GeminiDB Influx

User Guide

Issue	01
Date	2025-09-04

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 Service Overview

1.1 What Is GeminiDB Influx API?

GeminiDB Influx API is a cloud-native NoSQL time-series database with decoupled compute and storage and full compatibility with InfluxDB. This high availability database is secure and scalable, can be deployed, backed up, or restored quickly, and provides monitoring and alarm management. You can also expand storage or compute resources separately. It is widely used to monitor resources, services, IoT devices, and industrial production processes, evaluate production quality, and trace faults. GeminiDB Influx API meets the demand of high concurrent read and write, compressed storage, and SQL-like query, and supports multi-dimensional aggregation computing and visualized data analysis.

It provides high write performance, flexibility, high compression ratio, and high query performance.

- Efficient write

  Data is written in parallel, distributed mode, and up to trillions of data records can be written per day.
- Flexibility

  Compute nodes can be independently up or down scaled to meet service requirements, and data is not migrated during scale-out. Cluster nodes can be scaled in or out in minutes.
- High compression ratio

  Compared with open-source HBase, GeminiDB Influx API improves the compression ratio by 5 to 10 times based on the column-oriented storage and dedicated compression algorithm.
- Efficient query

  GeminiDB Influx API can easily handle a large number of analysis tasks by running multiple threads concurrently on multiple nodes.

Typical Application Scenarios

- IoT sensor time series data analysis

  IoT applications often require a high level of scale and reliability. GeminiDB Influx API can achieve very high throughput and concurrency, so it can handle

a large number of connections in a very short period of time, making it an excellent choice for IoT applications.

Advantages

Intensive write

In less write-intensive scenarios, the write performance is 4.5 times that of the open source version. When write demands are more intensive, the write performance is 3.3 times that of the open source version.

Elastic scalability

Thanks to a distributed architecture with decoupled compute and storage, compute nodes can be expanded in minutes to handle with service peaks.

- Securities and cryptocurrency transactions

  GeminiDB Influx API stores user bank statements and builds an anti-fraud system for risk control in banks.

  Advantages

  Efficient query

  GeminiDB Influx API can be deployed in a region close to your users, so they can enjoy faster processing and response.

  Real-time analysis

  The series data can be synchronized to the cloud to be analyzed in real-time.

- Real-time monitoring with hardware and software

  GeminiDB Influx API can store user behavior data to support precision marketing and user profiling.

  Advantages

  Efficient write and query

  GeminiDB Influx API can handle trillions of data records per day. It supports multi-node and multi-thread parallel queries.

  Real-time analysis

  The series data can be synchronized to the cloud to be analyzed in real-time.

- Environmental protection industry

  GeminiDB Influx API supports the writing of massive amounts of time series data, making it stable and reliable for environmental protection data collection.

  Advantages

  Efficient write and query

  Vectorized query APIs and efficient time series data query operators such as aggregation and convolution can process a large number of concurrent data writes and queries.

1.2 Compatible API and Versions

GeminiDB Influx instances support the following types: cluster, cluster (performance-enhanced), and single node.

- These types cannot be converted or upgraded to each other.

- Cluster (performance-enhanced) instances do not support the Flux syntax. To use Flux, you are advised to buy cluster instances.

Type	Compatible Version	Scenario
Cluster	InfluxQL 1.7/1.8 Flux	One cluster consists of at least three nodes. Nodes can be added to a cluster, which is unsuitable to cope with the ever growing data volume.
(Recommended) Cluster (performance-enhanced)	InfluxQL 1.7/1.8	Compared with cluster instances, instances in a performance-enhanced cluster support a larger scale and higher read/write performance.
Single node	InfluxQL 1.7/1.8 Flux	A single-node instance cannot ensure the SLA. You are advised to use it only for tests and function verification.

1.3 Instance Specifications

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

This section describes specifications supported by a GeminiDB Influx instance. The instance specifications depend on the selected CPU model.

Table 1-1 GeminiDB Influx cluster instance specifications

Data Node Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)	Default Maximum Connections per Node	Time Series per Node (unit: 10,000)	Max. Databases per Cluster	Max. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb.large.4	2	8	100	96,000	250	4	16	40	1,000	5,000

Data Node Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)	Default Maximum Connections per Node	Time Series per Node (unit: 10,000)	Max. Databases per Cluster	Max. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb.xlarge.4	4	16	100	96,000	500	16	16	40	2,000	20,000
geminidb.influxdb.2xlarge.4	8	32	100	96,000	1,000	64	32	80	4,000	80,000
geminidb.influxdb.4xlarge.4	16	64	100	96,000	2,000	256	32	160	8,000	320,000
geminidb.influxdb.8xlarge.4	32	128	100	192,000	4,000	1,024	64	320	16,000	1,280,000

Table 1-2 Specifications of a GeminiDB Influx instance with cloud native storage

Data Node Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)	Default Maximum Connections per Node	Time Series per Node (unit: 10,000)	Max. Databases per Cluster	Max. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb-geminifs.large.4	2	8	100	64000	250	4	16	40	1,000	5,000
geminidb.influxdb-geminifs.xlarge.4	4	16	100	64000	500	16	16	40	2,000	20,000
geminidb.influxdb-geminifs.2xlarge.4	8	32	100	64000	1,000	64	32	80	4,000	80,000
geminidb.influxdb-geminifs.4xlarge.4	16	64	100	64000	2,000	256	32	160	8,000	320,000
geminidb.influxdb-geminifs.8xlarge.4	32	128	100	64000	4,000	1,024	64	320	16,000	1,280,000

Table 1-3 Specifications of GeminiDB Influx single-node instances

Data Node Flavor	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)	Default Maximum Connections per Node	Time Series per Node (unit: 10,000)	Max. Databases per Cluster	Max. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb.single.xlarge.2	4	8	100	1,000	250	3	16	40	1,000	3,500
geminidb.influxdb.single.2xlarge.2	8	16	100	2,000	500	12	16	40	2,000	14,000
geminidb.influxdb.single.4xlarge.2	16	32	100	4,000	1,000	48	32	80	4,000	56,000
geminidb.influxdb.single.8xlarge.2	32	64	100	8,000	2,000	192	32	160	8,000	112,000

Table 1-4 Specifications of a GeminiDB Influx cluster (performance-enhanced) instance

Specification	vCPUs	Memory (GB)	Min. Storage Space (GB)	Max. Storage Space (GB)	Default Maximum Connections per Node	Time Series per Node (unit: 10,000)	Max. Databases per Cluster	Max. RPs per Cluster
geminidb.influxdb.sqlstore.large.4	2	8	100	96,000	250	4	16	40
geminidb.influxdb.sqlstore.xlarge.4	4	16	100	96,000	500	16	16	40
geminidb.influxdb.sqlstore.2xlarge.4	8	32	100	96,000	1,000	64	32	80
geminidb.influxdb.sqlstore.4xlarge.4	16	64	100	96,000	2,000	256	32	160
geminidb.influxdb.sqlstore.8xlarge.4	32	128	100	30,000	4,000	1,024	64	320

When the memory usage of a GeminiDB Influx instance node reaches:

- 90% or higher: Queries running the longest are killed and new queries are not allowed.
- 80% or higher: The speed for handling new I/O requests is slowed down.

A GeminiDB Influx single-node instance (including read replicas) is deployed on a single server. Therefore, SLA cannot be guaranteed. You are advised to use it for testing and function verification. When the timeline scale exceeds twice the time series scale supported by a single node, data cannot be written to the single-node instance.

Table 1-5 Requests per second on nodes of different specifications and memory usages

Memory Usage (Unit: %)	2 vCPUs | 8 GB		4 vCPUs | 16 GB		8 vCPUs | 32 GB		16 vCPUs | 64 GB		32 vCPUs | 128 GB	
-	Read	Write	Read	Write	Read	Write	Read	Write	Read	Write

Memory Usage (Unit: %)	2 vCPUs | 8 GB		4 vCPUs | 16 GB		8 vCPUs | 32 GB		16 vCPUs | 64 GB		32 vCPUs | 128 GB	
80 ≤ Memory usage < 85	100	300	100	300	180	480	280	750	470	1200
85 ≤ Memory usage < 90	66	200	66	200	120	320	186	500	313	800
90 ≤ Memory usage < 95	50	150	50	150	90	240	140	375	235	600
95 ≤ Memory usage < 100	40	120	40	120	72	192	112	300	188	480

1.4 DB Instance Statuses

The status of a DB instance indicates the health of the instance. You can view the DB instance statuses on the management console.

Table 1-6 DB instance statuses

Status	Description
Available	DB instance is available.
Abnormal	DB instance is faulty.
Creating	DB instance is being created.
Creation failed	DB instance creation fails.
Restarting	DB instance is being restarted.
Resetting password	Administrator password is being reset.
Adding node	Nodes are being added to a DB instance.
Deleting node	Nodes are being deleted from a DB instance.
Scaling up	The storage space of the DB instance is being expanded.
Changing instance class	The CPU or memory of a DB instance is being changed.
Uploading backup	The backup file is being uploaded.
Backing up	Backup is being created.
Checking restoration	The backup of the current DB instance is being restored to a new DB instance.

Status	Description
Changing to yearly/monthly	The billing mode is being changed from pay-per-use to yearly/monthly.
Changing to pay-per-use	The billing mode is being changed from yearly/monthly to pay-per-use.
Creating cold storage	Cold storage is being created.
Scaling up cold storage	Cold storage is being scaled up.
Configuring SSL	SSL is being enabled or disabled.
Frozen	The instance is frozen because your balance drops to or below zero.
Unfreezing	DB instance is unfrozen after the overdue payments are cleared.
Checking changes	The yearly/monthly instance is pending check when its billing mode is changed.
Storage full	An instance will be set as read-only and its status will change to Storage full in the following circumstances: • Storage space ≥ 600 GB; Available space < 18 GB • Storage space < 600 GB; Space usage ≥ 97% The instance will become normal in the following circumstances: • Storage space ≥ 600 GB; Available space ≥ 90 GB • Storage space < 600 GB; Space usage ≤ 85%

1.5 Usage Specifications and Suggestions

This section describes the GeminiDB Influx instance specifications and provides suggestions for using GeminiDB Influx from the aspects of naming, TAG, FIELD, and query to solve common problems such as incorrect usage, low efficiency, and difficult maintenance.

Terms and Definition

- Rule: a convention that must be followed when you use GeminiDB Influx API.
- Suggestion: a convention that must be considered when you use GeminiDB Influx API.

Description

- Retention Policy (RP): includes information such as the data retention period and number of backups.

- Data objects: database, RP, MEASUREMENT, TAG, and FIELD

Naming

- **Rules**
  a. The name of a database object must start with a lowercase letter and consist of letters or digits. The length of the name cannot exceed 32 bytes.
  b. The name of a database object contains a maximum of 120 characters in the format of *<Database name>.<RP name>.<MEASUREMENT name>*.
  c. The name of the database object cannot use the system reserved keyword.

     The system reserved keywords include: ALL,ALTER,ANY,AS,ASC,BEGIN,BY,CREATE,CONTINUOUS,DATABASE,DATABASES,DEFAULT,DELETE,DESC,DESTINATIONS,DIAGNOSTICS,DISTINCT,DROP,DURATION,END,EVERY,EXPLAIN,FIELD,FOR,FROM,GRANT,GRANTS,GROUP,GROUPS,IN,INF,INSERT,INTO,KEY,KEYS,KILL,LIMIT,SHOW,MEASUREMENT,MEASUREMENTS,NAME,OFFSET,ON,ORDER,PASSWORD,POLICY,POLICIES,PRIVILEGES,QUERIES,QUERY,READ,REPLICATION,RESAMPLE,RETENTION,REVOKE,SELECT,SERIES,SET,SHARD,SHARDS,SLIMIT,SOFFSET,STATS,SUBSCRIPTION,SUBSCRIPTIONS,TAG,TO,USER,USERS,VALUES,WHERE,WITH,WRITE,WARM
  d. The name of a database object cannot contain Chinese characters or the following special characters: ["].$,/\0*?~#:|'
  e. The database name cannot be the same as the database name used by systems such as _internal, _kapacitor, _heimdall, _vision and opentsdb.
  f. TAG names cannot be updated or renamed.
- **Suggestions**
  a. Shorter TAG names can save more resources because each tag name has an index which is stored in the memory.
  b. The names of TAG KEY and FIELD KEY cannot be the same.

TAG

- **Rules**
  a. Fields that use the InfluxQL function (such as MAX, MIN, and COUNT) are stored as FIELD.
  b. TAG supports only the character string type. If the stored value is not of the character string type, the value is stored as FIELD.
- **Suggestions**
  a. TAG can distinguish data better than the MEASUREMENT name does.
  b. Design the TIME precision as required. Lower precision can bring better performance.
  c. The field often used as a search criterion is stored as a TAG.
  d. The field that uses GROUP BY is stored as a TAG.

FIELD

- **Rule**: The type of each field must be the same.
- **Suggestion**: Do not set too many fields. Each field is calculated separately. If there are too many fields, the fuzzy query will fail.

Query

- **Rules**
  a. Do not run SELECT * FROM to query data.
  b. The query statement must contain the time range restriction.
  c. Before bringing a service online, perform a load test to measure the performance of the database in peak hours.
- **Suggestions**
  a. During the query, select only the fields that need to be returned.
  b. Shorter time range can bring better query performance.
  c. The more accurate the TAG value is, the better the query performance is. Use a single time series for query, that is, specify all TAG values or more TAG values.
  d. Add **fill(none)** after **group by time intervals** in queries. The function of **fill(none)** is that no timestamp or value is returned for an interval without data points. If there is sparse data, the number of returned query results can be greatly reduced.
  e. If nested queries are used, place the filter for querying time range in the outermost query.

DELETE

Suggestion: Do not execute the DELETE statement to delete data. Set a retention period so that data can be automatically deleted.

Others

- **Rule**: Select instance specifications based on the service time series scale, number of client connections, and number of retention policies. For details, see **Instance Specifications**.

  If the database load exceeds the specification limit, unpredictable problems may occur. In severe cases, the database may be unavailable.
- **Suggestion**: Use a load balancer address to connect to the database. For details, see **Method 1: Using a Load Balancer Address over SSL Connections (Recommended)**.
- **Note**: If cold storage is enabled, cold data cannot be written.

1.6 Constraints

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

Specifications

Table 1-7 Specifications

Resource Type	Specifications	Description
CPU and memory	GeminiDB Influx API supports cluster (performance-enhanced), cluster, and single-node instances.	• 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**.	Disk capacity of GeminiDB Influx instances can be scaled up. For details, see **Manually Scaling Up Storage Space of a GeminiDB Influx Instance**.

Quotas

Table 1-8 Quotas

Resource Type	Constraint	Description
Tag	A maximum of 20 tags can be added for each instance.	For more information, see **Managing Tags**.
Free backup space	GeminiDB Influx 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 1-9 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.

Item	Description
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 1-10 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 **Changing the Administrator Password of a GeminiDB Influx Database**. • Keep your password secure. The system cannot retrieve it if it is lost.
Database port	Port for accessing a database. The default value is **8635** and cannot be changed.
VPC	After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed.
Security group	A security group controls access between GeminiDB Influx API 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.
Access control	A load balancer address does not support security groups. After an instance is created, configure IP address access control. If no whitelist is configured, all IP addresses that can communicate with the VPC can access the instance.

Instance Operations

Table 1-11 Instance operations

Function	Constraint
Database access	• If remote access is not enabled, GeminiDB Influx instances and their associated ECSs must be in the same VPC subnet. • The security group must allow access from the associated ECS. By default, a GeminiDB Influx 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 the GeminiDB Influx instance is **8635** and cannot be changed.
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 Influx instance	• GeminiDB Influx 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 Influx instance backups	GeminiDB Influx instance backups are stored in OBS buckets and are invisible to you.
Changing the CPU or memory of a GeminiDB Influx 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 storage space of an instance is full, data cannot be written to databases. You are advised to periodically check the storage space.

Function	Constraint
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 recycling policy 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 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 instances are billed, how you can renew subscriptions and manage costs, and what happens if your account goes into arrears.

- **Billing Modes**

  There are yearly/monthly and pay-per-use billing modes. Each one has different advantages and disadvantages.

  – Yearly/Monthly: You pay upfront for the amount of time you expect to use the instance for. You will need to make sure you have a top-up account with a sufficient balance or have a valid payment method configured first.

  – Pay-per-use: You can start using the GeminiDB instance first and then pay as you go.

  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 space, backup space, and EIP bandwidths. For details about the billing factors and formulas for each billed item, see **Billing Items**.

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

- **Renewing Subscriptions**

  If you want to continue using an instance after it expires, you need to renew the instance subscription within the specified period. Otherwise, resources, such as compute and storage, will be automatically released, and data may be lost.

  You can renew your subscription manually or automatically. For details, see **Overview**.

- **Viewing Bills**

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

- **Arrears**

  If there is not a sufficient account balance to pay for your bill and there is no other payment method configured, your account will go into arrears. If you want to continue using your cloud services, you will need to top up your account in a timely manner. For details, see **Arrears**.

- **Stopping Billing**

  If you no longer need to use your GeminiDB Influx instance, you can unsubscribe from or delete it to stop the billing. For details, see **Billing Termination**.

- **Managing Costs**

  GeminiDB Influx costs include resource costs and O&M costs. You can allocate, analyze, and optimize GeminiDB costs to save more money. For details, see **Cost Management**.

2.2 Billing Modes

2.2.1 Overview

There are yearly/monthly and pay-per-use billing modes. Each one has different advantages and disadvantages.

- 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 differences between the two billing modes.

Table 2-1 Differences between billing modes

Billing Mode	Yearly/Monthly	Pay-per-use
Payment	Prepaid Billed by the subscription term you purchase	Postpaid Billed for what you use
Billing Method	Billed by the subscription term you purchase	Calculated by the second but billed every hour
Billing Items	Instance specifications (vCPUs and memory), storage space, backup space, and EIPs	Instance specifications (vCPUs and memory), storage space, backup space, and EIPs

Changing the Billing Mode	Yearly/Monthly can be changed to pay-per-use. The change takes effect only after the yearly/monthly subscription expires. For details, see **Changing a Yearly/Monthly Instance to Pay-per-Use**.	Pay-per-use can be changed to yearly/monthly. For details, see **Changing a Pay-per-Use Instance to Yearly/Monthly**.
Changing the Specifications	Supported	Supported
Application Scenarios	Recommended for resources expected to be in use long term. A cost-effective option for scenarios where the resource usage duration is predictable.	Recommended when the resource demands are likely to fluctuate and you want more flexibility.

2.2.2 Yearly/Monthly Billing

If you expect to use resources for a longer period, you can save money by selecting yearly/monthly billing. This section describes billing rules for yearly/monthly GeminiDB Influx resources.

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 specifications	Instance specifications, including vCPUs and memory.
Storage space	If the actual storage usage exceeds your purchased storage, you will be billed for additional storage on a pay-per-use basis.
Backup space	GeminiDB Influx provides backup storage up to 100% of your provisioned database storage at no additional charge. After the free backup storage is used up, additional usage will be priced by the hour based on the backup storage pricing details. If it has been used less than one hour, you will be billed based on the actual duration.
(Optional) EIP bandwidth	GeminiDB Influx instances are accessible from public networks, and you are billed for the generated public network traffic, but not for private network traffic.

If you want to purchase a 3-node (specifications of each node: 4 vCPUs | 16 GB) GeminiDB Influx instance with 100 GB of storage space. At the bottom of the instance buying page, price details (excluding the backup space fee) will be displayed.

Figure 2-1 Example price

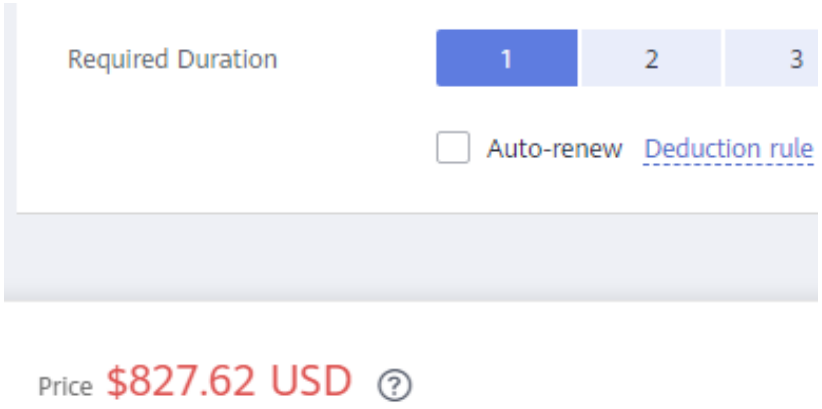

You are billed for:

- Selected specifications for your instance
- Storage space

NOTE

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

Backup Storage Space

DB Instance Type	Hourly	Currency
Cluster	0.00004	Price per GB

Billing Cycle

A yearly/monthly GeminiDB Influx instance is billed for the purchased duration (UTC+8). The billing starts when you activated or renewed the subscription, and ends at 23:59:59 of the expiry date.

For example, if you bought a GeminiDB Influx 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 Influx instance with 4 vCPUs, 16 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 one more month before it expired. The billing items include instance specifications (vCPUs, memory, and nodes), storage, backup storage, and EIP bandwidth.

- The first billing cycle is from 15:50:04 on March 8, 2023 to 23:59:59 on April 8, 2023.
- The second billing cycle is from 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 billed for 168 hours.

You need to pay in advance for each billing cycle. Each resource is billed separately.

Table 2-3 Billing formulas

Resource	Formula	Unit Price
Instance specifications (including vCPUs and memory)	Unit price of the instance specifications x Required duration x Number of nodes	For details about the unit price, see **Cluster CPU/Memory** on **Product Pricing Details**

Resource	Formula	Unit Price
Storage	Storage unit price x Required duration x Storage (GB)	For details about the unit price, see **Storage Space** on **Product Pricing Details**.
Backup storage	Backup storage unit price x Required duration x (Backup storage – Storage) (GB) **NOTE** The billing duration indicates how long the storage exceeding a free quota was used.	For details about the unit price, see **Backup Storage Space** on **Product Pricing Details**.
EIP bandwidth	Billed by fixed bandwidth	For details, see **Product Pricing Details**.

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

NOTE

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

Figure 2-2 Total price for a yearly/monthly GeminiDB Influx instance

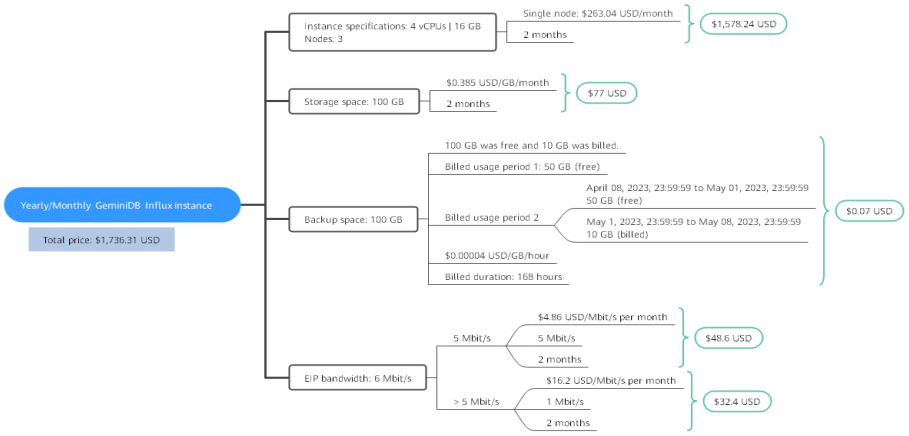

Price Change After Specification Change

If the specifications of a yearly/monthly GeminiDB Influx instance no longer meet your needs, you can change the specifications on the console. The system will recalculate the price and either bill or refund you the difference.

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

Decreasing instance specifications will affect instance performance. You are not advised to do so. Assume that you bought a yearly/monthly three-node GeminiDB

Influx instance with 4 vCPU and 16 GB of memory for one month on April 8, 2023 and increased its specifications to 8 vCPUs and 32 GB of memory on April 18, 2023. The old specifications cost USD827.62/month and the new specifications USD1616.74/month. The calculation formula is as follows:

Price difference = Price of new specifications x Remaining period – Price of old specifications x Remaining period

The remaining period is the remaining days of each calendar month divided by the maximum number of days in each calendar month. In this example, the remaining period is 0.6581 (12/30 + 8/31). The fee for increasing specifications is USD519.32 (1616.74 × 0.6581 – 827.62 × 0.6581).

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

Impact of Expiration

Figure 2-3 shows the statuses a yearly/monthly GeminiDB Influx instance can go through throughout its lifecycle. After a GeminiDB Influx 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 Influx instance

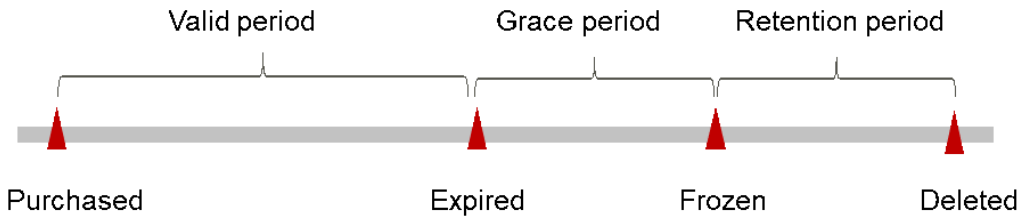

Expiration Reminder

The system will send you a reminder (by email, SMS, or in-app message) before a yearly/monthly GeminiDB Influx instance expires to the Huawei Cloud account creator.

- Expiration notifications will be sent 30 days, 15 days, 7 days, 3 days, and 1 day before yearly resources expire.
- Expiration notifications will be sent 15 days, 7 days, 3 days, and 1 day before monthly resources expire.

Impact of Expiration

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

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

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

NOTE

- For details about renewals, 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 Influx 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 specifications	vCPUs and memory
Storage	Instance storage space, which is billed hourly on a pay-per-use basis.
Backup storage	GeminiDB Influx API provides free backup storage equal to the amount of storage you purchased. After the free backup storage is used up, additional usage will be priced by the hour based on the backup storage pricing details. If it has been used less than one hour, you will be billed based on the actual duration.
(Optional) EIP bandwidth	GeminiDB Influx instances are accessible from public networks, and you are billed for the generated public network traffic, but not for private network traffic.

If you want to purchase a pay-per-use 3-node (specifications of each node: 4 vCPUs | 16 GB) GeminiDB Influx instance with 100 GB of storage space. At the

bottom of the instance buying page, price details (excluding the backup space fee) will be displayed.

Figure 2-4 Example price

You are billed for:

- 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**.

Backup Storage Space

DB Instance Type	Hourly	Currency
Cluster	0.00004	Price per GB

Billing Cycle

Pay-per-use instance usage is calculated by the second and priced by the hour. The billing starts when an instance is created and ends when the instance 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 Influx instance at 8:45:30 and delete 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 4 vCPUs, 16 GB of memory, 100 GB of storage, and 110 GB of backup storage (100 GB for free) at 09:59:30 on April 18 and deleted the instance at 10:45:46 on April 18, 2023. 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.

The price displayed in the pricing details is per hour, so you need to divide it by 3,600 to obtain the price for each second and then multiply the per-second price by the total number of seconds. GeminiDB Influx instances are billed individually as follows.

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 billing duration indicates how long the storage exceeding a free quota was used.	See the estimated price of a cluster instance with specified backup storage in **GeminiDB Price Calculator**.
Public network traffic	Tiered billing by fixed bandwidth • 0 Mbit/s to 5 Mbit/s (included): billed at a fixed unit price per Mbit/s • Greater than 5 Mbit/s: billed at a different price per Mbit/s	For details, see the estimated bandwidth price in **ECS Price Calculator** or **EIP Price Calculator**.

Figure 2-5 shows how the total price is calculated.

NOTE

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

If the price is not an integer, it is rounded off to the nearest two decimal places. If the rounded price is less than USD0.01, USD0.01 will be displayed.

Figure 2-5 Total price for a pay-per-use GeminiDB Influx instance

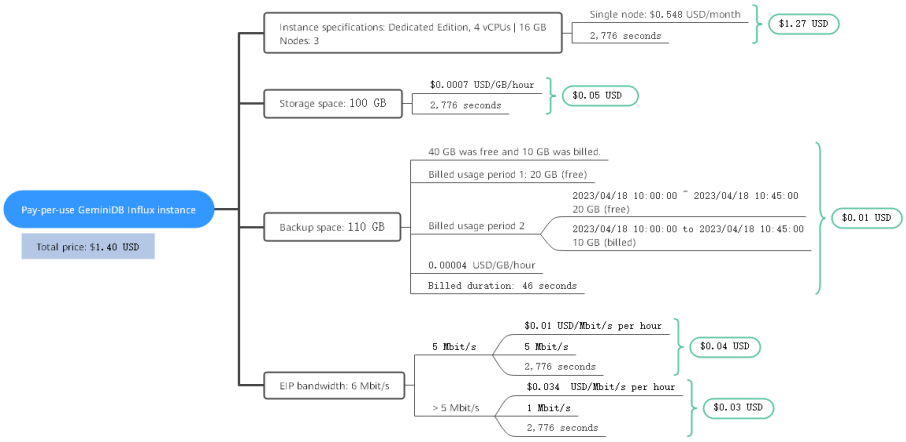

Impact on Billing After Specification Changes

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 4 vCPUs and 16 GB of memory at 9:00:00 and increased its specifications to 8 vCPUs and 32 GB of memory at 9:30:00, two billing records are generated between 9:00:00 and 10:00:00:

- 4 vCPUs and 16 GB of memory from 9:00:00 to 9:30:00
- 8 vCPUs and 32 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 Influx instance throughout its lifecycle. After a GeminiDB Influx instance is purchased, it enters the valid period and runs normally during this period. If your account goes into arrears, the instance enters a grace period and then a retention period.

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

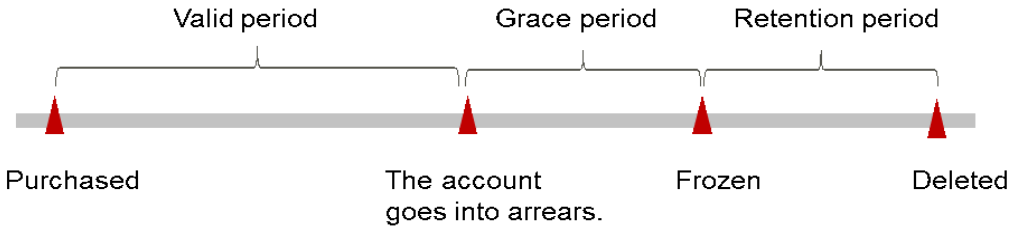

Arrears reminder

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

If your account is in arrears due to automated deduction for pay-per-use GeminiDB Influx instances, the instances are not immediately stopped but given a 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.

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 the 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 space, backup space, and public network traffic. For details, see **Table 2-6**.

NOTE

The billed items marked with asterisks (*) are mandatory.

Table 2-6 Billing Items of a GeminiDB Influx Instance

Billing Item	Description	Billing Mode	Formula
* Specifications	Billed by instance specifications, including vCPUs and memory. Computing and storage capabilities vary by the number of vCPUs and memory size.	Yearly/ Monthly and pay-per-use	**Unit price x Required duration** For details about the unit price, see **Cluster CPU/Memory** on **Product Pricing Details**.

Billing Item	Description	Billing Mode	Formula
* Storage space	Billed based on unified standards.	Yearly/ Monthly and pay-per-use	**Unit price x Storage space x Required duration** For details about the unit price, see **Storage Space** on **Product Pricing Details**.
Backup space	Billed based on unified standards.	Pay-per-use	**Unit price x Billed backup space x Required duration** For details about the unit price, see **Backup Storage Space** on **Product Pricing Details**. **NOTE** The billing duration indicates how long the storage exceeding a free quota was used.
Public network traffic	An EIP is required if an instance needs to access the Internet. Billing factors: bandwidth, traffic, and IP reservation • EIP for a yearly/monthly GeminiDB Influx instance: billed by bandwidth. • EIP for a pay-per-use instance: billed by bandwidth, traffic, or shared bandwidth. You are also billed for IP reservation.	Yearly/ Monthly and pay-per-use You can purchase a bandwidth add-on package or a shared traffic package.	**Tiered pricing based on fixed bandwidth**. • 0 Mbit/s to 5 Mbit/s (included): billed at a fixed unit price per Mbit/s. • Greater than 5 Mbit/s: billed at a different price per Mbit/s. For details about the unit price, see **Bandwidth Price** on **Product Pricing Details** or **Product Pricing Details**.

Billing Examples

Assume that you bought a three-node GeminiDB Influx instance with 4 vCPUs, 16 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 one more month before it expired. The billing items include instance specifications (vCPUs, memory, and nodes), storage, backup storage, and EIP bandwidth.

- The first billing cycle is from 15:50:04 on March 8, 2023 to 23:59:59 on April 8, 2023.
- The second billing cycle is from 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 billed for 168 hours.

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

NOTE

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

Figure 2-7 Total price for a yearly/monthly GeminiDB Influx instance

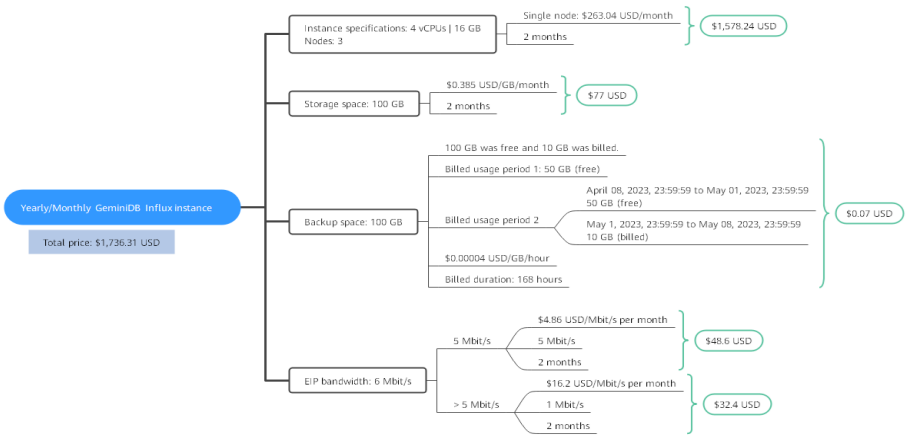

For more billing examples of a pay-per-use GeminiDB Influx instance, see **Billing Examples**.

