	Disable: Data is not encrypted.		
	• Enable : If you select this option, your data will be encrypted on disks and stored in ciphertext after you create an instance. When you download encrypted objects, the ciphertext will be decrypted into plain text and then sent to you. Disk encryption can improve data security and may have slight impacts on database writes and reads.		
	- Key Name : Select an existing key or create one.		
	 To use a shared key, ensure that you have created an agency. For details, see Creating an Agency (by a Delegating Party). Select another account from the drop-down list to share the key of the current account. VPC owners can share the keys with one or multiple accounts through Resource Access Manager (RAM). For details, see Creating a Resource Share. 		
	 Enter a key ID. The key must be in the current region. 		
	NOTE		
	 This function is now in CBT. To use it, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service. 		
	 After an instance is created, the disk encryption status and the key cannot be changed. 		
	 The key cannot be disabled, deleted, or frozen when being used. Otherwise, the database becomes unavailable. 		
	 For details about how to create a key, see "Creating a CMK" in Data Encryption Workshop User Guide. 		

Figure 4-5 Network configuration

WC (8) 61 MM 64 V C (9 MM 745

MA (9 CMC) Finds a shade of NTC week in shade and to drop of Dates and an elected for NTC 119 densities represent the proof of Dates and an elected for NTC 119 densities represent the proof of Dates and an elected for NTC 119 densities represent the proof of Dates and an elected for NTC 119 densities represent the proof of Dates and an elected for NTC elected for N

Table 4-4 Network configuration

Parameter	Description			
VPC	Virtual private network where your instances are located. A VPC isolates networks for different services. You can select an existing VPC or create a VPC.			
	If there are no VPCs available, the system automatically allocates a VPC to you.			
	For details, see "Creating a VPC" in the <i>Virtual Private Cloud User Guide</i> .			
	With VPC sharing, you can also use a VPC and subnet shared by another account.			
	VPC owners can share the subnets in a VPC with one or multiple accounts through Resource Access Manager (RAM). This allows for more efficient use of network resources and reduces O&M costs.			
	For more information about VPC subnet sharing, see VPC Sharing in Virtual Private Cloud User Guide.			
	NOTE			
	After the instance is created, its VPC cannot be changed.			
	 To connect an instance to an ECS over a private network, ensure they are in the same VPC. If they are not, create a VPC peering connection between them. 			
Subnet	A subnet where your instance is created. The subnet provides dedicated and isolated networks, improving network security. NOTE An IPv6 subnet cannot be associated with your instance. Select an IPv4 subnet.			
Security Group	A security group controls access between instances and other services. Ensure that the security group you selected allows your client to access the instance.			
	If no security group is available, the system creates one for you.			




Figure 4-6 Database configuration

Table 4-5 Database configuration

Parameter	Description	
Administrator	Username of the administrator account. The default value is rwuser.	
Administrator	Password of the administrator account. The password:	
Password	Can include 8 to 32 characters.	
	 Must contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!#%^*=+? 	
	Cannot contain @ or /	
	For security reasons, set a strong password. The system will verify the password strength.	
	Keep your password secure. The system cannot retrieve it if it is lost.	
Confirm Password	This password must be consistent with administrator password.	
Enterprise	This parameter is provided for enterprise users.	
Project	An enterprise project groups cloud resources, so you can manage resources and members by project. The default project is default .	
	Select an enterprise project from the drop-down list. For more information about enterprise projects, see <i>Enterprise Management User Guide</i> .	

Parameter	Description
Parameter Template	A parameter template contains the configuration values of a database engine. The parameters can be applied to one or more instances of the same type. After creating an instance, you can modify the parameter template.
	You can modify the instance parameters as required after the instance is created.
SSL	A security protocol. Secure Sockets Layer (SSL) certificates set up encrypted connections between clients and servers, preventing data from being tampered with or stolen during transmission.
	You are advised to enable SSL to improve data security.
	NOTE If SSL is not enabled when you create an instance, you can enable it after the instance is created. For details, see How Can I Connect to a GeminiDB HBase Instance over TLS (SSL)?

Table 4-6 Tags

Parameter	Description		
Tags	(Optional) Identifier of a GeminiDB HBase instance. Adding tags helps you better identify and manage your instance.		
	A maximum of 20 tags can be added for each instance.		
	If your organization has configured a tag policy, add a tag to the instance based on the tag policy. If the tag does not comply with the tag policy, the instance may fail to be creat Contact the organization administrator to learn details about the tag policy.		
	A tag consists of a tag key and a tag value.		
	 A tag key is mandatory if the instance is going to be tagged. Each tag key is unique for each instance. It can contain a maximum of 128 characters, cannot start with _sys_, and cannot start or end with a space. Only letters, digits, spaces, and the following special characters are allowed:@::/+= 		
	A tag value is optional if the instance is going to be tagged. The value can be empty.		
	The value can contain a maximum of 255 characters. Only letters, digits, spaces, and the following special characters are allowed::+=@/		

Step 4 On the displayed page, confirm instance details.

- To modify the configurations, click **Previous**.
- If no modification is required, read and agree to the service agreement and click **Submit**.

Step 5 On the **Instances** page, view and manage the created instance.

- The instance status is displayed as Creating.
- After the instance is created, its status becomes **Available**.
 - You can click in the upper right corner to refresh the instance status.
- An automated backup policy is enabled by default during instance creation.
 After the instance is created, a full backup is created.

----End


4.3 Instance Connection and Management

4.3.1 Connecting to a GeminiDB HBase Instance

You can connect to a GeminiDB HBase instance over a private network or using Java.

Figure 4-7 shows how to connect to a GeminiDB HBase instance.

Figure 4-7 Connection Methods



- 1 A GeminiDB HBase instance is connected over a private network (An ECS and a GeminiDB HBase instance are in the same security group).
- A GeminiDB HBase instance is connected over a private network (An ECS and a GeminiDB HBase instance are in different security groups).

Table 4-7 Connection methods

Met hod	Scenario	De fau lt	Description
		Por t	
Con nect ing to a Ge mini DB HBa se Inst anc e over a Priv ate Net wor k	Your applications are deployed on an ECS that is in the same region and VPC as your instances.	21 81	 High security and performance If the ECS and GeminiDB HBase instance are in the same security group, they can communicate with each other by default. No security group rule needs to be configured. If they are in different security groups, configure security group rules for them, separately. Configure inbound rules of a security group for the GeminiDB HBase instance by following Setting Security Group Rules for a GeminiDB HBase Instance. The default security group rule allows all outbound data packets, so you do not need to configure a security rule for the ECS. If not all access from the ECS is allowed, you need to configure an outbound rule for the ECS.
Con nect ing to a Ge mini DB HBa se Inst anc e Usin g Java	An example of connecting to a GeminiDB HBase instance using Java	21 81	

4.3.2 Connecting to a GeminiDB HBase Instance over a Private Network

You can install the HBase client on an ECS and access a GeminiDB HBase instance using a private IP address.

Prerequisites

- A GeminiDB HBase instance has been created and is running properly. For details about how to create a GeminiDB HBase instance, see Getting to Know GeminiDB HBase API.
- For details about how to create an ECS, see Purchasing an ECS in Getting Started with Elastic Cloud Server.
- JDK has been installed on the ECS.
- Download the **HBase client**. Click a directory of the latest version 2.6. X and download **hbase-2.6.X-client-bin.tar.gz**. For example, if the latest version is 2.6.1, click that directory and download **hbase-2.6.1-client-bin.tar.gz**. HBase 1. X is not recommended due to compatibility issues.
- **Step 1** Log in to the ECS.

For details, see Logging In to an ECS in Getting Started with Elastic Cloud Server.

- **Step 2** Upload the **HBase client** installation package to the ECS.
- **Step 3** Run the following command to decompress the client package. Replace the package name with the actual one.

tar -xvf hbase-2.6.1-client-bin.tar.gz

Step 4 Add the following configurations to **conf/hbase-site.xml** in the client directory and set **value** to the IP address of your instance. Use commas (,) to separate multiple IP addresses. By default, you do not need to set the port. You can obtain the private IP address by following **Viewing the Instance IP Address**.

Step 5 Go to the bin directory of the decompressed client and run the following command to connect to the instance: Replace YOUR_USERNAME and YOUR_PASSWORD with the user password set during instance creation. The username is fixed to rwuser. If you use a stress test tool to connect to the instance, you also need to set the username and password.

```
export HADOOP_PROXY_USER="YOUR_USERNAME"
export HADOOP_USER_NAME="YOUR_PASSWORD"
./hbase shell
```

Step 6 Check the results. If the following is displayed, the connection is successful. hbase:001:0>

----End

4.3.3 Connecting to a GeminiDB HBase Instance Using Program Code

4.3.3.1 Connecting to a GeminiDB HBase Instance Using Java

This section describes how to connect to a GeminiDB HBase instance using Java.

Prerequisites

- A GeminiDB HBase instance has been created and is running properly. For details about how to create a GeminiDB HBase instance, see Getting to Know GeminiDB HBase API.
- For details about how to create an ECS, see **Purchasing an ECS** in *Getting Started with Elastic Cloud Server*.
- JDK has been installed on the ECS.
- Download the **HBase client**. Click a directory of the latest version 2.6. X and download **hbase-2.6.X-client-bin.tar.gz**. For example, if the latest version is 2.6.1, click that directory and download **hbase-2.6.1-client-bin.tar.gz**. HBase 1. X is not recommended due to compatibility issues.
- **Step 1** Obtain the private IP address and port of the GeminiDB HBase instance.

For details about how to obtain the private IP address and port, see **Viewing the Instance IP Address**.

- **Step 2** Log in to the ECS. For details, see **Logging In to an ECS** in *Getting Started with Elastic Cloud Server*.
- **Step 3** Add the following Maven dependencies to the **pom.xml** file of your project. HBase 1.X is not recommended due to compatibility issues causing read and write failures. You are advised to use HBase 2.2.3 or later. SSL is supported only in HBase 2.6.0 or later.

```
<dependency>
    <groupld>org.apache.hbase</groupld>
    <artifactld>hbase-client</artifactld>
    <version>2.6.1</version>
</dependency>
```

Step 4 Edit the code for connecting to the GeminiDB HBase instance. Replace your_hbase_instance_quorum with the cluster IP address, your_user_name with the username (rwuser by default) you set when creating the cluster, and your_password with the password you set when creating the cluster. Before running the code, ensure that an HBase table has been created in the instance, and replace your_table_name with the table name.

```
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.*;
import org.apache.hadoop.hbase.util.Bytes;

import java.io.IOException;

public class HBaseExample {
    public static void main(String[] args) throws IOException {
        // Creates a configuration object and sets HBase connection parameters.
        Configuration config = HBaseConfiguration.create();
```

```
config.set("hbase.zookeeper.quorum", "your_hbase_instance_quorum");
     config.set("hbase.zookeeper.property.clientPort", "2181");
     // Enters a username and password.
     UserGroupInformation ugi = UserGroupInformation.createProxyUser("your_user_name",
UserGroupInformation.createRemoteUser("your_password"));
     // Establishes a connection to the HBase instance.
     Connection connection = ConnectionFactory.createConnection(config, User.create(ugi));
     try {
        // Obtains table objects.
        TableName tableName = TableName.valueOf("your table name");
       Table table = connection.getTable(tableName);
       // Inserts data.
        Put put = new Put(Bytes.toBytes("row_key"));
       put.addColumn(Bytes.toBytes("cf"), Bytes.toBytes("col"), Bytes.toBytes("value"));
       table.put(put);
       // Obtains a single row of data.
        Get get = new Get(Bytes.toBytes("row_key"));
       Result result = table.get(get);
       byte[] value = result.getValue(Bytes.toBytes("cf"), Bytes.toBytes("col"));
       System.out.println("Success: " + Bytes.toString(value));
     } finally {
       // Closes the connection.
       connection.close();
  }
```

Step 5 Run the sample code to check whether the result is normal. If it is, "Success: value" is displayed.

----End

4.3.4 How Can I Connect to a GeminiDB HBase Instance over TLS (SSL)?

Preparations

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance name. The **Basic Information** page is displayed.
- **Step 4** In the **DB Information** area, toggle on next to **SSL** to download the **ca.cert** file.



- **Step 5** Upload the **ca.cert** file to the ECS.
- **Step 6** Run the following command on the ECS to add the server certificate file to the truststore. The default certificate password is **PASSWORD**. You can change

PASSWORD in the following command. The generated truststore file will be used in subsequent connection method examples.

keytool -importcert -alias hw -file ca.cert -keystore truststore.jks -storepass PASSWORD

----End

Establishing a TLS Connection Using HBase Shell

Add the following configuration items to the **hbase-site.xml** file on the client:

- The value of **hbase.rpc.tls.truststore.location** is the path of the **truststore.jks** file generated in **Step 6**.
- Set **hbase.rpc.tls.truststore.password** to the password set in **Step 6**. The default password is **PASSWORD**.

Start HBase Shell to check whether the connection is successful.