2.4 Billing Examples

Billing Scenario

A user purchased a pay-per-use GeminiDB Influx instance at 15:30:00 on March 18, 2023. The instance configuration is as follows:

- Specifications: 4 vCPUs and 16 GB
- Nodes: 3
- EIP bandwidth: 6 Mbit/s

After a period, the user found that the current instance specifications could not meet service requirements and increased the specifications to 8 vCPUs and 32GB at 9:00:00 on March 20, 2023. Since the user wanted to use the instance long term, the user then changed the instance to yearly/monthly billing with a one-month duration at 10:30:00 on the same day. So how much will the user be billed for this GeminiDB Influx instance in March and April?

Billing Analysis

The total price of this GeminiDB Influx instance involves both pay-per-use and yearly/monthly usage:

- 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: 4 vCPUs and 16 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: 8 vCPUs and 32 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: 8 vCPUs and 32 GB of memory
  - Nodes: 3
  - Storage: 200 GB
  - Backup storage: 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
  - Billing duration: one month

NOTE

Unit prices in this example are used for reference only, and the calculated prices are only estimates. The actual unit price and cost may vary. 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 4 vCPUs and 16 GB was used for 41.5 hours, so the price is calculated as follows.

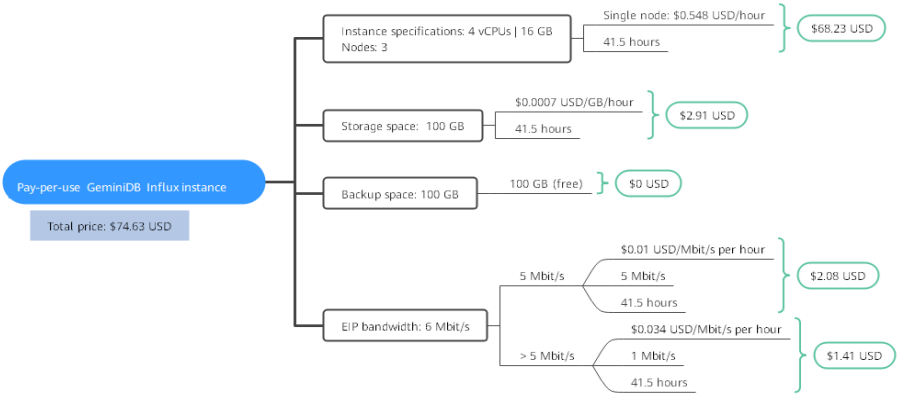

From 09:00:00 on March 20, 2023 to 10:30:00 on March 20, 2023, an instance with 8 vCPUs and 32 GB was used for 1.5 hours, so the price is calculated as follows.

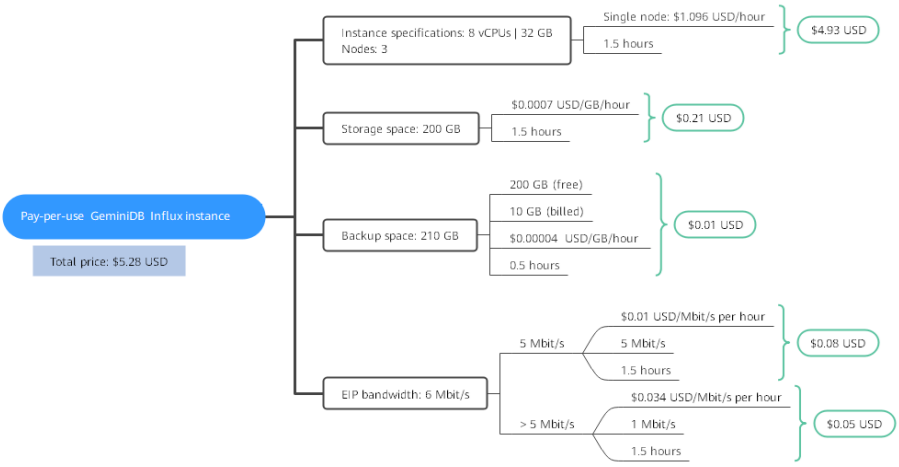

Yearly/Monthly

From 10:30:00 on March 20, 2023 to 23:59:59 on April 20, 2023, the yearly/monthly instance was used for one month. The price is calculated as follows.

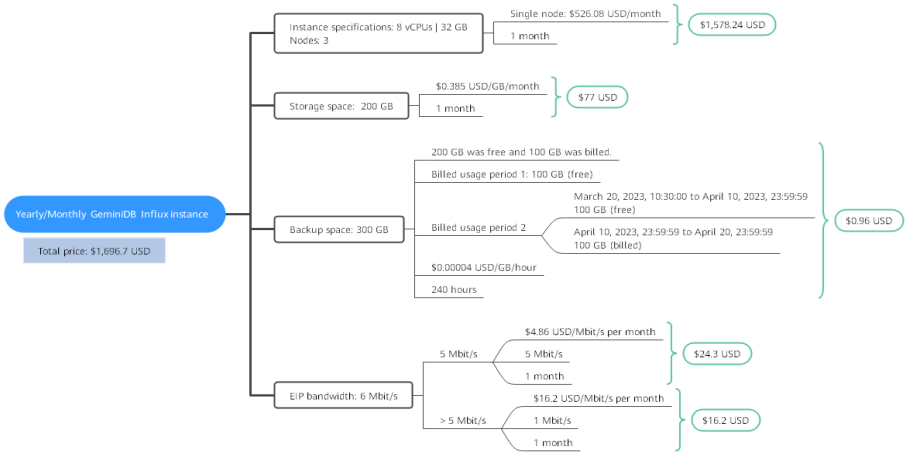

From March to April, the total price of this instance is USD1776.61 (74.63 + 5.28 + 1696.7).

2.5 Billing Mode Changes

2.5.1 Overview

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

Table 2-7 Changeable billing items of the GeminiDB Influx instance

Billing Item	Change Description	Reference
Instance specifications (vCPUs and nodes)	Changing the billing mode of a GeminiDB Influx instance includes the changes to compute resources (vCPUs and nodes). • Change from pay-per-use to yearly/monthly to enjoy lower prices. • Change from yearly/monthly to pay-per-use to use the GeminiDB Influx instance more flexibly **NOTE** Such a change takes effect only after the yearly/monthly subscription ends.	• **Changing a Pay-per-Use Instance to Yearly/Monthly** • **Changing a Yearly/Monthly Instance to Pay-per-Use**
EIP	• A yearly/monthly EIP can be changed to a pay-per-use EIP billed by bandwidth after the yearly/monthly subscription ends. • A pay-per-use EIP billed by bandwidth can be changed to a yearly/monthly EIP. • Pay-per-use EIPs billed by bandwidth can be changed to pay-per-use EIPs billed by traffic, and pay-per-use EIPs billed by traffic can be changed to pay-per-use EIPs billed by bandwidth. For details, see **Figure 2-8**.	• **Changing a Pay-per-Use Instance to Yearly/Monthly** • **Changing a Yearly/Monthly Instance to Pay-per-Use**

Figure 2-8 EIP billing mode change

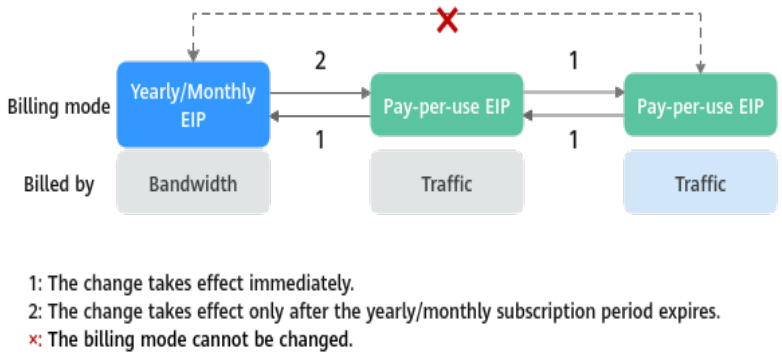

2.5.2 Changing a Pay-per-Use Instance to Yearly/Monthly

If you have a pay-per-use GeminiDB Influx instance that you expect to use for a long time, you can change it to yearly/monthly billing to reduce costs. 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 Influx instance at 15:29:16 on April 18, 2023 and changed it to yearly/monthly billing 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

Constraints

Resources such as EIPs that are used by an instance may not support the change with this instance. For details about their billing mode change rules and handling methods, see **Table 2-8**.

Table 2-8 EIP billing mode change rules

Resource	Billing Mode	Billed By	Bandwidth Type	Changed to Yearly/ Monthly Billing with the GeminiDB Influx Instance	Handling Measure
EIP	Pay-per-use	Bandwidth	Dedicated	Supported	Change the EIP to yearly/ monthly billing on the EIP console. For details, see **Changing EIP Billing Mode**.
EIP	Pay-per-use	Traffic	Dedicated	Not supported	An EIP that is billed by traffic on a pay-per-use basis cannot be directly changed to be billed on a yearly/monthly basis. To change this: 1. Change the EIP to be billed by bandwidth on a pay-per-use basis. 2. Change the EIP to be billed on a yearly/ monthly basis. For details, see **Changing EIP Billing Mode**.

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 2-9 Changing a pay-per-use instance to yearly/monthly

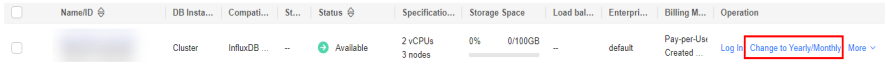

> **NOTE**
>
> The billing mode of multiple instances can be changed in batches. Perform the following steps:
>
> 1. Select the instances whose billing mode you want to change.
> 2. Click **Change to Yearly/Monthly** above the instance list.

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 results on the **Instances** page.

In the upper right corner of the instance list, click 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

After creating a yearly/monthly GeminiDB Influx instance, you can change it to pay-per-use for more flexibility, and you can recoup part of what you paid for the subscription.

Suppose you bought a yearly/monthly GeminiDB Influx instance at 15:29:16 on April 18, 2023 and changed it to pay-per-use billing 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.

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

Constraints

Resources such as EIPs that are used by an instance may not support the change with this instance. For details about their billing mode change rules and handling methods, see **Table 2-9**.

Table 2-9 EIP billing mode change rules

Resource	Billing Mode	Billed By	Bandwidth Type	Change to Pay-per-Use Billing with the GeminiDB Influx Instance	Handling Measure
EIP	Yearly/Monthly	Bandwidth	Dedicated	Not supported	Change the EIP to yearly/monthly billing on the EIP console. For details, see **Changing EIP Billing Mode**.
EIP	Yearly/Monthly	Traffic	Dedicated	Not supported	An EIP billed on a yearly/monthly basis cannot be directly changed to be billed by traffic on a pay-per-use basis. To change this: 1. Change the EIP to be billed by bandwidth on a pay-per-use basis. 2. Change the EIP to be billed by traffic on a pay-per-use basis. For details, see **Changing EIP Billing Mode**.

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 **More** > **Change to Pay-per-Use** in the **Operation** column.

Figure 2-10 Change to Pay-per-Use

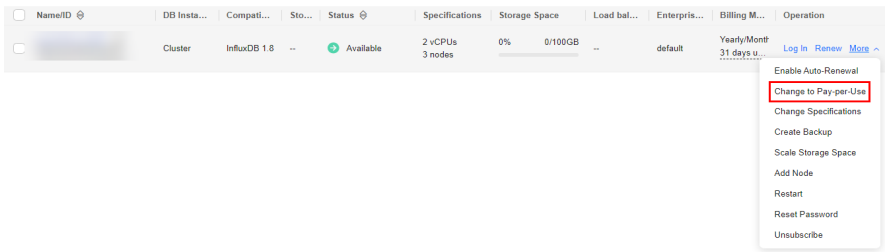

NOTE

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 Billing Center. On the **Renewals** page, locate the instance and click **More** > **Cancel Change to Pay-per-Use**.

Step 7 In the displayed dialog box, click **Yes**.

----End

2.6 Renewing Subscriptions

2.6.1 Overview

When to Renew Subscriptions

If a yearly/monthly instance is about to expire but you want to continue using it, you need to renew the instance subscription within a specified period, or 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 after instances have expired and the associated impacts, see **Impact of Expiration**.

How to Renew Subscriptions

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

Table 2-10 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-renewing an Instance	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 select a method to renew a yearly/monthly instance based on the phase the instance is currently in.

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

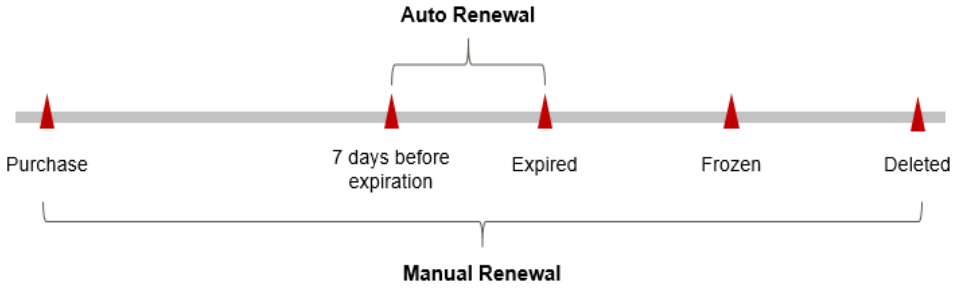

- An instance is in the **Provisioned** state after it is provisioned.
- When an instance subscription expires, the status will change from **Provisioned** to **Expired**.
- If an expired instance is not renewed, it enters a grace period. If it is not renewed by the time the grace period expires, the 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

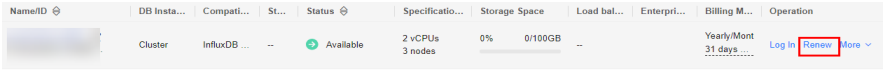

Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 2-13 Renewing an instance

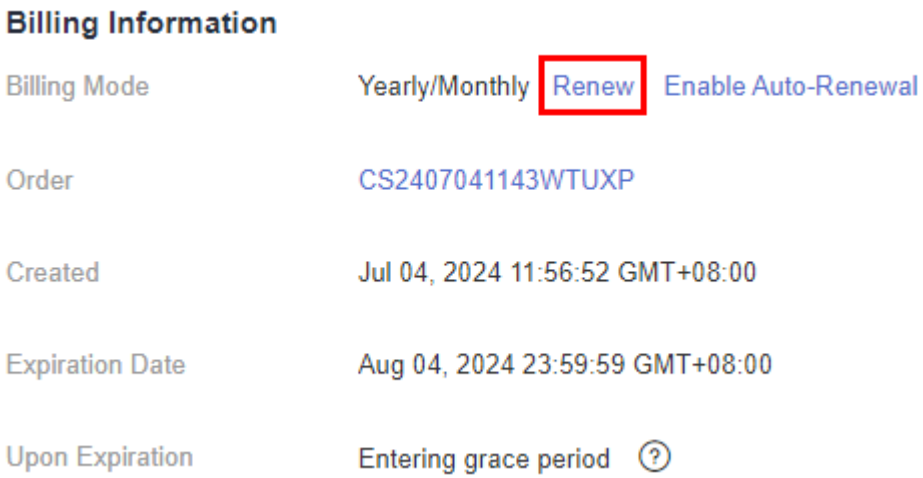

NOTE

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 On the top menu bar, choose **Billing** > **Renewal**.

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

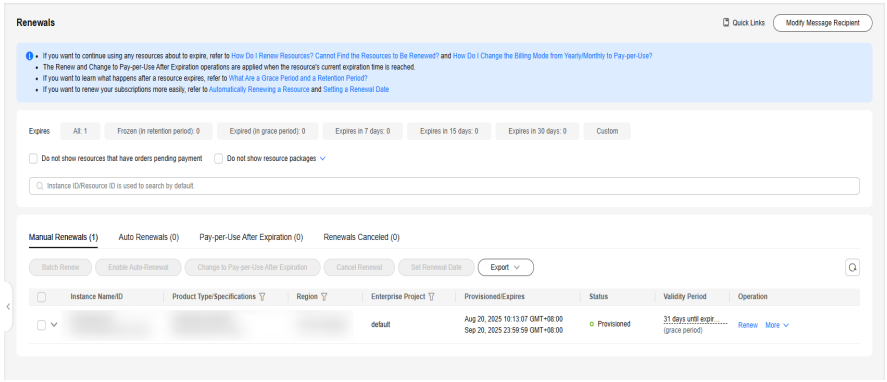

You can move all resources to be manually renewed to the **Manual Renewals** tab. For details, see **Restoring to Manual Renewal**.

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

  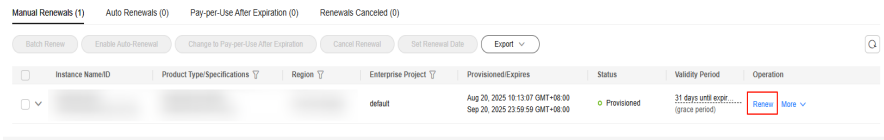

- Batch renewal: Select multiple instances that you want to renew and click **Batch Renew** in the upper left corner.

  Figure 2-16 Batch renewal

  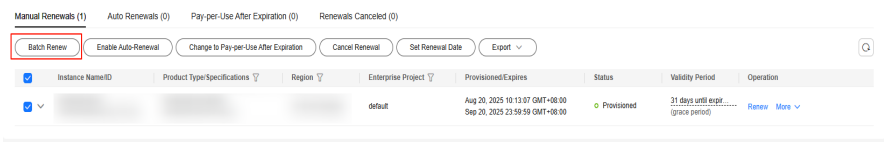

Step 5 Select a renewal duration and optionally select **Renew on the standard renewal date**. For details, see **Setting the Same Renewal Day for Yearly/Monthly Resources**. Confirm the price and click **Pay**.

Figure 2-17 Confirming renewal

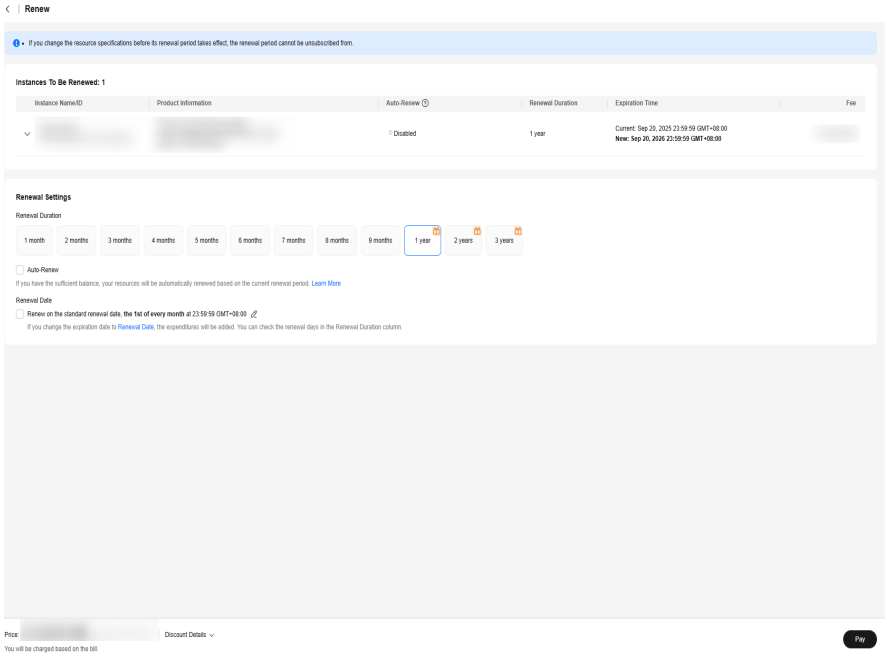

Step 6 Select a payment method and make your payment. Once the order is paid for, the renewal is complete.

----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.

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

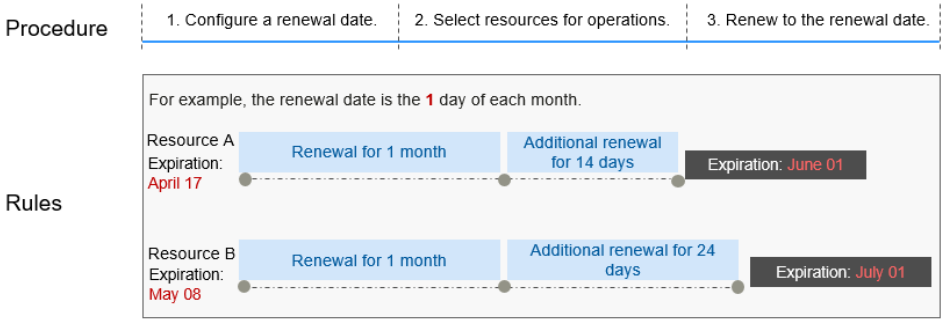

For more details, see **Setting a Renewal Date**.

2.6.3 Auto-renewing an Instance

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 is subject to the renewal duration you select.
  - Your monthly subscription will be renewed each month.
  - Your yearly subscription will be renewed 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 During Purchase

You can enable auto-renewal on the instance purchase page, as shown in **Figure 2-19**. For details, see **Buying an Instance**.

Figure 2-19 Enabling auto-renewal

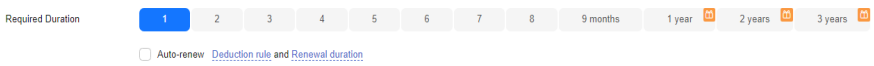

Enabling Auto-Renewal on the Renewals Page

Step 1 **Log in to the Huawei Cloud console**.

Step 2 On the top navigation bar, choose **Billing** > **Renewal**.

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**, **Pay-per-Use After Expiration**, and **Renewals Canceled** pages.

Figure 2-20 Renewal management

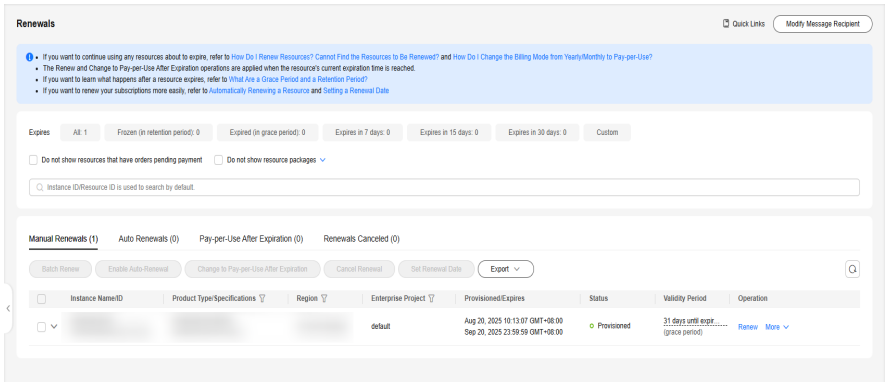

Step 4 Enable auto-renewal for yearly/monthly resources.

- 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

  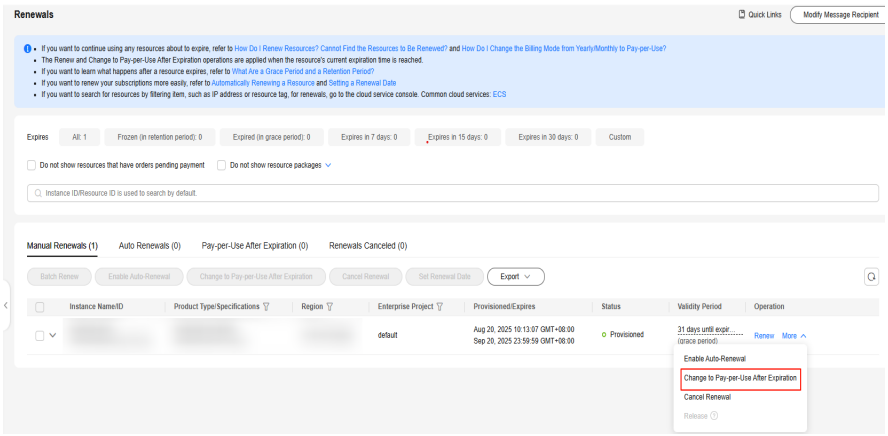

- Enabling auto-renewal for multiple instances at a time: Select the instances that you want to enable auto-renewal for and click **Enable Auto-Renew** above the list.

  Figure 2-22 Enabling auto-renewal for multiple instances

  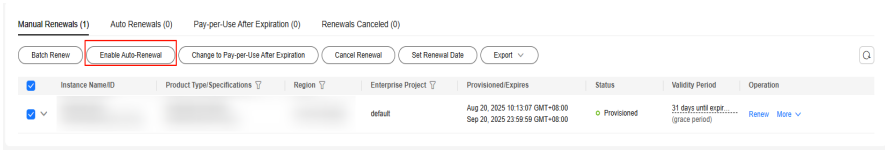

Step 5 Select a renewal period, specify the auto-renewal times, and click **Pay**.

Figure 2-23 Enabling auto-renewal

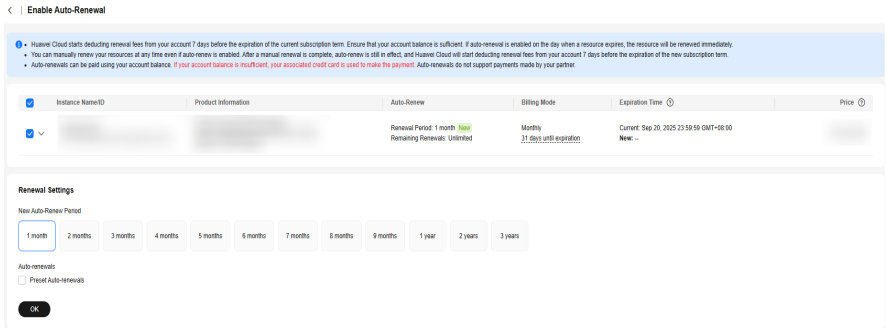

----End

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. A pay-per-use resource is billed by the hour, day, or month, depending on the resource's usage type. The GeminiDB Influx instance usage is billed by the hour. For details, see **Bill Run for Pay-per-Use Resources**.

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 Influx API**.

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

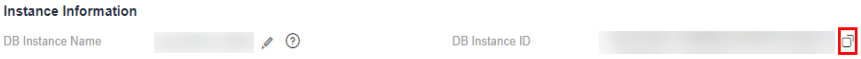

Step 4 On the top menu bar, choose **Billing** > **Bills**.

The **Bills** page is displayed.

Step 5 In the navigation pane, choose **Billing** > **Expenditure Details**. Select **Resource ID** as the filter, enter the resource ID, and click to search for the resource bills.

Figure 2-25 Querying resource bills

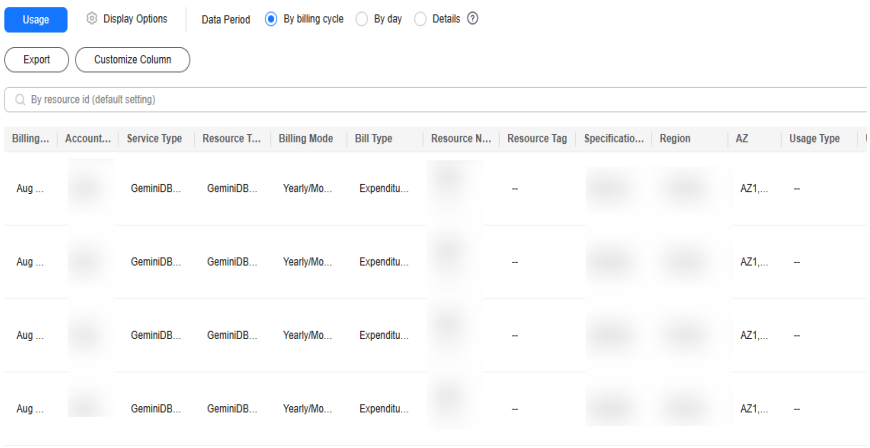

By default, the bill details are displayed by usage and billing cycle. You can choose other display options. For details, see **Bills**.

----End

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

Step 1 **Log in to the Huawei Cloud console**. Choose **Databases** > **GeminiDB Influx API**.

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

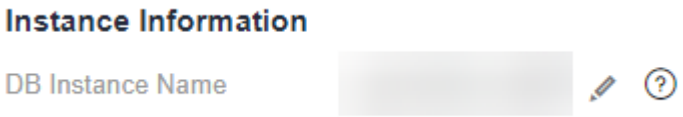

Step 4 On the top menu bar, choose **Billing** > **Bills**.

The **Bills** page is displayed.

Step 5 In the navigation pane, choose **Billing** > **Expenditure Details**. Select **Resource Name** as the filter, enter the resource name, and click to search for the resource bills.

Figure 2-27 Querying resource bills

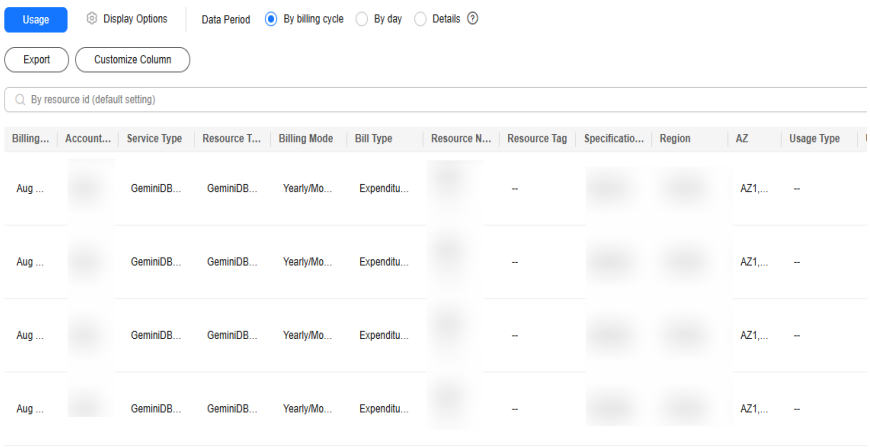

By default, the bill details are displayed by usage and billing cycle. You can choose other display options. For details, see **Bills**.

----End

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

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

- 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. The resources are billed separately. **Table 2-11** uses storage as an example.

  Table 2-11 Transaction records

Product Type	GeminiDB Influx
Resource Type	Storage
Billing Mode	Pay-per-use
Expenditure 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) × 0.0007 × 40. You can also calculate the list price in the other periods.
Discounted Amount	You can enjoy discounts on cloud services, such as business, partner-authorized, and promotional discounts. The discounts are calculated based on the list price.
Truncated Amount	Billing of Huawei Cloud is calculated to the 8th decimal place. 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.
Amount 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).

- Bill details of the GeminiDB Influx instance

  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 2-12** against the actual usage.

  Table 2-12 Bill details

Product Type	GeminiDB Influx
Resource Type	Storage
Billing Mode	Pay-per-use
Resource Name/ID	Name and ID Example: nosql-b388 and 21e8811a64bf4de88bc2e2556da17983in12
Specifications	Storage
Usage Type	Duration of a pay-per-use GeminiDB Influx instance

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 check the unit price of a pay-per-use 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 billed 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 × 0.0007 × 40).
Discounted Amount	You can enjoy discounts on cloud services, such as business, partner-authorized, and promotional discounts. The discounts are calculated based on the list price.
Amount 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.

Arrears Impact

- Yearly/Monthly

  This is a pre-paid billing mode, so you can continue using yearly/monthly GeminiDB Influx resources even if your account is in arrears. However, you cannot perform operations such as purchasing GeminiDB Influx instances, upgrading instance specifications, and renewing subscriptions, because they will generate new expenditures.

- Pay-per-Use

If your configured payment method is unable to pay a bill for pay-per-use resources, the resources enter a grace period. After you top up your account, Huawei Cloud will bill you for expenditures generated by the resources during the grace period. You can view the expenditures on the **Overview** page of the Billing Center.

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

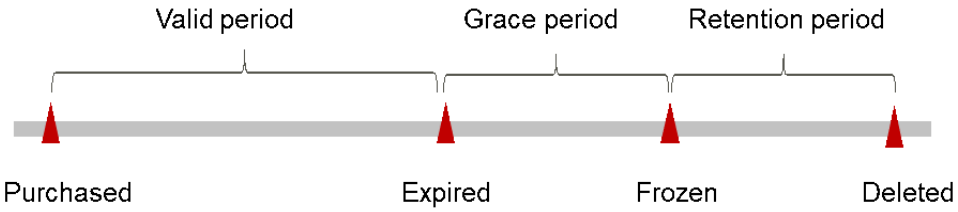

NOTE

The grace period and retention period are both 15 days.

Avoiding and Handling Arrears

Make sure you have a valid payment method configured as soon as possible after your account is in arrears. For details, see **Topping Up an Account**.

If a GeminiDB Influx instance is no longer used, you can delete it to avoid generating further expenditures.

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, any time 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

When you purchase a yearly/monthly resource, such as a yearly/monthly GeminiDB Influx instance, you make a one-time up-front payment. By default, the billing automatically stops when the purchased subscription expires.

- 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 pay-per-use resources, such as pay-per-use GeminiDB Influx instances, are no longer required, delete them in a timely manner.

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.

[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** > **Bills**.

The **Bills** page is displayed.

Step 2 In the navigation pane, choose **Billing** > **Expenditure Details**. Click the icon shown in the following figure to copy the resource ID.

Figure 2-29 Copying the resource ID

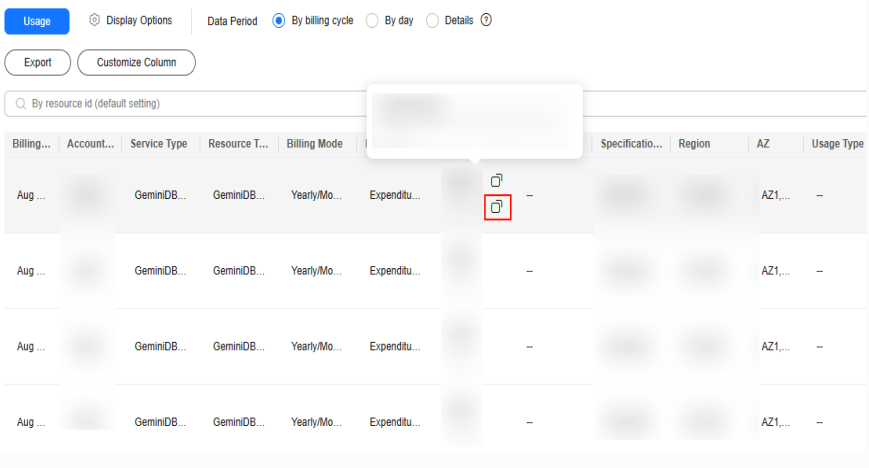

Step 3 Log in to the Huawei Cloud console. Choose **Databases** > **GeminiDB Influx API**.

Step 4 Select the region where the resource is located. Select **Instance ID**, enter the resource ID copied from **Step 2**, and click to search for the resource.