Establishing a TLS Connection Using a Java Application

Modify the connection address and certificate file directory in the following Java code:

```
package com.huawei;
import java.util.ArrayList;
import java.util.List;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.Admin;
import\ or g. apache. hadoop. hbase. client. Column Family Descriptor;
import org.apache.hadoop.hbase.client.ColumnFamilyDescriptorBuilder;
import org.apache.hadoop.hbase.client.Connection;
import org.apache.hadoop.hbase.client.ConnectionFactory;
import org.apache.hadoop.hbase.client.TableDescriptor;
import org.apache.hadoop.hbase.client.TableDescriptorBuilder;
import static org.apache.hadoop.hbase.io.crypto.tls.X509Util.HBASE_CLIENT_NETTY_TLS_ENABLED;
import static
org.apache.hadoop.hbase.io.crypto.tls.X509Util.HBASE_CLIENT_NETTY_TLS_VERIFY_SERVER_HOSTNAME;
import static org.apache.hadoop.hbase.io.crypto.tls.X509Util.TLS_CONFIG_TRUSTSTORE_LOCATION;
import static org.apache.hadoop.hbase.io.crypto.tls.X509Util.TLS_CONFIG_TRUSTSTORE_PASSWORD;
public class ExampleTlsConnection
  public static void main(String[] args) throws Throwable
     Configuration conf = HBaseConfiguration.create();
     // todo: change connect address
     conf.set("hbase.zookeeper.quorum", "127.0.0.1");
     conf.set("hbase.zookeeper.property.clientPort", "2181");
```

```
// todo: change those two strings.
     String path = "/absolute/path/to/your/truststore.jks";
     String password = "your_truststore_pass_word";
     conf.setBoolean(HBASE_CLIENT_NETTY_TLS_ENABLED, true);
     conf.setBoolean(HBASE_CLIENT_NETTY_TLS_VERIFY_SERVER_HOSTNAME, false);
     conf.set(TLS_CONFIG_TRUSTSTORE_LOCATION, path);
     conf.set(TLS_CONFIG_TRUSTSTORE_PASSWORD, password);
     // Enters a username and password.
     UserGroupInformation ugi = UserGroupInformation.createProxyUser("your_user_name",
UserGroupInformation.createRemoteUser("your password"));
     try (Connection connection = ConnectionFactory.createConnection(conf, User.create(ugi)))
       Admin admin = connection.getAdmin();
       TableName tb = TableName.valueOf("test");
       List<ColumnFamilyDescriptor> cfs = new ArrayList<>();
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf1".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf2".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf3".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf4".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf5".getBytes()).build());
       TableDescriptor tableDescriptor =
TableDescriptorBuilder.newBuilder(tb).setColumnFamilies(cfs).build();
       admin.createTable(tableDescriptor);
  }
```

4.3.5 Connection Information Management

4.3.5.1 Setting Security Group Rules for a GeminiDB HBase Instance

A security group is a collection of access control rules for ECSs and GeminiDB HBase instances that have the same security protection requirements and are mutually trusted in a VPC.

To ensure database security and reliability, configure security group rules to allow specific IP addresses and ports to access GeminiDB HBase instances.

This section describes how to configure security group rules for a GeminiDB HBase instance connected over a private network.

Usage Notes

- By default, a tenant can create a maximum of 500 security group rules.
- Too many security group rules will increase the first packet latency. You are advised to create a maximum of 50 rules for each security group.
- Currently, a GeminiDB HBase instance can be associated with only one security group.
- For details about the security group rules for connecting to an instance over a private network, see **Table 4-8**.

Table 4-8 Security group rules

Scenario	Description	
Connecting to an instance over a private network	 Configure security group rules as follows: If the ECS and GeminiDB HBase instance are in the same security group, they can communicate with each other by default. No security group rule needs to be configured. If they are in different security groups, configure security group rules for them, separately. 	
	 Configure inbound rules for the security group associated with the GeminiDB HBase instance. For details, see Procedure. 	
	 The default security group rule allows all outbound data packets, so you do not need to set a security rule for the ECS. If not all outbound traffic is allowed in the security group, set an outbound rule for the ECS. 	

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- Step 2 In the service list, choose Databases > GeminiDB.
- **Step 3** On the **Instances** page, click the target instance go to the **Basic Information** page.
- **Step 4** Set security group rules.

Method 1:

In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 4-8 Security group



Method 2


On the **Basic Information** page, choose **Connections** in the navigation pane on the left. In the **Security Group** area on the right, click the name of the security group. The **Security Group** page is displayed.



Step 5 Add an inbound rule.


1. Click the **Inbound Rules** tab.

Figure 4-9 Inbound rule



2. Click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Figure 4-10 Adding a rule



3. Add a security group rule as prompted.

Table 4-9 Inbound rule settings

Parame ter	Description	Example Value
Protoco l & Port	 Protocol: Currently, GeminiDB HBase instances can be accessed only over TCP. Port: The port (1 to 65535) for accessing the ECS. 	TCP
Туре	IP address type. This parameter is available after IPv6 is enabled. - IPv4	IPv4

Parame ter	Description	Example Value
Source	Source: The source can be an IP address, a security group, or an IP address group which allows access from IP addresses or instances in other security groups. For example:	
	- xxx.xxx.xxx/32 (IPv4 address)	
	- xxx.xxx.xxx.0/24 (subnet)	
	- 0.0.0.0/0 (any IP address)	
	– sg-abc (security group)	
Descrip tion	(Optional) Provides supplementary information about the security group rule.	-
	The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	

Step 6 Click OK.

----End

4.3.5.2 Viewing the IP Address and Port of a GeminiDB HBase Instance

This section describes how to query the IP address and port of a GeminiDB HBase instance on the console.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance.

Method 1:

In the **Node Information** area on the **Basic Information** page, view the private IP address or EIP of each instance node.

Figure 4-11 Viewing the IP address



In the **Network Information** area, view the instance port. The default port is **2181**.

Figure 4-12 Viewing the port



Method 2

In the navigation pane, choose **Connections** to view the private IP address, port, and EIP bound to the instance.

Figure 4-13 Viewing the IP address and port



----End

4.3.5.3 Changing the Security Group of a GeminiDB HBase Instance

You can change the security group of a GeminiDB HBase instance.

Usage Notes

If you are adding nodes to an instance, the security group cannot be changed.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance whose security group you want to change and click its name.
- **Step 4** In the navigation pane, choose **Connections**.
- **Step 5** In the **Security Group** area, click next to the security group name and select the security group of the instance.
 - Click to submit the modification. This process takes about 1 to 3 minutes.

Click X to cancel the modification.

Step 6 View the modification result.

----End

4.4 Data Migration

4.4.1 Solution Overview

If you have any migration questions about GeminiDB HBase API, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the console.

4.5 Instance Lifecycle Management

The status of an instance indicates its health. You can view the status on the console.

Table 4-10 Instance statuses

Status	Description
Available	The DB instance is available.
Abnormal	The instance is abnormal.
Creating	The instance is being created.
Creation failed	DB instance creation fails.
Restarting	The instance is being restarted.
Resetting password	The administrator password is being reset.
Adding node	Nodes are being added to an instance.
Deleting node	Nodes are being deleted from an instance.
Scaling storage space	The storage space of an instance is being scaled up.
Changing specifications	The vCPUs and memory of an instance are being changed.
Uploading backup	The backup file is being uploaded.
Backing up	A database backup is being created.
Checking restoration	The backup of the instance is being restored to a new instance.
Changing to yearly/monthly	The billing mode is being changed from pay-per-use to yearly/monthly.

Status	Description
Changing to pay- per-use	The billing mode is being changed from yearly/monthly to pay-per-use.
Creating a DR cluster	A DR instance is being created.
Canceling DR relationship	A DR instance is being deleted.
Configuring SSL	SSL is being enabled or disabled.
Frozen	The instance is frozen because your balance drops to or below zero.
Unfreezing	Overdue payments are cleared, and the DB instance is being unfrozen.
Checking changes	The yearly/monthly instance is pending check when its billing mode is changed.

4.5.1 Restarting a GeminiDB HBase instance

You may need to restart an instance for routine maintenance.

Usage Notes

- Only instances in states **Available**, **Abnormal**, or **Checking restoration** can be restarted
- Restarting an instance will interrupt services, so off-peak hours are the best time. Ensure that your application can be reconnected.
- After you restart an instance, all nodes in the instance are also restarted.
- If you enable operation protection, two-factor authentication is required for sensitive operations to secure your account and cloud products. For details about how to enable operation protection, see *Identity and Access Management User Guide*.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Restart** in the **Operation** column.
 - You can also click the instance and click **Restart** in the upper right corner of the displayed page.
- **Step 4** If you have enabled operation protection, click **Start Verification** in the **Restart DB Instance** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

Step 5 In the displayed dialog box, click **Yes**.

----End

4.5.2 Exporting Instance Information

Scenarios

You can export information about all or selected instances to view and analyze instance information.

Exporting All Instance Information

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3 On the Instances page, click in the upper right corner. By default, information about all instances are exported. In the displayed dialog box, you can select the items to be exported and click **Export**.
- **Step 4** After the export task is complete, check an XLS file is generated locally.

----End

Exporting Information About Selected Instances

- Step 1 On the Instances page, select the target instances or search for required instances by project, compatible API, name, ID, or tag and click in the upper right corner. In the displayed dialog box, select the items to be exported and click Export.
- **Step 2** After the export task is complete, check an XLS file is generated locally.

----End

4.5.3 Deleting a Pay-per-Use Instance

You can manually delete a pay-per-use instance on the **Instances** page. Before deleting a yearly/monthly instance, you need to unsubscribe from it. For details, see **How Do I Unsubscribe from a Yearly/Monthly Instance?**.

Usage Notes

- Instances that an operation is being performed on cannot be deleted. They can be deleted only after the operations are complete.
- If a pay-per-use instance is deleted, its automated backups will also be deleted and you will no longer be billed for them. Manual backups, however, will be retained and generate additional costs.
- After an instance is deleted, all its data and automated backups are automatically deleted as well and cannot be recovered. You are advised to create a backup before deleting an instance. For details, see Creating a Manual Backup.

- After you delete an instance, all of its nodes are deleted.
- A deleted instance will be retained in the recycle bin for a period of time after being released, so you can rebuild the instance and restore data from it.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Delete** in the **Operation** column.
- **Step 4** If you have enabled operation protection, click **Start Verification** in the **Delete DB Instance** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

□ NOTE

If you enable operation protection, two-factor authentication is required for sensitive operations to secure your account and cloud products. For details about how to enable operation protection, see *Identity and Access Management User Guide*.

Step 5 In the displayed dialog box, click **Yes**.

Deleted instances are not displayed in the instance list any longer.

----End

4.5.4 Recycling an Instance

You can move unsubscribed yearly/monthly and deleted pay-per-use GeminiDB HBase instances to the recycle bin and rebuild them if necessary.

You can move deleted GeminiDB HBase instances to the recycle bin and rebuild them if necessary.


Usage Notes

- The recycling bin is enabled by default and cannot be disabled. Instances in the recycle bin are retained for 7 days by default, and this will not incur any charges.
- Currently, you can put a maximum of 100 instances into the recycle bin.
- If you delete an instance of full storage, the deleted instance will not be moved to the recycle bin.
- You can modify the retention period, and the changes only apply to the DB instances deleted after the changes, so exercise caution when performing this operation.
- After an instance is deleted, the most recent automated full backup (if no automated full backup is available one day ago, the latest one is retained) is retained and a full backup is performed. You can select any backup file to restore the instance data.

Modifying the Recycling Policy

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Recycling Bin** page, click **Modify Recycling Policy**. In the displayed dialog box, set the retention period from 1 day to 7 days. Then, click **OK**.

Figure 4-14 Modifying a recycling policy



----End

Rebuilding an Instance

You can rebuild DB instances from the recycle bin within the retention period to restore data.

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Recycling Bin** page, locate the target instance and click **Rebuild** in the **Operation** column.

Figure 4-15 Rebuilding an instance



Step 4 On the displayed page, configure required parameters and submit the task.

----End

4.6 Instance Modifications

4.6.1 Upgrading a Minor Version

GeminiDB HBase API can be upgraded by installing patches to improve performance, release new features, or fix bugs.

If a new patch is released, you can upgrade your instance by clicking the upgrade button in the **Compatible API** column on the **Instances** page.

Figure 4-16 Patch installation



If the kernel version of your instance has potential risks or major defects, has expired, or has been brought offline, the system will notify you by SMS message or email and deliver an upgrade task during maintenance.

Usage Notes

- Upgrade your instance once there are new patches released.
- If the database version is a risky version, the system prompts you to upgrade the database patch.
- The instance will be restarted and services may be interrupted during the upgrade. The interruption duration depends on services, quantity of nodes, and the amount of service data. Upgrade your instance during off-peak hours.
- When you upgrade a cluster, services may be interrupted a number of times
 equal to the number of nodes in the cluster plus one. Each interruption will
 last for no more than 60s and will only affect the services on that node. If
 your cluster has more than one node, the upgrade duration is as follows:
 - $600 + (N \times 60) \le Total upgrade duration (s) \le 600 + (N \times 120)$

For example, if there are 9 nodes in a cluster instance, the upgrade duration is 19 to 28 minutes.

The upgrade duration of most instances is close to 600+ (N x 60). If there are too many tokens on a single node, the upgrade duration may be increased.

 Before you upgrade a DR instance, upgrade the corresponding standby instance first and then the primary instance afterwards.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance you want to upgrade and click **Upgrade Minor Version** in the **Compatible API** column.

Figure 4-17 Patch installation




Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Minor Version** in the **Compatible API** field.

Figure 4-18 Patch installation



Step 4 In the displayed dialog box, click **OK**.

Figure 4-19 Confirming information



Step 5 View the upgrade result on the **Instances** page.

- When the upgrade is ongoing, the instance status is **Upgrading minor** version.
- After the upgrade is complete, the instance status changes **Available**.

----End

4.6.2 Changing an Instance Name

This section describes how to modify the name of a GeminiDB HBase instance.

Method 1

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click $\stackrel{\checkmark}{=}$ next to the target instance name and change it.
 - To submit the change, click **OK**.
 - To cancel the change, click **Cancel**.

Ⅲ NOTE

The instance name:

- Can be the same as an existing instance name.
- Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).

Step 4 View the results on the **Instances** page.

----End

Method 2

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the **Instance Information** area on the **Basic Information** page, click next to **DB Instance Name** and change the instance name.
 - Click to submit the modification.
 - Click X to cancel the modification.

The instance name:

- Can be the same as an existing instance name.
- Can include 4 to 64 bytes and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).

Step 5 On the **Basic Information** page, check the new name.

----End

4.6.3 Resetting the Administrator Password

For security reasons, regularly change your administrator password.

Usage Notes

- If the instance status is Available, Backing up, Checking restoration, Scaling up or certain nodes become abnormal, you can reset the administrator password.
- The administrator password takes effect immediately after being reset.
- For two instances with an intra-region DR or cross-region dual-active relationship, make sure that they have the same administrator passwords.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see *Identity* and Access Management User Guide.



You are advised to change the password during off-peak hours to avoid service interruption.

Method 1

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance whose password you want to reset and choose **More** > **Reset Password** in the **Operation** column.
- **Step 4** Enter and confirm the new administrator password and click **OK**.

The password must contain 8 to 32 characters. It can include uppercase letters, lowercase letters, digits, and any of the following special characters: $\sim !@#\%^*-=+?$

Step 5 If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

----End

Method 2

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance whose password you want to reset to go to the **Basic Information** page.
- Step 4 In the DB Information area, click Reset Password in the Administrator field.
- **Step 5** Enter and confirm the new administrator password and click **OK**.