Figure 2-30 Searching for resources

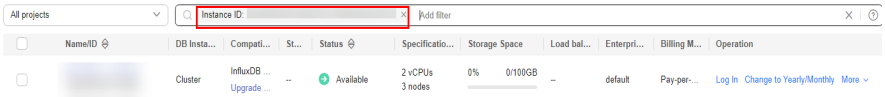

Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

> **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** > **Bills**.

The **Bills** page is displayed.

Step 2 In the navigation pane, choose **Billing** > **Expenditure Details**. Click the icon shown in the following figure to copy the resource name.

Figure 2-31 Copying the resource name

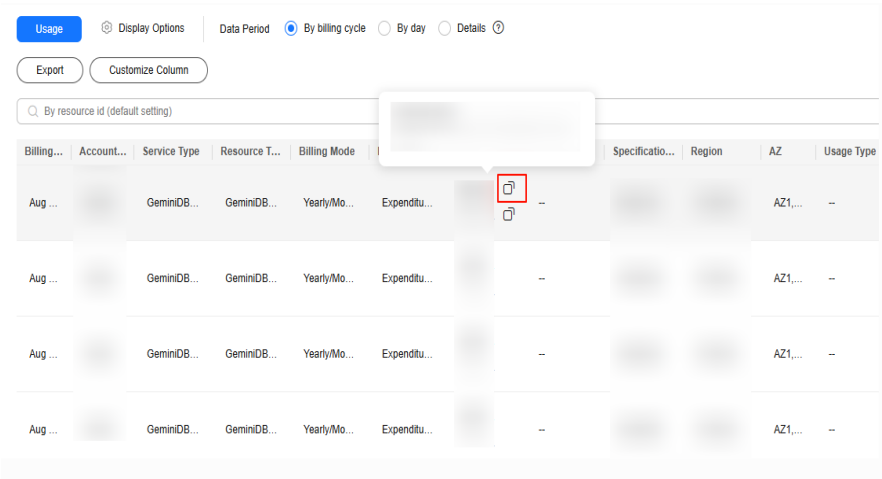

Step 3 Log in to the Huawei Cloud console. Choose **Databases** > **GeminiDB Influx API**.

Step 4 Enter the instance name copied from **Step 2** in the search box and click .

Figure 2-32 Searching for resources

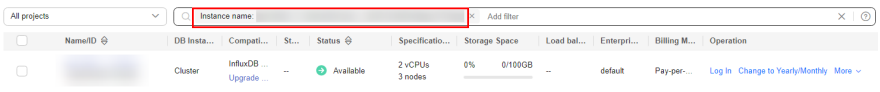

Step 5 Locate the instance you want to delete and click **More** > **Delete** in the **Operation** column. Ensure that the resource is not found in the list.

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

2.10 Cost Management

2.10.1 Cost Composition

GeminiDB Influx costs consist of two parts:

- Resource costs: costs of compute and storage resources. For details, see **Billing Modes**.
- O&M costs: labor costs incurred during the use of GeminiDB Influx.

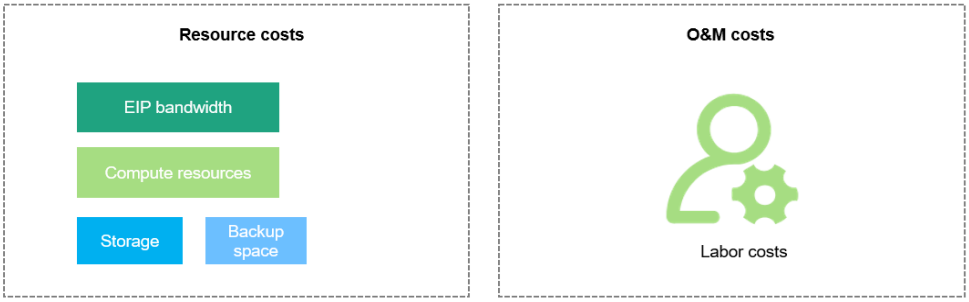

2.10.2 Cost Allocation

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

  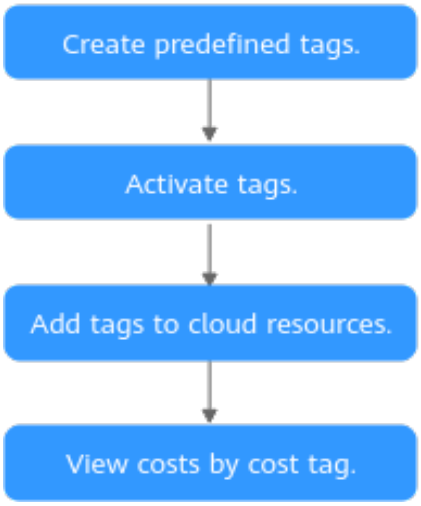

  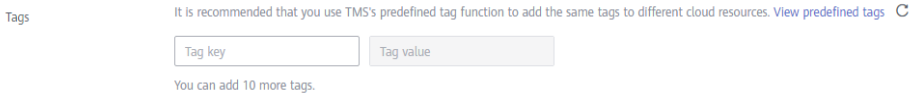

  For details, see **Viewing Costs by Cost Tag**.

- **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 more information, see **Allocating Costs 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 rightsizing

- View GeminiDB Influx monitoring metrics on Cloud Eye, such as the CPU, memory, and disk usage. If the current configuration is too high, you can reduce the configuration by changing specifications.
- Monitor idle GeminiDB Influx resources and delete idle instances 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 billing mode in which resources are billed based on the service duration. This cost-effective mode is ideal when the duration of resource usage is predictable. It is recommended for long-term users.

Pay-per-use billing is a postpaid payment mode. This billing mode allows you to make or cancel subscriptions at any time. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.

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

You can change the billing mode of your instance from yearly/monthly to pay-per-use or vice versa.

- For details about how to change the billing mode from yearly/monthly to a pay-per-use, see **Changing a Yearly/Monthly Instance to Pay-per-Use**.
- For details about how to change the billing mode from pay-per-use 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 Influx instances.

Usage Notes

- Pay-per-use GeminiDB Influx instances do not support this function.
- This function is only available for cluster instances.

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 instance that you want to renew and click **Renew** in the **Operation** column.

Figure 2-35 Renewal

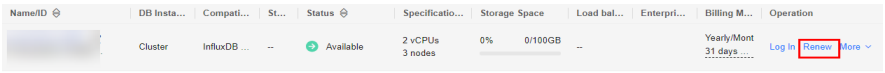

Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 2-36 Renewal

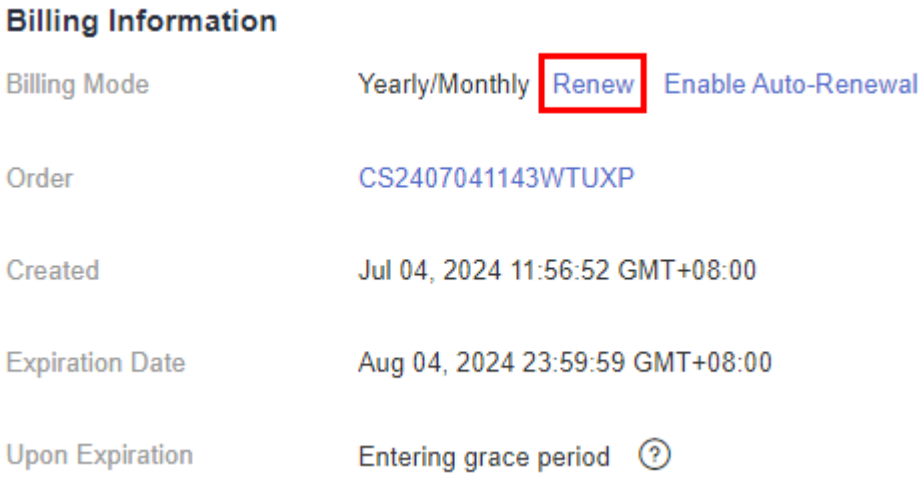

Step 4 On the displayed page, renew the instance.

----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

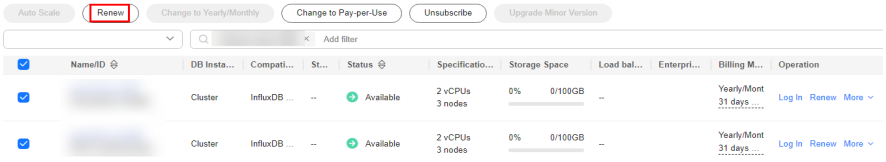

Step 4 In the displayed dialog box, click **Yes**.

----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. To retain data, back it up before submitting the unsubscription request.

- This function is only available for cluster instances.

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 instance you want to unsubscribe from and click **Unsubscribe** or choose **More** > **Unsubscribe** in the **Operation** column.

Figure 2-38 Unsubscribe

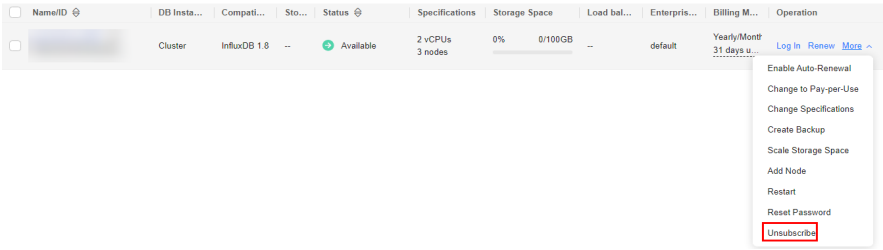

Step 4 In the displayed dialog box, click **Yes**.

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**.

NOTE

1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
2. To retain data, back it up 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

Batch Unsubscribing from Yearly/Monthly Instances

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

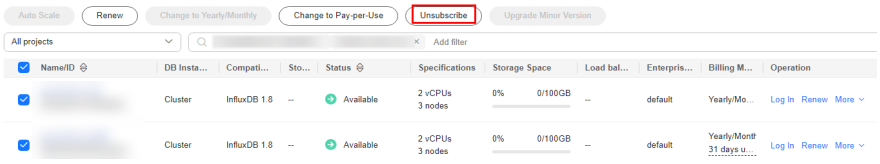

Step 4 In the displayed dialog box, click **Yes**.

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**.

NOTE

1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
2. To retain data, back it up 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 Influx API

3.1 Getting to Know GeminiDB Influx API

This section describes GeminiDB Influx instance types and instructs you to quickly create and connect to a GeminiDB Influx instance.

Table 3-1 Instance types

Instance Type	Scenario	Reference
Cluster	One cluster consists of at least three nodes. Nodes can be added to a cluster, which is unsuitable to cope with the ever growing data volume.	**Buying and Connecting to an Instance**
(Recommended) Cluster (performance-enhanced)	Compared with cluster instances, instances in a performance-enhanced cluster support a larger scale and higher read/write performance.	**Buying and Connecting to an Instance**
Single node	A single-node instance cannot ensure the SLA. You are advised to use it only for testing and function verification.	**Buying and Connecting to a Single-Node Instance**

Connection Methods

DAS enables you to manage instances on a web-based console, simplifying database management and improving working efficiency. You can connect and

manage instances through DAS. By default, you have the permission of remote login. DAS is secure and convenient for connecting to GeminiDB Influx instances.

Table 3-2 Connection on DAS

Method	Scenario	Remarks
DAS	You can connect to a GeminiDB Influx instance on a web-based console.	• Easy to use, secure, advanced, and intelligent • By default, you have the permission of remote login. DAS is secure and convenient for connecting to instances.

More Connection Operations

- See **Connecting to a GeminiDB Influx Instance**.

3.2 Buying and Connecting to a Cluster or Cluster (Performance-Enhanced) Instance

This section describes how to buy and connect to a GeminiDB Influx cluster or cluster (performance-enhanced) instance on the GeminiDB console.

- Cluster: One cluster consists of at least three nodes. Nodes can be added to a cluster, which is unsuitable to cope with the ever-growing data volume.
- (Recommended) Cluster (performance-enhanced): Compared with cluster instances, instances in a performance-enhanced cluster support a larger scale and higher read/write performance.

Each tenant can create a maximum of 50 GeminiDB Influx instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

- **Step 1: Buying a Cluster or Cluster (Performance-Enhanced) Instance**
- **Step 2: Connecting to an Instance Through DAS**

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

Step 1: Buying a Cluster or Cluster (Performance-Enhanced) Instance

For details, see **Buying a GeminiDB Influx Cluster Instance** or **(Recommended) Buying a GeminiDB Influx Cluster (Performance-Enhanced) 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 specifications, and click **Next**.

The following parameters are for reference only. Select proper specifications as needed. **Table 4-8** lists details about the parameters.

Figure 3-1 Billing mode and basic information of a cluster instance

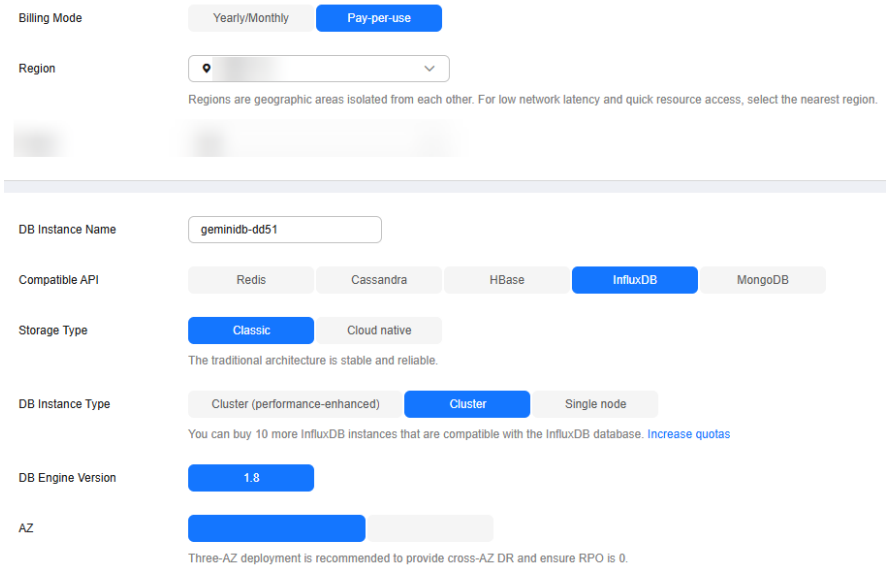

Figure 3-2 Billing mode and basic information of a cluster (performance-enhanced) instance

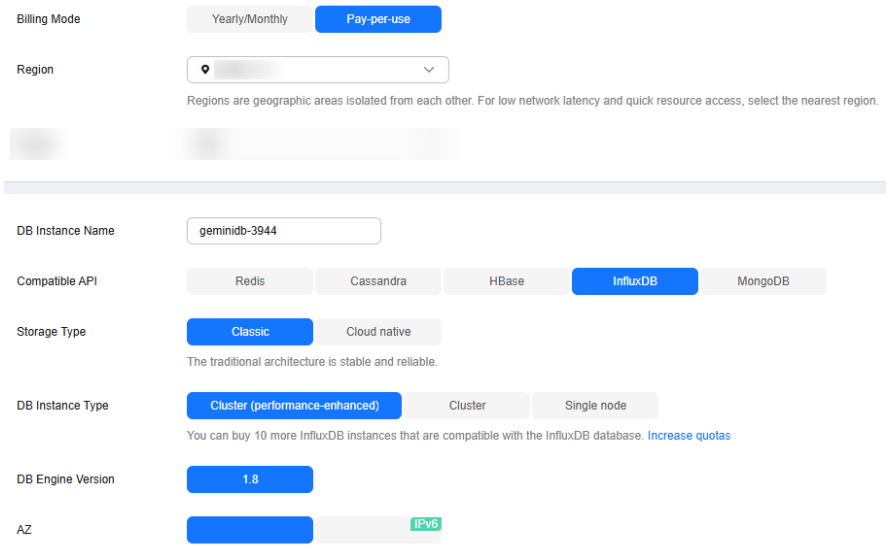

Parameter	Example Value	Description
Billing Mode	Pay-per-use	Billing mode of an instance • **Yearly/Monthly** is a prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. The longer the subscription term, the bigger the discount. This mode is a good option for long-term stable services. • **Pay-per-use** is a postpaid billing mode. You are billed based on how long you have actually used GeminiDB in 1-second increments and settled by the hour. This mode allows you to adjust resource usage easily. You neither need to prepare for resources in advance, nor end up with excessive or insufficient preset resources.
Region	Select **CN-Hong Kong**.	The region where the tenant is located. It can be changed in the upper left corner. **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 through a private network. After you buy an instance, you cannot change its region. • Cluster (performance-enhanced) instances are available only in the following region: AP-Bangkok
DB Instance Name	User-defined	The instance name: • Can be the same as an existing instance name. • Can contain 4 to 64 characters and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).
Compatible API	InfluxDB	GeminiDB is compatible with mainstream NoSQL APIs, including Redis, DynamoDB, Cassandra, HBase, MongoDB, and InfluxDB. You can select GeminiDB APIs by following **How Do I Select an API?**

Parameter	Example Value	Description
Storage Type	Classic	• **Classic**: classic architecture with decoupled storage and compute • **Cloud native**: more flexible, new-gen version with support for more AZs **NOTE** • Cloud native storage is only available for cluster (performance-enhanced) instances. • 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.
DB Instance Type	Cluster	• One cluster consists of at least three nodes. Nodes can be added to a cluster, which is unsuitable to cope with the ever growing data volume. • Compared with cluster instances, instances in a performance-enhanced cluster support a larger scale and higher read/write performance.
DB Engine Version	1.8	• If **Storage Type** is set to **Classic**, the version is fixed at **1.8**. • If **Storage Type** is set to **Cloud native**, the version is fixed at **1.7**.
AZ	AZ 1, AZ 2, and AZ 3	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. Instances can be deployed in a single AZ or three AZs. • To deploy instances in a single AZ, select one AZ. • If you want to deploy your instance across AZs for disaster recovery, select three AZs. Nodes of the instance are evenly distributed across the three AZs.

Figure 3-3 Storage and specifications

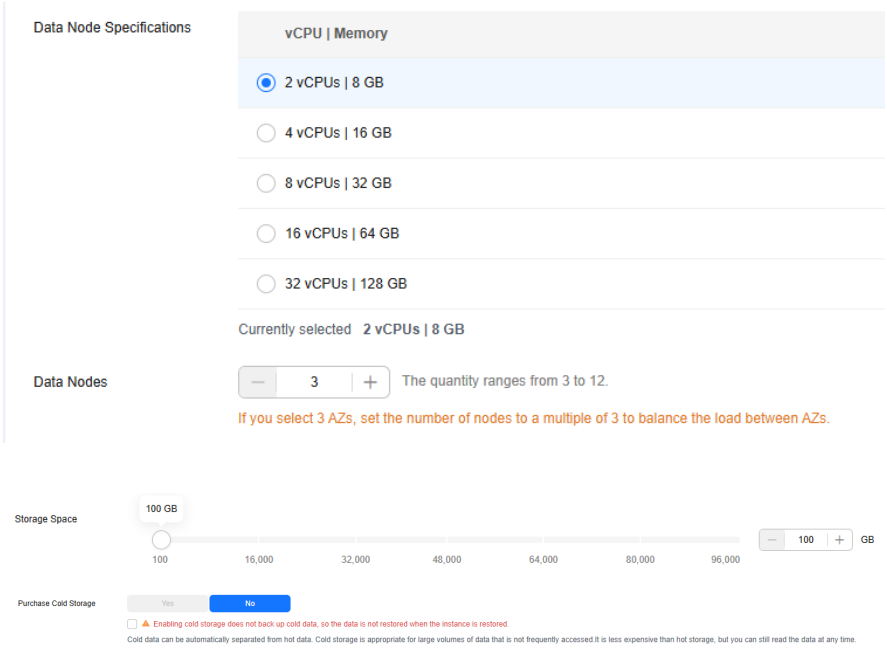

Parameter	Example Value	Description
Data Node Specifications	2U8GB	Data nodes provide read and write capabilities for time series databases. The specifications depend on configurations of the DFV shared resource pool and memory. Select specifications based on service requirements. For details about supported specifications, see **Instance Specifications**.
Data Nodes	3	Select the number of data nodes based on service requirements. After an instance is created, you can add nodes. For details, see **Adding Instance Nodes**. Currently, a maximum of 12 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.
Storage Space	100 GB	The storage is an integer and the minimum storage is 100 GB. You can add a minimum of 1 GB at a time.

Parameter	Example Value	Description
Purchase Cold Storage	**No**	Do not purchase cold storage. If you do not enable cold storage when creating an instance, you can enable it later based on service requirements. For details, see **Enabling Cold Storage**. **NOTE** Cold storage cannot be disabled after being enabled.

Figure 3-4 Network and database configurations

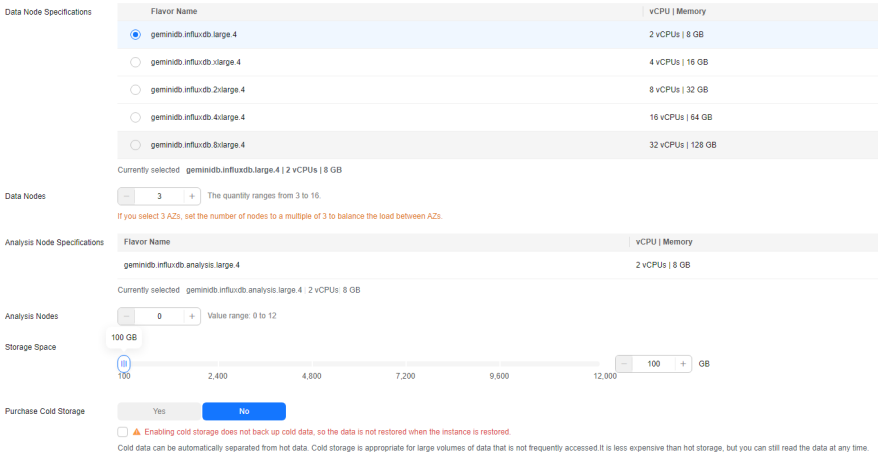

Parameter	Example Value	Description
VPC	default_vpc	Virtual private network where your instances are located. A VPC isolates networks for different services. You can select an existing VPC or create a VPC. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • To connect a GeminiDB Influx 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.

Parameter	Example Value	Description
Subnet	default_subnet	A subnet provides dedicated network resources that are logically isolated from other networks for security purposes.
Security Group	default	A security group controls access between GeminiDB Influx 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.
Administrator Password	Configured based on the password policy	Password of the administrator account. The password: • Can contain 8 to 32 characters. • Can include uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+? • 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.
Parameter Template	Default-InfluxDB-1.8	A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances. After an instance is created, you can modify its parameters to better meet your service requirements. For details, see **Modifying Parameters of GeminiDB Influx Instances**.

Parameter	Example Value	Description
Enterprise Project	default	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 ***Enterprise Management User Guide***.

Retain the default values for other parameters.

5. On the order confirmation page, check the instance information. If you need to modify the information, click **Previous**. If no modification is required, read and agree to the service agreement and click **Submit**.
6. Click **Back to Instance Management** to go to the instance list.
7. On the **Instances** page, view and manage the created instance.
   - Creating an instance takes about 5 to 9 minutes. During the process, the instance status displayed in the DB instance list is **Creating**.
   - After the instance is created, its status becomes **Available**.

Figure 3-5 Available cluster instance

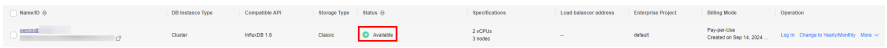

Figure 3-6 Available cluster (performance-enhanced) instance

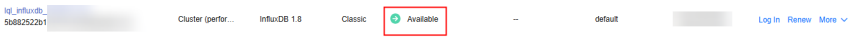

Step 2: Connecting to an Instance Through DAS

1. Log in to the Huawei Cloud console.
2. In the service list, choose **Databases** > **GeminiDB**.
3. In the instance list, locate a target instance and click **Log In** in the **Operation** column.

Figure 3-7 Connecting to a GeminiDB Influx Instance

Alternatively, click the instance name to go to the **Basic Information** page. Click **Log In** in the upper right corner of the page.

Figure 3-8 Connecting to a GeminiDB Influx Instance

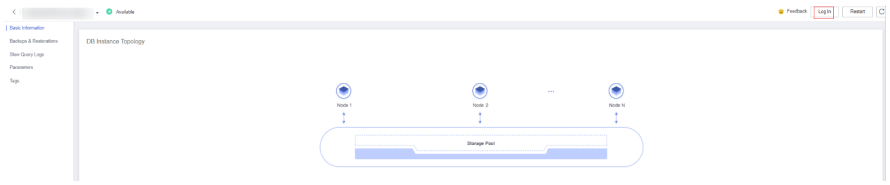

4. Enter a password for logging in to the instance.

Figure 3-9 Logging in to the GeminiDB Influx instance

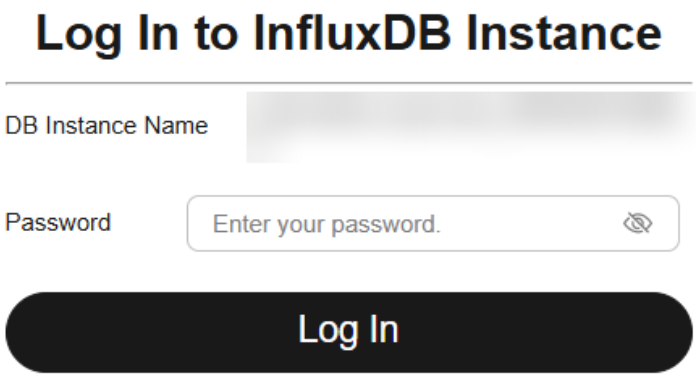

If you need to log in again after the password is reset, click **Re-login** in the upper right corner and use the new password.

Figure 3-10 Re-login

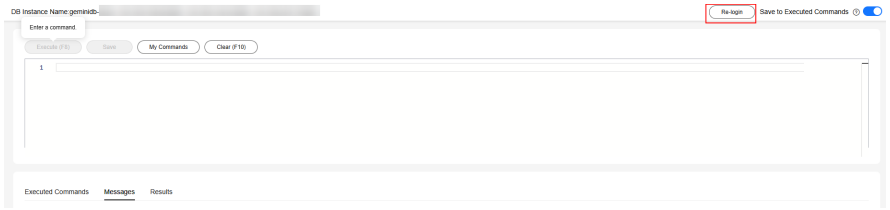

5. Manage relevant databases.

Figure 3-11 Instance homepage

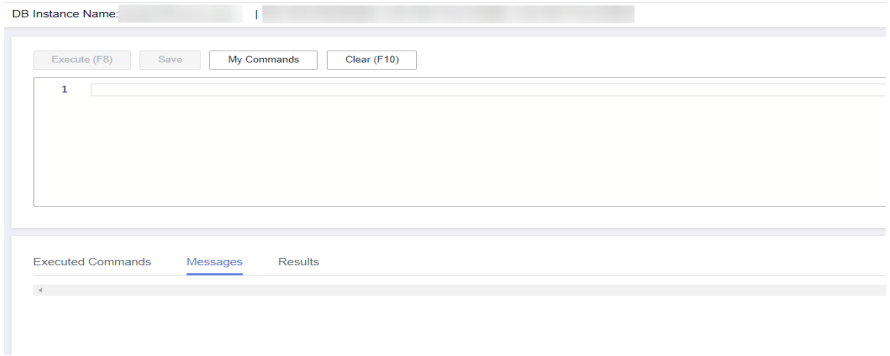

- Save commands to the execution record.

  This function is enabled by default to save the recently executed commands for your later query.

  Then you can click the **Executed Commands** tab on the lower page to view historical commands.

  > **NOTE**
  >
  > Commands with passwords are not displayed on the **Executed Commands** tab page.

  Figure 3-12 Executed commands

  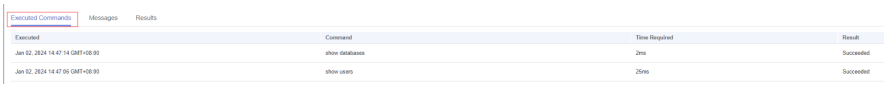

  If this function is disabled, the commands executed subsequently are not displayed. You can click next to **Save Executed SQL Statements** in the upper right corner to disable this function.

- Execute a command.

  Enter a command in the command window and click **Execute** or **F8**.

  Figure 3-13 Execute a command.

  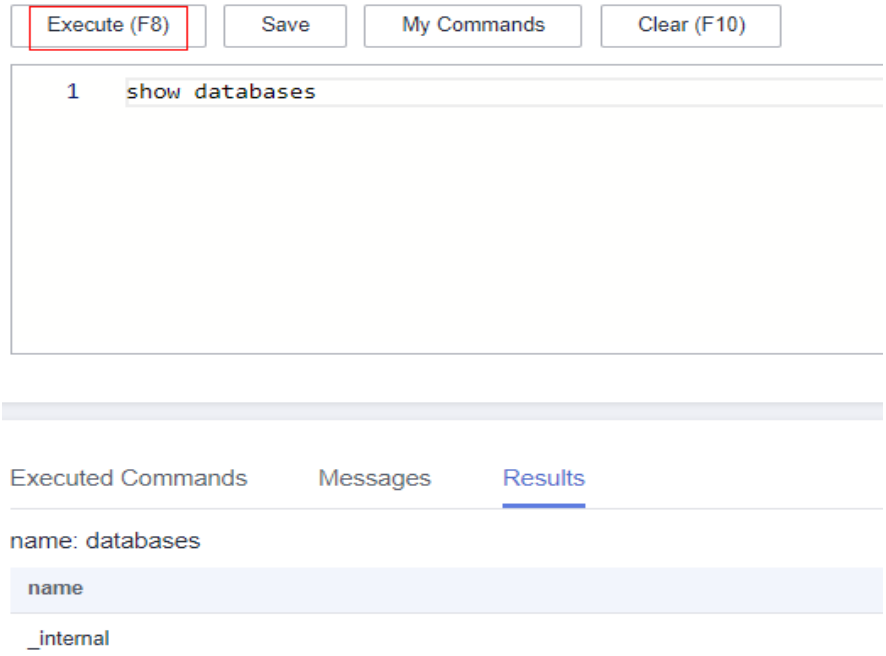

  After a command is executed, you can view the execution result on the **Results** page.

- Save a command.

  You can save a command to all instances or the current instance. Then you can view details in **My Commands**.

NOTE

Commands with passwords cannot be saved to **My Commands**.

Figure 3-14 Save a command.

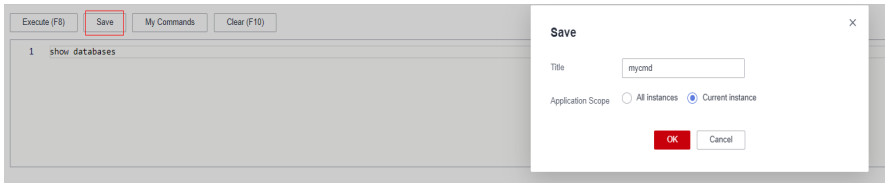

- View my commands.

  Common commands are displayed the **My Commands** page.

  You can set a filter to narrow the scope of commands. If you select **All**, all commands saved in the current account are displayed.

Figure 3-15 Filtering commands

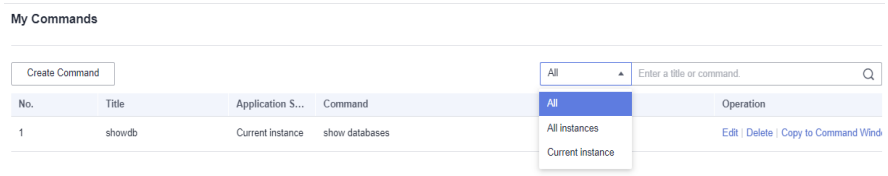

Alternatively, you can enter a command title or statement in the search box to search for the corresponding command.

Figure 3-16 Searching for a command

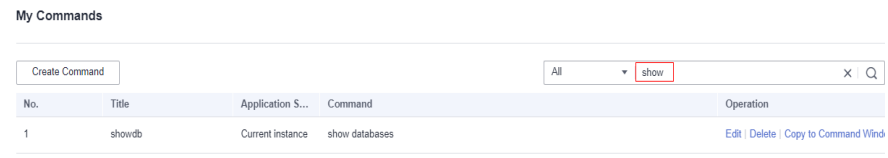

On the **My Commands** page, you can also create, edit, and delete a command or copy it to the command window.

Figure 3-17 Managing a command

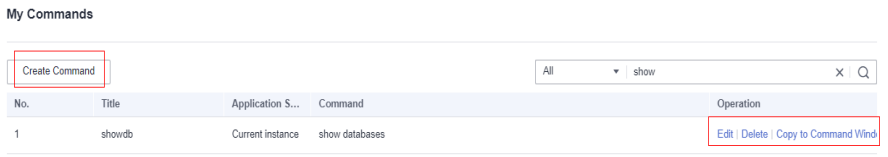

- Clear a command.

  You can also press **F10** to clear the command in the command window.

FAQs

Question: What should I do if the DAS console cannot be redirected after I click **Log In** in the **Operation** column in the instance list or click **Log In** on the **Basic Information** page?

Solution: Set your browser to allow pop-ups and try again.

3.3 Buying and Connecting to a Single-Node Instance

This section describes how to buy and connect to a single-node GeminiDB Influx instance on the GeminiDB console.

A single-node instance cannot ensure the SLA. You are advised to use it only for testing and function verification.

Each tenant can create a maximum of 50 GeminiDB Influx instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

- **Step 1: Buying a Single-Node Instance**
- **Step 2: Connecting to an Instance Through DAS**

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

Step 1: Buying a Single-Node Instance

For details, see **Buying a Single-Node GeminiDB Influx 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 specifications, and click **Next**.

   The following parameters are for reference only. Select proper specifications as needed. **Table 4-15** lists details about the parameters.

   Figure 3-18 Billing mode and basic information

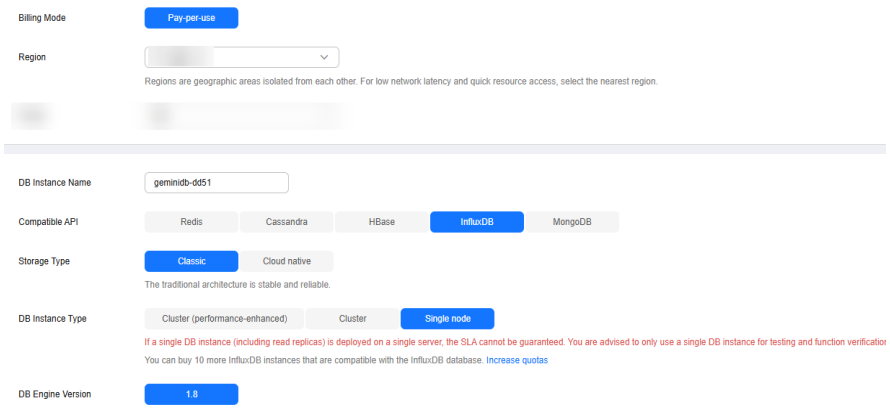

Parameter	Example Value	Description
Billing Mode	Pay-per-use	Billing mode of an instance • **Yearly/Monthly**: A prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. The longer the subscription term, the bigger the discount. This mode is a good option for long-term stable services. • **Pay-per-use** is a postpaid mode. You are billed based on how long you have actually used GeminiDB. Pay-per-use instance usage is calculated by the second and priced by the hour. This mode allows you to adjust resource usage easily. You neither need to prepare for resources in advance, nor end up with excessive or insufficient preset resources.
Region	Select **CN-Hong Kong**.	The region where the tenant is located. It can be changed in the upper left corner. **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	User-defined	The instance name: • Can be the same as an existing instance name. • Can contain 4 to 64 characters and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_).
Compatible API	InfluxDB	GeminiDB is compatible with mainstream NoSQL APIs, including Redis, DynamoDB, Cassandra, HBase, MongoDB, and InfluxDB. You can select GeminiDB APIs by following **How Do I Select an API?**

Parameter	Example Value	Description
Storage Type	Classic	• **Classic**: classic architecture with decoupled storage and compute • **Cloud native**: more flexible, new-gen version with support for more AZs **NOTE** • Cloud native storage is only available for cluster (performance-enhanced) instances. • 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.
DB Instance Type	Single node	A single-node instance cannot ensure the SLA. You are advised to use it only for testing and function verification.
DB Engine Version	1.8	1.8
AZ	AZ 1, AZ 2, and AZ 3	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. A single-node instance can be deployed in one AZ.

Figure 3-19 Specifications and storage

Parameter	Example Value	Description
Instance Specifications	2U8GB	Data nodes provide read and write capabilities for time series databases. The specifications depend on configurations of the DFV shared resource pool and memory. Select specifications based on service requirements. For details about supported specifications, see **Instance Specifications**.
Nodes	1	A single-node instance can have only one node.
Storage Space	100 GB	The minimum storage space is 100 GB and must be an integer. You can add at least 1 GB each time you scale up storage space.
Purchase Cold Storage	No	Do not purchase cold storage. If you disable cold storage when creating an instance, you can enable it later based on service requirements. For details, see **Enabling Cold Storage**. **NOTE** Cold storage cannot be disabled after being enabled.

Figure 3-20 Network and database configurations

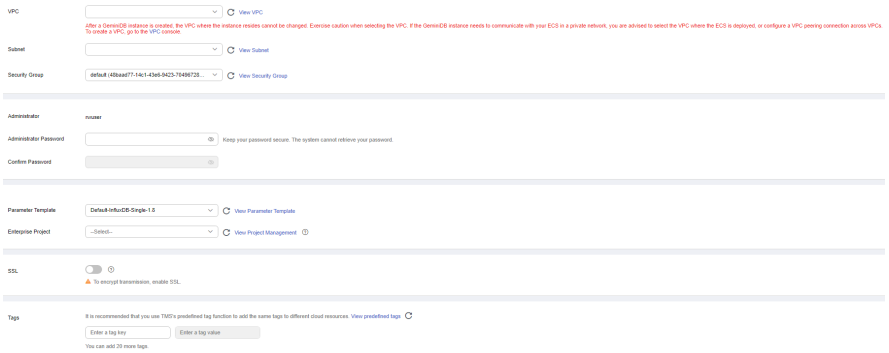

Parameter	Example Value	Description
VPC	default_vpc	Virtual private network where your instances are located. A VPC isolates networks for different services. You can select an existing VPC or create a VPC. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • To connect a GeminiDB Influx 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	default_subnet	A subnet provides dedicated network resources that are logically isolated from other networks for security purposes.
Security group	default	A security group controls access between GeminiDB Influx 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.

Parameter	Example Value	Description
Administrator Password	Configured based on the password policy	Password of the administrator account. The password: • Can include 8 to 32 characters. • Can include uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+? • 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.
Parameter Template	Default-InfluxDB-1.8	A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances. After an instance is created, you can modify its parameters to better meet your service requirements. For details, see **Modifying Parameters of GeminiDB Influx Instances**.
Enterprise Project	default	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 ***Enterprise Management User Guide***.

Retain the default values for other parameters.

5. On the displayed page, confirm instance details. To modify the configurations, click **Previous**.

If you do not need to modify the settings, read and agree to the service agreement and click **Submit**.

6. Click **Back to Instance Management** to go to the instance list.
7. On the **Instances** page, view and manage the created instance.
   - Creating an instance takes about 5 to 9 minutes. During the process, the instance status displayed in the instance list is **Creating**.
   - After the instance is created, its status becomes **Available**.

Figure 3-21 Available instance

Step 2: Connecting to an Instance Through DAS

1. Log in to the Huawei Cloud console.
2. In the service list, choose **Databases** > **GeminiDB**.
3. In the instance list, locate a target instance and click **Log In** in the **Operation** column.

Figure 3-22 Connecting to a GeminiDB Influx Instance

Alternatively, click the instance name to go to the **Basic Information** page. Click **Log In** in the upper right corner of the page.

Figure 3-23 Connecting to a GeminiDB Influx Instance

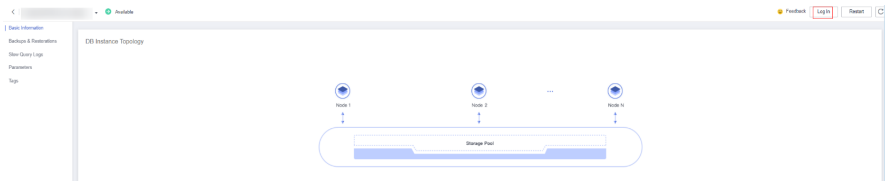

4. Enter a password for logging in to the instance.

Figure 3-24 Logging in to the GeminiDB Influx instance

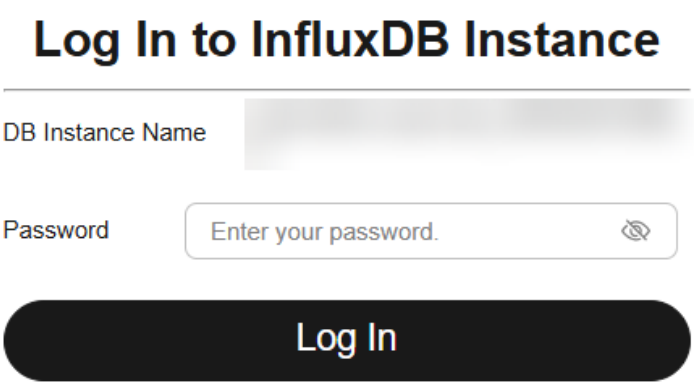

If you need to log in again after the password is reset, click **Re-login** in the upper right corner and use the new password.

Figure 3-25 Re-login

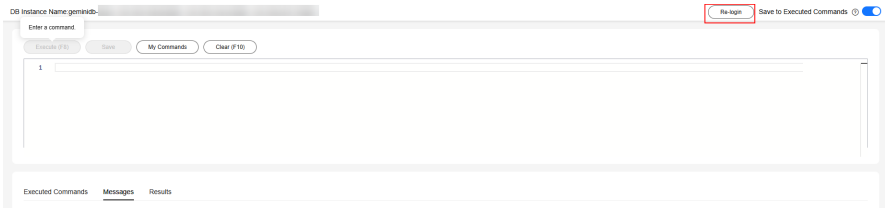

5. Manage relevant databases.

Figure 3-26 Instance homepage

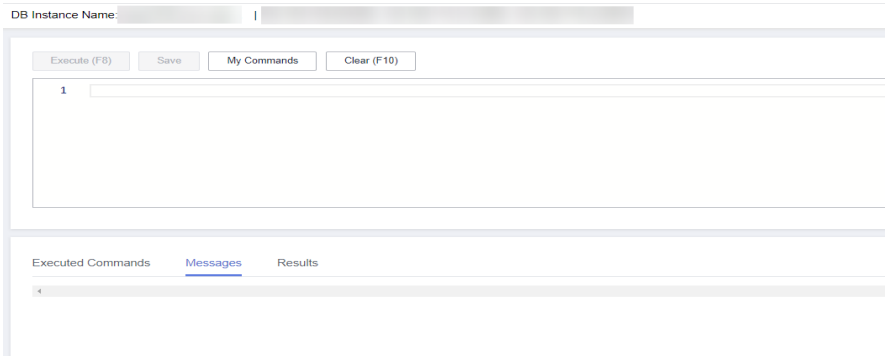

- Save commands to the execution record.

  This function is enabled by default to save the recently executed commands for your later query.

  Then you can click the **Executed Commands** tab on the lower page to view historical commands.

  NOTE

  Commands with passwords are not displayed on the **Executed Commands** tab page.

  Figure 3-27 Executed commands

  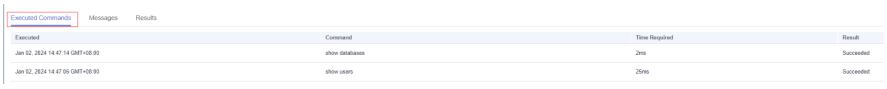

  If this function is disabled, the commands executed subsequently are not displayed. You can click next to **Save Executed SQL Statements** in the upper right corner to disable this function.

- Execute a command.

  Enter a command in the command window and click **Execute** or **F8**.

Figure 3-28 Executing a command

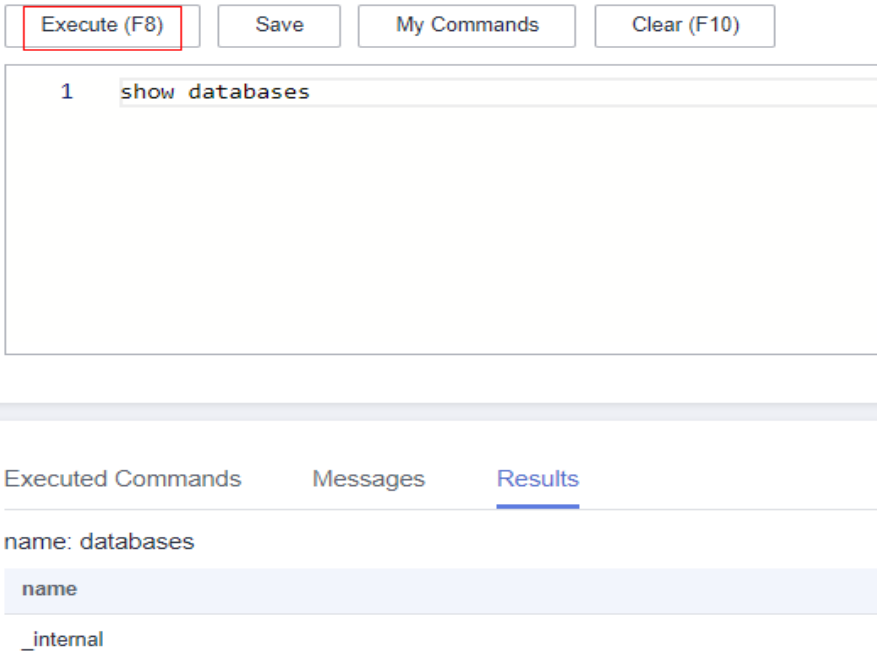

After a command is executed, you can view the execution result on the **Results** page.

- Save a command.

  You can save a command to all instances or the current instance. Then you can view details in **My Commands**.

  NOTE

  Commands with passwords cannot be saved to **My Commands**.

Figure 3-29 Saving a command

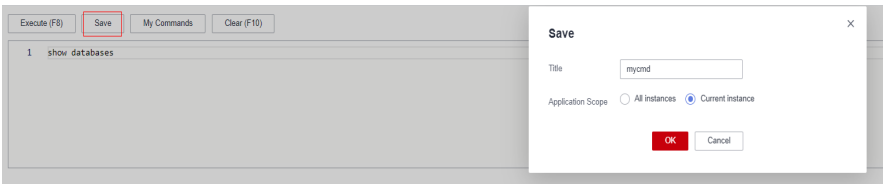

- View my commands.

  Common commands are displayed the **My Commands** page.

  You can set a filter to narrow the scope of commands. If you select **All**, all commands saved in the current account are displayed.

Figure 3-30 Filtering commands

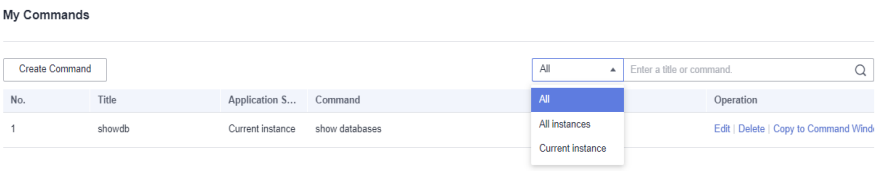

Alternatively, you can enter a command title or statement in the search box to search for the corresponding command.

Figure 3-31 Searching for a command

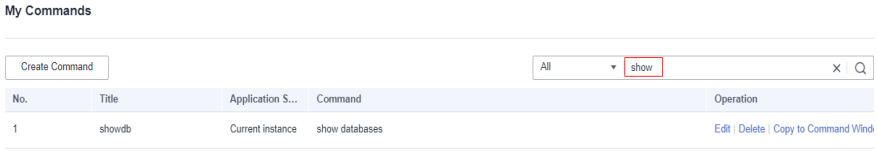

On the **My Commands** page, you can also create, edit, and delete a command or copy it to the command window.

Figure 3-32 Managing a command

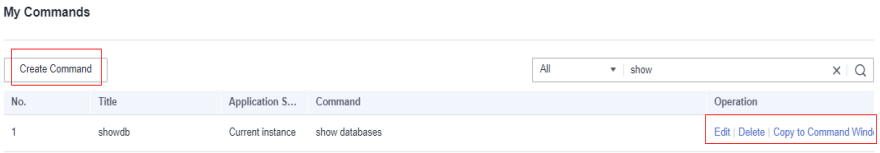

- Clear a command.

  You can also press **F10** to clear the command in the command window.

FAQs

Question: What should I do if the DAS console cannot be redirected after I click **Log In** in the **Operation** column in the instance list or click **Log In** on the **Basic Information** page?

Solution: Set your browser to allow pop-ups and try again.

3.4 Getting Started with Common Practices

After purchasing and connecting to a GeminiDB Influx instance, you can view common practices to better use the instance.

Table 3-3 Common practices

Practice		Description
Usage rules	Usage Specifications and Suggestions	Describes rules and suggestions for using GeminiDB Influx instances in the aspects of naming, TAG, FIELD, and query to solve common problems such as incorrect usage, low efficiency, and difficult maintenance.

Practice		Description
Instance modifications	**Changing a GeminiDB Influx Instance Name**	Describes how to change the name of a GeminiDB Influx instance to help you identify different instances.
	Changing the Administrator Password of a GeminiDB Influx Database	Describes how to change your administrator password. For security reasons, change it periodically.
	Changing vCPUs and Memory	Describes how to change the CPU or memory of your instance to suit your service requirements.
Data backup	**Managing Automated Backups**	This practice describes how GeminiDB Influx API automatically creates backups for an instance during a backup window and saves the backups based on the configured retention period.
	Managing Manual Backups	Describes how to create manual backups for a DB instance. These backups can be used to restore data for improved reliability.
Data restoration	**Restoring Data to a New Instance**	Describes how to restore an existing automated or manual backup to a new instance. The restored data is the same as the backup data.

4 Working with GeminiDB Influx API

4.1 Using IAM to Grant Access to GeminiDB Influx API

4.1.1 Creating a User Group and Assigning Permissions

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 your Huawei Cloud account does not require individual IAM users, skip this section.

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

Prerequisites

Learn about the permissions supported by GeminiDB and choose policies or roles based on your requirements. For details about the permissions, see **Permissions Management**. For system policies of other services, see **System Permissions**.

Process Flow

Figure 4-1 Process of granting GeminiDB permissions

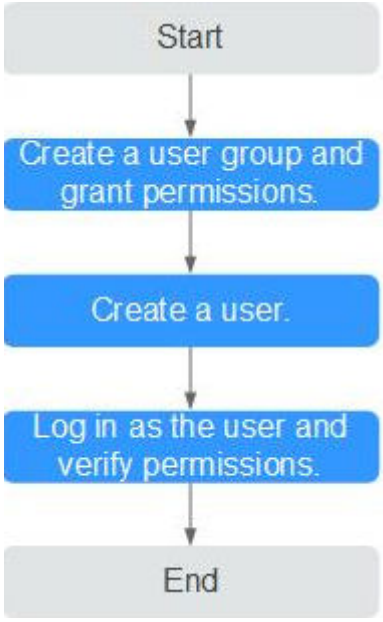

1. **Create a user group and assign permissions** to it.

   Create a user group on the IAM console and attach the **GeminiDB FullAccess** policy to the group.

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 that the user only has read 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

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

Example Custom Policy

- Example 1: Allowing users to create GeminiDB instances

```
{
   "Version": "1.1",
   "Statement": [
      {
         "Effect": "Allow",
         "Action": [
            "nosql:instance:create"
         ]
      }
   ]
}
```

- Example 2: Deny users the permission 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:

```
{
   "Version": "1.1",
   "Statement": [
      {
        "Effect": "Deny"
        "Action": [
            "nosql:instance:delete"
         ],
        }
   ]
}
```

- 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:

```
{
      "Version": "1.1",
      "Statement": [
            {
                  "Action": [
                        "nosql:instance:create",
                        "nosql:instance:rename",
                        "nosql:instance:delete",
                        "vpc:publicIps:list",
                        "vpc:publicIps:update"
                  ],
                  "Effect": "Allow"
            }
      ]
}
```


4.2 Buying a GeminiDB Influx Instance

4.2.1 (Recommended) Buying a GeminiDB Influx Cluster (Performance-Enhanced) Instance

This section describes how to buy a GeminiDB Influx instance on the GeminiDB console.

Each tenant can create a maximum of 50 GeminiDB Influx instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Prerequisites

- You have created a Huawei Cloud account.

Procedure

Step 1 Log in to the Huawei Cloud console.

Step 2 Choose **Databases** > **GeminiDB** from the service list.

Step 3 On the **Instances** page, click **Buy DB Instance**.

Step 4 On the displayed page, select a billing mode, select instance specifications and click **Next**.

Figure 4-2 Billing mode and basic information

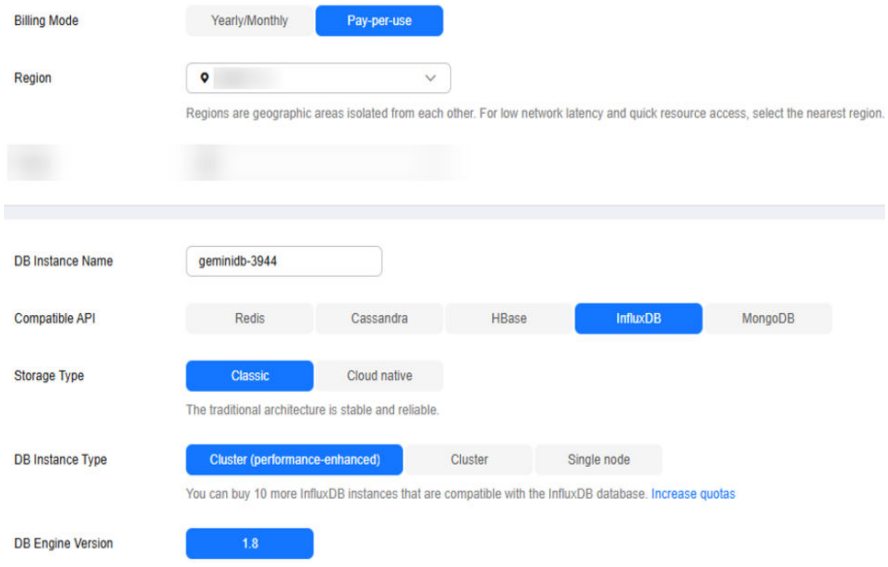

Table 4-1 Billing Mode

Parameter	Description
Billing Mode	• **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. It can be changed in the upper left corner. **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 through a private network. After you buy an instance, you cannot change its region. • Cluster (performance-enhanced) instances are available only in the following region: AP-Bangkok

Parameter	Description
DB Instance Name	The instance name: • Can be the same as an existing instance name. • Can contain 4 to 64 characters and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_). You can change the name of an instance after it is created. For details, see **Changing a GeminiDB Influx Instance Name**.
Compatible API	InfluxDB GeminiDB is compatible with mainstream NoSQL APIs, 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** – Cloud native storage is only available for cluster (performance-enhanced) instances. – 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.
DB Instance Type	Cluster (performance-enhanced) Compared with cluster instances, instances in a performance-enhanced cluster support a larger scale and higher read/write performance.
Version	• If **Storage Type** is set to **Classic**, the version is fixed at **1.8**. • If **Storage Type** is set to **Cloud native**, the version is fixed at **1.7**.
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. Instances can be deployed in a single AZ or three AZs. • To deploy instances in a single AZ, select one AZ. • To deploy instances across AZs for disaster recovery, select three AZs, where the instance nodes will be evenly distributed.

Figure 4-3 Specifications and storage

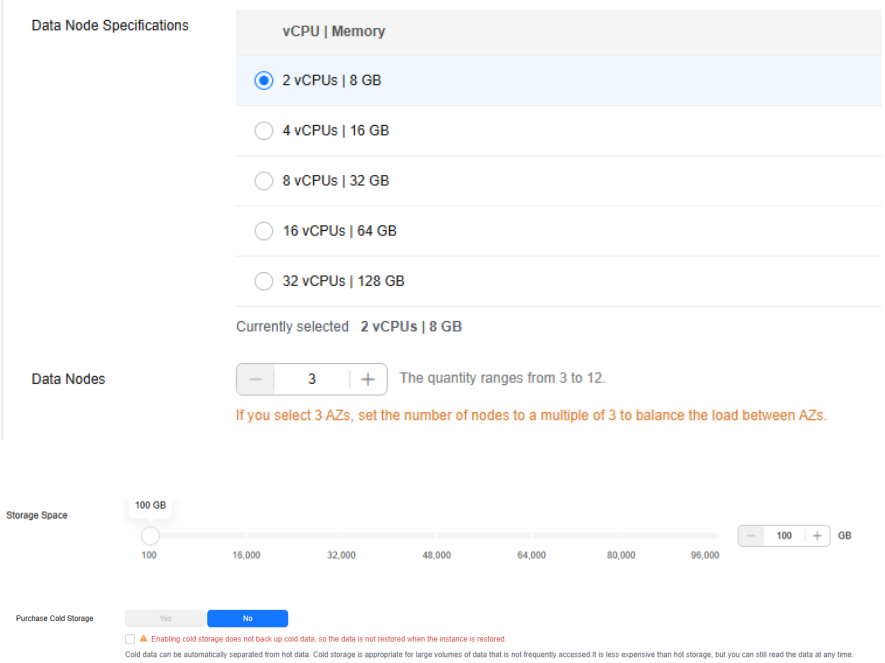

Table 4-3 Specifications and storage

Parameter	Description
Data Node Specifications	Data nodes provide read and write capabilities for time series databases. The specifications depend on configurations of the DFV shared resource pool and memory. Select specifications based on service requirements. For details about supported specifications, see **Instance Specifications**.
Data Nodes	Select the number of data nodes based on service requirements. After an instance is created, you can add nodes. For details, see **Adding Instance Nodes**.
Storage Space	The storage is an integer, and the minimum storage is 100 GB. You can add at least 1 GB each time.
Auto Scale	You can determine whether to enable it based on the site requirements. 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. • **Trigger If Available Storage Drops To**: When the available storage drops to or below the specified threshold, autoscaling will be triggered. The value can be **10%**, **15%**, and **20%**. • **Increase By**: percentage that your instance storage will be scaled up at. The value can be **10%**, **15%**, or **20%**. • **Storage Limit**: limit of storage that can be automatically scaled up to

Parameter	Description
Purchase Cold Storage	Cold storage stores historical data that is not frequently queried. When purchasing a GeminiDB Influx instance, you can purchase cold storage and configure the retention policy to specify the retention period of hot data. In this way, hot data will be automatically archived in cold storage after the retention period expires, reducing storage costs. The value can be: • **Yes** Set the cold storage capacity to suit your service requirements. • **No** You do not want to purchase cold storage. For more information about cold and hot data separation, see **Cold and Hot Data Separation**. If you do not enable cold storage when creating an instance, you can enable it later based on service requirements. For details, see **Enabling Cold Storage**. **NOTE** Cold storage cannot be disabled after being enabled.
Cold Storage	The cold storage space is an integer from 500 GB to 1,024,000 GB. You can add at least a multiple of 10 GB each time. After an instance is created, you can scale up its cold storage. For details, see **Scaling Up Cold Storage**.
Disk Encryption	You need to set this parameter if **Storage Type** is set to **Cloud native**. • **Disable**: Data is not encrypted. • **Enable**: Your data will be encrypted on disks and stored in ciphertext after you create an instance. **Key Name**: Select an existing key or create one. **NOTE** – 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. – This parameter is only available for instances with cloud native storage. – 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 Key**" in *Data Encryption Workshop User Guide*.

Figure 4-4 Network and database configurations

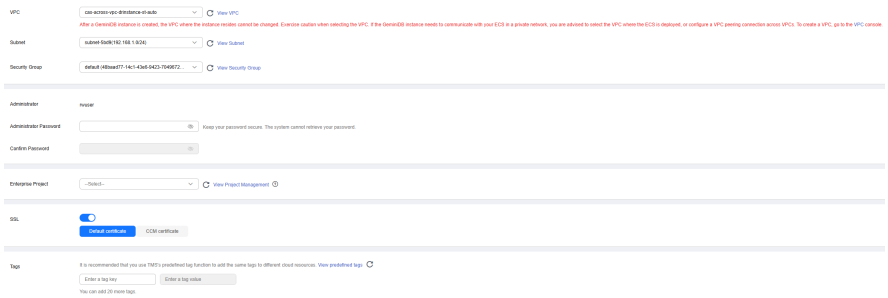

Table 4-4 Network

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. For details about how to create a VPC, see "Creating a VPC" in *Virtual Private Cloud User Guide*. If there are no VPCs available, the system allocates resources to you by default. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • To connect a GeminiDB Influx 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 GeminiDB Influx instances and other services. When you select a security group, you must ensure that it allows your client to access your instances. If there are no security groups available, the system allocates resources to you by default.

Table 4-5 Database configuration

Parameter	Description
Administrator	The default administrator account is **rwuser**.

Parameter	Description
Administrator Password	Password of the administrator account. The password: • Must be 8 to 32 characters long. • The password must contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+? • 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 ***Enterprise Management User Guide***.

Table 4-6 Tags

Parameter	Description
Tag	Tags a GeminiDB Influx instance. This parameter is optional. Adding tags helps you better identify and manage your instances. A maximum of 20 tags can be added for each instance. 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 1 to 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: _.-:+=@/ After a DB instance is created, you can view its tag details on the **Tags** tab. In addition, you can add, modify, and delete tags of an existing instance. For details, see **Managing Tags**.

Table 4-7 Required duration

Parameter	Description
Required duration	The length of your subscription if you select **Yearly/Monthly** billing. Subscription lengths range from one month to three years.
Auto-renewing an Instance	• By default, this option is not selected. • If you select this option, the auto-renew cycle is determined by the selected required duration.

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

- Yearly/Monthly
  - To modify the configurations, click **Previous**.
  - If no modification is required, read and agree to the service agreement, click **Pay Now**, and complete the payment.
- Pay-per-use
  - To modify the configurations, click **Previous**.
  - If no modification is required, read and agree to the service agreement and click **Submit**.

Step 6 Click **Back to Instance Management** to go to the instance list.

Step 7 On the **Instances** page, view and manage your instances.

- Creating an instance takes about 5 to 9 minutes. During the process, the instance status becomes **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. A full backup is automatically triggered after a DB instance is created.
- The default database port of the instance is **8635** and cannot be changed.

Figure 4-5 Available instance

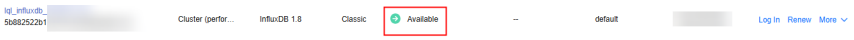

----End

4.2.2 Buying a GeminiDB Influx Cluster Instance

This section describes how to buy a GeminiDB Influx cluster instance on the GeminiDB console.

Each tenant can create a maximum of 50 GeminiDB Influx instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Prerequisites

- You have created a Huawei Cloud account.

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 **Buy DB Instance**.

Step 4 On the displayed page, specify a billing mode and instance specifications and click **Next**.

Figure 4-6 Billing mode and basic information

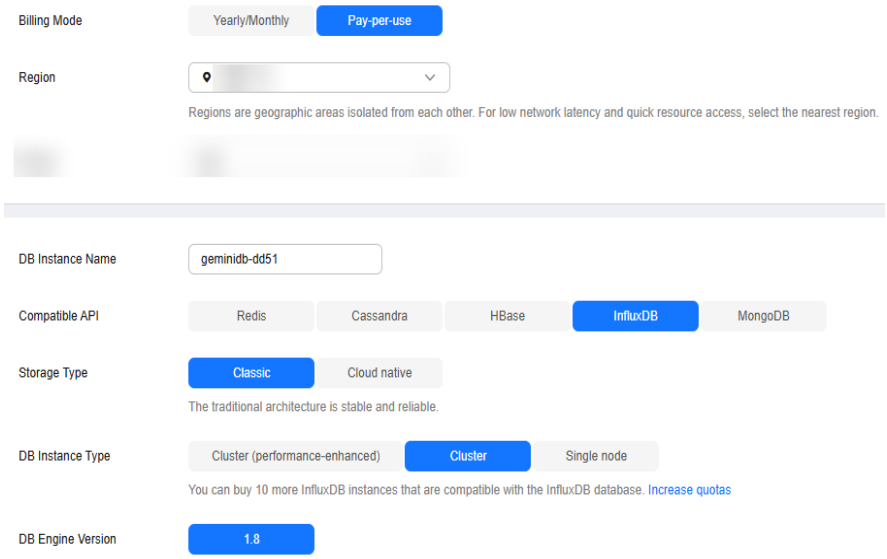

Table 4-8 Billing mode

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?**. • 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 4-9 Basic information

Parameter	Description
Region	The region where the tenant is located. It can be changed in the upper left corner. **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 contain 4 to 64 characters and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_). You can change the name of an instance after it is created. For details, see **Changing a GeminiDB Influx Instance Name**.
Compatible API	InfluxDB GeminiDB is compatible with mainstream NoSQL APIs, including Redis, DynamoDB, Cassandra, HBase, MongoDB, and InfluxDB. You can select GeminiDB APIs by following **How Do I Select an API?**

Parameter	Description
Storage Type	• **Classic**: classic architecture with decoupled storage and compute • **Cloud native**: more flexible, new-gen version with support for more AZs **NOTE** – Cloud native storage is only available for cluster (performance-enhanced) instances. – 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.
DB Instance Type	Cluster One cluster consists of at least three nodes. Nodes can be added to a cluster, which is unsuitable to cope with the ever-growing data volume.
DB Engine Version	1.8
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. Instances can be deployed in a single AZ or three AZs. • To deploy instances in a single AZ, select one AZ. • If you want to deploy your instance across AZs for disaster recovery, select three AZs. Nodes of the instance are evenly distributed across the three AZs.

Figure 4-7 Specifications and storage

Table 4-10 Specifications and storage

Parameter	Description
Data Node Specifications	Data nodes provide read and write capabilities for time series databases. The specifications depend on configurations of the DFV shared resource pool and memory. Select specifications based on service requirements. For details about supported specifications, see **Instance Specifications**.
Data Nodes	Select the number of data nodes based on service requirements. After an instance is created, you can add nodes. For details, see **Adding Instance Nodes**. Currently, a maximum of 12 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.
Storage Space	The storage is an integer and the minimum storage is 100 GB. You can add a minimum of 1 GB at a time.
Purchase Cold Storage	Cold storage is used to store historical data that is not frequently queried. When purchasing a GeminiDB Influx instance, you can purchase cold storage and configure the retention policy to specify the retention period of hot data. In this way, hot data will be automatically archived in cold storage after the retention period expires, reducing storage costs. The value can be: • **Yes** Set the cold storage capacity to suit your service requirements. • **No** Do not purchase cold storage. For more information about cold and hot data separation, see **Cold and Hot Data Separation**. If you do not enable cold storage when creating an instance, you can enable it later based on service requirements. For details, see **Enabling Cold Storage**. **NOTE** Cold storage cannot be disabled after being enabled.
Cold Storage	The cold storage is an integer from 500 GB to 100,000 GB. You can add a minimum of 1 GB each time you scale up storage space. After an instance is created, you can scale up its cold storage. For details, see **Scaling Up Cold Storage**.

Figure 4-8 Network and database configurations

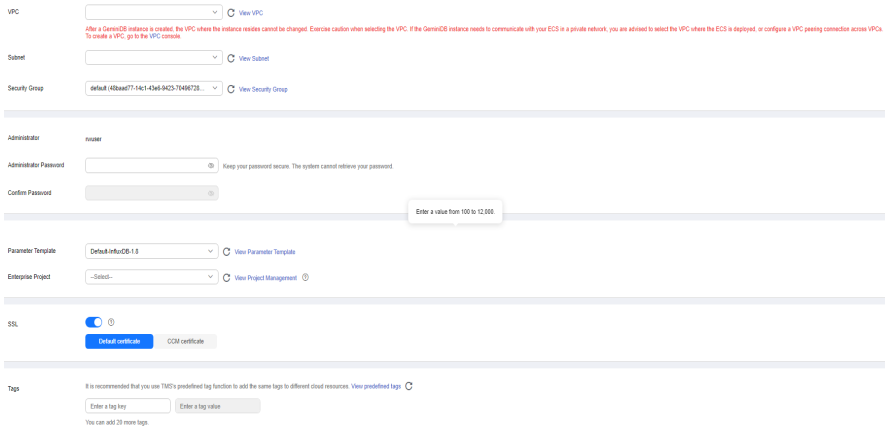

Table 4-11 Network configurations

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 one. For details about how to create a VPC, see "Creating a VPC" in *Virtual Private Cloud User Guide*. 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*. If there are no VPCs available, the system automatically allocates a VPC to you. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • To connect a GeminiDB Influx 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.

Parameter	Description
Security Group	A security group controls access between GeminiDB Influx 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.