The password must contain 8 to 32 characters. It can include uppercase letters, lowercase letters, digits, and any of the following special characters: $\sim !@#\%^*-=+?$

Step 6 If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

----End

4.6.4 Changing vCPUs and Memory

This section describes how to change your instance vCPUs and memory to suit your service requirements.

Usage Notes

- You can increase or decrease the vCPUs and memory.
- If one instance has multiple nodes, the change will be performed on the nodes one by one. It takes about 5 to 10 minutes for each node, and the total time required depends on the number of the nodes.
- For a node whose specifications are being changed, its computing tasks are handed over to other nodes. Change specifications of nodes during off-peak hours to prevent the instance from overload.

• Do not perform DDL operations when you change the instance specifications.

◯ NOTE

A data definition language (DDL) is a language for defining data structures and database objects. Common examples of DDL statements are CREATE, ALTER, and DROP. Data Definition Language (DDL) is used to create, modify, and delete database objects, such as tables, indexes, views, functions, stored procedures, and triggers.

- vCPU and memory changes are applied on all nodes in sequence. During this process, temporary I/O disruptions or increased latency may occur. You are advised to perform this operation during off-peak hours.
- If you forcibly change the specifications of an instance when the instance is abnormal, services may be affected in seconds.

Method 1


- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance whose vCPUs and memory you want to change and click its name.
- **Step 4** In the in **DB Information** area on the **Basic Information** page, click **Change** next to **Specifications**.

Figure 4-20 Changing specifications



Step 5 On the displayed page, select the required specifications and click **Next**.

Figure 4-21 Changing specifications



- **Step 6** On the displayed page, confirm the instance specifications.
 - Yearly/Monthly
 - If you need to modify your settings, click Previous.
 - If you do not need to modify your settings, click **Submit**. If you are scaling up the instance specifications, go to the payment page, select a payment method, and complete the payment.
 - Pay-per-use
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.

Step 7 View the change results.

In the **DB Information** area on the **Basic Information** page, you can see the new specifications.

----End

Method 2


- **Step 1** Log in to the Huawei Cloud console.
- Step 2 In the service list, choose Databases > GeminiDB.
- **Step 3** On the **Instances** page, locate the instance whose specifications you want to change and choose **More** > **Change Specifications** in the **Operation** column.

Figure 4-22 Changing specifications



Step 4 On the displayed page, select the required specifications and click **Next**.

Figure 4-23 Changing specifications



Step 5 On the displayed page, confirm the instance specifications.

- Yearly/Monthly
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click **Submit**. If you are scaling up the instance specifications, go to the payment page, select a payment method, and complete the payment.
- Pay-per-use
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.

Step 6 View the change results.

In the **DB Information** area on the **Basic Information** page, you can see the new specifications.

----End

4.6.5 Setting a Maintenance Window

The default maintenance window is 10:00–14:00 (GMT+08:00) but you can change it if needed. To prevent service interruption, set the maintenance window to offpeak hours. Before calling this API:


Usage Notes

- You can configure a maintenance window only for restarting a DB instance, changing an instance class, or upgrading the minor version of a DB instance.
- The specification change and patch upgrade that have been performed during the maintenance period cannot be performed immediately. The instance can be restarted immediately.
- You can cancel a task to be executed.
- Changing the maintenance window will not affect the timing that has already been scheduled.
- The maintenance window cannot overlap the time window configured for backups. Otherwise, scheduled tasks may fail.
- During the maintenance window, the scheduled task is scanned and executed every 10 minutes. If the task is delivered near the end of the maintenance period, the task may fail to be scanned and the execution is canceled.

Setting a Maintenance Window

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the **DB Information** area, locate **Maintenance Window** and click **Change**.


Figure 4-24 The change button



Step 5 On the **Change Maintainable Window** page, select the maintenance time period as needed, and then click **OK**.

Supported time periods: 02:00-06:00, 06:00-10:00, 10:00-14:00, 14:00-18:00, 18:00-22:00, and 22:00-02:00

Figure 4-25 Changing a maintenance window



Step 6 Check the result.

On the **Basic Information** page, you can view the changed maintenance window.

----End

Querying an Executed Task

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Task Center** page, click the **Instant Tasks** or **Scheduled Tasks** tab to view a task.

Figure 4-26 Querying a task



Canceling a Scheduled Task

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Task Center** page, locate a scheduled task, and click **Cancel** in the **Operation** column.

Figure 4-27 Canceling a task



Step 4 Check the result.

On the **Task Center** page, you can view the result. After the task is cancelled, its status changes to **Cancelled**.

Figure 4-28 Checking cancelled tasks



----End

4.6.6 Adding and Deleting Instance Nodes


4.6.6.1 Overview

After you buy a GeminiDB HBase instance, resource requirements may change along with workload volumes. You can scale your instance nodes in the following ways.

Manually Adding Instance Nodes

For example, if three nodes have been deployed and two more nodes need to be added, there will be five nodes in total. For details, see **Manually Adding Instance Nodes**.


Figure 4-29 Adding instance nodes



Automatically Adding Instance Nodes

For example, if three nodes have been deployed and two more nodes need to be added, there will be five nodes in total. For details, see **Automatically Adding Instance Nodes**.


Figure 4-30 Adding instance nodes



Manually Deleting Instance Nodes

For example, if five nodes have been deployed and two of them need to be deleted, three nodes will be left. For details, see **Manually Deleting Instance Nodes**.

Figure 4-31 Deleting instance shards



4.6.6.2 Manually Adding Instance Nodes

This section describes how to add nodes to an instance to suit your service requirements.

Usage Notes

- Adding nodes may lead to the decrease of operations per second (OPS).
 Perform this operation during off-peak hours.
- You can only add nodes when the instance status is Available or Checking restoration.
- Instances that one or more nodes are added to cannot be deleted.
- After the nodes are added, you can delete them by following Manually Deleting Instance Nodes.
- Currently, a maximum of 60 nodes are supported. To add more, choose
 Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.

Method 1


- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance that you want to add nodes to and click its name.
- **Step 4** In the **Node Information** area on the **Basic Information** page, click **Add Node**.

Figure 4-32 Node information



Step 5 Specify **Add Nodes** and click **Next**.

Figure 4-33 Adding nodes



■ NOTE

- New nodes are of the same specifications as existing nodes. Once a new node is added, its specifications cannot be changed.
- New nodes and the instance can be in different subnets of the same VPC.

Step 6 On the displayed page, confirm the node configurations.

- Yearly/Monthly
 - If you need to modify your settings, click Previous.
 - If you do not need to modify your settings, click Next and complete the payment.
- Pay-per-use
 - If you need to modify your settings, click Previous.
 - If you do not need to modify your settings, click **Submit**.

Step 7 View the results.

- When new nodes are being added, the instance status is **Adding node**.
- After the nodes are added, the instance status becomes **Available**.
- Click the instance name. In the **Node Information** area on the **Basic Information** page, view information about the new nodes.

----End

Method 2


- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Add Node** in the **Operation** column.

Figure 4-34 Adding nodes



Step 4 Specify Add Nodes and click Next.

Figure 4-35 Adding nodes



□ NOTE

- New nodes are of the same specifications as existing nodes. Once a new node is added, its specifications cannot be changed.
- New nodes and the instance can be in different subnets of the same VPC.

Step 5 On the displayed page, confirm the node configurations.

- Yearly/Monthly
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Next and complete the payment.
- Pay-per-use
 - If you need to modify your settings, click **Previous**.
 - If you do not need to modify your settings, click Submit.

Step 6 View the results.

- When new nodes are being added, the instance status is **Adding node**.
- After the nodes are added, the instance status becomes **Available**.
- Click the instance name. In the **Node Information** area on the **Basic Information** page, view information about the new nodes.

----End

4.6.6.3 Automatically Adding Instance Nodes

When an autoscaling threshold is met, GeminiDB HBase instance nodes can be automatically added to reduce I/O pressure.

□ NOTE

- If you enable **Auto Scale** using a Huawei Cloud account, no additional configuration is required.
- If you enable Auto Scale as an IAM user first time, you need to obtain the permission to create an agency.

Configuring Permissions

If you are using an IAM user, perform the following operations to configure GeminiDB and IAM permissions before you enable storage autoscaling:

- 1. Configure the GeminiDB FullAccess permission.
- 2. Configure fine-grained permissions for IAM.

For details about how to configure IAM permissions, see **Creating a Custom Policy**.

Custom policy in JSON format:

```
"iam:agencies:listAgencies",
    "iam:roles:listRoles",
    "iam:roles:createRole"
]
}
]
}
```

3. Create a user group and assign permissions.

You can create a user group on the IAM console and grant it custom permissions created in 2 and the security administrator role.

4. Create an IAM user and add it to a user group.

Log in to the IAM console using a Huawei Cloud account or as an IAM user, locate the IAM user that the target instance belongs to, and add it to the user group created in 3. The IAM user will inherit permissions of the user group.


Usage Notes

- This function is now in OBT. To use it, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.
- If resources in the current region are insufficient, nodes may fail to be added.
- Autoscaling is available only when your account balance is sufficient.
- The instance is in the **Available** status.
- Once autoscaling is enabled, an agency will be created and fees will be automatically deducted.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the **Node Information** area on the **Basic Information** page, click **Auto Scale**.

Figure 4-36 Auto Scale



Step 5 Set autoscaling triggers and thresholds listed in **Table 4-11**.




Table 4-11 Description

Parameter	Description		
Auto Scale	Whether autoscaling is enabled		
Percentage of Overloaded Nodes ≥	Percentage of overloaded nodes that need to be scaled out		
Memory Usage ≥	Memory usage of nodes for which autoscaling is triggered		
CPU Usage ≥	CPU usage of nodes for which autoscaling is triggered		
Increase By	Number of nodes to be added each time		
Maximum Nodes	Maximum number of nodes that can be automatically added		

----End

4.6.6.4 Manually Deleting Instance Nodes

You can delete nodes that are no longer used to release resources.

Usage Notes

- Deleted nodes cannot be recovered. Exercise caution when performing this operation.
- Only pay-per-use instances can be deleted.
- Deleting nodes will cause the OPS to decrease for a short period of time. Deleting nodes during off-peak hours.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see *Identity* and Access Management User Guide.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- Step 2 In the service list, choose Databases > GeminiDB.
- **Step 3** On the **Instances** page, locate the instance that you want to delete nodes from and click its name.
- **Step 4** In the **Node Information** area on the **Basic Information** page, locate the target node and click **Delete** in the **Operation** column.
- **Step 5** If you have enabled operation protection, click **Start Verification** in the **Delete Node** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.
- **Step 6** In the displayed dialog box, click **Yes**.
 - When the node is being deleted, the instance status is **Deleting node**.
 - After the node is deleted, the instance status becomes **Available**.

----End

4.6.7 Scaling Storage Space

4.6.7.1 Overview

As more data is added, you may run out of storage. This section describes how to scale up storage space of your instance. As data volumes decrease, you can scale down storage to avoid low database node utilization and resource waste. **Table 4-12** lists the scaling methods supported by GeminiDB HBase instances.

Table 4-12 Scaling methods


Method	Description
Manually Scaling Up Storage Space	You can specify how much storage space needs to be added. The added value must be a multiple of 1 (GB). The total storage space cannot exceed the upper limit defined by your instance specifications.

Method	Description	
Automatica lly Scaling	If storage usage exceeds the configured threshold, autoscaling will be triggered.	
Up Storage Space	The storage is scaled up by a percentage you specify. The added storage space is the current storage space multiplied by the scaling increment.	
Manually Scaling Down Storage Space	You can specify how much storage space needs to be reduced. The storage space to be reduced must be an integer multiple of 1 GB and greater than or equal to 125% of the used storage space. The value is rounded up.	

Manually Scaling Up Storage Space

For example, if the storage space of a cluster instance is 24 GB and is increased by 8 GB, the storage space will become 32 GB.


Figure 4-37 Manually scaling up storage space



Automatically Scaling Up Storage Space

For example, the storage space of a cluster instance is 24 GB before scale-up, the storage usage threshold for triggering autoscaling is set to 80%, and the total storage needs to be automatically scaled up by 10%. If the storage usage is greater than or equal to 80%, the storage space is automatically scaled up by 2.4 GB (24 x 10%), which is rounded up to 3 GB. In this case, the total storage space becomes 27 GB (24 + 3).


Figure 4-38 Automatically scaling up storage space



Manually Scaling Down Storage Space

For example, if the storage space of a cluster instance is 32 GB and is decreased by 8 GB, the storage space will become 24 GB.

Figure 4-39 Manually scaling down storage space



4.6.7.2 Manually Scaling Up Storage Space

This section describes how to scale up storage of an instance to suit your service requirements.

Usage Notes

- Scaling up storage does not interrupt your services. After storage scale-up is complete, you do not need to restart your instance.
- If your yearly/monthly instance is running out of storage, additional usage will be billed on a pay-per-use basis. To avoid these extra costs and maintain the benefits of your yearly/monthly subscription, you are advised to scale up storage.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Scale Storage Space** in the **Operation** column.

Click the instance name. In the **Storage Space** area on the **Basic Information** page, click **Scale**.

Step 4 On the displayed page, specify new storage and click **Next**.

Figure 4-40 Scaling up storage space



- To scale up classic storage, you need to add at least 1 GB each time. The value must be an integer.
- To scale up cloud native storage, you need to add at least 10 GB each time. The value must be an integer multiple of 10.

Step 5 On the displayed page, confirm the storage space.

- Yearly/Monthly
 - If you need to modify your settings, click Previous.
 - If you do not need to modify your settings, click Next and complete the payment.
- Pay-per-use
 - If you need to modify your settings, click Previous.
 - If you do not need to modify your settings, click Submit.

Step 6 Check the results.

- When the scaling task is ongoing, the instance status is **Scaling storage space**.
- After the scaling process, the instance status becomes **Available**.
- Click the instance name. In the **Storage Space** area on the **Basic Information** page, check the new storage space.

----End

4.6.7.3 Automatically Scaling Up Storage Space

You can enable storage autoscaling for GeminiDB HBase instances. When storage usage reaches the limit, autoscaling is triggered.

You can enable Auto Scale:

- When you create an instance. For details, see Getting to Know GeminiDB HBase API.
- 2. After you create an instance

This section describes how to configure **Auto Scale** after an instance is created.

□ NOTE

- If you enable Auto Scale using a Huawei Cloud account, no additional configuration is required.
- If you enable **Auto Scale** as an IAM user first time, you need to obtain the permission to create an agency.

Configuring Permissions

If you are using an IAM user, perform the following operations to configure GeminiDB and IAM permissions before you enable storage autoscaling:

- 1. Configure the GeminiDB FullAccess permission.
- Configure fine-grained permissions for IAM.
 For details about how to configure IAM permissions, see Creating a Custom Policy.

If you use the JSON view to configure a custom policy, the policy content is as follows:

3. Create a user group and assign permissions.

You can create a user group on the IAM console and grant it custom permissions created in 2 and the security administrator role.

4. Create an IAM user and add it to a user group.

Log in to the IAM console using a Huawei Cloud account or as an IAM user, locate the IAM user that the target instance belongs to, and add it to the user group created in 3. The IAM user will inherit permissions of the user group.


Usage Notes

- Autoscaling is available only when your account balance is sufficient.
- The instance is in the **Available** status.
- Once **Auto Scale** is enabled, an agency will be created and fees will be automatically deducted.

Automatically Scaling Up Storage of a Single Instance

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the **Storage Space** area, click **Auto Scale**.

Figure 4-41 Auto Scale



Step 5 Toggle on **Auto Scale** and specify the parameters below.

Figure 4-42 Configuring autoscaling parameters




Table 4-13 Description

Parameter	Description	
Auto Scale	If you toggle on this switch, autoscaling is enabled.	
Trigger If Available Storage Drops To	When the available storage usage drops to a specified threshold or the available storage drops to 10 GB, autoscaling is triggered.	
Increase By	Percentage that your instance storage will be scaled up at. The value can be 10% , 15% , or 20% . If the value is not a multiple of 10, it is rounded up. At least 100 GB is added each time.	
Storage Limit	Limit of storage (GB) that can be automatically scaled up to. The value must be no less than the storage of your instance and cannot exceed the maximum storage supported by your instance.	


Step 6 Click OK.

----End

Automatically Scaling Up Storage of Multiple Instances In Batches

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- Step 3 Select instances and click Auto Scale.

Figure 4-43 Auto Scale



Step 4 Select an instance, toggle on **Auto Scale**, and specify the parameters below.

Figure 4-44 Batch Auto Scale




Table 4-14 Description

Parameter	Description
Auto Scale	If you toggle on this switch, autoscaling is enabled.
Trigger If Available Storage Drops To	When the available storage usage drops to a specified threshold or the available storage drops to 10 GB, autoscaling is triggered.

Parameter	Description
Increase By	Percentage that your instance storage will be scaled up at. The value can be 10% , 15% , or 20% . If the value is not a multiple of 10, it is rounded up. At least 100 GB is added each time.
Storage Limit	The value cannot be specified. By default, the storage is scaled up to the maximum defined by your instance specifications.

Step 5 Click OK.

----End

4.6.7.4 Manually Scaling Down Storage Space

As data volumes decrease, you can scale down storage space to avoid low database node utilization and resource waste.

Usage Notes

- To scale down storage, ensure the new storage space is at least 1.25 times more than the used space and rounded up.
- Scaling down storage does not interrupt your services, and you do not need to restart your instance.
- If your yearly/monthly instance is running out of storage, additional usage will be billed on a pay-per-use basis. To avoid these extra costs and maintain the benefits of your yearly/monthly subscription, you are advised to scale up storage.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the target instance and choose **More** > **Scale Storage Space** in the **Operation** column.

Click the instance name. In the **Storage Space** area on the **Basic Information** page, click **Scale**.

Step 4 On the displayed page, specify new storage and click **Next**.

Figure 4-45 Scaling down storage space scale Storage Space ⊙



Select at least 1 GB each time, and the value must be an integer.

Step 5 On the displayed page, confirm the storage space.

- Yearly/Monthly
 - To modify your settings, click **Previous** to go back to the page where you specify details.
 - If you do not need to modify your settings, click Next and complete the payment.
- Pay-per-use
 - To modify your settings, click **Previous** to go back to the page where you specify details.
 - If you do not need to modify your settings, click Submit.

Step 6 Check the results.

- During the scale-down process, the instance status becomes Scaling storage space.
- After the scaling process, the instance status becomes **Available**.
- Click the instance name. In the **Storage Space** area on the **Basic Information** page, check the new storage space.

----End

4.7 Data Backup

4.7.1 Overview

You can create backups for GeminiDB HBase instances to ensure data reliability. After an instance is deleted, the manual backup data is retained. Automated backup data is released together with instances. Backup data cannot be downloaded or exported.

Usage Notes

Backing up data consumes a few CPUs. Uploading backup files to OBS occupies bandwidth of compute nodes, causing slight latency and jitter.

Backup Methods

You can create backups manually or automatically.

Automated backup

You can **modify a backup policy** on the GeminiDB console, and the system will automatically back up your instance data based on the time window and backup cycle you configured in the backup policy and will store the data based on the specified retention period.

Automated backups cannot be manually deleted. You can adjust their retention period by following **Modifying an Automated Backup Policy** . Expired backups will be automatically deleted.

Manual backup

A manual backup is a full backup of a DB instance and can be retained until you manually delete it. Manual backup can be triggered at any time to meet your service requirements.

Regularly backing up your database is recommended. If your database becomes faulty or data is corrupted, you can restore it from backups.

Table 4-15 Comparison between automated backup and manual backup

Backup Method	Scenario
Managing Automated Backups	After you set a backup policy, the system automatically backs up your database based on the policy. You can also modify the policy based on service requirements. Either incremental or full backup is supported.
Managing Manual Backups	You can enable full backup for your instance based on service requirements.

Cross-region and table-level backups are supported based on application scenarios.

Table 4-16 Application scenarios

Method	Scenario
Managing Cross-Region Backups	Backups can be stored in the destination region. Then for disaster recovery, you can restore the backups to a new instance in another region. Only an automated full backup is supported.
Managing Table-level Backups	If a database or table is deleted maliciously or accidentally, you can use backups to restore data. Manual and automated backups are supported.

Full and incremental backups are created based on data volumes.

Table 4-17 Comparison between full and incremental backups

Bac kup Typ e	Full backup	Incremental backup
Des crip tion	All data in an instance is backed up.	Only data that has changed within a certain period is backed up.

Ena bled by Def ault	Yes	Yes
Ret enti on Dur atio n	 You can specify how many days automated backups can be retained for. If you shorten the retention duration, the new backup policy takes effect for existing backups. Manual backups are always retained even though a GeminiDB HBase instance is deleted. They can only be deleted manually. 	Incremental backups will be deleted along with automated full backups.
Feat ure	 All data of your instance is backed up in the current point of time. You can use a full backup to restore all data generated when its backup was created. Full backups can be created automatically or manually. 	 Incremental data in your instance is backed up since the last full backup. When you use an incremental backup for restoration, the last full backup data and the incremental data generated since then are downloaded. Incremental backups can be created automatically only.
Ho w to Vie w	Click an instance name. On the Backups & Restorations page, click the Instance-level Backups and Table-level Backups tabs to view the backup size.	Click an instance name. On the Backups & Restorations page, click the Incremental Backup tab to view the backup size.

How Data Is Backed Up

GeminiDB HBase API provides a seed node dedicated to backup management. As shown in the following figure, the seed node backs up data of GeminiDB HBase cluster instances. The node takes snapshots in seconds and then stores them as compressed backups in OBS buckets, without occupying storage of your instance. The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.




Figure 4-46 Backup process

Backup Storage

Backups are stored in OBS buckets to provide disaster recovery without occupying storage of your instance.

After you buy instance storage, an equivalent amount of backup storage will be provided free of charge. For example, if you purchase an instance of 100 GB, you will obtain additional backup storage of 100 GB free of charge. If the backup data size stays within 100 GB, it is stored in OBS at no cost. If it exceeds 100 GB, additional data will be billed based on the billing rules.

4.7.2 Managing Automated Backups

You can create automated backups of GeminiDB HBase instances to ensure data reliability. If a database or table is deleted maliciously or accidentally, backups can help recover your data.

Usage Notes

 Backup files are saved as packages in OBS buckets. Uploading backup files and reading service data both consume bandwidth, so the upload bandwidth of OBS is limited. The upload bandwidth of a single node ranges from 20 MB/s to 70 MB/s.

You need to specify appropriate nodes based on the backup bandwidth limit to achieve better performance.

- The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.
- While backups are uploaded, the memory usage does not increase significantly. Generally, the memory usage is about 300 MB. How much memory usage actually increases depends on the instance data volume. The increased memory mainly caches data during backup upload and read. After the backups are uploaded, the memory usage becomes normal.
- You can manually modify incremental backups of a GeminiDB HBase instance.
- To enable the incremental backup function, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.
- After the incremental backup function is enabled, differential backup is selected by default. To enable PITR, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.

Configuring an Automated Backup Policy

Automated backups are generated according to a backup policy and saved as packages in OBS buckets to ensure data confidentiality and durability. You are advised to regularly back up your database, in case it becomes faulty or damaged. Backing up data affects the database read and write performance so you are advised to set the automated backup time window to off-peak hours.


When you create an instance, automated backup is enabled by default.

Modify Backup Policy Automated Backup Incremental Backup Incremental Backup Interval Create a backup immediately after the incremental backup policy is modified. Retention Period 7 + days Enter an integer from 1 to 3660 GMT+08:00 Time Zone 02:00-03:00 Time Window Backup Cycle All Monday Tuesdav Wednesday Thursday Friday Saturday Sunday A minimum of one day must be selected OΚ Cancel

Figure 4-47 Enabling the automated backup policy

- Incremental Backup is enabled by default. You can click to manually enable or disable it. After it is enabled, the system stores backup data in OBS.
 Select an incremental backup type. Differential backup is selected by default.
 - **Differential backup**: Data can be restored to a specified point in time.
 - **PITR backup**: Data can be restored to any point in time.

Figure 4-48 Selecting an incremental backup type



Enabling Incremental Backup will take effect in the next full backup. You are advised to select Create a backup immediately after the incremental backup policy is modified.

- If you select it, the full backup request is delivered immediately, and the incremental backup takes effect.
- If you do not select it, the incremental backup will take effect in the next full backup.




Figure 4-49 Selecting Create a backup immediately after the incremental backup policy is modified

- Incremental Backup Interval: Incremental backups are generated every 15 minutes.
- **Retention Period**: Automated backup files are saved for seven days by default. The retention period ranges from 1 to 3660 days. Full backups are retained till the retention period ends. However, even if the retention period has ended, the most recent backup will be retained.
 - Extending the retention period improves data reliability. You can extend the retention period as needed.
 - If you shorten the retention period, the new backup policy takes effect for existing backups. Any automated backups (including full and incremental backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.

□ NOTE

- If the retention period is less than seven days, the system automatically backs up data daily.
- The system checks existing automated backups and deletes any backups that exceed the backup retention period you configured.
- **Time Window**: A one-hour period the backup will be scheduled within 24 hours, such as 00:00–01:00. The backup time is displayed in GMT. After the DST or standard time is switched, the backup time segment changes with the time zone.

If **Retention Period** is set to **2**, full and incremental backups that have been stored for more than two days will be automatically deleted. For instance, a backup generated on Monday will be deleted on Wednesday; or a backup generated on Tuesday will be deleted on Thursday.

Policy for automatically deleting full backups:

To ensure data integrity, even after the retention period ends, the most recent backup will be retained:

If **Backup Cycle** was set to **Monday** and **Tuesday** and the **Retention Period** was set to **2**:

- The full backup generated on Monday will be automatically deleted on Thursday. The reasons are as follows:
 - The backup generated on Monday expires on Wednesday, but it is the last backup, so it will be retained until a new backup expires. The next backup will be generated on Tuesday and will expire on Thursday. So the full backup generated on Monday will not be automatically deleted until Thursday.
- A full backup generated on Tuesday will be automatically deleted on the following Wednesday. The reasons are as follows:
 - The backup generated on Tuesday will expire on Thursday, but as it is the last backup, so it will be retained until a new backup expires. The next backup will be generated on the following Monday and will expire on the following Wednesday. So the full backup generated on Tuesday will not be automatically deleted until the following Wednesday.
- Backup Cycle: All options are selected by default.
 - All: Each day of the week is selected. The system automatically backs up data every day.
 - You can select one or more days in a week. The system automatically backs up data at the specified time.

Ⅲ NOTE

A full backup starts within one hour of the time you specify. The amount of time required for the backup depends on the amount of data to be backed up. The more data has to be backed up, the longer it will take.


- After the DB instance is created, you can modify the automated backup policy as needed. You can change the time window after the DB instance is created. The system backs up data based on an automated backup policy you configure.
- If **Automated Backup** is disabled, any automated backups in progress stop immediately.

Modifying an Automated Backup Policy

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance you want to back up.
- **Step 4** Choose **Backups & Restorations** in the navigation pane one the left, and click **Modify Backup Policy**. In the displayed dialog box, configure the backup policy. Click **OK**.

For details about how to set a backup policy, see **Configuring an Automated Backup Policy**.

Figure 4-50 Modifying the backup policy



Step 5 Check or manage the generated backups on the **Backups** or **Backups & Restorations** page.

----End


Viewing Incremental Backups

You can view incremental backups and their size of a GeminiDB HBase instance.

- To view the size and records of incremental backups, choose Service Tickets >
 Create Service Ticket in the upper right corner of the console and contact
 the customer service.
- You can view incremental backups and their size only after you enable Incremental Backup, or no data is displayed.
- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance that you want to connect to and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.

Step 5 On the **Backups & Restorations** page, click **Incremental Backup**.

Figure 4-51 Incremental backup



Step 6 View incremental backups and their size.

----End

Disabling Incremental Backup

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance you want to back up.
- Step 4 In the navigation pane, choose Backups & Restorations. On the displayed page, click Modify Backup Policy and click next to Incremental Backup.




Figure 4-52 Disabling Incremental Backup

----End

Disabling Automated Backup

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance you want to back up.
- **Step 4** Choose **Backups & Restorations** in the navigation pane one the left, and click **Modify Backup Policy**.
- Step 5 In the displayed dialog box, click and click OK.




Figure 4-53 Disabling backup policies

When your disable **Automated Backup**, specify whether to delete the automated backups:

- If you select **Delete automated backups**, all backup files within the retention period will be deleted. There are no automated backups displayed until you enable automated backup again.
- If you do not select **Delete automated backups**, backup files within the retention period will be retained, but you can still manually delete them later if needed. For details, see **Deleting an Automated Backup**.