Table 4-12 Database configurations

Parameter	Description
Administrator	Username of the administrator account. The default value is **rwuser**.
Administrator Password	Password of the administrator account. The password: • Can contain 8 to 32 characters. • Can include uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+? • 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 the administrator password.
Parameter Template	A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances. After an instance is created, you can modify its parameters to better meet your service requirements. For details, see **Modifying Parameters of GeminiDB Influx Instances**.
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 ***Enterprise Management User Guide***.

Parameter	Description
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 can enable SSL to improve data security. After an instance is created, you can connect to it using SSL. **Figure 4-9** Enabling SSL 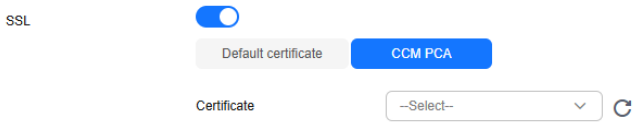 After SSL is enabled, you can select the default certificate or the certificate issued by the CCM service. **NOTE** • If SSL is not enabled when you create an instance, you can enable SSL after the instance is created. For details, see **Encrypting Data over SSL for a GeminiDB Influx Instance**. • For details about how to disable SSL, see **Encrypting Data over SSL for a GeminiDB Influx Instance**.

Table 4-13 Tags

Parameter	Description
Tags	The setting is optional. Adding tags helps you better identify and manage your instances. A maximum of 20 tags can be added for each instance. If your organization has configured a tag policy for a GeminiDB Influx instance, you need to 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 created. 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 will be tagged. Each tag key is unique for each instance. It can contain 1 to 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 will 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: _.-:+=@/ After an instance is created, you can view its tag details on the **Tags** tab. In addition, you can add, modify, and delete tags of an existing instance. For details, see **Managing Tags**.

Table 4-14 Required duration

Parameter	Description
Required Duration	The length of your subscription if you select **Yearly/Monthly** billing. Subscription lengths range from one month to three years.
Auto-renew	• This option is not selected by default. • If you select this option, the instance is automatically renewed based on the subscription duration.

Step 5 On the displayed page, confirm instance details.

- Yearly/Monthly
  - To modify the configurations, click **Previous**.
  - If no modification is required, read and agree to the service agreement, click **Pay Now**, and complete the payment.
- Pay-per-use
  - To modify the configurations, click **Previous**.

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

Step 6 Click **Back to Instance Management** to go to the instance list.

Step 7 On the **Instances** page, view and manage the created instance.

- Creating an instance takes about 5 to 9 minutes. During the process, the instance status displayed in the DB instance list is **Creating**.
- After the instance is created, its status becomes **Available**.

  You can click  in the upper right corner of the page to refresh the instance status.
- An automated backup policy is enabled by default during instance creation. A full backup is automatically triggered after an instance is created.
- The default database port of the instance is **8635** and cannot be changed.

Figure 4-10 Available instance

----End

4.2.3 Buying a Single-Node GeminiDB Influx Instance

This section describes how to buy a single-node GeminiDB Influx instance on the GeminiDB console.

Each tenant can create a maximum of 50 GeminiDB Influx instances by default. To request a higher quota, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Prerequisites

- You have created a Huawei Cloud account.

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 **Buy DB Instance**.

Step 4 On the displayed page, select a billing mode, select instance specifications and click **Next**.

Figure 4-11 Billing mode and basic information

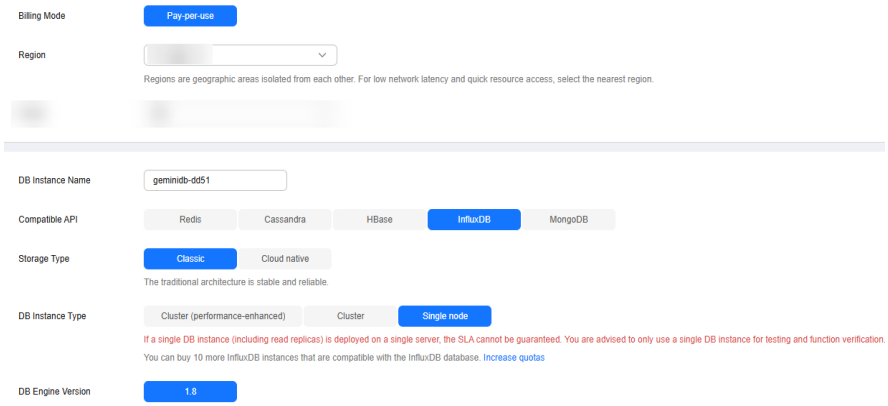

Table 4-15 Billing mode

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?**. • 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 4-16 Basic information

Parameter	Description
Region	The region where the tenant is located. It can be changed in the upper left corner. **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 through 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 contain 4 to 64 characters and must start with a letter. It is case-sensitive and allows only letters, digits, hyphens (-), and underscores (_). You can change the name of an instance after it is created. For details, see **Changing a GeminiDB Influx Instance Name**.
Compatible API	InfluxDB GeminiDB is compatible with mainstream NoSQL APIs, 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** – Cloud native storage is only available for cluster (performance-enhanced) instances. – 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.
DB Instance Type	Single node A single-node instance cannot ensure the SLA. You are advised to use it only for tests and function verification.
DB Engine Version	1.8
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. A single-node instance can be deployed in one AZ.

Figure 4-12 Specifications and storage

Table 4-17 Specifications and storage

Parameter	Description
Instance Specifications	The specifications depend on configurations of the DFV shared resource pool and memory. Select specifications based on service requirements. For details about supported specifications, see **Instance Specifications**.
Nodes	A single-node instance can have only one node.
Storage Space	The storage is an integer and the minimum storage is 100 GB. You can add a minimum of 1 GB at a time.

Figure 4-13 Network and database configurations

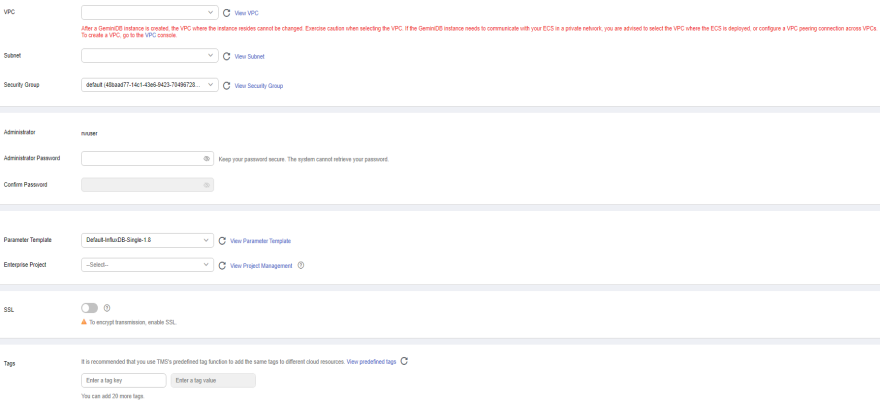

Table 4-18 Network configurations

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 one. For details about how to create a VPC, see "Creating a VPC" in *Virtual Private Cloud User Guide*. 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*. If there are no VPCs available, the system automatically allocates a VPC to you. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • To connect a GeminiDB Influx 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 GeminiDB Influx 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.

Table 4-19 Database configurations

Parameter	Description
Administrator	Username of the administrator account. The default value is **rwuser**.

Parameter	Description
Administrator Password	Password of the administrator account. The password: • Can include 8 to 32 characters. • Can include uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+? • 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 the administrator password.
Parameter Template	A template of parameters for creating an instance. The template contains API configuration values that are applied to one or more instances. After an instance is created, you can modify its parameters to better meet your service requirements. For details, see **Modifying Parameters of GeminiDB Influx Instances**.
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 application information about enterprise projects, see ***Enterprise Management User Guide***.

Table 4-20 Tags

Parameter	Description
Tags	The setting is optional. Adding tags helps you better identify and manage your GeminiDB Influx instances. A maximum of 20 tags can be added for each instance. If your organization has configured a tag policy for a GeminiDB Influx instance, you need to 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 created. 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 will be tagged. Each tag key is unique for each instance. It can contain 1 to 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 will 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: _.-:+=@/ After an instance is created, you can view its tag details on the **Tags** tab. In addition, you can add, modify, and delete tags of an existing instance. For details, see **Managing Tags**.

Table 4-21 Required duration

Parameter	Description
Required Duration	The length of your subscription if you select **Yearly/Monthly** billing. Subscription lengths range from one month to three years.
Auto-renew	• This option is not selected by default. • If you select this option, the instance is automatically renewed based on the subscription duration.

Step 5 On the displayed page, confirm instance details.

- Yearly/Monthly
  - To modify the configurations, click **Previous**.
  - If no modification is required, read and agree to the service agreement, click **Pay Now**, and complete the payment.
- Pay-per-use
  - To modify the configurations, click **Previous**.

- If you do not need to modify the settings, read and agree to the service agreement and click **Submit**.

Step 6 Click **Back to Instance Management** to go to the instance list.

Step 7 On the **Instances** page, view and manage the created instance.

- Creating an instance takes about 5 to 9 minutes. During the process, the instance status displayed in the instance list is **Creating**.
- After the instance is created, its status becomes **Available**.

  You can click  in the upper right corner of the page to refresh the instance status.
- An automated backup policy is enabled by default during instance creation. A full backup is automatically triggered after an instance is created.
- The default database port of the instance is **8635** and cannot be changed.

Figure 4-14 Available instance

----End

4.3 Instance Connection and Management

4.3.1 Connecting to a GeminiDB Influx Instance

You can connect to a GeminiDB Influx instance over a private network, public network, load balancer IP address, or program code.

Figure 4-15 shows the process of connecting to a GeminiDB Influx instance.

Figure 4-15 Connection Methods

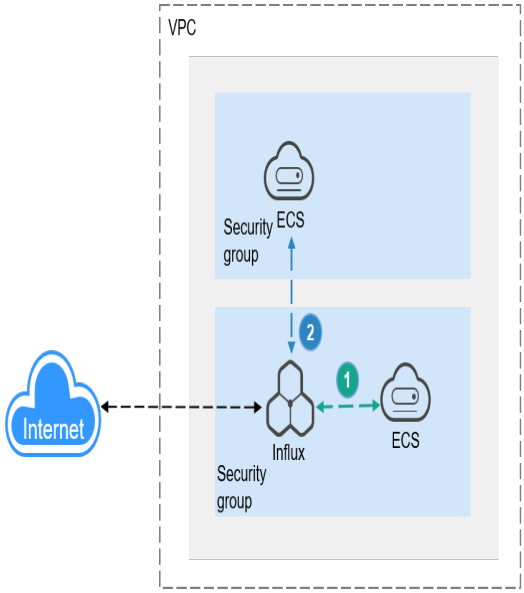

1 A GeminiDB Influx instance is connected over a private network (An ECS and a GeminiDB Influx instance are in the same security group).

2 A GeminiDB Influx instance is connected over a private network (An ECS and a GeminiDB Influx instance are in different security groups).

Table 4-22 Connection methods

Method	Scenario	Default Port	Description
DAS	You can connect to a GeminiDB Influx instance on a web-based console.	-	• Easy to use, secure, advanced, and intelligent • By default, you have the permission of remote login. DAS is secure and convenient for connecting to instances.

Method	Scenario	Default Port	Description
Private network	Connect to a GeminiDB Influx instance using a **private IP address** or **load balancer address**. This method is suitable when your application is deployed on an ECS that is in the same region and VPC as your instance.	8635	• To improve connection reliability and eliminate the impact of a single point of failure, the load balancer address is recommended. • High security and performance • If the ECS and GeminiDB Influx 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 GeminiDB Influx instances by following **Setting Security Group Rules for a GeminiDB Influx 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.

Method	Scenario	Default Port	Description
Public network	You can connect to a GeminiDB Influx instance through an EIP. This method is suitable when DB instances cannot be accessed over a private network. You can bind an EIP to an ECS (or a server on the public network) to access the instance.	8635	• Low security • For faster transmission and improved security, migrate your applications to an ECS that is in the same subnet as your instance and use a private IP address to access the instance. • For EIP pricing details, see **EIP billing details**.
Program code	Connect to a GeminiDB Influx instance using **Go**, **Java**, or **Python**.	8635	-

4.3.2 Connecting to a GeminiDB Influx Instance on the DAS Console

This section describes how to connect to a GeminiDB Influx instance on the console.

Prerequisites

A GeminiDB Influx instance has been created and is running properly.

Usage Notes

- SELECT query commands are supported.
- INSERT commands for writing data are supported.
- Commands for database operations (including creating, deleting, and displaying databases) are supported.
- Commands for user operations (including creating, deleting, displaying, and authorizing users, and changing user passwords) are supported.
- Commands of retention policies (including creating, deleting, displaying, and modifying retention policies) are supported.
- CONTINUOUS QUERY commands (including CREATE CONTINUOUS QUERY, DROP CONTINUOUS QUERY, and SHOW CONTINUOUS QUERY) are supported.

Procedure

Step 1 Log in to the Huawei Cloud console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 In the instance list, locate a target instance and click **Log In** in the **Operation** column.

Figure 4-16 Connecting to a GeminiDB Influx Instance

Alternatively, click the instance name to go to the **Basic Information** page. Click **Log In** in the upper right corner of the page.

Figure 4-17 Connecting to a GeminiDB Influx Instance

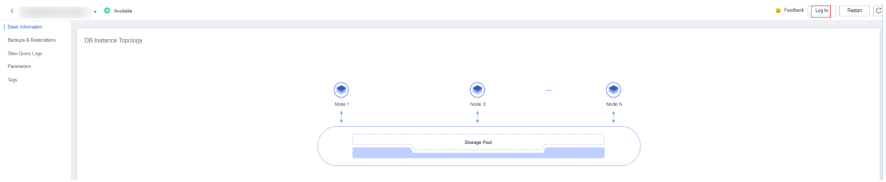

Step 4 Enter the password for logging in to the instance.

Figure 4-18 Logging in to the GeminiDB Influx instance

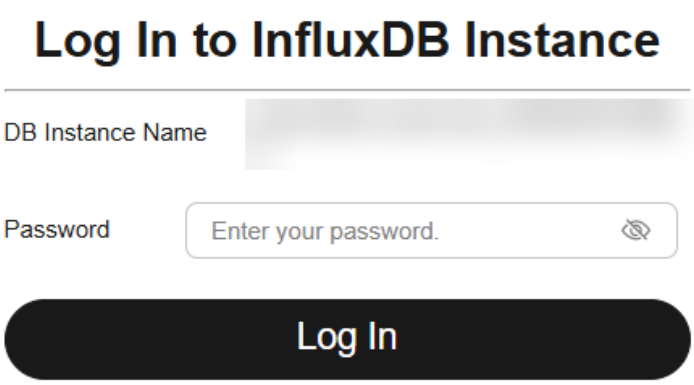

If you need to log in again after the password is reset, click **Re-login** in the upper right corner and use the new password.

Figure 4-19 Re-login

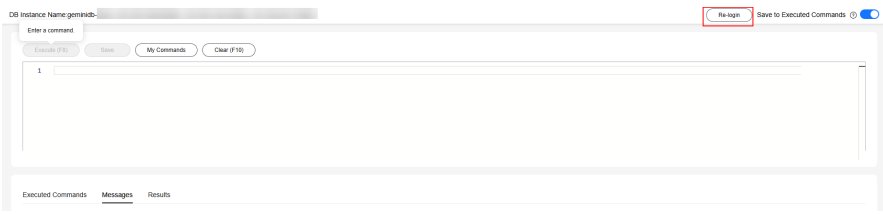

Step 5 Manage relevant databases.

Figure 4-20 Instance homepage

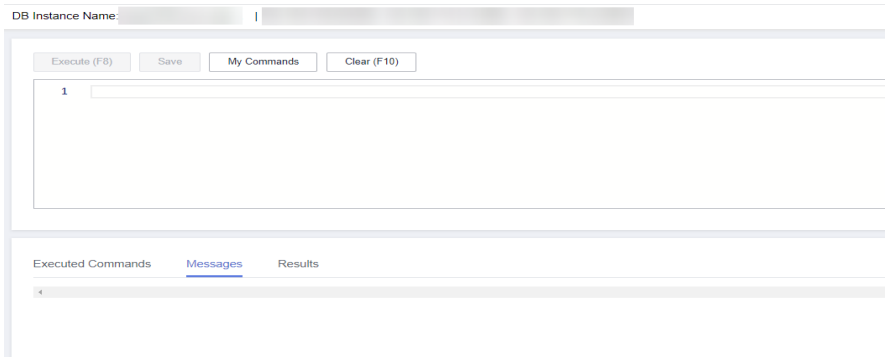

- Save commands to executed commands.

  This function is enabled by default to save the recently executed commands for your later query.

  Then you can click the **Executed Commands** tab on the lower page to view historical commands.

  NOTE

  Commands with passwords are not displayed on the **Executed Commands** tab page.

  Figure 4-21 Executed commands

  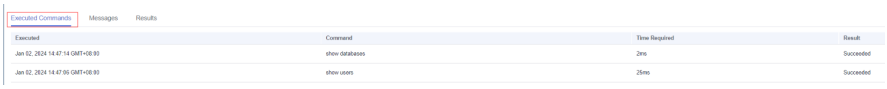

  If this function is disabled, the commands executed subsequently are not displayed any longer. You can click next to **Save Executed SQL Statements** in the upper right corner to disable this function.

- Execute a command.

  You can enter a command in the command window and click **Execute** or **F8**.

Figure 4-22 Execute a command.

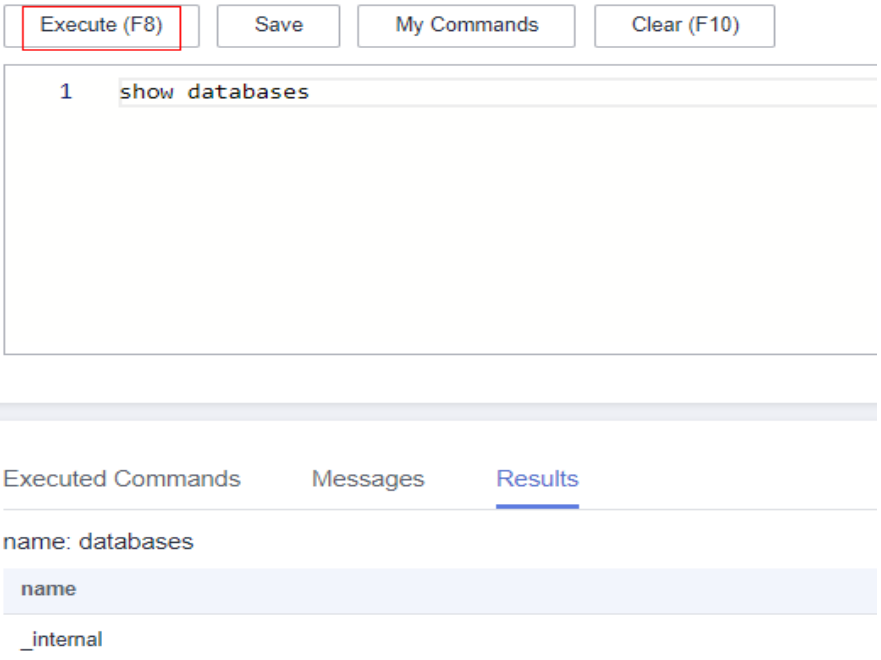

After a command is executed, you can view the execution result on the **Results** page.

- Save a command.

  You can save a command to all instances or the current instance. Then you can view details in **My Commands**.

  NOTE

  Commands with passwords cannot be saved to **My Commands**.

Figure 4-23 Save a command.

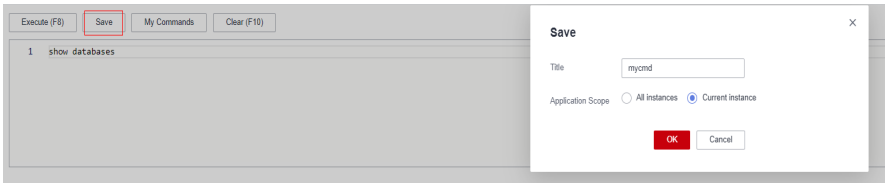

- View my commands.

  Common commands are displayed the **My Commands** page.

  You can set a filter to narrow the scope of commands. If you select **All**, all commands saved in the current account are displayed.

Figure 4-24 Filtering commands

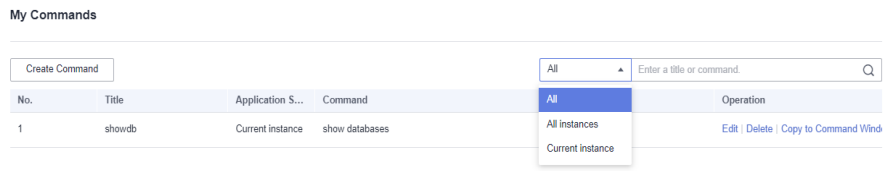

Alternatively, you can enter a command title or statement in the search box to search for the corresponding command.

Figure 4-25 Searching for a command

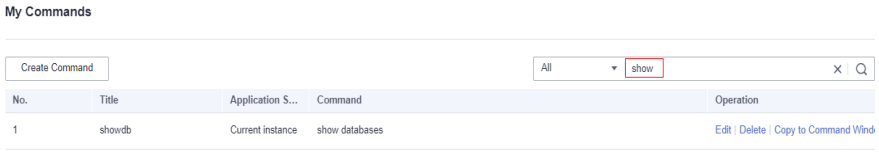

On the **My Commands** page, you can also create, edit, and delete a command or copy it to the command window.

Figure 4-26 Managing a command

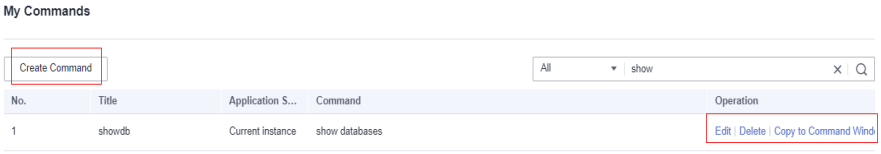

- Clear commands.

  You can also press **F10** to clear the command in the command window.

----End

FAQs

Question: What should I do if the DAS console cannot be redirected after I click **Log In** in the **Operation** column in the instance list or click **Log In** on the **Basic Information** page?

Solution: Set your browser to allow pop-ups and try again.

4.3.3 Connecting to a GeminiDB Influx Instance Using the InfluxDB CLI over a Private Network

Scenarios

This section uses Linux as an example to describe how to connect to a GeminiDB Influx instance using a load balancer address or the private IP address of an ECS.

Usage Notes

- The instance and ECS must be in the same VPC and subnet.
- The ECS must be allowed by the security group to access the instance.
  - If the instance is associated with the default security group, you do not need to set security group rules.
  - If the instance is not associated with the default security group, check whether the security group rules allow the ECS to access the instance. For details, see **Setting Security Group Rules for a GeminiDB Influx Instance**.

Prerequisites

- An ECS has been created. The following uses a Linux ECS as an example. For details, see **Purchasing an ECS** in *Getting Started with Elastic Cloud Server*.
- Download the **x86 client** or **Arm client** of InfluxDB. The following uses the Linux 64-bit client as an example.

Method 1: Using a Load Balancer Address over SSL Connections (Recommended)

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 InfluxDB client installation package to the ECS using file transfer tools like XFTP.

Step 3 Decompress the client (for example, x86 client) tool package.

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 4 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:

   ```
   cd influxdb-1.8.10-1/usr/bin
   ```

2. Connect to the GeminiDB Influx instance.

   ```
   ./influx -ssl -unsafeSsl -username '<DB_USER>' -password '<DB_PWD>' -host <DB_HOST> -port <DB_PORT>
   ```

   Example:

   ```
   ./influx -ssl -unsafeSsl -username 'rwuser' -password '<DB_PWD>' -host 192.xx.xx.xx -port 8635
   ```

   Table 4-23 Parameters

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, you can see the administrator username.
<DB_PWD>	Administrator password

Parameter	Description
<DB_HOST>	Load balancer address of the instance to be connected. Connecting to an instance using a load balancer address 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. Scenario 1: If you have obtained a load balancer address before creating an instance, the load balancer address is selected by default on the instance creation page. After the instance is created, click its name to go to the **Basic Information** page and obtain the load balancer IP address in the **Network Information** area. Scenario 2: To use a load balancer address after the instance is created, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service. Then you can click the instance name to view the load balancer address in the **Network Information** area on the **Basic Information** page.
<DB_PORT>	Instance port Click the instance name to go to the **Basic Information** page. In the **Network Information** area, you can see the instance port.

Step 5 Check the results. If information similar to the following is displayed, the connection is successful.

```
Connected to https://host:port version x.x.x
InfluxDB shell version 1.8.10
>
```

----End

Method 2: Using a Load Balancer Address over Non-SSL Connections

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 InfluxDB client installation package to the ECS using file transfer tools like XFTP.

Step 3 Decompress the client (for example, x86 client) tool package.

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 4 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:

```
cd influxdb-1.8.10-1/usr/bin
```


2. Connect to the GeminiDB Influx instance.
   ./influx -username '*<DB_USER>*' **-password** '*<DB_PWD>*' **-host** *<DB_HOST>* **-port** *<DB_PORT>*

   Example:

   ./influx -username 'rwuser' -password '*<DB_PWD>*' **-host 192.xx.xx.xx -port 8635**

   Table 4-24 Parameters

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, you can see the administrator username.
<DB_PWD>	Administrator password
<DB_HOST>	Load balancer address of the instance to be connected. Connecting to an instance using a load balancer address 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 you have enabled the load balancer address before creating an instance, you can view that the load balancer address is selected by default on the instance creation page. After the instance is created, click the instance name to go to the **Basic Information** page and obtain the load balancer address in the **Network Information** area. – If you have already created an instance and enabled the load balancer address, you can click the instance name and view the address in the **Network Information** area on the **Basic Information** page.
<DB_PORT>	Instance port Click the instance name to go to the **Basic Information** page. In the **Network Information** area, you can see the instance port.

Step 5 Check the results. If information similar to the following is displayed, the connection is successful.

```
Connected to https://host:port version x.x.x
InfluxDB shell version: 1.8.10
>
```

----End

Method 3: Using a Private IP Address over SSL Connections

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 InfluxDB client installation package to the ECS using file transfer tools like XFTP.

Step 3 Decompress the client (for example, x86 client) tool package.

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 4 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:

   ```
   cd influxdb-1.8.10-1/usr/bin
   ```

2. Connect to the GeminiDB Influx instance.

   – Use the default certificate.

   ```
   ./influx -ssl -unsafeSsl -host <DB_HOST> -port <DB_PORT>
   ```

   Example:

   ```
   ./influx -ssl -unsafeSsl -host 192.xx.xx.xx -port 8635
   ```

   Table 4-25 Parameters

Parameter	Description
<DB_HOST>	Private IP address of a node to be connected On the **Instances** page, click the target instance name. You can see the private IP address in the **Private IP Address** column in the **Node Information** area. If the instance you purchased has multiple nodes, select the private IP address of any node.
<DB_PORT>	Port of an instance to be connected. The default port is 8635 and cannot be changed. Click the GeminiDB Influx instance to go to the **Basic Information** page. In the **Network Information** area, you can see the port.

3. Run the **auth** command to authenticate the user.

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

   Table 4-26 Parameters

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, you can see the administrator username.
<DB_PWD>	Administrator password

Step 5 After the authentication is successful, run the **show databases** command.

show databases

If the following information is displayed, the connection is successful.

```
name: databases
name
----
_internal
```

----End

Method 4: Using a Private IP Address over Non-SSL Connections

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 InfluxDB client installation package to the ECS using file transfer tools like XFTP.

Step 3 Decompress the client (for example, x86 client) tool package.

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 4 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:
   ```
   cd influxdb-1.8.10-1/usr/bin
   ```
2. Connect to the GeminiDB Influx instance.
   ```
   ./influx -host <DB_HOST> -port <DB_PORT>
   ```
   Example:
   ```
   ./influx -host 192.xx.xx.xx -port 8635
   ```

   Table 4-27 Parameters

Parameter	Description
<DB_HOST>	Private IP address of a node to be connected On the **Instances** page, click the target instance name. You can see the private IP address in the **Private IP Address** column in the **Node Information** area. If the instance you purchased has multiple nodes, select the private IP address of any node.
<DB_PORT>	Port of an instance to be connected. The default port is 8635 and cannot be changed. Click the GeminiDB Influx instance to go to the **Basic Information** page. In the **Network Information** area, you can see the port.

3. Run the **auth** command to authenticate the user.

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

Table 4-28 Parameters

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, you can see the administrator username.
<DB_PWD>	Administrator password

Step 5 After the authentication is successful, run the **show databases** command.

show databases

If the following information is displayed, the connection is successful.

```
name: databases
name
----
_internal
```

----End

4.3.4 Connecting to a GeminiDB Influx Instance Using the InfluxDB CLI over a Public Network

This section uses the Linux operating system as an example to describe how to connect an ECS to a GeminiDB Influx instance over a public network.

Prerequisites

- Bind an EIP to the GeminiDB Influx instance and configure security group rules to ensure that the instance is accessible from ECSs through the EIP. For details, see **Binding an EIP to a GeminiDB Influx Instance Node** and **Setting Security Group Rules for a GeminiDB Influx Instance**.
- An ECS has been created. The following uses a Linux ECS as an example. For details, see **Purchasing an ECS** in *Getting Started with Elastic Cloud Server*.
- Download the **x86 client** or **Arm client** of InfluxDB. The following uses the Linux 64-bit client as an example.

Procedure

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 InfluxDB client installation package to the ECS using file transfer tools like XFTP.

Step 3 Decompress the client (for example, x86 client) tool package.

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 4 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:
   ```
   cd influxdb-1.8.10-1/usr/bin
   ```
2. Connect to the GeminiDB Influx instance.
   - Use SSL to connect to a database.
     ```
     ./influx -ssl -unsafeSsl -host <DB_HOST> -port <DB_PORT>
     ```
     Example:
     ```
     ./influx -ssl -unsafeSsl -host 10.xx.xx.xx -port 8635
     ```
   - Use a non-SSL connection to access a database.
     ```
     ./influx -host <DB_HOST> -port <DB_PORT>
     ```
     Example:
     ```
     ./influx -host 10.xx.xx.xx -port 8635
     ```

Table 4-29 Parameters

Parameter	Description
<DB_HOST>	EIP of the node to be connected On the **Instances** page, click the target instance name. You can see the EIP in the **EIP** column in the **Node Information** area. If the instance you purchased has multiple nodes, select the EIP of any node. If no EIP has been bound to the current node, bind an EIP to the instance by following **Binding an EIP to a GeminiDB Influx Instance Node**.
<DB_PORT>	Port of an instance to be connected. The default port is 8635 and cannot be changed. Click the GeminiDB Influx instance to go to the **Basic Information** page. In the **Network Information** area, you can see the port.

3. Run the **auth** command to authenticate the user.

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

Table 4-30 Parameters

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, you can see the administrator username.
<DB_PWD>	Administrator password

Step 5 After the authentication is successful, run the **show databases** command.

show database

If the following information is displayed, the connection is successful.

```
name: databases
name
----
_internal
```

----End

Follow-up Operations

After you log in to the instance, you can create databases or data retention policies. For details, see **Buying and Connecting to a GeminiDB Influx Instance**.

4.3.5 Connecting to a GeminiDB Influx Instance Using Programming Languages

4.3.5.1 Connecting to a GeminiDB Influx Instance Using Go

This section describes how to connect to a GeminiDB Influx instance using the Go programming language.

Prerequisites

- You have downloaded the client code from the **InfluxDB open-source project website**.

Example Code for Accessing an Instance Using a Non-SSL Connection

```
package main

import (
    "fmt"
    _ "github.com/influxdata/influxdb1-client" // this is important because of the bug in go mod
    client "github.com/influxdata/influxdb1-client/v2"
    "os"
)

func main(){
    c, err := client.NewHTTPClient(client.HTTPConfig{
        Addr: "http://ip:port",
        // There will be security risks if the username and password used for authentication are directly written into code. Store the username and password in ciphertext in the configuration file or environment variables.
        // In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
        username = os.Getenv("EXAMPLE_USERNAME_ENV"),
        password = os.Getenv("EXAMPLE_PASSWORD_ENV"),
        Username: username,
        Password: password,
    })
    if err != nil {
        fmt.Println("Error creating InfluxDB Client: ", err.Error())
    }
```



```
    q := client.NewQuery("select * from cpu","db0","ns")
    if response, err := c.Query(q); err == nil && response.Error() == nil {
        fmt.Println("the result is: ",response.Results)
    }
}
```

4.3.5.2 Connecting to a GeminiDB Influx Instance Using Java

This section describes how to connect to a GeminiDB Influx instance using the Java programming language.