If **Automated Backup** is disabled, any automated backups in progress stop immediately.

----End

Deleting an Automated Backup

If automated backup is disabled, you can delete stored automated backups to free up storage space.

If automated backup is enabled, the system will delete automated backups when they expire. You cannot delete them manually.

NOTICE

Deleted backups cannot be recovered. Exercise caution when performing this operation.

Method 1

- a. Log in to the Huawei Cloud console.
- b. In the service list, choose **Databases** > **GeminiDB**.
- c. On the **Instances** page, click the instance you want to back up.
- d. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup you want to delete, and click **Delete** in the **Operation** column.
- e. In the displayed dialog box, confirm the backup details and click Yes.

Method 2

- a. Log in to the Huawei Cloud console.
- b. In the service list, choose **Databases** > **GeminiDB**.
- c. On the **Backups** page, locate the backup that you want to delete and click **Delete**.
- d. In the displayed dialog box, confirm the backup details and click **Yes**.

4.7.3 Managing Manual Backups

To ensure data reliability, you can manually back up GeminiDB HBase instances in the **Available** status. If a database or table is deleted, maliciously or accidentally, backups can help recover your data.

Usage Notes

- Manual backups are full backups.
- Backup files are saved as packages in OBS buckets. Uploading backup files and reading service data both consume bandwidth, so the upload bandwidth of OBS is limited. The upload bandwidth of a single node ranges from 20 MB/s to 70 MB/s.
 - For better performance, you need to specify appropriate nodes for an instance and take into account the bandwidth for uploading backups.
- The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.
- While backups are uploaded, the memory usage does not increase significantly. Generally, the memory usage is about 300 MB. How much memory usage actually increases depends on the instance data volume. The increased memory mainly caches data during backup upload and service read. After the backup upload is complete, the memory recovers.
- Manual backups are charged for instances with cloud native storage during OBT.

Creating a Manual Backup

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** Create a manual backup.

Method 1

On the **Instances** page, locate the instance that you want to create a backup for and choose **More** > **Create Backup** in the **Operation** column.

Method 2

- On the **Instances** page, click the instance that you want to create a backup for and click its name.
- Choose Backups & Restorations in the navigation pane on the left, click Create Backup.

Method 3

In the navigation pane on the left, choose **Backups** and click **Create Backup**.

Step 4 In the displayed dialog box, enter a backup name and description and click **OK**.




Figure 4-54 Creating a backup

Table 4-18 Parameter description

Parameter	Description	
DB Instance Name	Must be the name of the DB instance to be backed up and cannot be modified.	
Backup Name	Must be 4 to 64 characters in length and start with a letter. It is case-insensitive and contains only letters, digits, hyphens (-), and underscores (_).	
Description	Can contain a maximum of 256 characters and cannot contain line breaks or special characters >!<"&'=	

Step 5 View the backup status.

When the backup is being created, query the backup status on the **Backups** or Backups & Restorations page. The backup status is Backing up.

After the backup is created, the backup status is Completed.

You can create manual table-level backups by following **Creating and Managing Table-level Backups**.

----End

Deleting a Manual Backup

If you do not need a manual backup any longer, delete it on the **Backups** or **Backups & Restorations** page.

Deleted backups are not displayed in the backup list.

NOTICE

Deleted backups cannot be recovered. Exercise caution when performing this operation.

Method 1

- 1. Log in to the Huawei Cloud console.
- 2. In the service list, choose **Databases** > **GeminiDB**.
- 3. On the **Instances** page, locate the instance whose backup you want to delete and click its name.
- 4. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup you want to delete and click **Delete** in the **Operation** column.
- 5. In the displayed dialog box, confirm the backup details and click **Yes**.

Method 2

- 1. Log in to the Huawei Cloud console.
- 2. In the service list, choose **Databases** > **GeminiDB**.
- 3. On the **Backups** page, locate the backup you want to delete and click **Delete** in the **Operation** column.
- 4. In the displayed dialog box, confirm the backup details and click **Yes**.

4.7.4 Managing Cross-Region Backups

GeminiDB HBase instance backups can be stored in the destination region or OBS buckets, so you can use the backups to restore data to a new instance.

After a cross-region backup policy is set for an instance, the system will synchronize backups of the instance to the destination region you specified. You can manage cross-region backup files on the **Backups** page.

Usage Notes

 To enable the cross-region backup function, choose Service Tickets > Create Service Ticket in the upper right corner of the console and contact the customer service.

- Before you configure a cross-region backup policy, make sure to enable automated backup first. Otherwise, the cross-region backup policy cannot take effect. For details, see **Modifying an Automated Backup Policy**.
- Only automated full backups can be created across regions.

Billing

Table 4-19 Billing

Flavor	Billing Item	Unit Price
geminidb.cassandra.cross reg.backup.space.dfv	Storage space	CNY0.0009/GB/hour
geminidb.cassandra.cross reg.backup.flow	Cross-region backup traffic	CNY0.5/GB

Setting or Modifying a Cross-Region Backup Policy

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance that you want to connect to and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Set Cross-Region Backup Policy**.
- **Step 6** In the displayed dialog box, set required parameters.

Figure 4-55 Setting a cross-region backup policy

All cross-region backups of your DB instances are stored in the region you specify. Only automated full backups will be replicated to the target region. Cross-Region Full Backup Region Retention Period I + Enter an integer from 1 to 3660.

Table 4-20 Description

Parameter	Description
Cross-Region Full Backup	If you enable Cross-Region Full Backup , automated full backup files of the instance will be stored in the region you specify.
Region	You can select the region for storing backups based on service requirements.
Retention Period	Number of days that cross-region backups are kept. The value ranges from 1 to 1825 . You can increase the retention period to improve data reliability.

■ NOTE

- Only new backups generated after you set a cross-region backup policy will be stored in the region you specify.
- All cross-region backups of your DB instances are stored in the same region you specify.
- Cross-region backups are synchronized to the destination region you specify only after your instance is backed up locally.
- Only automated full backups are replicated to the destination region.


Step 7 Click OK.

----End

Managing Cross-Region Backups

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** In the navigation pane, choose **Backups**. Click the **Cross-Region Backups** tab.

Figure 4-56 Cross-region backups



All cross-region backups are displayed by default.

• To modify the cross-region backup policy, click **Set Cross-Region Backup** in the **Operation** column.

To view all cross-region backups, click View Cross-Region Backup. To restore
a backup to a new instance, click Restore in the Operation column. For
details, see Restoring a Backup to a New Instance.

Figure 4-57 Restoring a cross-region backup



□ NOTE

- Cross-region full backup replicates only automated full backups to another region for full restoration. Any of the full backups can be restored to a new instance that has no relationships with the original instance.
- The new instance uses the same parameter group as the original instance.
- During the instance restoration, backup files are downloaded from OBS buckets to the data directory of the new instance. The download bandwidth of OBS is 40 MB/s.
- If the original instance has cross-region backup disabled, the restoration may fail.
- Cross-region backup is not supported for instances that has disk encryption enabled.

----End

Disabling Cross-Region Backup

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance that you want to connect to and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Set Cross-Region Backup Policy**.
- **Step 6** In the displayed dialog box, disable **Cross-Region Full Backup**.

Set Cross-Region Backup Policy

1 All cross-region backups of your DB instances are stored in the region you specify.
Only automated full backups will be replicated to the target region.

Cross-Region Full Backup

If the cross-region backup policy is disabled, the cross-region backup task will be stopped immediately, and all cross-region backups will be immediately deleted. Operations related to cross-region backup may fail.

Region

Retention Period

OK

Cancel

Figure 4-58 Disabling cross-region backup

□ NOTE

- After cross-region backup is disabled, the cross-region backup task is stopped and all cross-region backups are deleted immediately. As a result, operations using cross-region backups will fail.
- If an instance with cross-region backup enabled is deleted, its cross-region backups will be retained. The retention period depends on settings of the cross-region backup policy.

Step 7 Click OK.

----End

4.7.5 Managing Table-level Backups

You can manually create table-level backups. If a database or table is deleted maliciously or accidentally, you can use backups to restore data.

Usage Notes

- To enable the table-level backup, choose Service Tickets > Create Service
 Ticket in the upper right corner of the console and contact the customer
 service.
- Table-level backups can be created automatically or manually.

Enabling or Modifying a Table-level Backup Policy

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance name to access the **Basic Information** page.

- **Step 4** In the navigation pane, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Configure Table-level Backup Policy**.
- **Step 6** On the displayed page, configure related parameters.

Figure 4-59 Setting a table-level backup policy



Table 4-21 Parameters

Parameter	Description
Automated Backup	After this function is enabled, database table data is backed up based on the backup policy.
Retention Period	Automated backup files are stored for 7 days by default. The retention period ranges from 1 to 732 days.
	 If the retention period is shorter than seven days, the system automatically backs up data daily.
	The system automatically checks existing backup files and deletes files that exceed the retention period you set.
Time Window	A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00 or 12:00-13:00. The backup time is in GMT format. If the DST or standard time is switched, the time window changes with the time zone.
Backup Cycle	All: Each day of the week is selected. This option is selected by default. The system automatically backs up data every day.
	 You can select one or more days in a week. The system automatically backs up data on the specified days.
Databases and Tables	Select tables to be backed up.

A full backup starts within one hour of the time you specify. How long the backup takes depends on the data volume.

Step 7 Click OK.

----End

Creating and Managing Table-level Backups

Creating a Manual Table-level Backup

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance that you want to create a table-level backup for and click its name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Create Table-level Backup**.
- **Step 6** On the displayed page, configure related parameters.

Figure 4-60 Creating a table-level backup



Table 4-22 Parameters

Parameter	Description
Backup Name	Can contain 4 to 64 characters and must start with a letter. The name is case-sensitive and can contain only letters, digits, hyphens (-), or underscores (_).
Description	The description can contain a maximum of 256 characters and cannot contain line breaks or special characters >! <"&'=
Databases and Tables	You can select the databases and specify tables therein that you want to back up.

- Step 7 Click OK.
- **Step 8** Choose **Backups and Restorations** > **Table-level Backup** and manage the created backup.

Alternatively, click **Backups** in the navigation pane on the left, choose **Intraregion Backups** > **Table-level Backups**, and manage the created backup.

Figure 4-61 Managing the created table-level backup



- Click View Tables to view tables contained in the backup file.
- Click **Restore** in the **Operation** column to restore the backup to a new instance. For details, see **Restoring a Backup to a New Instance**.
- Click **Delete** in the **Operation** column to delete the created backup.

NOTICE


Deleted backups cannot be recovered.

----End

Disabling a Table-level Backup Policy

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance name to access the **Basic Information** page.
- **Step 4** In the navigation pane, choose **Backups & Restorations**.
- **Step 5** On the displayed page, click **Configure Table-level Backup Policy**.
- **Step 6** On the displayed page, configure related parameters.

Figure 4-62 Disabling a table-level backup policy



□ NOTE

After the table-level backup policy is disabled, any table-level backup task in progress stops immediately, and all table-level backups of the instance are retained. The retention duration depends on **Retention Period** specified when you enabled the table-level backup policy.

Step 7 Click OK.

----End

4.8 Data Restoration

4.8.1 Restoration Methods

You can select a proper method to restore GeminiDB HBase instance data.

Table 4-23 Restoration Methods

Method	Scenario
Rebuilding an Instance	If an instance is deleted by mistake, you can rebuild it within a retention period in the recycle bin.
Restoring a Backup to a New Instance	You can restore an existing backup file to a new instance.
Restoring a Backup to a Specified Point in Time	You can use an automated backup to restore an instance to a specified point in time.

4.8.2 Restoring a Backup to a New Instance

You can restore an existing backup to a new GeminiDB HBase instance.

Usage Notes

- The new instances must have at least as many nodes as the original instance.
- The new instance must have at least as much storage as the original instance.
- Incremental backup and PITR are not supported.
- Restoration to the current instance is not supported.
- You can scale in the memory, but the memory decrease cannot become less than the actual memory used during the backup.
- The restored instance uses the same parameter group as the original instance.
- During the instance restoration, backups are downloaded from OBS buckets to the data directory of the restored instance. The download bandwidth of OBS is 40 MB/s.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** Restore an instance from the backup.

Method 1

- 1. On the **Instances** page, locate the instance whose backup you want to restore and click its name.
- 2. Choose **Backups & Restorations** in the navigation pane on the left, locate the backup that you want to restore and click **Restore** in the **Operation** column.

Figure 4-63 Backups and restorations



Method 2


On the **Backups** page, locate the backup that you want to restore and click **Restore** in the **Operation** column.

Figure 4-64 Backup management



Step 4 In the displayed dialog box, confirm the current instance details and restoration method and click **OK**.

Figure 4-65 Restoring data to a new instance



- The default API type and DB engine version are the same as those of the original instance and cannot be changed.
- The system automatically calculates the minimum storage required for restoring data to the new instance based on the size of the selected backup file. You need to select an integer for the storage.
- The administrator password needs to be reset.
- To modify other parameters, see the description of buying instances of other DB engines in *Getting Started*.

Step 5 View the results.

A new instance is created using the backup data. The status of the new instance changes from **Creating** to **Available**.

After the restoration, the system will perform a full backup.

The new instance is independent from the original one.

----End

4.8.3 Restoring a Backup to a Specified Point in Time

Existing automated backups can be restored to a specified point in time on a GeminiDB HBase instance.

The most recent full backup will be downloaded from OBS for restoration. After the restoration is complete, incremental backups will be replayed to the specified point in time. The time required depends on the amount of data to be restored.

Usage Notes

- Data can only be restored to a new instance.
- After automated backup is enabled, the system performs an incremental backup based on the preset incremental backup interval. The incremental backup is stored in OBS.
- Keep your account balance above zero so that backup data can be restored to a new instance.
- Data can be restored to a specified time point only after the automated backup policy is enabled.
- During the instance restoration, backup files are downloaded from OBS buckets to the data directory of the restored instance. The download bandwidth of OBS is 40 MB/s.

Procedure


- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the instance name.
- **Step 4** In the navigation pane on the left, choose **Backups & Restorations**.
- **Step 5** On the **Backups & Restorations** page, click **Restore to Point in Time**.

Figure 4-66 Restoring data to a point in time



Step 6 Select a restoration date and a time point and click **OK**.