Dependencies on the pom File

```
<dependency>
  <groupId>org.influxdb</groupId>
  <artifactId>influxdb-java</artifactId>
  <version>2.21</version>
</dependency>
```

Example Code for Connecting to an Instance Using SSL

```
package influxdb;

 import java.security.SecureRandom;
 import java.security.cert.X509Certificate;
 import java.util.concurrent.TimeUnit;
 import javax.net.ssl.SSLContext;

 import okhttp3.OkHttpClient;
 import org.influxdb.InfluxDB;
 import org.influxdb.InfluxDBFactory;
 import org.influxdb.dto.Point;
 import org.influxdb.dto.Query;
 import org.influxdb.dto.QueryResult;

 import org.apache.http.ssl.SSLContexts;
 import javax.net.ssl.*;

 public class demo {
     public static void main(String[] args) {
         OkHttpClient.Builder client = new OkHttpClient.Builder()
             .connectTimeout(10, TimeUnit.SECONDS)
             .writeTimeout(10, TimeUnit.SECONDS)
             .readTimeout(10, TimeUnit.SECONDS)
             .retryOnConnectionFailure(true);

         client.sslSocketFactory(defaultSslSocketFactory(), defaultTrustManager());
         client.hostnameVerifier(noopHostnameVerifier());

        // There will be security risks if the username and password used for authentication are directly written into code. Store the username and password in ciphertext in the configuration file or environment variables.
        // In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
         String username = System.getenv("EXAMPLE_USERNAME_ENV");
         String password = System.getenv("EXAMPLE_PASSWORD_ENV");
         final String serverURL = "https://127.0.0.1:8086", username = username, password = password;

         InfluxDB influxdb = InfluxDBFactory.connect(serverURL, username, password, client);

         // Create a database...
         String databaseName = "foo";
         influxdb.query(new Query("CREATE DATABASE " + databaseName, databaseName));
         influxdb.setDatabase(databaseName);

         // Write points to influxdb.
```



```
        influxdb.write(Point.measurement("bar")
            .time(System.currentTimeMillis(), TimeUnit.MILLISECONDS)
            .tag("location", "chengdu")
            .addField("temperature", 22)
            .build());

        // Query your data using InfluxQL.
        QueryResult queryResult = influxdb.query(new Query("SELECT * FROM bar", databaseName));

        // Close it if your application is terminating or you are not using it anymore.
        influxdb.close();
    }

    private static X509TrustManager defaultTrustManager() {
        return new X509TrustManager() {
            public X509Certificate[] getAcceptedIssuers() {
                return new X509Certificate[0];
            }

            public void checkClientTrusted(X509Certificate[] certs, String authType) {
            }

            public void checkServerTrusted(X509Certificate[] certs, String authType) {
            }
        };
    }

    private static SSLSocketFactory defaultSslSocketFactory() {
        try {
            SSLContext sslContext = SSLContexts.createDefault();

            sslContext.init(null, new TrustManager[] {
                defaultTrustManager()
            }, new SecureRandom());
            return sslContext.getSocketFactory();
        } catch (Exception e) {
            throw new RuntimeException(e);
        }

    }

    private static HostnameVerifier noopHostnameVerifier() {
        return new HostnameVerifier() {
            @Override
            public boolean verify(final String s, final SSLSession sslSession) {
                return true; //true indicates that SSL is enabled but the SSL certificate is not verified. This mode
is recommended.
            }
        };
    }
}
```

Example Java Code for Connecting to an Instance Using an Unencrypted Connection

```
package influxdb;

import okhttp3.OkHttpClient;
import org.influxdb.InfluxDB;
import org.influxdb.InfluxDBFactory;
import org.influxdb.dto.Point;
import org.influxdb.dto.Query;
import org.influxdb.dto.QueryResult;

import java.util.concurrent.TimeUnit;

public class demoNoSSL {
    public static void main(String[] args) {
        OkHttpClient.Builder client = new OkHttpClient.Builder()
```



```
            .connectTimeout(10, TimeUnit.SECONDS)
            .writeTimeout(10, TimeUnit.SECONDS)
            .readTimeout(10, TimeUnit.SECONDS)
            .retryOnConnectionFailure(true);

        // There will be security risks if the username and password used for authentication are directly written into code. Store the username and password in ciphertext in the configuration file or environment variables.
        // In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
        String username = System.getenv("EXAMPLE_USERNAME_ENV");
        String password = System.getenv("EXAMPLE_PASSWORD_ENV");
        final String serverURL = "http://127.0.0.1:8086", username = username, password = password;
        InfluxDB influxdb = InfluxDBFactory.connect(serverURL, username, password, client);

        // Create a database...
        String databaseName = "foo";

        influxdb.query(new Query("CREATE DATABASE " + databaseName, databaseName));
        influxdb.setDatabase(databaseName);

        // Write points to influxdb.
        influxdb.write(Point.measurement("bar")
            .time(System.currentTimeMillis(), TimeUnit.MILLISECONDS)
            .tag("location", "chengdu")
            .addField("temperature", 22)
            .build());

        // Query your data using InfluxQL.
        QueryResult queryResult = influxdb.query(new Query("SELECT * FROM bar", databaseName));

        // Close it if your application is terminating or you are not using it anymore.
        influxdb.close();
    }
}
```

Example Java Code for Connecting to an Instance Using the Connection Pool

```
package influxdb;

import okhttp3.ConnectionPool;
import okhttp3.OkHttpClient;
import org.influxdb.InfluxDB;
import org.influxdb.InfluxDBFactory;
import org.influxdb.dto.Point;
import org.influxdb.dto.Query;
import org.influxdb.dto.QueryResult;

import java.util.concurrent.TimeUnit;

public class demoConnectionPool {
    public static void main(String[] args) {
        // The client connection pool is based on OkHttpClient.
        OkHttpClient.Builder client = new OkHttpClient().newBuilder();
        client.connectTimeout(10, TimeUnit.SECONDS);
        client.readTimeout(10, TimeUnit.SECONDS);
        client.writeTimeout(10, TimeUnit.SECONDS);
        // Set this parameter to true to mask some connection errors so that the system automatically retries.
        client.retryOnConnectionFailure(true);
        // Maximum number of idle connections in the connection pool. The default value is 5.
        // The connection that stays idle longer than the threshold will be disabled by the connection pool. Then sockets enter into the TIME_WAIT status for the system to reclaim. Set parameter new ConnectionPool based on the number of the idle connections.
        client.connectionPool(new ConnectionPool(5, 30, TimeUnit.SECONDS));

        // There will be security risks if the username and password used for authentication are directly written into code. Store the username and password in ciphertext in the configuration file or environment variables.
```



```
        // In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
        String username = System.getenv("EXAMPLE_USERNAME_ENV");
        String password = System.getenv("EXAMPLE_PASSWORD_ENV");
        final String serverURL = "http://127.0.0.1:8086", username = username, password = password;
        InfluxDB influxdb = InfluxDBFactory.connect(serverURL, username, password, client);

        // Create a database...
        String databaseName = "foo";

        influxdb.query(new Query("CREATE DATABASE " + databaseName, databaseName));
        influxdb.setDatabase(databaseName);

        // Write points to influxdb.
        influxdb.write(Point.measurement("bar")
                .time(System.currentTimeMillis(), TimeUnit.MILLISECONDS)
                .tag("location", "chengdu")
                .addField("temperature", 22)
                .build());

        // Query your data using InfluxQL.
        QueryResult queryResult = influxdb.query(new Query("SELECT * FROM bar", databaseName));

        // Close it if your application is terminating or you are not using it anymore.
        influxdb.close();
    }
}
```

Example Java Code for Connecting to an Instance Using a Short Connection

```
/**
        Scenarios:
        * * When the ELB connection is used, the client sends multiple query requests at a time.
        * If HTTP persistent connections are used, most query requests are sent to one InfluxDB node, causing load imbalance.
        HTTP short connections (The value of Connection is close in the request header) can be used to achieve load balancing among InfluxDB nodes.
        */
/**
        In this mode, only part of the code is displayed.
        */
OkHttpClient.Builder client = new OkHttpClient.Builder()
                .connectTimeout(10, TimeUnit.SECONDS)
                .writeTimeout(10, TimeUnit.SECONDS)
                .readTimeout(10, TimeUnit.SECONDS)
                .retryOnConnectionFailure(true)
                .addNetworkInterceptor(chain -> {
                    Request newRequest = chain.request().newBuilder().header("Connection", "close").build();
                    return chain.proceed(newRequest);
                });
```

4.3.5.3 Connecting to a GeminiDB Influx Instance Using Python

This section describes how to connect to a GeminiDB Influx instance using the Python programming language.

Prerequisites

The Python client of InfluxDB has been installed.

Example Code for Accessing an Instance Using a Non-SSL Connection

```
from influxdb import InfluxDBClient

# There will be security risks if the username and password used for authentication are directly written into
```



```
code. Store the username and password in ciphertext in the configuration file or environment variables.
# In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
username = os.getenv('EXAMPLE_USERNAME_ENV')
password = os.getenv('EXAMPLE_PASSWORD_ENV')
client = InfluxDBClient(host=IP, port=****, username=username, password=password, ssl=False)
client.get_list_database()
```

NOTE

Replace **host** and **port** with the actual values.

Example Code for Accessing an Instance Using an SSL Connection

```
from influxdb import InfluxDBClient

# There will be security risks if the username and password used for authentication are directly written into code. Store the username and password in ciphertext in the configuration file or environment variables.
# In this example, the username and password are stored in the environment variables. Before running this example, set environment variables EXAMPLE_USERNAME_ENV and EXAMPLE_PASSWORD_ENV as needed.
username = os.getenv('EXAMPLE_USERNAME_ENV')
password = os.getenv('EXAMPLE_PASSWORD_ENV')
client = InfluxDBClient(host=IP, port=****, username=username, password=password, ssl=True)
client.get_list_database()
```

NOTE

- Replace **host** and **port** with the actual values.
- The value of **ssl** must be **True**.
- If SSL is not set or is set to **False**, the following error information is displayed:

  InfluxDBClientError: 400: Client sent an HTTP request to an HTTPS server.

4.3.6 Connection Information Management

4.3.6.1 Setting Security Group Rules for a GeminiDB Influx Instance

A security group is a collection of access control rules for ECS, , and GeminiDB Influx 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 the GeminiDB Influx instances.

This section describes how to configure security group rules for a GeminiDB Influx instance that is connected through a private or a public network.

Usage Notes

- By default, you can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- One security group can be associated with only one GeminiDB Influx instance.
- For details about security group rules, see **Table 4-31**.

Table 4-31 Parameter description

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

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.

Step 4 Configure security group rules.

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

Figure 4-27 Security group

Step 5 Add an inbound rule.

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

Figure 4-28 Inbound rules

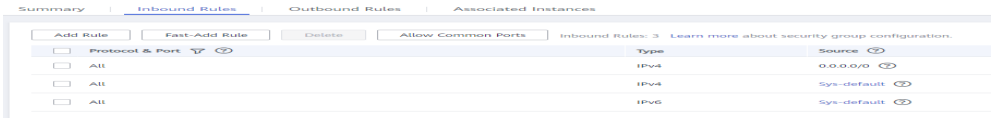

2. Click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Figure 4-29 Adding a rule

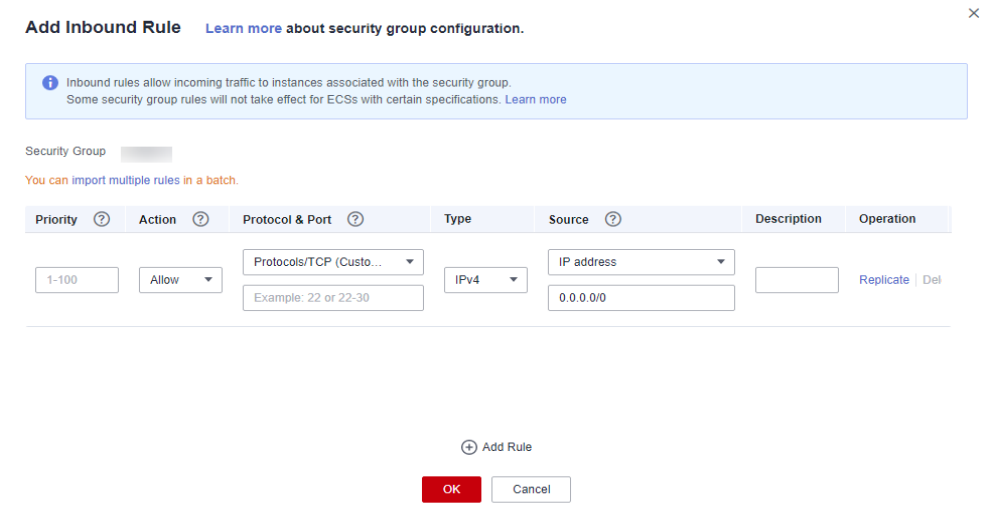

3. In the displayed dialog box, set required parameters.

Table 4-32 Inbound rule settings

Parameter	Description	Example Value
Protocol & Port	– The network protocol required for access. Currently, GeminiDB Influx instances can be accessed only over TCP. – **Port**: The port (1 to 65535) for accessing the ECS.	TCP
Type	IP address type. This parameter is available after IPv6 is enabled. – IPv4 – IPv6	IPv4
Source	The IP address, IP address group, or security group that the rule applies to, which allows access from IP addresses or instances in other security group. Example: – Single IP address: xxx.xxx.xxx.xxx/32 (IPv4) – Subnet: xxx.xxx.xxx.0/24 – All IP addresses: 0.0.0.0/0 – sg-abc (security group)	0.0.0.0/0

Parameter	Description	Example Value
Description	(Optional) Provides supplementary information about the security group rule. The description can contain up to 255 characters and cannot contain angle brackets (<>).	-

Step 6 Click **OK**.

----End

4.3.6.2 Binding an EIP to a GeminiDB Influx Instance Node

Scenarios

An EIP provides independent public IP addresses and bandwidth for Internet access. After you create a GeminiDB Influx instance, you can bind an EIP to its node to allow external access. If later you want to prohibit external access, you can also unbind the EIP.

Usage Notes

- Configure security group rules and enable specific IP addresses and ports to access the target DB instance. Before accessing a database, apply for an EIP on the VPC console. Then, add an inbound rule to allow the IP addresses or IP address ranges of ECSs. For details, see **Setting Security Group Rules for a GeminiDB Influx Instance**.
- To change the EIP that has been bound to a node, unbind it from the node first.

Binding an EIP

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 to which you want to bind an EIP to and click its name.

Step 4 On the **Basic Information** page, in the **Node Information** area, locate the target node and click **Bind EIP** in the **Operation** column.

Figure 4-30 Binding an EIP

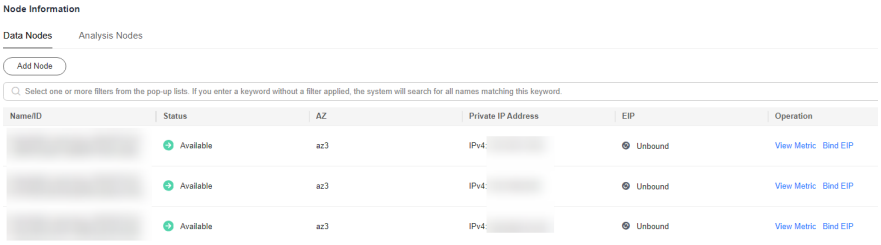

Step 5 In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **Yes**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.

Figure 4-31 Selecting an EIP

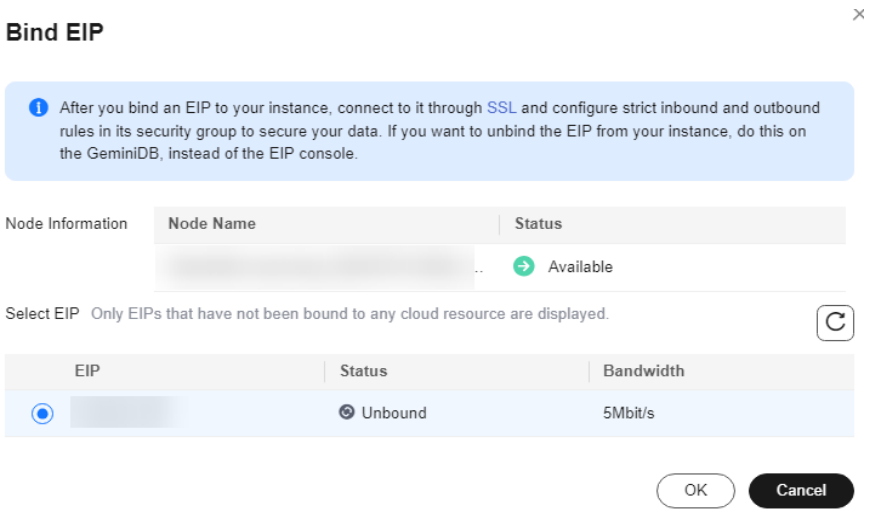

Step 6 In the **EIP** column, view the EIP that is successfully bound.

To unbind the EIP from the DB instance, see **Unbinding an EIP**.

----End

Unbinding an EIP

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 from which you wish to unbind an EIP.

Step 4 On the **Basic Information** page, in the **Node Information** area, locate the target node and click **Unbind EIP** in the **Operation** column.

Figure 4-32 Unbinding an EIP

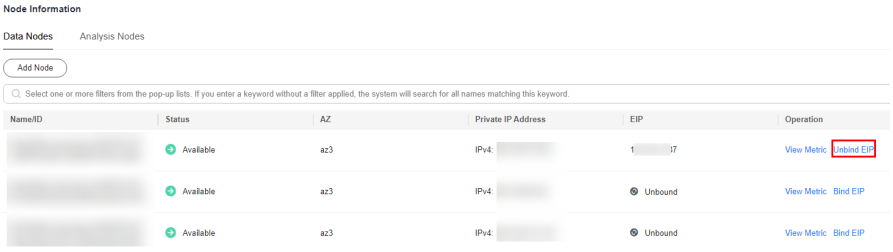

Step 5 In the displayed dialog box, click **Yes**.

To bind an EIP to the DB instance again, see **Binding an EIP**.

----End

4.3.6.3 Changing the Security Group of a GeminiDB Influx Instance

Scenarios

You can change security groups of GeminiDB Influx instances.

Usage Notes

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

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 DB instance.

Step 4 In the navigation pane on the left, choose **Connections**.

Step 5 In the **Security Group** area, click  beside the security group name and select the required security group.

- To submit the change, click  . This process takes about 1 to 3 minutes.
- To cancel the change, click  .

Step 6 View the modification result.

----End

4.3.6.4 Encrypting Data over SSL for a GeminiDB Influx Instance

After a GeminiDB Influx instance is created, you can enable or disable SSL.

Usage Notes

- After enabling or disabling SSL, restart the DB instance for the change to take effect.

Enabling SSL

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. The **Basic Information** page is displayed.

Step 4 In the **DB Information** area, click  to enable SSL.

Figure 4-33 Enabling SSL

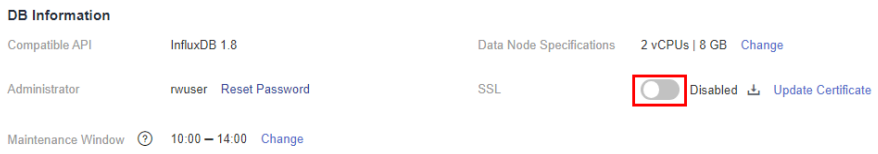

----End

Disabling SSL

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. The **Basic Information** page is displayed.

Step 4 In the **DB Information** area, click  to disable SSL.

Figure 4-34 Disabling SSL

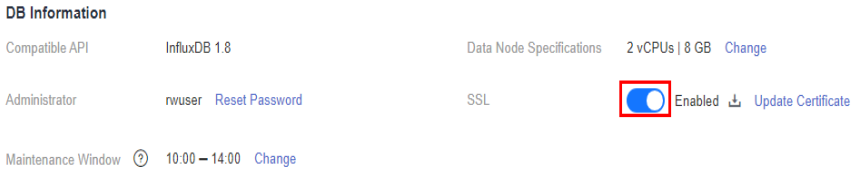

----End

4.3.6.5 Accessing a GeminiDB Influx Instance Using a Load Balancer Address

Scenarios

You can access a GeminiDB Influx instance using a load balancer address.

Usage Notes

- A load balancer address does not support security groups. After an instance is created, configure IP address access control. If no whitelist is configured, all IP addresses that can communicate with the VPC can access the instance.
- To use the access control function, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service to grant required permissions.

Enabling a Blacklist/Whitelist for a Load Balancer IP Address

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.

Step 4 In the **Network Information** area, click next to **Access Control**.

Figure 4-35 Enabling access through a load balancer

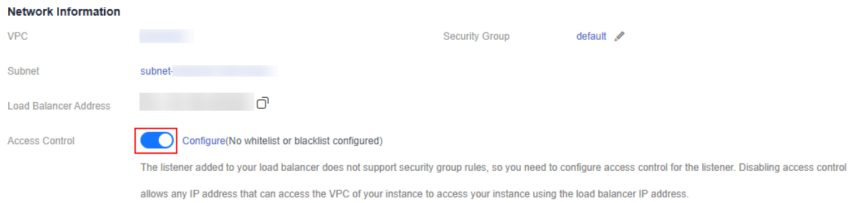

Step 5 Select **Blacklist** or **Whitelist** and specify IP addresses in that list.

Figure 4-36 Configuring access control

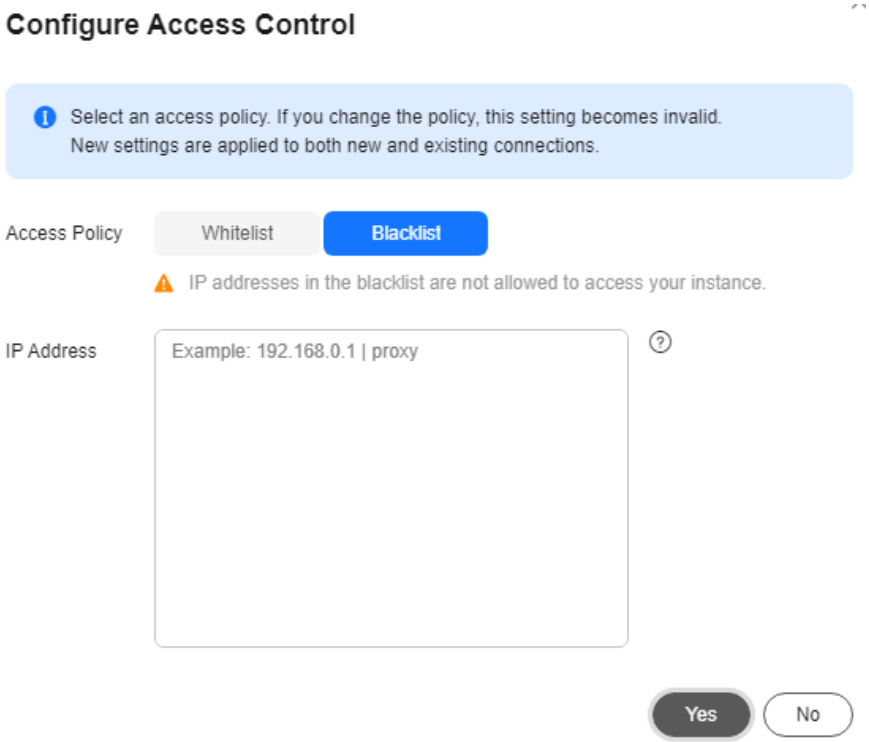

- **Blacklist**: The whitelist and blacklist cannot be configured at the same time. If you switch between lists, your previous settings will be lost. IP addresses in the blacklist cannot be accessed.
- **Whitelist**: The whitelist and blacklist cannot be configured at the same time. If you switch between lists, your previous settings will be lost. Only IP addresses in the whitelist are allowed to access the system.

----End

Disabling Access Through a Load Balancer

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.

Step 4 In the **Network Information** area, click next to **Access Control**. In the displayed dialog box, click **Yes** to disable access through a load balancer.

Figure 4-37 Disabling access through a load balancer

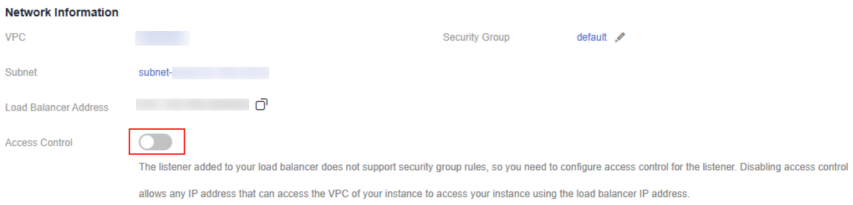

Step 5 Check whether the settings take effect.

----End

4.4 Migrating Data

InfluxDB Community Edition is a popular time series database that focuses on high-performance query and storage of time series data.

GeminiDB Influx API is a cloud-native NoSQL time-series database with a decoupled compute and storage architecture developed by Huawei and is compatible with InfluxDB. This high availability database is secure and scalable, can be deployed, backed up, or restored quickly, and offers monitoring and alarm management capabilities. You can also add storage or compute resources separately. GeminiDB Influx API has better query, write, and data compression performance than InfluxDB Community Edition.

This section describes how to migrate data from InfluxDB Community Edition to GeminiDB Influx API.

Migration Principles

Use open-source migration tool **data-migration-tools** to parse the tsm and wal files of the InfluxDB Community Edition and write the files to a line protocol file. Then, the line protocol file data is parsed and migrated to the destination.

The migration process is divided into two phases:

- Export: tsm files of InfluxDB Community Edition are concurrently parsed, and the parsed data is written into memory.
- Import: The read data is sent to the GeminiDB Influx cluster.

You can specify a migration period while the migration tool is running.

NOTE

Download and decompress the release package of **data-migration-tools**.

Usage Notes

- Deploy the migration tool on the same server as InfluxDB Community Edition and prepare a configuration file.
- The migration tool needs to extract data from tsm to the local line protocol file, obtain data from the line protocol file, and send the data to the destination GeminiDB Influx database. This process may affect the performance of the source side. You are advised to run the migration tool during off-peak hours.
- The migration tool supports only InfluxDB 1.*X* Community Edition.

Prerequisites

- Ensure that the network connection between the source and destination is normal.
- The corresponding database has been created and the retention policy (RP) has been configured in the destination GeminiDB Influx.

Procedure

For details about how to migrate data from InfluxDB Community Edition to GeminiDB Influx API, see **Data Migration Tool Usage Guide**.

Migration Performance Reference

- Migration environment
  - Source: Deploy InfluxDB and the migration tool on an ECS with 4 vCPUs and 16 GB of memory.
  - Destination: three-node GeminiDB Influx instance with 4 vCPUs and 16 GB of memory
- Migration performance
  - The data migration rate of a single process on the source database is 1 GB/min.

4.5 Converting Data into a Parquet file and Exporting the Data to OBS

Scenarios

New data inserted by users on GeminiDB Influx instances in a performance-enhanced cluster can be converted into a Parquet file and automatically uploaded to a specified OBS bucket. You can access or download a Parquet file on OBS.

Usage Notes

- To export data, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

- Data can be converted into a Parquet file by table. By default, data in all tables is converted. To export data in a specified table, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.
- This function is available only for GeminiDB Influx instances with classic storage in a performance-enhanced cluster.
- To use this function, you need to enable **Export Data**.
- OBS supports parallel file systems and standard buckets.
- You can specify a target folder, which must be created in advance.
- When you export data, a policy named data-dump-access is created on OBS. This policy provides only the PUT permission and applies only to the selected folder.
- You need to set the shard duration in the retention policy to one day.
- Conversion into a Parquet file is an asynchronous process. The duration depends on multiple factors such as the concurrency level and system load. The real-time performance cannot be guaranteed.
- The converted Parquet file is uploaded on the hour, for example, at 02:00, every two hours.
- Historical data cannot be converted into a Parquet file or exported to OBS.

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** > **Export Data** in the **Operation** column.

Figure 4-38 Export Data

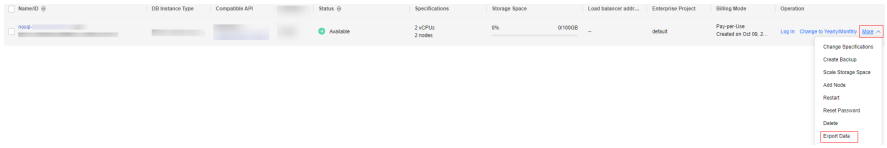

You can also click the target instance to go to the basic information page. Click **Export Data** next to **Bucket Configuration** in the **DB Information** area.

Figure 4-39 Export Data

Step 4 On the **Export Data** page, enable **Export Data**.

Figure 4-40 Export Data

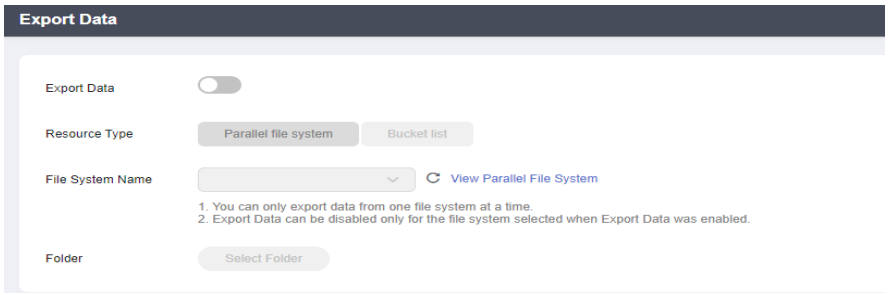

Step 5 Select **Parallel file system** or **Bucket list** for **Resource Type** and select a file system name or an OBS bucket name.

Step 6 Click **Select Folder**.

Figure 4-41 Select Folder

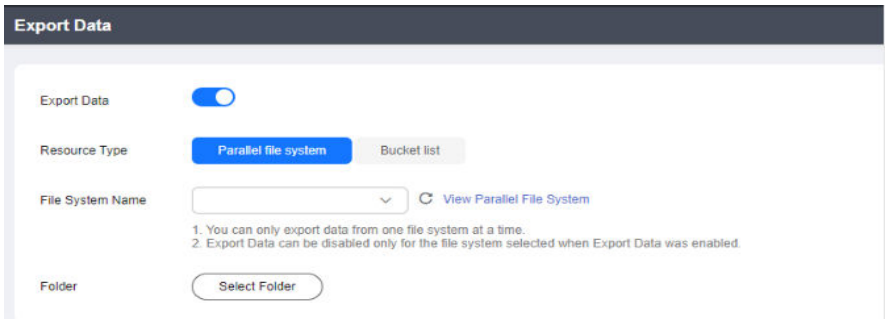

Step 7 Click **OK**.

----End

4.6 Instance Lifecycle Management

4.6.1 Restarting a GeminiDB Influx Instance

Scenarios

You may need to restart an instance for routine maintenance.

Usage Notes

- If the instance status is **Available**, **Abnormal**, or **Checking restoration**, you can restart the instance.
- Restarting an instance will interrupt services. Exercise caution when performing this operation.
- If you restart an instance, all nodes in the instance are also restarted.
- 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 you wat to restart and choose **More** > **Restart** in the **Operation** column.

Alternatively, click the name of the instance, and on the displayed **Basic Information** page, click **Restart** in the upper right corner.

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**.

For a GeminiDB Influx cluster instance, you can restart nodes one by one or all at once.

Figure 4-42 Restarting a GeminiDB Influx cluster instance

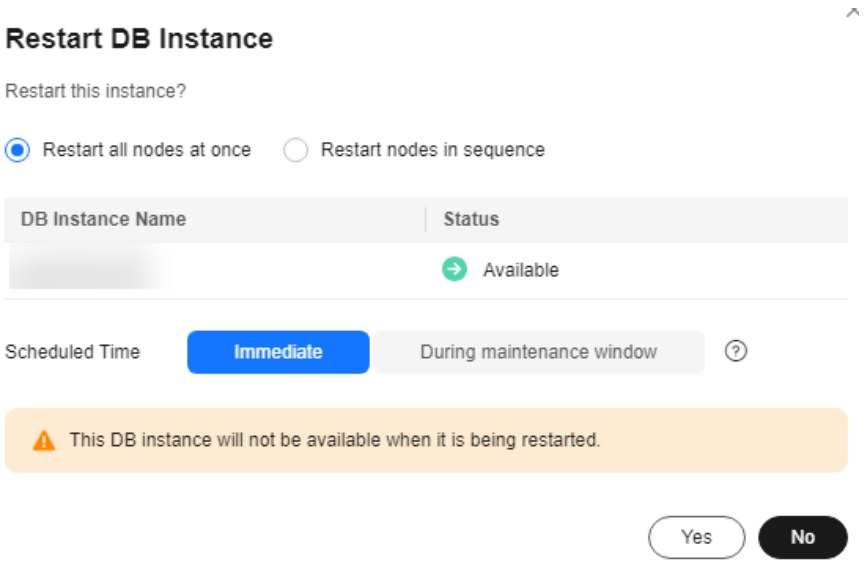

For a single-node GeminiDB Influx instance, click **Yes** or **Immediate**.

Figure 4-43 Restarting a single-node GeminiDB Influx instance

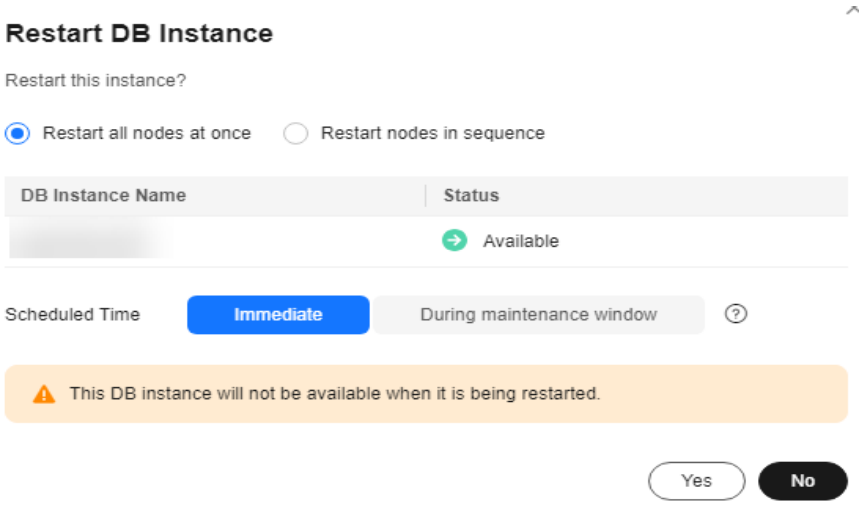

- Instance with classic storage

  For a GeminiDB Influx cluster instance, you can restart nodes one by one or all at once.

  Figure 4-44 Restarting a GeminiDB Influx cluster instance

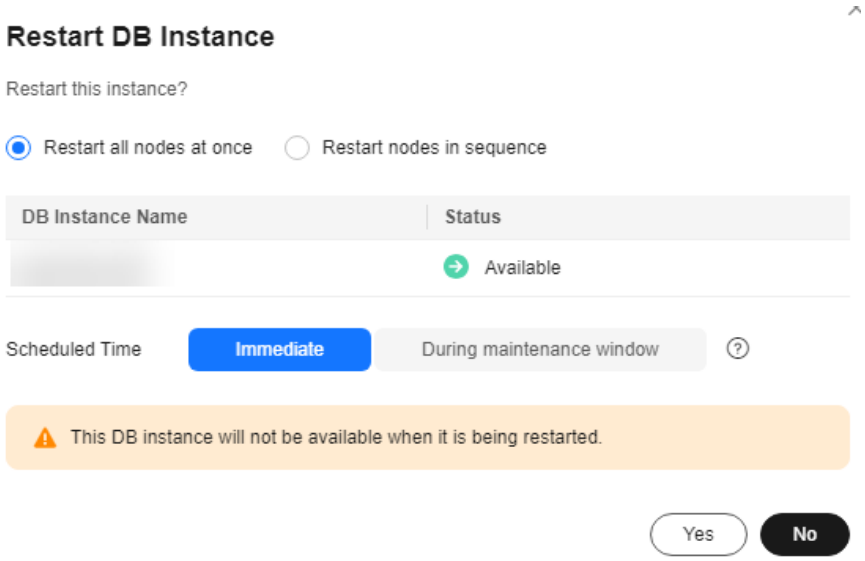

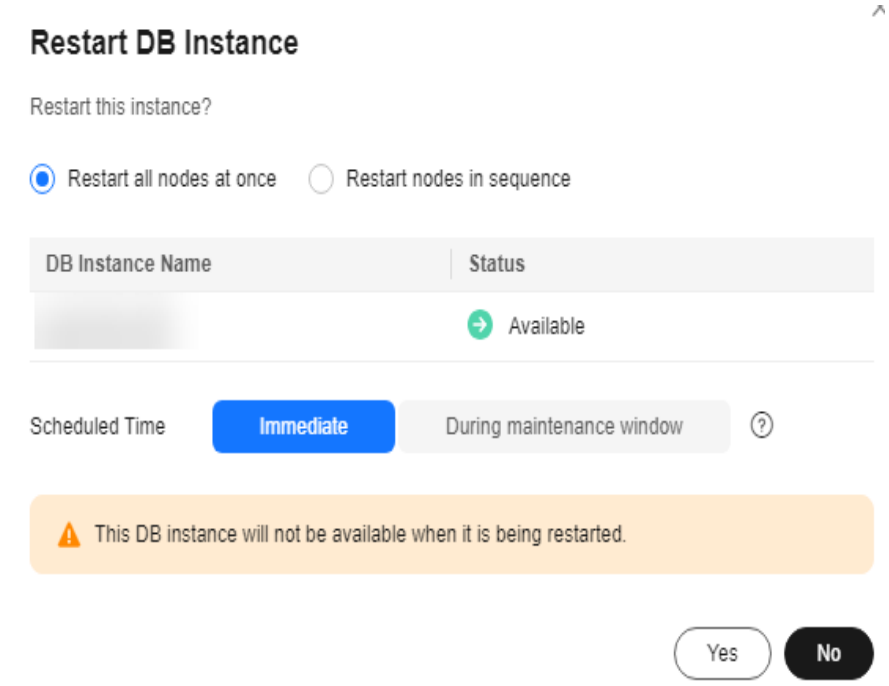

For a single-node GeminiDB Influx instance, click **Yes** or **Immediate**.