Figure 4-67 Restore to Point in Time



- **Step 7** On the **Create New Instance** page, create an instance of the same specifications as the instance to be restored. The new instance is independent from the original one.
 - The new instance should be deployed in a different AZ to ensure that your applications will not be affected by SPOFs.
 - The API compatibility, instance type, version, and CPU type are the same as those of the original instance and cannot be changed.
 - Other settings are the same as those of the original instance by default but can be modified. For details, see Getting to Know GeminiDB HBase API.

----End

4.9 Logs and Audit

This function is unavailable to external systems.

4.10 Viewing Metrics and Configuring Alarms

4.10.1 Supported Metrics

This section describes GeminiDB HBase API metrics reported to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics and alarms.

□ NOTE

Because GeminiDB HBase API is based on GeminiDB Cassandra API, they share some metrics.

Namespace

SYS.NoSQL

Metrics

₩ NOTE

You can view metrics on instance nodes by following Viewing Metrics.

Table 4-24 GeminiDB HBase API metrics

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
nosql005 _disk_usa ge	Storage Space Usage	Storage usage of the current instance.	0- 100	%	N/ A	GeminiDB HBase instance	1 minut e
nosql006 _disk_tot al_size	Total Storage Space	Total storage space of the current instance.	≥ 0	GB	102 4(I EC)	GeminiDB HBase instance	1 minut e
nosql007 _disk_use d_size	Storage Space Usage	Storage space usage of the current instance.	≥ 0	GB	102 4(I EC)	GeminiDB HBase instance	1 minut e
nosql009 _dfv_writ e_delay	Storage Write Latency	Average delay of writing data to the storage layer in a specified period	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
nosql010 _dfv_rea d_delay	Storage Read Latency	Average latency of reading data from the storage layer in a specified period	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a001_cpu _usage	CPU Usage	CPU usage of an instance	0- 100	%	N/ A	GeminiDB HBase instance node	1 minut e

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
cassandr a002_me m_usage	Memory Usage	Memory usage of the instance	0- 100	%	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a003_byt es_out	Network Output Through put	Outgoing traffic in bytes per second	≥ 0 Bytes /s	By tes /s	102 4(I EC)	GeminiDB HBase instance node	1 minut e
cassandr a004_byt es_in	Network Input Through put	Incoming traffic in bytes per second	≥ 0	By tes /s	102 4(I EC)	GeminiDB HBase instance node	1 minut e
cassandr a014_con nections	Active Node Connecti ons	Active connections to the current GeminiDB Cassandra instance node	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a015_rea d_latency	Average Read Latency	Average amount of time consumed by read requests	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a016_wri te_latenc y	Average Write Latency	Average amount of time consumed by write requests	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a037_pe nding_wr ite	Suspende d Write Tasks	Number of write tasks waiting in the queue	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a038_pe nding_re ad	Suspende d Read Tasks	Number of read tasks waiting in the queue.	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a044_ran ge_slice_l atency	Scan Duration	Average time consumed by scan operations	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
cassandr a049_dro pped_mu tation	Dropped Writes	Average number of dropped writes	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a052_dro pped_rea d	Dropped Reads	Average number of dropped reads	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a092_loa d_info	Data Volume on a Node	Data volume on a node	≥ 0 Byte	By te	102 4(I EC)	GeminiDB HBase instance node	1 minut e
cassandr a093_wri te_count _latency	Accumul ated Write Requests	Number of write requests initiated by a node	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a094_wri te_1min_ rate	Average Write Rate in the Last Minute	Average write rate in the last minute	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a095_wri te_p75_la tency	p75 Write Latency	p75 write latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a096_wri te_p95_la tency	p95 Write Latency	p95 write latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a097_wri te_p99_la tency	p99 Write Latency	p99 write latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a098_rea d_count_l atency	Accumul ated Read Requests	Number of read requests initiated by a node	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
cassandr a099_rea d_1min_r ate	Average Read Rate in the Last Minute	Average read rate in the last minute	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a100_rea d_p75_la tency	p75 Read Latency	p75 read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a101_rea d_p95_la tency	p95 Read Latency	p95 read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a102_rea d_p99_la tency	p99 Read Latency	p99 read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a103_ran ge_slice_ count_lat ency	Accumul ated Range Read Requests	Number of range read requests	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a104_ran ge_slice_ 1min_rat e	Average Range Read Rate in the Last Minute	Average range read rate in the last minute	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a105_ran ge_slice_ p75_late ncy	p75 Range Read Latency	p75 range read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a106_ran ge_slice_ p95_late ncy	p95 Range Read Latency	p95 range read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
cassandr a107_ran ge_slice_ p99_late ncy	p99 Range Read Latency	p99 range read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a163_wri te_p999_l atency	p999 Write Latency	p999 write latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a164_rea d_p999_l atency	p999 Read Latency	p999 read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a165_lar ge_partiti on_num	Big Keys	Number of big keys on the current node	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a166_wri te_max_l atency	Maximu m Write Latency	Maximum write latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a167_rea d_max_la tency	Maximu m Read Latency	Maximum read latency	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a168_im balance_t able_nu m	Tables with Uneven Data Distributi on	Number of tables whose data is not evenly distributed	≥ 0	Co un ts	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a169_mo dify_requ est_size_ mean	Average Write Request Size	Average write request size	≥ 0	By tes	102 4(I EC)	GeminiDB HBase instance node	1 minut e

Metric ID	Name	Description	Valu e Rang e	Un it	Nu mb er Sys te m	Monitore d Object	Moni torin g Perio d (Raw Data)
cassandr a170_qu ery_respo nse_size_ mean	Average Query Response Size	Average size of query requests	≥ 0	By tes	102 4(I EC)	GeminiDB HBase instance node	1 minut e
cassandr a173_lim it_diff_co unt_mea n	Mean of limit Value and Returned Rows	Mean of limit difference	≥ 0	-	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a174_to mbstone _query_r ate	Tombsto ne Query Requests per Second	Rate of tombstone query requests	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a175_sin gle_delet e_rate	Row Delete Requests per Second	Rate at which a single row is deleted	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a176_ran ge_delete _rate	Range Delete Requests per Second	Range deletion rate	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a177_lar ge_row_c ount	Large Rows per Second	Number of large rows	≥ 0	Co un ts/ s	N/ A	GeminiDB HBase instance node	1 minut e
cassandr a174_to mbstone _query_r ate	Maximu m sync delay	Maximum synchronization latency between the primary and standby instances	≥ 0	ms	N/ A	GeminiDB HBase instance node	1 minut e

Dimensions

Key	Value
cassandra_cluster_id	Cluster ID of a GeminiDB Cassandra instance
cassandra_node_id	ID of a GeminiDB Cassandra instance node

4.10.2 Configuring Alarm Rules

Setting alarm rules allows you to customize objects to be monitored and notification policies so that you can closely monitor your instances.

Alarm rules include the alarm rule name, instance, metric, threshold, monitoring interval, and whether to send notifications. This section describes how to set alarm rules.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Management & Governance** > **Cloud Eye**.
- **Step 3** In the navigation pane on the left, choose **Alarm Management > Alarm Rules**.
- **Step 4** On the **Alarm Rules** page, click **Create Alarm Rule**.

Figure 4-68 Creating an alarm rule



Step 5 Set alarm parameters.

1. Configure basic alarm information.

Figure 4-69 Configuring basic information for an alarm rule



Table 4-25 Basic alarm rule information

Parameter	Description	Example Value
Name	Name of the rule. The system generates a random name and you can modify it.	alarm-cag2
Description	(Optional) Alarm rule description.	-

2. Select objects to be monitored and specify the monitoring scope.

Figure 4-70 Configuring objects to be monitored



Table 4-26 Parameter description

Parameter	Description	Example Value
Alarm Type	Alarm type that the alarm rule is created for. The value can be Metric or Event .	Metric
Resource Type	Type of the resource the alarm rule is created for. Select GeminiDB .	-
Dimension	Metric dimension of the alarm rule. Select Cassandra - Cassandra Nodes. You can select either Cassandra or HBase.	-

Parameter	Description	Example Value
Monitoring Scope	Monitoring scope the alarm rule applies to. NOTE	All resources
	 If you select All resources, an alarm notification will be sent when any instance meets an alarm policy, and existing alarm rules will be automatically applied for newly purchased resources. 	
	 If you select Resource groups and any resource in the group meets the alarm policy, an alarm notification will be sent. 	
	 If you select Specific resources, click Select Specified Resources, select one or more resources, and click OK. 	
Group	This parameter is mandatory when Monitoring Scope is set to Resource groups.	-

3. Configure an alarm policy.

Figure 4-71 Configuring an alarm policy



Table 4-27 Parameter description

Parameter	Description	Example Value
Method	Select Associate template, Use existing template, or Configure manually. NOTE If you set Monitoring Scope to Specific resources, you can set Method to Use existing template.	Configure manually
Template	Select the template to be used. This parameter is available only when you set Method to Use existing template .	-

Parameter	Description	Example Value
Alarm Policy	Policy for triggering an alarm. You can configure the threshold, consecutive periods, alarm interval, and alarm severity based on service requirements.	Take the CPU usage as an example. The alarm policy
	 Metric Name: metric that an alarm rule is created for. The following metrics are recommended: 	configured in Figure 4-71 indicates that
	Storage Space Usage:	a major alarm
	Storage usage of a GeminiDB HBase instance. If the storage usage is greater than 80%, scale up the storage in a timely manner by following Manually Scaling Up Storage Space.	notification will be sent to users every 10 minutes if the original
	CPU Usage and Memory Usage:	CPU usage
	Compute resource usage of each GeminiDB HBase instance node. If the CPU usage or memory usage is greater than 80%, you are advised to add nodes or increase node specifications.	reaches 80% or above for three consecutive periods.
	For more metrics, see Supported Metrics .	
	 Alarm Severity: specifies the severity of the alarm. Valid values are Critical, Major, Minor, and Informational. 	
	NOTE A maximum of 50 alarm policies can be added to an alarm rule. If any one of these alarm policies is met, an alarm is triggered.	

4. Configure alarm notification information.

Figure 4-72 Configuring alarm notification information



Table 4-28 Parameter description

Parameter	Description	Example Value
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message. Enabling alarm notification is recommended. When the metric data reaches the threshold set in the alarm rule, Cloud Eye immediately notifies you through SMN that an exception has occurred.	Enabled Alarm Notification.
Notification Recipient	Select Notification group or Topic subscription .	-
Notification Group	Notification group the alarm notification is to be sent to.	-
Notification Object	Specifies the object that receives alarm notifications. You can select the account contact or a topic. - Account contact is the mobile phone number and email address provided for registration.	-
	 Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see Creating a Topic and Adding Subscriptions. 	
Notification Window	Cloud Eye sends notifications only within the notification window specified in the alarm rule. For example, if Notification Window is set to 00:00-8:00, Cloud Eye sends notifications only within 00:00-08:00.	-

Parameter	Description	Example Value
Trigger Condition	Condition for triggering an alarm notification. You can select Generated alarm (when an alarm is generated), Cleared alarm (when an alarm is cleared), or both.	-

5. Configure advanced settings.

Figure 4-73 Advanced settings




Table 4-29 Parameter description

Parameter	Description	Example Value
Enterprise Project	Enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule. For details about how to create an enterprise project, see Creating an Enterprise Project.	default
Tag	A tag is a key-value pair. Tags identify cloud resources so that you can easily categorize and search for your resources. You are advised to create predefined tags on TMS. For details about how to create predefined tags, see Creating Predefined Tags.	-
	 A key can contain a maximum of 128 characters, and a value can contain a maximum of 255 characters. A maximum of 20 tags can be added. 	

Step 6 After the configuration is complete, click **Create**.

When the metric data reaches the threshold set in the alarm rule, Cloud Eye immediately notifies you through SMN that an exception has occurred.

□ NOTE

For more information about alarm rules, see Cloud Eye User Guide.

----End

4.10.3 Viewing Metrics

Cloud Eye monitors the GeminiDB HBase instance status. You can view metrics on the console.

Monitored data requires a period of time for transmission and display. The status of the monitored object displayed on the Cloud Eye page is the status obtained 5 to 10 minutes before. You can view the monitored data of a newly created DB instance 5 to 10 minutes later.

Usage Notes

- The DB instance is running properly.
 Cloud Eye does not display the metrics of a faulty or deleted DB instance. You can view the monitoring information only after the instance is restarted or recovered.
- The DB instance has been properly running for at least 10 minutes.

 The monitoring data and graphics are available for a new DB instance after the instance runs for at least 10 minutes.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instance** page, click the instance whose metrics you want to view and click its name.
- **Step 4** In the **Node Information** area on the **Basic Information** page, click **View Metric** in the **Operation** column.

Figure 4-74 Viewing metrics



Step 5 In the monitoring area, select a time range to view monitoring data.

You can view the monitoring data in the last 1, 3, or 12 hours.

To view the graph in a longer time range, click 5.

----End

4.10.4 Event Monitoring

4.10.4.1 Introduction

Events and alarms will be reported on the **Event Monitoring** page. You can create alarm rules for both system and custom events. When a specific event occurs, Cloud Eye generates and sends an alarm for you.

Key operations on GeminiDB HBase instances are monitored and recorded by Cloud Eye as events. Events include operations performed by specific users on specific resources, such as changing instance names and specifications.

Event monitoring provides an API for reporting custom events, which helps you collect and report abnormal events or important change events generated by services to Cloud Eye.

Event monitoring is enabled by default. You can view monitoring details of system and custom events. For details about system events, see **Events Supported by Event Monitoring**.

□ NOTE

If you do not create an alarm rule, no alarm will be sent by default.

4.10.4.2 Viewing Event Monitoring Data

Scenarios

Events and alarms will be reported on the **Event Monitoring** page. You can create alarm rules for both system and custom events. When a specific event occurs, Cloud Eye generates and sends an alarm for you.

Event monitoring is enabled by default. You can view monitoring details about system events and custom events.

This topic describes how to view the event monitoring data.

Procedure

- Step 1 Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, locate the instance whose event monitoring data you want to view. In the **Node Information** area on the **Basic Information** page, click **View Metric** in the **Operation** column.
- **Step 4** Click in the upper part to return to the Cloud Eye console.
- **Step 5** In the navigation pane on the left, choose **Event Monitoring**.

On the displayed **Event Monitoring** page, all system events generated in the last 24 hours are displayed by default.

You can also click **1h**, **3h**, **12h**, **1d**, **7d**, or **30d** to view events generated in different time periods.