Figure 4-45 Restarting a single-node GeminiDB Influx instance

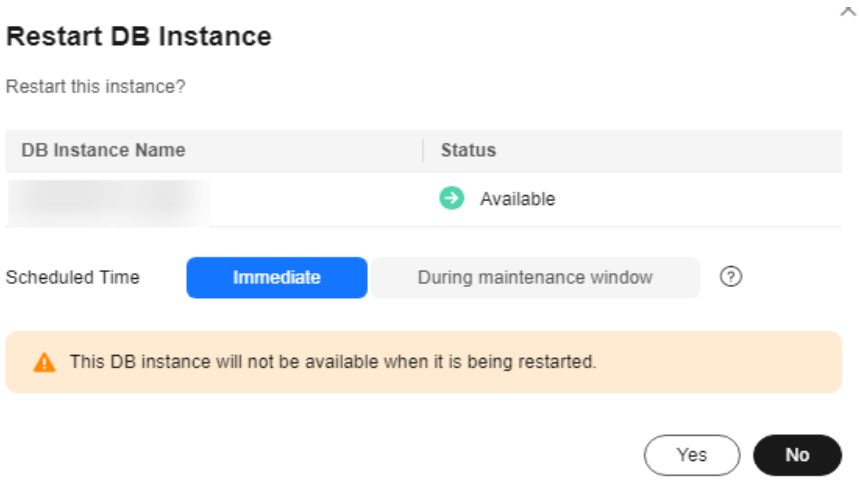

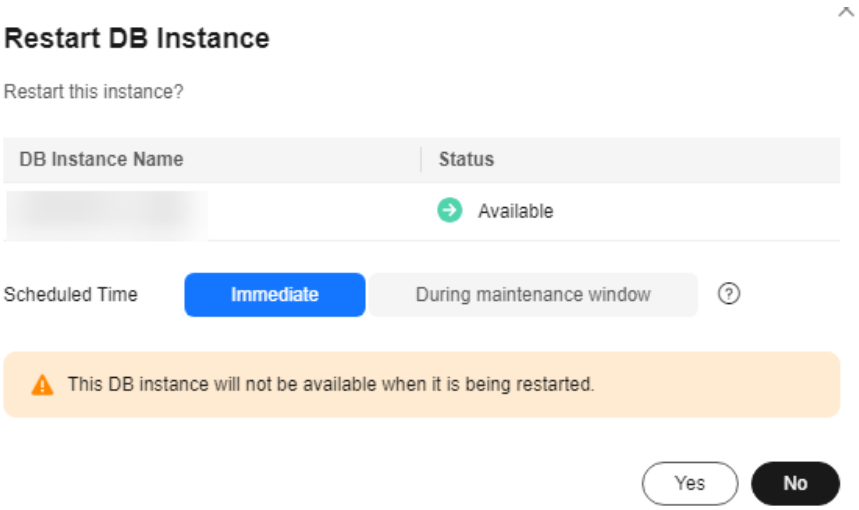

- Instance with cloud native storage

  For a GeminiDB Influx instance with cloud native storage, click **Yes** or **Immediate**.

  Figure 4-46 Restarting an instance

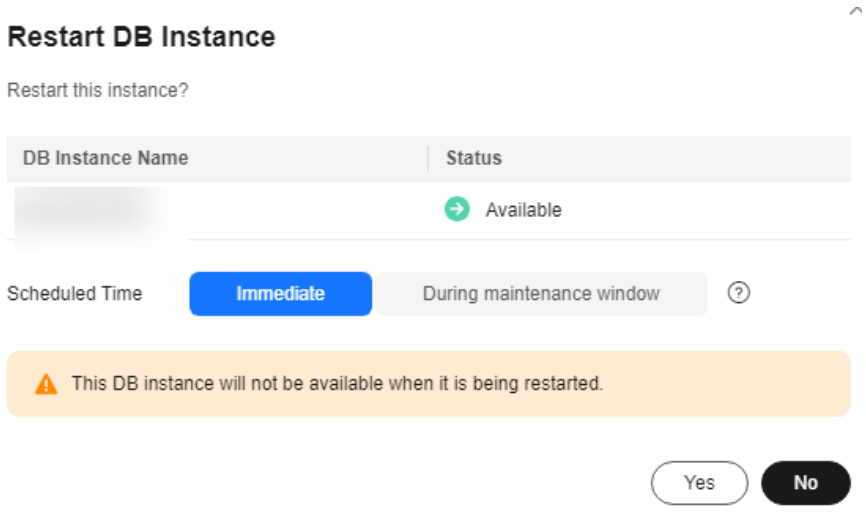

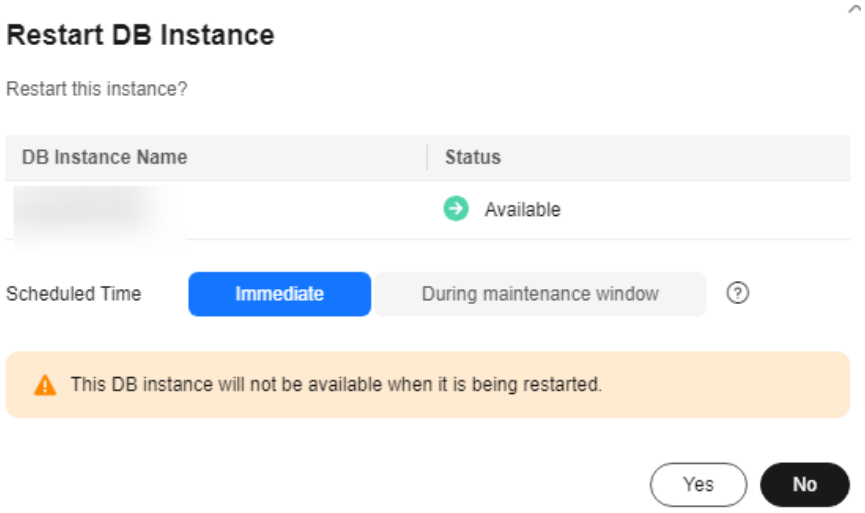

----End

4.6.2 Exporting Instance Information

Scenarios

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

Exporting Information About All Instances

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 on your local PC.

----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 on your local PC.

----End

4.6.3 Deleting a Pay-per-Use Instance

Scenarios

You can choose to delete a pay-per-use instance on the **Instances** page based on service requirements. To delete a yearly/monthly instance, unsubscribe from it. For details, see **How Do I Unsubscribe from a Yearly/Monthly Instance?**.

Precautions

- 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.

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 and in the **Operation** column choose **Delete** or **More** > **Delete**.

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.

----End

4.6.4 GeminiDB Influx Instance Recycle Bin

You can move unsubscribed yearly/monthly and deleted pay-per-use GeminiDB Influx instances to the recycle bin. You can also 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.

- You can put up to 100 instances into the recycle bin. If the maximum number of instances is reached, you cannot put instances into the recycle bin anymore.
- If you delete an instance running out of storage, it will not be moved to the recycle bin.
- Data of a single-node instance cannot be restored after it is moved to the recycle bin.
- 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

NOTE

You can modify the retention period, and the modifications are only applied to instances deleted after the retention period is modified, so exercise caution when performing this operation.

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 4-47 Modifying a recycling policy

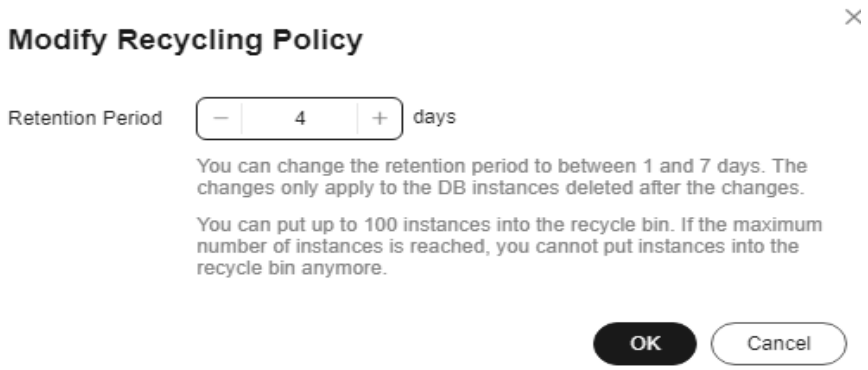

----End

Rebuilding an Instance

You can rebuild instances from the recycle bin within the retention period to restore data (Only cluster instances can be rebuilt.).

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 4-48 Rebuilding an instance

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

----End

4.7 Instance Modifications

4.7.1 Upgrading a Minor Version

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

After a new patch version involving performance improvement, new functions, or problem rectification is released, you can upgrade your instance to the latest version at a proper time based on service requirements.

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 4-49 Patch installation

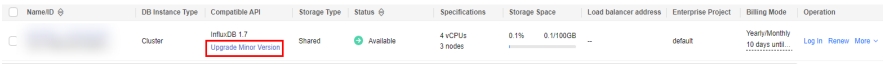

If the kernel version of your instance has potential risks or major defects, has expired, or has been brought offline, the system will notify you by SMS message or email and deliver an upgrade task during maintenance.

Precautions

- Upgrade your instance once there is a new patch released.
- If the database version is a risky version, the system prompts you to upgrade the database patch.
- Upgrading the minor version of an instance will restart each node of the instance in sequence. When a node is being restarted, its services will be taken over by another node. Each takeover will interrupt services for 5 to 10 seconds. So, perform an upgrade during off-peak hours and enable automatic reconnection so that each node can be reconnected immediately after being restarted.
- Upgrading basic components takes about 15 minutes. The duration of upgrading data components depends on the number of nodes. The upgrade of each node takes about 1 to 2 minutes.
- Services are unavailable during the parallel upgrade. You are advised to perform the parallel upgrade during off-peak hours. The parallel upgrade takes about 17 to 20 minutes, regardless of the number of nodes.

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 4-50 Patch installation

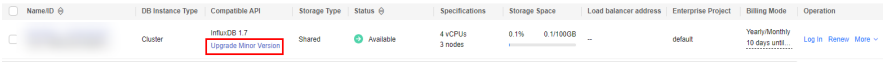

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Information** area, click **Upgrade Minor Version** in the **Compatible API** field.

Figure 4-51 Patch installation

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

- You can upgrade nodes one by one or all at once.
- You can select **Immediate** or **During maintenance window** for **Scheduled Time**. If **During maintenance window** is selected, the scheduled upgrade task will be executed in the next time window.

Figure 4-52 Confirming dialog box

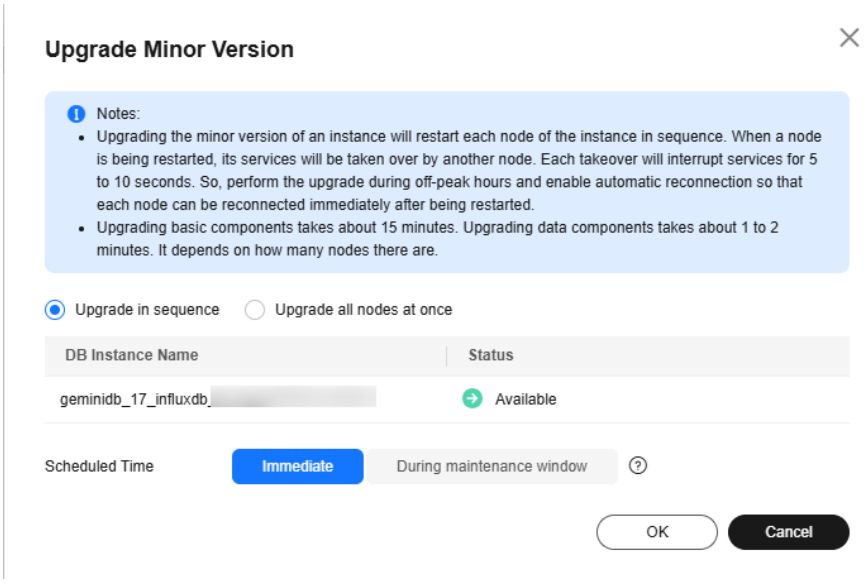

Step 5 View the upgrade result on the **Instances** page.

- When the upgrade is ongoing, the instance status is **Upgrading minor version**.

- After the upgrade is complete, the instance status changes **Available**.

----End

4.7.2 Changing a GeminiDB Influx Instance Name

Scenarios

This section describes how to change a GeminiDB Influx instance name to identify different instances.

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  next to the target instance name and change it.

- To submit the change, click **OK**.
- To cancel the change, click **Cancel**.

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 On the **Basic Information** page, click  next to **DB Instance Name** and change the instance name.

- To submit the change, click  .
- To cancel the change, click  .

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 5 View the results on the **Instances** page.

----End

4.7.3 Changing the Administrator Password of a GeminiDB Influx Database

Scenarios

For security reasons, regularly change your administrator password.

Usage Notes

- You can reset the administrator password only when your instance is in the **Available**, **Backing up**, **Checking restoration**, or **Scaling up** state. You can also choose to reset the password if an instance node becomes abnormal.
- 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***.

CAUTION

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 administrator 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 be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+?

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 target instance. The **Basic Information** page is displayed.

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 be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and any of the following special characters: ~!@#%^*-_=+?

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

4.7.4 Changing vCPUs and Memory

Scenarios

This section describes how to change instance specifications to suit your service requirements.

Usage Notes

- Instances can be scaled up or down by changing their specifications.
- 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.

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 specifications you want to change and click its name.

Step 4 In the **DB Information** area, click **Change** next to the data node specifications.

Figure 4-53 Changing specifications

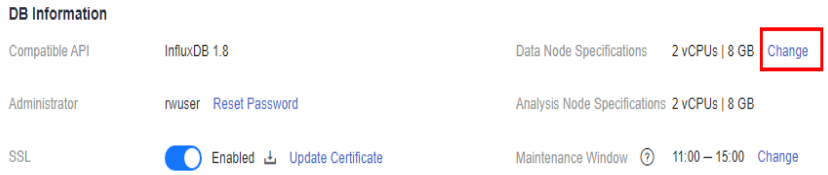

Step 5 On the displayed page, select new specifications and click **Next**.

Figure 4-54 Changing specifications

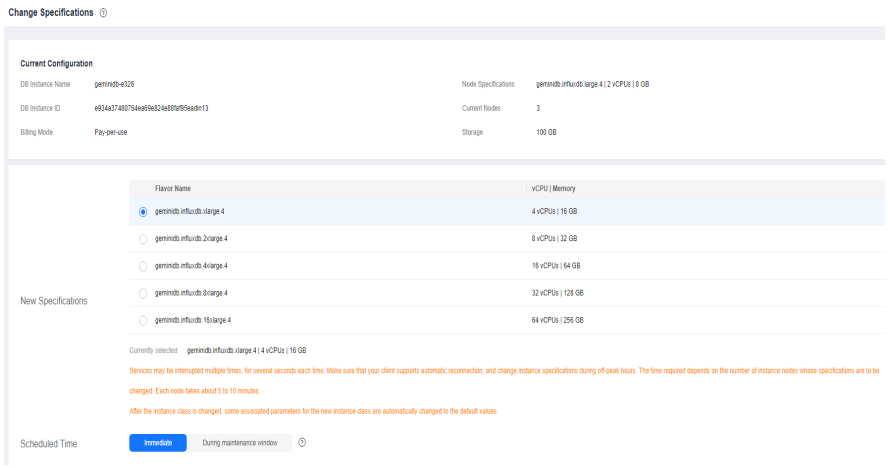

Step 6 On the displayed page, confirm the specifications.

- Yearly/Monthly
  - If you need to modify the settings, click **Previous**.
  - If you do not need to modify the 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 the settings, click **Previous**.
  - If you do not need to modify the 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 4-55 Changing specifications

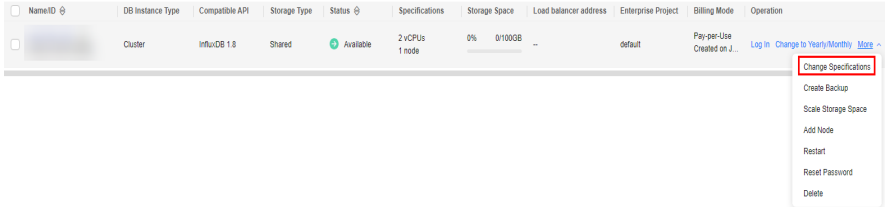

Step 4 On the displayed page, select new specifications and click **Next**.

Figure 4-56 Changing specifications

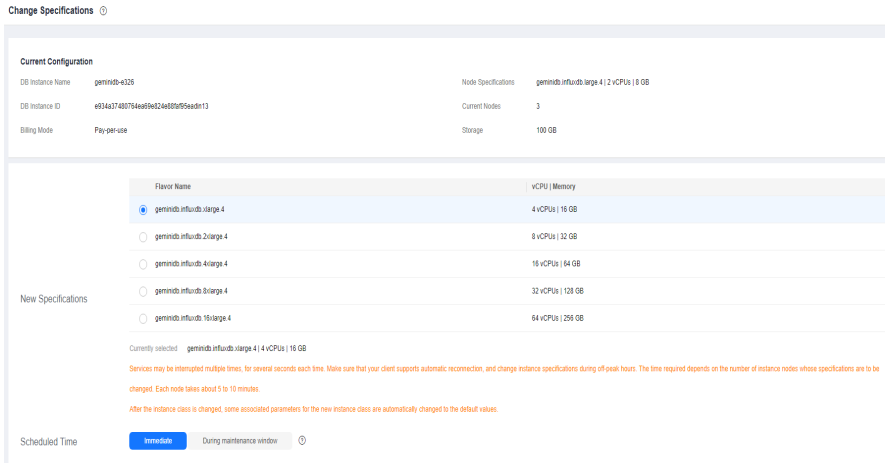

Step 5 On the displayed page, confirm the specifications.

- Yearly/Monthly
  - If you need to modify the settings, click **Previous**.
  - If you do not need to modify the 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 the settings, click **Previous**.
  - If you do not need to modify the 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

4.7.5 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 off-peak 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 On the **Basic Information** page, locate **Maintenance Window** and click **Change**.

Figure 4-57 Change

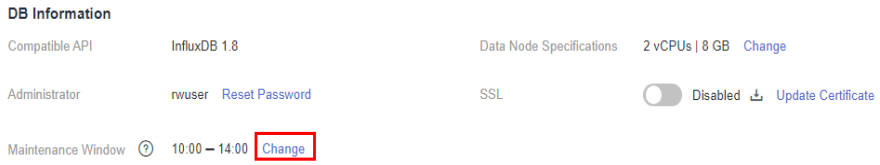

Step 5 On the **Change Maintainable Window** page, select the maintenance time period as needed, and then click **OK**.

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 4-58 Changing a maintenance window

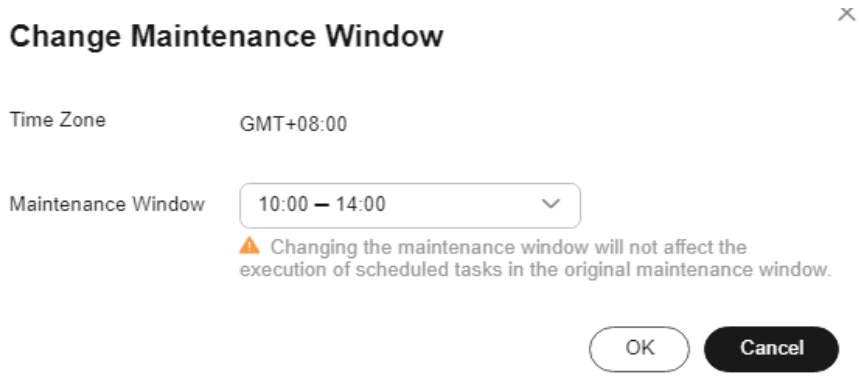

Step 6 Check the result.

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 4-59 Querying a task

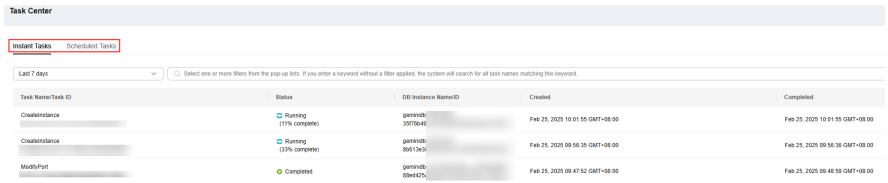

----End

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 4-60 Canceling a task

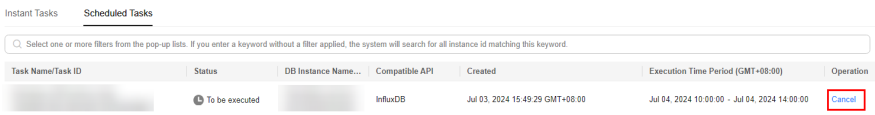

Step 4 Check the result.

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

Figure 4-61 Checking cancelled tasks

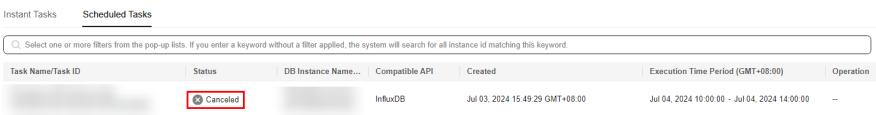

----End

4.7.6 Adding and Deleting Instance Nodes

4.7.6.1 Overview

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

Table 4-33 Scaling methods

Method	Supported Instance Type
Adding Instance Nodes	• Cluster • Cluster (performance-enhanced)
Deleting Instance Nodes	Cluster (performance-enhanced)

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 **Adding Instance Nodes**.

Figure 4-62 Adding instance nodes

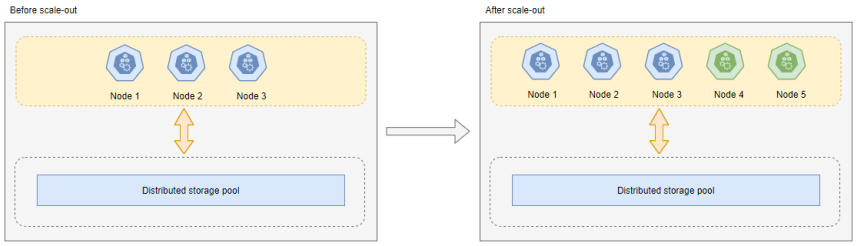

Deleting Instance Nodes

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

Figure 4-63 Deleting instance shards

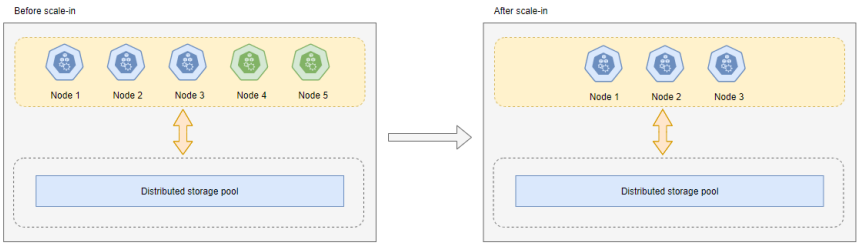

4.7.6.2 Adding Instance Nodes

Scenarios

This section describes how to add nodes to an instance to suit your service requirements. Nodes added for either a cluster or single-node instance cannot be deleted.

Usage Notes

- Adding nodes may lead to the decrease of OPS. Perform this operation during off-peak hours.
- You can only add nodes when the instance status is **Available** or **Checking restoration**.
- An instance cannot be deleted when one or more nodes are being added.
- Currently, this function is only available for cluster and cluster (performance-enhanced) instances.
- Currently, a maximum of 12 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 4-64 Adding a node

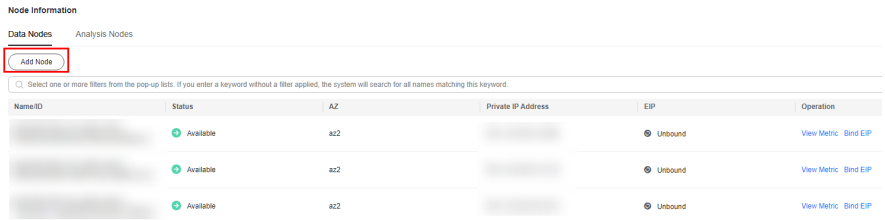

Step 5 Specify **Add Nodes** and click **Next**.

Figure 4-65 Adding a node

By default, specifications of the new node are the same as the instance specifications and cannot be modified.

Step 6 On the displayed page, confirm the node configuration details.

- 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 result of adding nodes.

- When new nodes are being added, the instance status is **Adding node**.
- After the nodes are added, the DB instance status becomes **Available**.
- Click the instance name. In the **Node Information** area on the **Basic Information** page, view the 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 instance you want to add nodes for and choose **More** > **Add Node** in the **Operation** column.

Figure 4-66 Adding a node

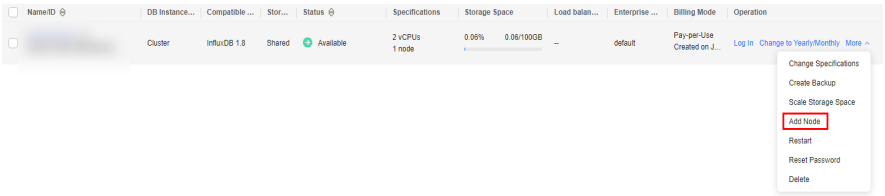

Step 4 Specify **Add Nodes** and click **Next**.

Figure 4-67 Adding a node

Step 5 On the displayed page, confirm the node configuration details.

- 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 result of adding nodes.

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

----End

4.7.6.3 Deleting Instance Nodes

Scenarios

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. Delete nodes during off-peak hours because OPS will decrease during the deletion.
- 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***.
- This function is available only to cluster (performance-enhanced) instances.

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 from which you want to delete nodes and click its name.

Step 4 In the **Node Information** area on the **Basic Information** page, locate the target node and click **Delete**.

- Yearly/Monthly
  - To delete a single node, click **Delete** in the **Operation** column.
  - To delete multiple nodes, select them and click **Delete Node**.

  Figure 4-68 Node information

  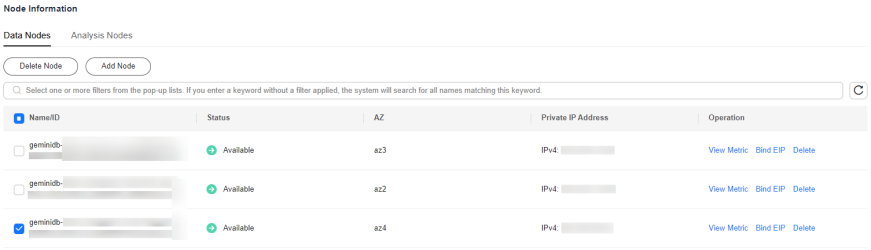

- Pay-per-use
  - To delete a single node, click **Delete** in the **Operation** column.
  - To delete multiple nodes, select them and click **Delete Node**.

Figure 4-69 Node information

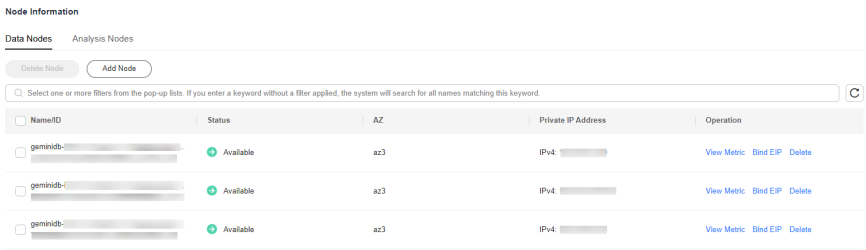

Step 5 If you have enabled operation protection, click **Start Verification** in the **Delete Node** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

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

4.7.7 Scaling Storage Space

4.7.7.1 Overview

As more data is added, you may run out of storage. This section describes how to scale up storage of your instance. As data volumes decrease, you can scale down disk space to avoid low database node utilization and resource waste. **Table 4-34** lists the scaling methods supported by GeminiDB Influx instances.

Table 4-34 Scaling methods

Method	Supported Instance Type	Description
Manually Scaling Up Storage Space of a GeminiDB Influx Instance	• Cluster (performance-enhanced) • Cluster • Single node	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.
Automatically Scaling Up Storage Space of a GeminiDB Influx Instance	Cluster (performance-enhanced)	When the available storage drops to or below the specified threshold, autoscaling will be triggered. The storage is scaled up by a percentage you specify. The added storage space is the current storage space multiplied by the scaling increment.

Method	Supported Instance Type	Description
Manually Scaling Down Storage Space of a GeminiDB Influx Instance	Cluster (performance-enhanced) instance with classic storage	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 100 GB and is increased by 50 GB, the storage space will become 150 GB.

Figure 4-70 Manually scaling up storage space

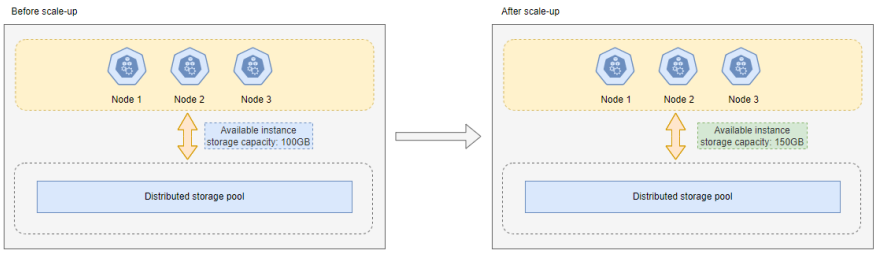

Automatically Scaling Up Storage Space

For example, the storage of a cluster (performance-enhanced) instance is 1250 GB, the storage usage threshold for triggering autoscaling is set to 10%, and the total storage needs to be automatically scaled up by 10%. If the available storage usage of an instance drops to or below 10%, the storage is automatically scaled up by 125 GB (1250 x 10%), which is rounded up to 130 GB. In this case, the total storage becomes 1380 GB (1250 + 130).

Figure 4-71 Automatically scaling up storage space

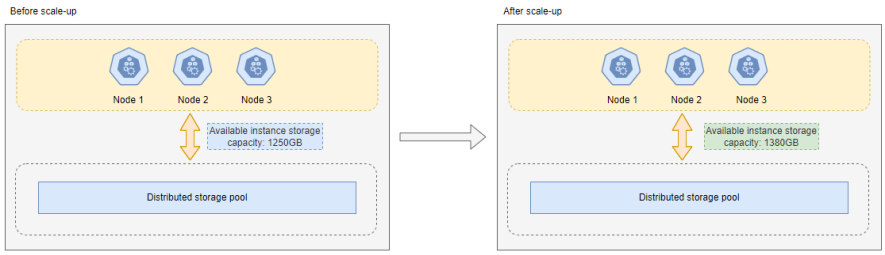

Manually Scaling Down Storage Space

For example, the storage of a cluster (performance-enhanced) instance is 200 GB. If the storage is scaled down by 10 GB, it becomes 190 GB.

Figure 4-72 Manually scaling down storage space

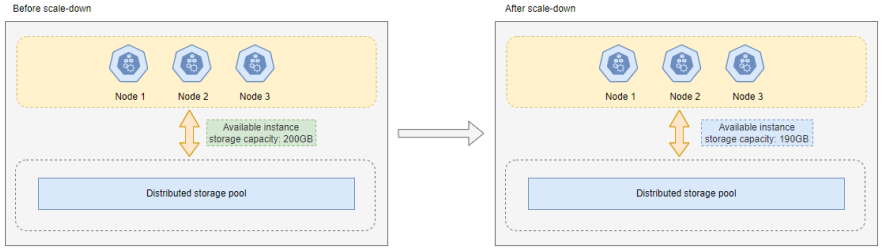

4.7.7.2 Manually Scaling Up Storage Space of a GeminiDB Influx Instance

Scenarios

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

Storage scaling does not interrupt your services. After storage scaling is complete, you do not need to restart your instance.

Usage Notes

- Currently, cloud native storage of cluster, single-node, or performance-enhanced cluster instances can only be scaled up.

Setting an Instance Status to Read-Only

To ensure that the GeminiDB Influx instance can still run properly when the storage space is about to be used up, the database is set to read-only, and data cannot be modified. If this happens, you can scale up the storage to restore the database status to read/write.

Table 4-35 Setting an instance status to read-only

Storage	Description
< 600 GB	• When the storage usage reaches 97%, the instance status is set to read-only. • When the storage usage decreases to 85%, the read-only status is automatically disabled for the instance.
≥ 600 GB	• If the remaining storage space is less than 18 GB, the instance status is set to read-only. • When the remaining storage space is greater than or equal to 90 GB, the read-only status is automatically disabled for the instance.

The kernel uses an LSM architecture. When written or deleted data reaches a certain amount, it will be merged. New data and old data to be deleted are stored together, and the disk usage increases temporarily based on the amount of

merged data. In this case, the read-only status may be triggered. You are advised to reserve sufficient disk space to prevent the read-only status.