Step 6 Locate an event and click **View Event** in the **Operation** column to view its details.

----End

4.10.4.3 Creating an Alarm Rule for Event Monitoring

Scenarios

This topic describes how to create an alarm rule for event monitoring.

Usage Notes

If you do not create an alarm rule, no alarm will be sent by default.

Procedure

- **Step 1** Log in to the Huawei Cloud console.
- Step 2 Click in the upper left corner of the page. Choose Management & Governance > Cloud Eye.
- **Step 3** In the navigation pane on the left, choose **Event Monitoring**.
- **Step 4** On the event list page, click **Create Alarm Rule** in the upper right corner.
- **Step 5** On the **Create Alarm Rule** page, configure the parameters.

Table 4-30 Parameter description

Parameter	Description
Name	Specifies the name of the alarm rule. The system generates a random name, but you can change it if needed.
Description	(Optional) Provides supplementary information about the alarm rule.
Enterprise Project	You can select an existing enterprise project or click Create Enterprise Project to create one.
Alarm Type	Specifies the alarm type corresponding to the alarm rule.
Event Type	Specifies the event type of the metric corresponding to the alarm rule.
Event Source	Specifies the service the event is generated for. Select GeminiDB.

Parameter	Description
Monitoring Scope	Specifies the monitoring scope for event monitoring.
Method	Specifies the event creation method.
Alarm Policy	Event Name indicates the instantaneous operations users performed on system resources, such as login and logout.
	For details about events supported by Event Monitoring, see Events Supported by Event Monitoring.
	You can select a trigger mode and alarm severity as needed.

Click to enable **Alarm Notification**. The validity period is 24 hours by default. If required topics are not displayed in the drop-down list, click **Create an SMN topic**.

Table 4-31 Alarm notification parameters

Parameter	Description
Alarm Notification	Whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/ HTTPS message.
Notification Object	Object an alarm notification is to be sent to. You can select the account contact or a topic.
	Account contact is the mobile phone number and email address provided for registration.
	Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see Creating a Topic and Adding Subscriptions.
Validity Period	Notification window which Cloud Eye only sends notifications within.
	If you set Validity Period to 08:00-20:00 , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Condition for triggering the alarm notification.

Step 6 After the configuration is complete, click **Create**.

----End

4.10.4.4 Events Supported by Event Monitoring

Table 4-32 Events supported by GeminiDB Event Monitoring

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
NoSQ L	Instance creation failure	NoSQL Createl nstance Failed	Maj or	The instance quota or underlying resources are insufficient.	Release unnecessary instances and try again. You can also choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to adjust the quota.	Instan ces fail to be create d.
	Specificati ons change failure	NoSQL Resizel nstance Failed	Maj or	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to the O&M personnel to coordinate resources in the background and change the specifications again.	Servic es are interr upted.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Node adding failure	NoSQL AddNo desFail ed	Maj or	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background, delete nodes that failed to be added, and add the nodes again.	None
	Node deletion failure	NoSQL Delete NodesF ailed	Maj or	Releasing underlying resources failed.	Delete the node again.	None
	Storage space scale-up failure	NoSQL ScaleU pStorag eFailed	Maj or	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background and scale up storage again.	Servic es may be interr upted.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Password resetting failure	NoSQL ResetPa ssword Failed	Maj or	Resetting the password times out.	Reset the password again.	None
	Parameter template change failure	NoSQL Updatel nstance Param GroupF ailed	Maj or	Changing a parameter template times out.	Change the parameter template again.	None
	Backup policy configurat ion failure	NoSQL SetBack upPolic yFailed	Maj or	The database connection is abnormal.	Configure the backup policy again.	None
	Manual backup creation failure	NoSQL Create Manual Backup Failed	Maj or	The backup files fail to be exported or uploaded.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	Data canno t be backe d up.
	Automate d backup creation failure	NoSQL CreateA utomat edBack upFaile d	Maj or	The backup files fail to be exported or uploaded.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	Data canno t be backe d up.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Instance status abnormal	NoSQL FaultyD BInstan ce	Maj or	This event is a key alarm event and is reported when an instance is faulty due to a disaster or a server failure.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The datab ase servic e may be unava ilable.
	Instance status recovery	NoSQL DBInsta nceRec overed	Maj or	If a disaster occurs, NoSQL provides an HA tool to automatically or manually rectify the fault. After the fault is rectified, this event is reported.	No further action is required.	None
	Node status abnormal	NoSQL FaultyD BNode	Maj or	This event is a key alarm event and is reported when a database node is faulty due to a disaster or a server failure.	Check whether the database service is functional. Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The datab ase servic e may be unava ilable.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Node status recovery	NoSQL DBNod eRecov ered	Maj or	If a disaster occurs, NoSQL provides an HA tool to automatically or manually rectify the fault. After the fault is rectified, this event is reported.	No further action is required.	None
	Primary/ standby switchove r or failover	NoSQL Primary Standb ySwitch ed	Maj or	This event is reported when a primary/ secondary switchover or a failover is triggered.	No further action is required.	None
	Occurrenc e of hotspot partitionin g keys	HotKey Occurs	Maj or	Hotspot data is stored in one partition because the primary key is improper. Improper application design causes frequent read and write operations on a key.	1. Choose a proper partition key. 2. Add service cache so that service applications read hotspot data from the cache first.	The servic e reque st succes s rate is affect ed, and the cluste r perfor manc e and stabili ty deteri orates .

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	BigKey occurrenc e	BigKey Occurs	Maj or	The primary key design is improper. There are too many records or too much data in a single partition, causing load imbalance on nodes.	1. Choose a proper partition key. 2. Add a new partition key for hashing data.	As more and more data is stored in the partiti on, cluste r stabili ty deteri orates .
	Insufficien t storage space	NoSQL RiskyDa taDiskU sage	Maj or	The storage space is insufficient.	Scale up storage space. For details, see section "Scaling Up Storage Space" in the user guide of GeminiDB.	The instan ce is set to readonly and data canno t be writte n to the instan ce.
	Data disk expanded and being writable	NoSQL DataDi skUsag eRecov ered	Maj or	The data disk has been expanded and becomes writable.	No further action is required.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Index creation failure	NoSQL Createl ndexFai led	Maj or	The service load exceeds what the instance specifications can take. In this case, creating indexes consumes more instance resources. As a result, the response is slow or even frame freezing occurs, and the creation times out.	1. Select matched instance specifications based on the service loads. Create indexes during offpeak hours. Create indexes in the background. Select indexes as required.	The index fails to be create d or is incom plete. Delet e the index and create a new one.
	Write speed decrease	NoSQL Stalling Occurs	Maj or	The write speed is close to the maximum write speed allowed by the cluster scale and instance specifications. As a result, the database flow control mechanism is triggered, and requests may fail.	1. Adjust the cluster scale or node specifications based on the maximum write rate of services. 2. Measure the maximum write request rate of services and distribute the peak write rate of services.	The succes s rate of servic e reque sts is affect ed.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Data write stopped	NoSQL Stoppin gOccur s	Maj or	The data write is too fast, reaching the maximum write capability allowed by the cluster scale and instance specifications. As a result, the database flow control mechanism is triggered, and requests may fail.	1. Adjust the cluster scale or node specifications based on the maximum write rate of services. 2. Measure the maximum write request rate of services and distribute the peak write rate of services.	The succes s rate of servic e reque sts is affect ed.
	Database restart failure	NoSQL Restart DBFaile d	Maj or	The instance status is abnormal.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The instan ce status may be abnor mal.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Restoratio n to new instance failure	NoSQL Restore ToNewl nstance Failed	Maj or	The underlying resources are insufficient.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console. Submit a service ticket to O&M personnel to coordinate resources in the background and add nodes again.	Data canno t be restor ed to a new instan ce.
	Restoratio n to existing instance failure	NoSQL Restore ToExistI nstance Failed	Maj or	The backup file fails to be downloaded or restored.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	The curren t instan ce may be unava ilable.
	Backup file deletion failure	NoSQL DeleteB ackupF ailed	Maj or	The backup files fail to be deleted from OBS.	Delete the backup files again.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Failure to display slow query logs in plaintext	NoSQL SwitchS lowlog PlainTe xtFailed	Maj or	The API does not support this function.	Refer to GeminiDB User Guide to ensure that the API supports slow query logs in plaintext. Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	None
	EIP binding failure	NoSQL BindEip Failed	Maj or	The node status is abnormal, an EIP has been bound to the node, or the EIP to be bound is invalid.	Check whether the node is normal and whether the EIP is valid.	The instan ce canno t be access ed from a public netwo rk.
	EIP unbinding failure	NoSQL Unbind EipFaile d	Maj or	The node status is abnormal or the EIP has been unbound from the node.	Check whether the node and EIP status are normal.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Parameter modificati on failure	NoSQL Modify Parame terFaile d	Maj or	The parameter value is invalid.	Check whether the parameter value is valid. Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	None
	Parameter template applicatio n failure	NoSQL ApplyP aramet erGrou pFailed	Maj or	The instance status is abnormal. So, the parameter template cannot be applied.	Choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket to O&M personnel.	None

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Enabling or disabling SSL failure	NoSQL SwitchS SLFaile d	Maj or	Enabling or disabling SSL times out.	Try again or choose Service Tickets > Create Service Ticket in the upper right corner of the console and submit a service ticket. Retain the SSL connection mode configured before the event occurred.	The SSL conne ction mode canno t be chang ed.
	Too much data in a single row	LargeR owOcc urs	Maj or	If there is too much data in a single row, queries may time out, causing faults like OOM error.	1. Limit the write length of each column and row so that the key and value length of each row does not exceed the preset threshold. 2. Check whether there are abnormal writes or coding, causing large rows.	If there are too many record s in a single row, cluste r stabili ty will deteri orate as the data volum e increa ses.

Event Sourc e	Event Name	Event ID	Eve nt Sev erit y	Description	Solution	Impa ct
	Schedule for deleting a KMS key	planDel eteKms Key	Maj or	The user plans to delete a KMS key.	Check whether the GeminiDB instance associated with the key has been deleted or is no longer used. Deleting the key will affect the instance services.	The key will be auto matic ally delete d after it expire s. Deleti ng the key will affect the instan ce servic es.
	Too many tombston es	TooMa nyQuer yTombs tones	Maj or	Querying too many tombstones may time out.	Use a proper query and deletion method to avoid batch range queries.	The query may time out.
	Ultra- large collection column	TooLar geColle ctionCo lumn	Maj or	If there are too many elements in the collection column, the query will fail.	Set a threshold for the number of elements in the collection column. Check whether there is an error while data is written and encoded.	The query on the collec tion colum n will fail.

4.11 Enterprise Project

4.11.1 Overview

An enterprise project facilitates project-level management and grouping of cloud resources and users. The default project is **default**.

You can also customize enterprise projects to meet your service requirements. For details, see **Enterprise Management User Guide**.

4.11.2 Quota Management

GeminiDB HBase API allows you to manage resources by controlling the number of resources in each enterprise project to ensure that resources can be used and managed properly.

This section describes how to query used resources in each enterprise project and its resource quotas.

This function is now in OBT. To use it, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Viewing Resource Quotas in Each Enterprise Project

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** In the navigation pane on the left, choose **My Quotas** to view quota details of the current enterprise project.

Figure 4-75 Quota management




Table 4-33 Parameter description

Parameter	Description
Enterprise Project	Enterprise project that an instance belongs to.

Parameter	Description
Used/Total DB Instances	Number of used instances in the current enterprise project
Used/Total vCPUs	vCPUs of all instances in the current enterprise project
Used/Total Memory (GB)	Memory of all instances in the current enterprise project

NOTE

If there are no resources in an enterprise project, the default quota is 0. Before creating an instance, you need to set quotas first by following **Modifying Resource Quotas of an Enterprise Project**.

----End

Modifying Resource Quotas of an Enterprise Project

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** In the navigation pane on the left, choose **My Quotas**. In the quota list, select the enterprise project you want to set quotas for and click **Modify** in the **Operation** column.

Figure 4-76 Managing quotas




Table 4-34 Quota management

Parameter	Value Range
DB Instances	0–5,000
vCPUs	0-8,000,000
Memory (GB)	0–16,000,000

----End

4.12 Tag Management

Tag Management Service (TMS) enables you to manage resources using tags on the management console. TMS works with other cloud services to manage tags. TMS manages tags globally while other cloud services manage their own tags.

You can tag your GeminiDB HBase instance to easily identify and manage its resources. An instance can be tagged when or after it is created.

After an instance is tagged, you can search for the tag key or value to quickly query the instance details.

Usage Notes

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
 For details about the naming rules of tag keys and values, see Table 4-35.
- A maximum of 20 tags can be added for each instance.
- The tag name must comply with the naming rules described in Table 4-35.

Table 4-35 Naming rules

Parameter	Requirement	Example Value
Tag key	Cannot be left blank.Must be unique for each instance.	Organization
	Can contain a maximum of 128 characters.	
	 Cannot start with _sys_ and cannot start or end with a space. Only letters, digits, spaces, and the following special characters are allowed:@.:/+= 	

Parameter	Requirement	Example Value
Tag value	 Can be left blank. Can contain a maximum of 255 characters. 	nosql_01
	Only letters, digits, spaces, and the following special characters are allowed:@.:/+=	

Adding a Tag

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the navigation pane on the left, click **Tags**.
- **Step 5** On the **Tags** page, click **Add Tag**. In the displayed dialog box, enter a tag key and value, and click **OK**.
- **Step 6** View and manage tags on the **Tags** page.

----End

Editing a Tag

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.
- **Step 4** In the navigation pane on the left, click **Tags**.
- **Step 5** On the **Tags** page, locate the tag that you want to edit and click **Edit** in the **Operation** column. In the displayed dialog box, change the tag value and click **OK**.

Only the tag value can be edited.

Step 6 View and manage tags on the **Tags** page.

----End

Deleting a Tag

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, click the target instance. The **Basic Information** page is displayed.


- **Step 4** In the navigation pane on the left, click **Tags**.
- **Step 5** On the **Tags** page, locate the tag that you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.
- **Step 6** Check whether the deleted tag is displayed on the **Tags** page.

----End

Searching by tag

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** On the **Instances** page, select **Tags** in the search box.

Figure 4-77 Selecting tags



Step 4 Select the tag to be queried and click **OK** to query information about instances associated with the tag.

Figure 4-78 Searching by tag

----End

4.13 User Resource Quotas

Scenarios


Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. Quotas limit the number or amount of resources available to users, for example, the maximum number of GeminiDB instances that you can create.

If a quota cannot meet your needs, apply for a higher quota.

Checking Quotas

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** Click on the upper left corner and select a region and project.
- **Step 4** In the upper right corner, choose **Resources** > **My Quotas**.

Figure 4-79 My quotas



Step 5 On the displayed page, check the used and total quotas of each type of GeminiDB instance resources.

----End

Increasing Quotas

- **Step 1** Log in to the Huawei Cloud console.
- **Step 2** In the service list, choose **Databases** > **GeminiDB**.
- **Step 3** Click on the upper left corner and select a region and project.
- **Step 4** In the upper right corner, choose **Resources** > **My Quotas**.
- **Step 5** In the upper right corner of the page, click **Increase Quota**.

Figure 4-80 Increasing quotas



- **Step 6** On the **Create Service Ticket** page, configure parameters.
 - In the **Problem Description** area, describe why you need the adjustment.
- **Step 7** After all mandatory parameters are configured, read and agree to the agreement and click **Submit**.

----End

5 Best Practices

5.1 How Do I Set Pre-partition Keys When Creating a Table on a GeminiDB HBase Instance?

This section describes how to set a pre-partition key when creating a table on a GeminiDB HBase instance.

What Is Pre-partitioning

On a GeminiDB HBase instance, data is stored in different data partitions. Row key prefixes uniquely identify entities within each partition. By evenly distributing the data across partitions, workloads can be distributed evenly, so that cluster resources can be efficiently used.

For example, if two pre-partition keys are set to [1111, 2222] during table creation, data is divided into three ranges. The partitions to which the data belongs are divided based on the lexicographic order of row keys and partition keys. If rowkey < '1111' is specified, data is stored in the first partition. If '1111' <= rowkey < '2222' is specified, data is stored in the second partition. If rowkey >='2222' is specified, data is stored in the third partition. Ideally, the three partitions belong to different nodes. If partition keys are not properly set, partitions may belong to one cluster node.

Designing Pre-partition Keys

Theoretically, customer's application data can be evenly distributed by prefix in each partition. On a GeminiDB HBase instance, the ideal data volume in a partition is about 100 GB. There is no upper limit on the data volume in a single partition. If there is more than 100 GB of data in a partition, the data will be automatically partitioned. You can choose **Service Tickets** > **Create Service Ticket** in the upper right corner of the console to disable automated partitioning.

• Example 1:

If the first digit of row key values are evenly distributed from **0** to **9**, 10 partition keys can be set: **[0, 1, 2, 3, 4, 5, 6, 7, 8, 9]**. Values starting with these digits belong to their own partitions.

Example 2:

If the first two digits of row key values are evenly distributed from **00** to **FF** and the estimated data volume in each partition is about 100 GB, 256 partition keys are recommended: **[00, 01, 02, ..., FD, FE, FF]**.

Specifying Pre-partitions During Table Creation

On a GeminiDB HBase instance, HBase Shell or Java code can be used to specify pre-partitions during table creation.

Specify pre-partitions using HBase Shell when creating a table.

```
create 'tb','cf1','cf2', 'cf3', SPLITS => ['1111', '2222', '3333']
```

You can replace '1111', '2222', and '3333' with other custom partition key values. Use commas (,) to separate multiple values.

Specify pre-partitions using Java code when creating a table.

```
import java.util.ArrayList;
import java.util.List;
import org.apache.hadoop.conf.Configuration;
import org.apache.hadoop.hbase.HBaseConfiguration;
import org.apache.hadoop.hbase.TableName;
import org.apache.hadoop.hbase.client.Admin;
import org.apache.hadoop.hbase.client.ColumnFamilyDescriptor;
import org.apache.hadoop.hbase.client.ColumnFamilyDescriptorBuilder;
import org.apache.hadoop.hbase.client.Connection;
import org.apache.hadoop.hbase.client.ConnectionFactory;
import org.apache.hadoop.hbase.client.TableDescriptor;
import org.apache.hadoop.hbase.client.TableDescriptorBuilder;
import org.apache.hadoop.hbase.security.User;
import org.apache.hadoop.security.UserGroupInformation;
public class ExampleCreateTable
  public static void main(String[] args) throws Throwable
     // Create HBase configuration
     Configuration hbaseConfig = HBaseConfiguration.create();
// Enters an IP address of the HBase cluster.
     hbaseConfig.set("hbase.zookeeper.quorum", "127.0.0.1");
// Sets the default port number to 2181.
     hbaseConfig.set("hbase.zookeeper.property.clientPort", "2181");
     TableName tableName = TableName.valueOf("default", "tb1");
     // Enters a username and password.
     UserGroupInformation ugi = UserGroupInformation.createProxyUser("your_user_name",
UserGroupInformation.createRemoteUser("your_password"));
     // Establishes a connection to the HBase instance.
     try (Connection connection = ConnectionFactory.createConnection(hbaseConfig, User.create(ugi)))
        Admin admin = connection.getAdmin();
        // provide your split key here
        byte[][] splitkey = new byte[][]{ "row1".getBytes(), "row2".getBytes()};
        // 5 column families
        List<ColumnFamilyDescriptor> cfs = new ArrayList<>();
        cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf1".getBytes()).build());
        cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf2".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf3".getBytes()).build());
        cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf4".getBytes()).build());
       cfs.add(ColumnFamilyDescriptorBuilder.newBuilder("cf5".getBytes()).build());
        TableDescriptor tableDescriptor =
TableDescriptorBuilder.newBuilder(tableName).setColumnFamilies(cfs).build();
        // create table
        admin.createTable(tableDescriptor, splitkey);
  }
```

6 Performance White Paper

6.1 Performance Test Methods

This section describes performance tests of GeminiDB HBase instances, including environment, procedure, and model tests.

Test Environment

- Region: CN North-Beijing4
- AZ: AZ 1
- Elastic Cloud Server (ECS): h3.4xlarge.2 with 16 vCPUs, 32 GB of memory, and CentOS 7.5 64-bit image
- Stress test data model: 20 columns in a single row and 20 bytes in each column
- Instance specifications: all specifications described in Table 6-1

Table 6-1 Instance specifications

No.	Specifications	
Cluster 1	16 vCPUs 64 GB	

Test Tool

YCSB is an open-source tool for testing database performance. In this test, YCSB 0.17.0 is used.

For details on how to use this tool, see YCSB.

Test Metrics

Operations per Second (OPS): operations executed by a database per second

Test Procedure

- 1. Configure the workload file.
 - Set readproportion, insertproportion, updateproportion, scanproportion, and readmodifywriteproportion in the workload file based on Table 6-2.
- 2. Use workload-insert-mostly as an example. Run the following command to prepare test data:
 - nohup ./bin/ycsb load hbase -s -P workloads/workload-insert-mostly -P hbase.properties -p operationcount=400000000 -p recordcount=400000000 -threads \${thread} -p maxexecutiontime= \${maxexecutiontime} -s 1> data load.log 2>&1 &
- 3. Use workload-insert-mostly as an example. Run the following command to test performance:
 - nohup ./bin/ycsb run hbase -s -P workloads/workload-insert-mostly -P hbase.properties -p operationcount=9000000 -p recordcount=9000000 -threads ${\text -s 1> workload-insert-mostly_run.log 2>&1 &}$

Test Models

Workload model

Table 6-2 Workload models

Workload Model	Test Models		
workload-read (single-row read)	100% read		
workload-insert (single-row write)	100% insert		
workload-range-read (range read)	100% range read		

Preset data volume

In this performance test, 2 billion data records are preset for instances of each specification. Each data record contains 20 fields, and the size of each field is 20 bytes.

6.2 Performance Test Data

OPS of instances with different specifications can be tested using different service models with the same preset data volume. For details, see the numbers in bold in **Table 6-3**.

Table 6-3 Test data

,		Service Performance - Throughput (QPS)	Service Performance - P99 Latency (ms)
---	--	---	---

		Single- Row Write	Single- Row Read	Range Read	Single- Row Write	Singl e- Row Read	Rang e Read
GeminiDB HBase instance	2 billion	258729	87581	754631	3.5	6	20
Open-source HBase instance	2 billion	139723	18931	63350	3.9	5	46

◯ NOTE

- Operations per Second (OPS): operations executed by a database per second
- Test Model No.: test model sequence number. Table 6-2 lists the test models.

7 FAQs

7.1 Product Consulting

7.1.1 What Are the Precautions for Using GeminiDB HBase API?

- 1. DB instance operating systems (OSs) are invisible to you. Your applications can access a database only through an IP address and a port.
- 2. Backup files stored in OBS and system containers used by GeminiDB HBase API are invisible to you. They are visible only to the background management system.
- 3. Precautions after purchasing DB instances:

After purchasing DB instances, you do not need to perform basic database O&M operations, such as applying HA and security patches, but you should still note:

- a. The CPU, input/output operations per second (IOPS), and space are sufficient for the DB instances.
- b. The DB instance has performance problems and whether optimization is required.

7.1.2 What Is GeminiDB HBase Instance Availability?

The formula for calculating the instance availability is as follows:

DB instance availability = (1 - Failure duration/Total service duration) × 100%

The failure duration refers to the total duration of faults that occur during the running of an instance after you buy the instance. The total service duration refers to the total running time of the instance.

7.2 Billing

7.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-Use Billing Modes?

Yearly/Monthly is a prepaid mode. 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 is a postpaid mode. 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.

7.2.2 Can I Switch Between Yearly/Monthly and Pay-per-use Billing Modes?

Yes.

- To change a yearly/monthly instance to pay-per-use, see Changing a Yearly/ Monthly Instance to Pay-per-Use.
- To change a pay-per-use instance to yearly/monthly, see Changing a Pay-per-Use Instance to Yearly/Monthly.

7.3 Database Connection

7.3.1 How Can I Create and Connect to an ECS?

- 1. To create an ECS, see Elastic Cloud Server User Guide.
 - A GeminiDB HBase instance can connect to an ECS only when they are in the same VPC and subnet.
 - Configure the security group rules to allow the ECS to access to the instance.
- 2. To connect to an ECS, see "Logging in to an ECS" *Getting Started with Elastic Cloud Server*.

7.3.2 Can I Change the VPC After Buying a GeminiDB HBase Instance?

After a GeminiDB HBase instance is created, the VPC cannot be changed on the console.

However, you can change a VPC by restoring the full backup of your instance to the VPC you want to use. For details, see **Restoring a Backup to a New Instance**.

7.4 Backup and Restoration

7.4.1 How Long Can GeminiDB HBase Instance Backups Be Stored?

Automated backups are kept based on the retention period you specified. There is no limit for the manual backup retention period. You can delete manually backup files as needed

7.5 Regions and AZs

7.5.1 What Is AZ and How Can I Select an AZ?

ΑZ

An AZ is a part of a physical region with its own independent power supply and network. An AZ is generally an independent physical equipment room, ensuring independence of the AZ.

Each region contains multiple AZs. If one AZ becomes faulty, the other AZs in the same region can continue to provide services.

AZs in the same region can communicate with each other over an intranet.

Selecting an AZ

You can deploy your instances in different AZs for high availability. If one of an AZ becomes faulty, databases in other AZs will not be affected. When selecting AZs:

- If only one AZ is available in a region, there is only one AZ in the region.
- The AZ of a purchased DB instance cannot be changed.
- The AZs in one region can communicate with each other over an intranet.

For more information, see Regions and AZs.

7.5.2 Can Different AZs Communicate with Each Other?

An AZ is a part of a physical region with its own independent power supply and network. An AZ is generally an independent physical equipment room, ensuring independence of the AZ.

Each region contains multiple AZs. If one AZ becomes faulty, the other AZs in the same region can continue to provide services normally.

By default, different AZs in the same VPC can communicate with each other through an internal network.

For more information, see Regions and AZs.

7.5.3 Can I Change the Region After Buying a GeminiDB HBase Instance?

No. After an instance is created, its region cannot be changed.

7.6 Instance Freezing, Release, Deletion, and Unsubscription

Why Is My GeminiDB HBase Instance Released?

If your subscriptions have expired but not been renewed, or you are in arrears due to insufficient balance, your instances enter a grace period. If you do not renew your subscription or top up your account before the grace period ends, your instance will enter a retention period, during which the resources will be suspended. If you still do not renew them or top up your account after the retention period ends, your instances will be released and your data stored will be deleted. For details, see **Resource Suspension and Release**.

Why Is My GeminiDB HBase Instance Frozen?

Your instances may be frozen for a variety of reasons. The most common reason is that you are in arrears.

Can I Back Up Data After My Instance Is Frozen?

No. If your instance is frozen due to arrears, you need to renew the subscription to unfreeze the instance before backing up data.

How Do I Unfreeze My Instance?

If your instance is frozen because your account is in arrears, renew the subscription or top up your account. If your GeminiDB HBase instance is frozen due to arrears, you can renew the subscription, release the instance, or delete the instance. If your yearly/monthly GeminiDB HBase instance has expired, you cannot unsubscribe from it. You can unsubscribe from a yearly/monthly GeminiDB HBase instance that has not expired.

What Happens After My Resources Are Frozen, Unfrozen, or Released?

- After an instance is frozen:
 - It cannot be accessed, and your services will be interrupted. For example, if a GeminiDB HBase instance is frozen, it cannot be connected.
 - If they are yearly/monthly resources, no changes can be made to them.
 - It can be unsubscribed from or deleted manually.
- After it is unfrozen, you can connect to it again.
- When resources are released, the instance is deleted. Before the deletion, the system determines whether to move the instance to the recycle bin based on the recycling policy you specified.

How Do I Renew My Instance?

After a yearly/monthly GeminiDB HBase instance expires, you can renew it on the **Renewal Management** page. For details, see **Renewal Management**.

Can I Restore Released or Unsubscribed Resources?

If a deleted instance is moved to the recycle bin, you can restore it by following **Recycling an Instance**. If the recycling policy is not enabled, you cannot restore it.

When you unsubscribe from an instance, confirm the instance information carefully. If you have unsubscribed from an instance by mistake, purchase a new one.

How Do I Delete a GeminiDB HBase Instance?

- To delete a pay-per-use instance, see **Deleting a Pay-per-Use Instance**.
- To delete a yearly/monthly instance, see How Do I Unsubscribe from a Yearly/Monthly Instance?.