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 storage space you want to scale up and click its name.

Step 4 In the **Storage Space** area on the **Basic Information** page, click **Scale**.

Figure 4-73 Scaling up storage space

Step 5 On the displayed page, specify new storage and click **Next**.

Step 6 On the displayed page, confirm the storage space.

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

Step 7 Check the results.

- When the scale-up task is ongoing, the instance status is **Scaling up**.
- After the scale-up task is complete, the instance status becomes **Available**.
- In the **Storage Space** area on the **Basic Information** page, check whether the scale-up is successful.

----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 storage you want to scale up and choose **More** > **Scale Storage Space** in the **Operation** column.

Figure 4-74 Scaling up storage space

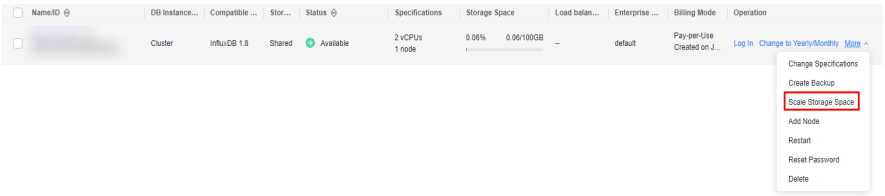

Step 4 On the displayed page, specify new storage and click **Next**.

Figure 4-75 Scaling up storage space

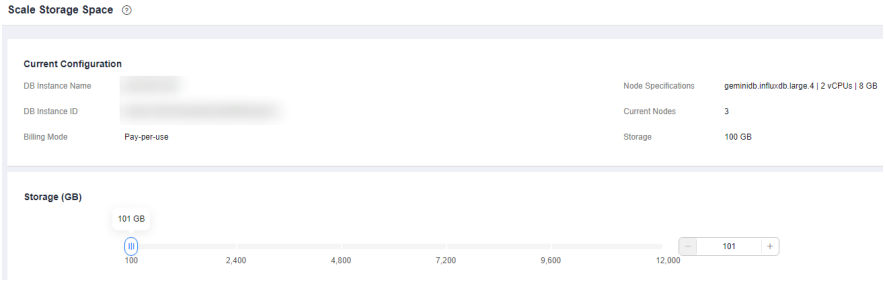

Select at least 1 GB each time, and the value must be an integer.

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

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

Step 6 Check the results.

- When the scale-up task is ongoing, the instance status is **Scaling up**.
- After the scale-up task is complete, the instance status becomes **Available**.
- In the **Storage Space** area on the **Basic Information** page, check whether the scale-up is successful.

----End

4.7.7.3 Automatically Scaling Up Storage Space of a GeminiDB Influx Instance

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

You can enable storage autoscaling when or after creating an instance. For details, see **Buying a GeminiDB Influx Instance**.

This section describes how to configure storage autoscaling after an instance is created.

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 fine-grained permissions for IAM and least permissions for GeminiDB.

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

   ```
   {
       "Version":"1.1",
       "Statement":[
           {
               "Effect":"Allow",
               "Action":[
                   "iam:permissions:listRolesForAgencyOnProject",
                   "iam:permissions:grantRoleToGroupOnProject",
                   "iam:agencies:createAgency",
                   "iam:agencies:listAgencies",
                   "iam:roles:listRoles",
                   "iam:roles:createRole"
               ]
           }
       ]
   }
   ```

2. **Create a user group and assign permissions to it.**

   You can create a user group on the IAM console and grant it custom permissions created in **1** and the security administrator role.

3. **Add a user to a user group.**

   Log in to the IAM console using a Huawei Cloud account or an IAM account, locate the IAM user that the target instance belongs to, and add it to the user group created in **2**. The IAM user will inherit permissions of the user group.

Usage Notes

- **Auto Scale** 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.
- Currently, only GeminiDB Influx instances in a performance-enhanced cluster support autoscaling.
- By default, when the storage is used up (the available storage space is less than or equal to 10% or the available storage is less than or equal to 10 GB), autoscaling is triggered.

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 4-76 Auto Scale

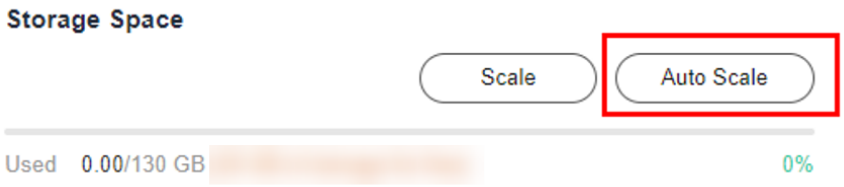

Step 5 Toggle on **Auto Scale** and specify the trigger condition and increment.

Figure 4-77 Configuring autoscaling parameters

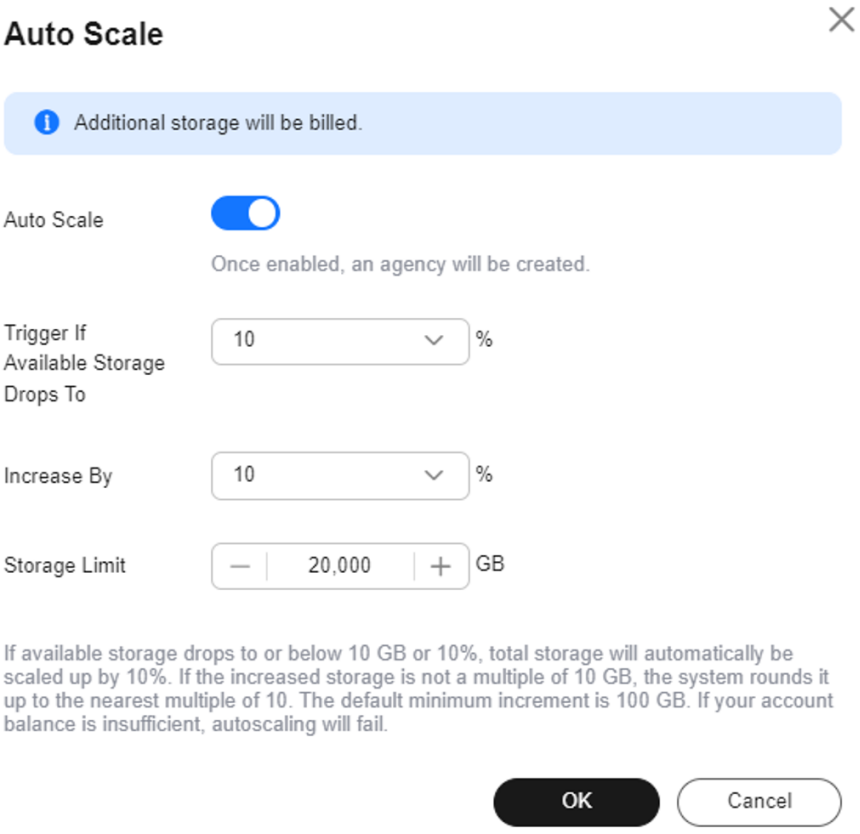

Table 4-36 Parameters

Parameter	Description
Auto Scale	If you toggle on this switch, autoscaling is enabled.
Trigger If Available Storage Drops To	If the available storage usage drops to a specified threshold (10%, 15%, or 20%), autoscaling is triggered.
Increase By	Percentage of the current storage to be automatically 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 that can be automatically scaled up to

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 In the instance list, select the target instances and click **Auto Scale**.

Figure 4-78 Auto Scale

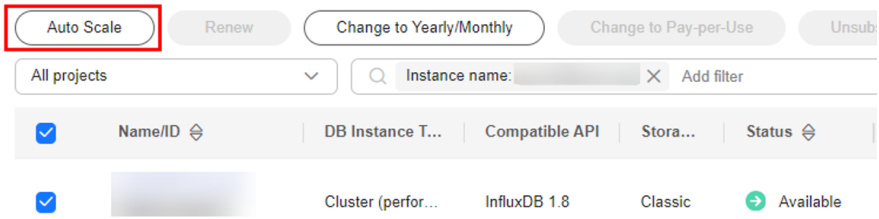

Step 4 Toggle on **Auto Scale** and specify trigger conditions.

Figure 4-79 Configuring autoscaling parameters

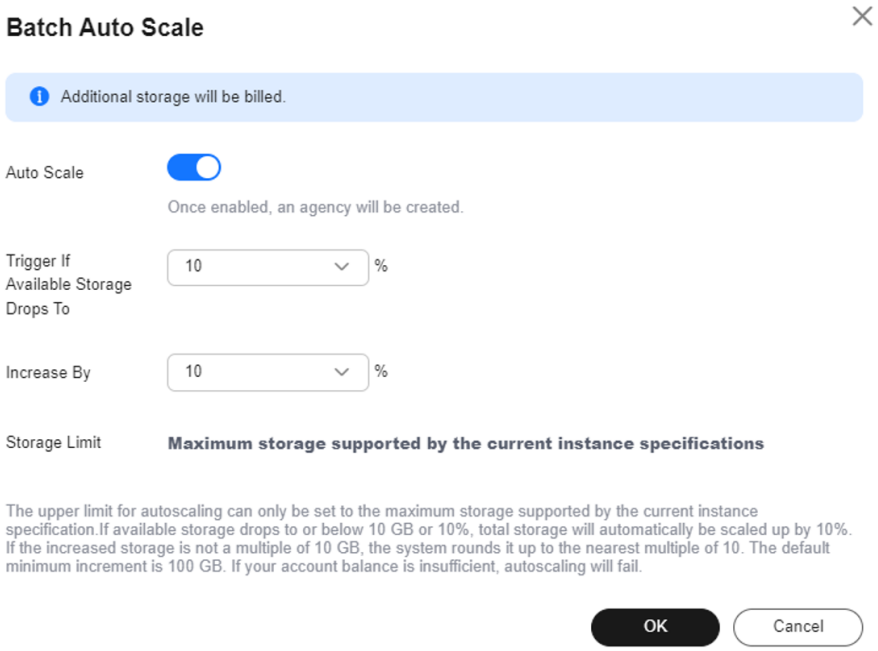

Table 4-37 Parameters

Parameter	Description
Auto Scale	If you toggle on this switch, autoscaling is enabled.
Trigger If Available Storage Drops To	If the available storage usage drops to a specified threshold (10%, 15%, or 20%), autoscaling is triggered.
Increase By	Percentage of the current storage to be automatically 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 that can be automatically scaled up to

Step 5 Click **OK**.

----End

4.7.7.4 Manually Scaling Down Storage Space of a GeminiDB Influx Instance

Scenarios

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.
- Storage scaling does not interrupt your services. After storage space is scaled, you do not need to restart your instance.
- You can only scale down classic storage of instances in a performance-enhanced cluster.

Setting an Instance Status to Read-only

To ensure that the GeminiDB Influx instance can still run properly when the storage space is about to be used up, the database is set to read-only, and data cannot be modified. If this happens, you can scale up the storage to restore the database status to read/write.

Table 4-38 Setting an Instance Status to Read-only

Storage	Description
< 600 GB	• When the storage usage reaches 97%, the instance is read-only. • When the storage usage decreases to 85%, the read-only status is automatically disabled for the instance.
≥ 600 GB	• When the remaining storage space is less than 18 GB, the instance is read-only. • When the remaining storage space is greater than or equal to 90 GB, the read-only status is automatically disabled for the 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, locate the target instance and choose **More** > **Scale Storage Space** in the **Operation** column.

Figure 4-80 Scaling storage space

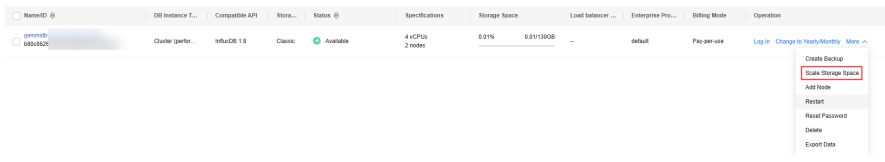

Step 4 On the displayed page, specify the new storage space and click **Next**.

Figure 4-81 Scaling storage space

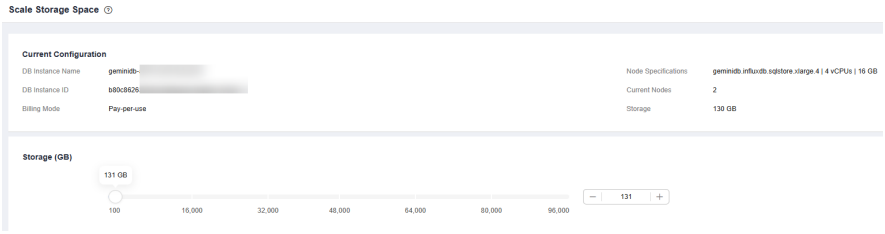

Select at least 1 GB each time, and the value must be an integer.

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

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

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.

Step 4 In the **Specification Information** area on the **Basic Information** page, click **Scale** .

Figure 4-82 Scaling storage space

Step 5 On the displayed page, specify the new storage space and click **Next**.

Figure 4-83 Scaling storage space

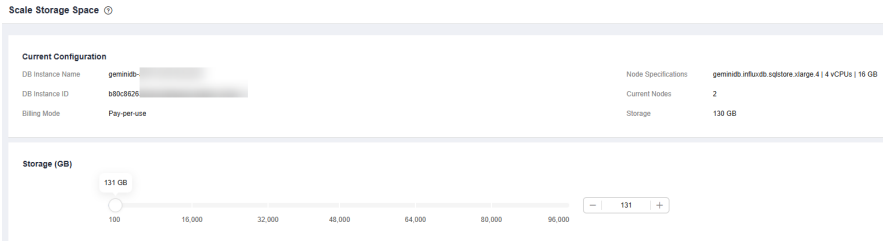

Select at least 1 GB each time, and the value must be an integer.

Step 6 On the displayed page, confirm the storage space.

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

Step 7 Check the results.

- During the scale-down process, the instance status becomes **Scaling storage space**.
- After the scaling task is complete, 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

4.8 Database Commands

4.8.1 Supported Commands

The following table lists the commands supported by GeminiDB Influx API.

NOTE

For GeminiDB Influx common commands, basic syntax, and examples, see **Buying and Connecting to a GeminiDB Influx Instance**.

Usage Notes

DROP MEASUREMENT can be used only for instances in a performance-enhanced cluster.

User Management

Table 4-39 Commands supported by user management

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
create user	√	√
show user	√	√
drop user	√	√
set password	√	√
grant	√	√
show grants	√	√
revoke	√	√

CLI Commands Used on an InfluxDB Client

Table 4-40 CLI commands used on an InfluxDB client

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
connect	√	√
auth	√	√
pretty	√	√
chunked	√	√
chunk size	√	√
use	√	√
format	√	√
precision	√	√
consistency	√	√
history	√	√
settings	√	√
clear	√	√
exit/quit/ctrl+d	√	√

Metadata Management

Table 4-41 Commands supported by metadata management

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
create database	√	√
show databases	√	√
drop database	√	√
show measurements	√	√
show measurement cardinality	√	√
show measurement exact cardinality	√	√
drop measurement	√	√
create retention policy	√	√
alter retention policy	√	√
drop retention policy	√	√
show retention policies	√	√
create continuous query	√	√
show continuous queries	√	√
drop continuous query	√	√
show series	√	√
show series cardinality	√	√
show series exact cardinality	√	√
drop series	×	×
show tag keys	√	√
show tag key cardinality	√	√
show tag key exact cardinality	√	√
show tag values	√	√
show tag values cardinality	√	√
show tag values exact cardinality	√	√

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
show field keys	√	√
show field key cardinality	√	√
show field key exact cardinality	√	√
show shards	√	√
show shard groups	√	√
drop shard	√	√

Monitoring and Management of Queries

Table 4-42 Commands for monitoring and management of queries

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
kill query	√	√
show queries	√	√

Querying, Writing, and Deleting Data Points

Table 4-43 Commands supported by data points

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
select	√	√
select xxx into	√	√
insert into	√	×
insert	√	×
limit	√	√
offset	√	√
delete	×	×
explain	√	√
explain analyze	√	√

Aggregate Functions

Table 4-44 Commands supported by aggregate functions

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
count	√	√
distinct	√	√
integral	√	√
mean	√	√
median	√	√
mode	√	√
spread	√	√
stddev	√	√
sum	√	√

SELECT Function

Table 4-45 Commands supported by the SELECT function

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
bottom	√	√
top	√	√
first	√	√
last	√	√
max	√	√
min	√	√
percentile	√	√
sample	√	√

Conversion Function

Table 4-46 Commands supported by the conversion function

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
abs	√	√
acos	√	√
asin	√	√
atan	√	√
atan2	√	√
ceil	√	√
cos	√	√
sin	√	√
tan	√	√
sqrt	√	√
round	√	√
floor	√	√
exp	√	√
ln	√	√
log2	√	√
log10	√	√
log	√	√
pow	√	√
cumulative_sum	√	√
difference	√	√
non_negative_difference	√	√
derivative	√	√
non_negative_derivative	√	√
elapsed	√	√
moving_average	√	√

NOTE

√ indicates that an item is supported, and × indicates that an item is not supported.

4.9 Cold and Hot Data Separation

4.9.1 Enabling Cold Storage

Cold storage is mainly used to store historical data with low query frequency. As the amount of historical data increases, the need to reduce storage costs becomes necessary. GeminiDB Influx provides cold storage to help you store cold data at low costs in just a few clicks.

In addition, GeminiDB Influx can separate cold data from hot data based on the retention policy. If you need to separate cold data from hot data, create cold storage and set the **time boundary between hot and cold data**. In this way, hot data will be automatically archived in cold storage after the retention period expires.

Both new and existing instances support cold storage. This section describes how to create cold storage.

Usage Notes

- Cold data cannot be written.
- Cold storage is supported only when the kernel version of an instance with classic storage in a performance-enhanced cluster is 1.0.0.240200, when the kernel version of an instance with cloud native storage in a performance-enhanced cluster is 1.7.4.250700, and when the kernel version of a cluster instance is 1.7.4.6. If the kernel version is earlier than 1.7.4.6, you can upgrade it by following **Upgrading a Minor Version**.
- Cold data of GeminiDB Influx instances will not be backed up.
- Cold storage cannot be disabled after being enabled.
- For instances with classic or cloud native storage in a performance-enhanced cluster and cluster instances, this function is now in OBT. You can choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Creating Cold Storage for a New Instance

You can specify **Purchase Cold Storage** on the page for purchasing an instance. For details, see **Buying a GeminiDB Influx Instance**.

Creating Cold Storage for an Existing Instance

If you select **No** for **Purchase Cold Storage** on the page for purchasing an instance. To create cold storage, you can perform the following steps:

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 cold storage for and click its name.

Step 4 In the **Cold Storage** area on the **Basic Information** page, click **Create**.

Figure 4-84 Creating cold storage

Step 5 On the displayed page, specify the amount of cold storage and click **Next**.

Figure 4-85 Specifying cold storage

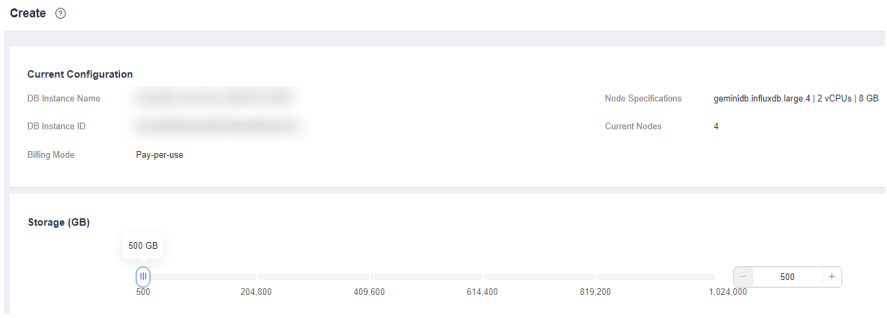

The cold storage is an integer from 500 GB to 1,024,000 GB. You can add a minimum of 1 GB each time you scale up storage space.

Step 6 On the displayed page, confirm the cold 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 the settings, click **Submit**.

Step 7 Check the results.

- When the cold storage is being created, the instance status is **Creating cold storage**.
- After the cold storage is created, the instance status becomes **Available**.

- Click the instance name. In the **Cold Storage** area on the **Basic Information** page, you can view the cold storage capacity after the cold storage is created.

----End

4.9.2 Cold and Hot Data Separation

GeminiDB Influx allows you to separate cold and hot data based on the retention policy (RP). You can configure data retention duration and number of backups, and then the system automatically archives hot data that meets the conditions to cold storage.

Background

In big data scenarios, cold data and hot data is distinguished. Historical time-series data is less likely to be queried and analyzed as time goes by. In addition, the historical data will take up space that may increase storage costs. Therefore, it is necessary for enterprises to reduce cold data storage costs. Cold data and hot data are separately stored on a GeminiDB Influx instance. You can store cold data on low-cost media with just a few clicks.

Cold and hot data separation is based on the RP. You need to set a time boundary between cold and hot data in the RP, and the system will automatically archives cold data to cold storage. When you query data, the system will automatically retrieve it from hot or cold data storage based on the time range you specify.

Principles

You can configure the retention period of hot data. When data is written, it is stored in the hot storage first. GeminiDB Influx determines whether the data is hot or cold based on the data timestamp. If the data timestamp is within the hot data storage duration, the data is still hot. Otherwise, the hot data will be automatically archived in cold storage.

Figure 4-86 Diagram

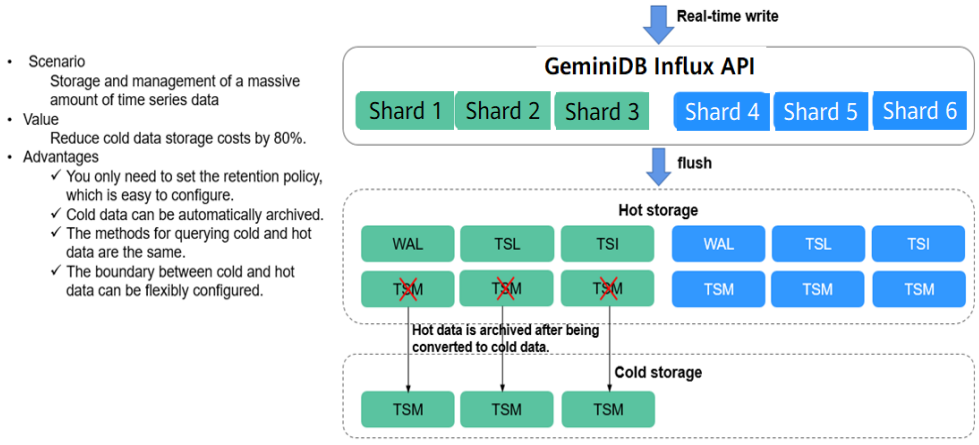

Basic Usage

1. Set the cold and hot time boundary.

Specify **WARM DURATION** in the RP. Data generated before the value of **WARM DURATION** is cold data.

To set **WARM DURATION**, perform the following steps:

```
//Create an RP named myrp for database named mydb. The value of WARM DURATION is 6d, indicating that data generated six days ago is cold data.
create retention policy myrp on mydb duration 30d replication 1 warm duration 6d shard duration 3d
//Create an RP named myrp for database mydb. If WARM DURATION is not specified, no cold data exists.
create retention policy myrp on mydb duration 30d replication 1 shard duration 3d
//Create a database named mydb with an RP named myrp. The value of WARM DURATION is 3d, indicating that data generated three days ago is cold data.
create database mydb with duration 6d warm duration 3d name myrp
//Change the value of WARM DURATION to 7d, indicating that data generated seven days ago is cold data.
alter retention policy myrp on mydb warm duration 7d
```

2. Write data to the storage.

   Hot and cold data is written in the same way. Data is first stored in the hot storage when being written. As time goes by, if the timestamp of the data in the hot storage exceeds the value of **WARM DURATION**, the system automatically archives the data to the cold storage. This process is completely transparent to the user.

3. Query data.

   The methods for querying hot and cold data are the same. During data query, the system automatically queries hot or cold storage based on the TimeRange condition in the query statement. This process is completely transparent to the user. The response to a cold data query is longer than that to a hot data query.

4. Check the status of hot and cold data.

```
> show shards
name: _internal
id database  retention_policy shard_group start_time           end_time
expiry_time          owners tier
-- --------  ---------------- ----------- ----------          --------            -----------          ------
----
1  _internal monitor          1           2021-06-29T00:00:00Z 2021-06-30T00:00:00Z
2021-07-07T00:00:00Z 4      warm
2  _internal monitor          1           2021-06-29T00:00:00Z 2021-06-30T00:00:00Z
2021-07-07T00:00:00Z 5      warm
3  _internal monitor          1           2021-06-29T00:00:00Z 2021-06-30T00:00:00Z
2021-07-07T00:00:00Z 7      warm
4  _internal monitor          1           2021-06-29T00:00:00Z 2021-06-30T00:00:00Z
2021-07-07T00:00:00Z 6      warm

name: hsdb
id database retention_policy shard_group start_time           end_time
expiry_time          owners tier
-- -------- ---------------- ----------- ----------          --------            -----------          ------
----
5  hsdb     myrp             2           2019-08-12T00:00:00Z 2019-08-19T00:00:00Z
2019-08-19T00:00:00Z 4      cold
6  hsdb     myrp             2           2019-08-12T00:00:00Z 2019-08-19T00:00:00Z
2019-08-19T00:00:00Z 5      moving
7  hsdb     myrp             2           2019-08-12T00:00:00Z 2019-08-19T00:00:00Z
2019-08-19T00:00:00Z 6      warm
8  hsdb     myrp             2           2019-08-12T00:00:00Z 2019-08-19T00:00:00Z
2019-08-19T00:00:00Z 7      cold
```


- If the **tier** value is **cold**, the current shard stores cold data.
- If the **tier** value is **warm**, the current shard store hot data.
- If the **tier** value is **moving**, the current shard is being changed from hot data to cold data.
- The process of changing hot data to cold data involves only the transfer of TSM files from hot storage to cold storage. Other files of the shard are still stored in hot storage and do not need to be moved.

4.9.3 Scaling Up Cold Storage

Scenarios

If the existing cold storage cannot meet your service requirements, scale up it.

Usage Notes

- Scaling down cold storage does not interrupt your services, and you do not need to restart your instances.
- Cluster instance storage can only be scaled up. Scale-down is not supported.

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 whose cold storage you want to scale up and click its name.

Step 4 In the **Cold Storage** area on the **Basic Information** page, click **Scale** for an instance.

Figure 4-87 Scaling up cold storage of cluster and single-node instances

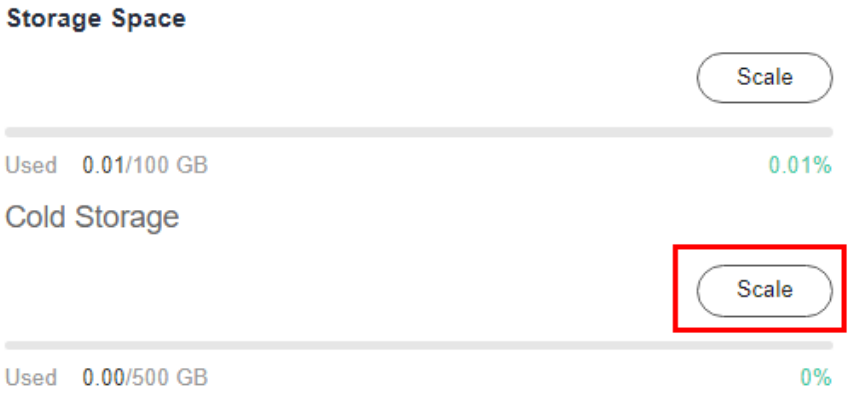

Figure 4-88 Scaling up cold storage of cluster (performance-enhanced) instances

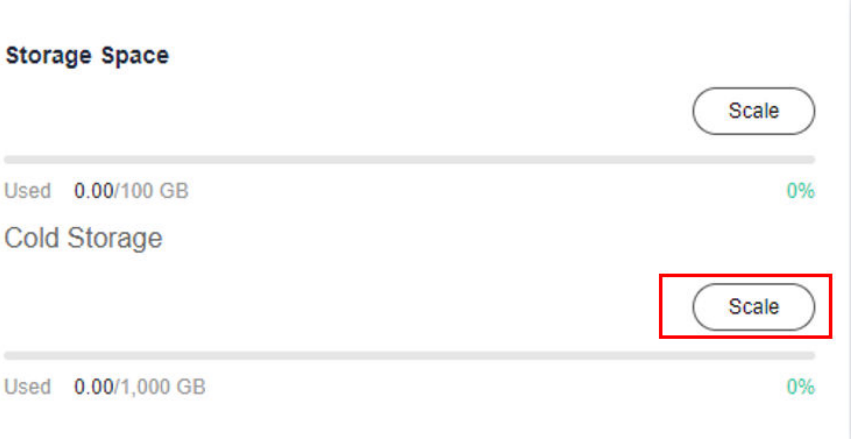

Step 5 On the displayed page, specify desired cold storage space and click **Next**.

Figure 4-89 Scaling up cold storage of cluster and single-node instances

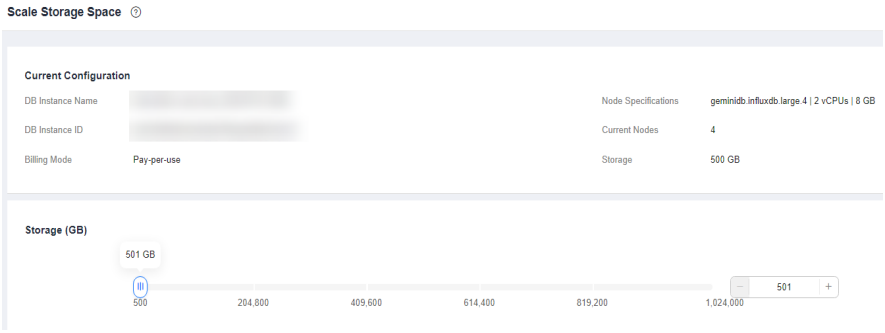

Figure 4-90 Scaling up cold storage of cluster (performance-enhanced) instances

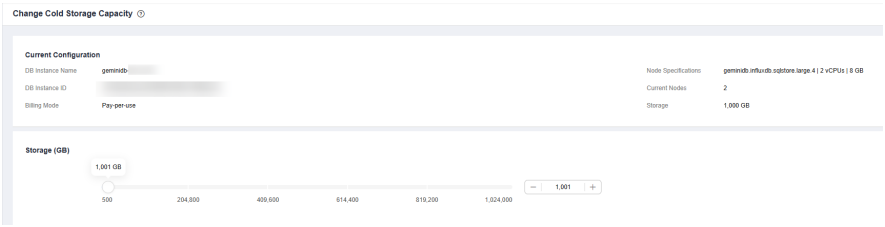

Select at least 1 GB each time, and the value must be an integer.

Step 6 On the displayed page, confirm the cold storage space.

- For yearly/monthly instances
  - If you need to modify your settings, click **Previous**.
  - If you do not need to modify your settings, click **Next** and complete the payment.
- For pay-per-use instances
  - If you need to modify your settings, click **Previous**.
  - If you do not need to modify your settings, click **Submit**.

Step 7 Check the scale-up result.

- During the scale-up, the instance status becomes **Scaling up Cold storage** or **Changing cold storage capacity**.
- After the scale-up is complete, the instance status becomes **Available**.
- Click the instance name. In the **Cold Storage** area on the **Basic Information** page, you can view the new cold storage.

----End

4.9.4 Scaling Down Cold Storage

Scenarios

Cold storage of GeminiDB Influx cluster (performance-enhanced) instances can be scaled down when the requested cold storage space needs to be released.

Usage Notes

- Scaling down cold storage does not interrupt your services, and you do not need to restart your instances.
- Currently, only cluster (performance-enhanced) instances support this function.

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 name.

Step 4 In the **Cold Storage** area on the **Basic Information** page, click **Scale**.

Figure 4-91 Changing the cold storage capacity

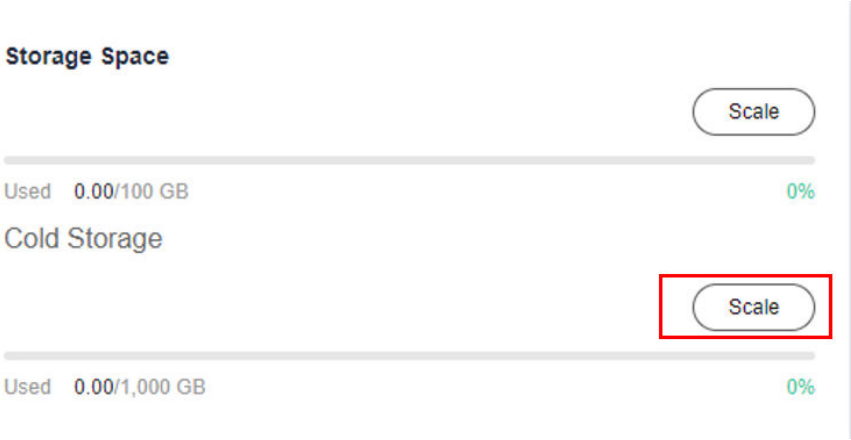

Step 5 On the displayed page, specify the cold storage space and click **Next**.

Figure 4-92 Scaling down cold storage

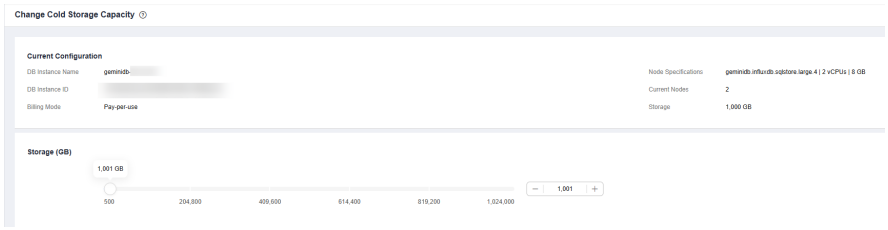

Step 6 On the displayed page, confirm the cold storage space.

- Yearly/Monthly
  - To modify your settings, click **Previous**.
  - If you do not need to modify your settings, click **Submit**.
- Pay-per-use
  - To modify your settings, click **Previous**.
  - If you do not need to modify your settings, click **Submit**.

Step 7 Check the scale-down result.

- During the scale-down, the instance status becomes **Changing cold storage capacity**.
- After the scale-down is complete, the instance status becomes **Available**.
- Click the instance name. In the **Cold Storage** area on the **Basic Information** page, you can view the new cold storage.

----End

4.10 Certificate Management

4.10.1 Downloading the Default SSL Certificate

Scenarios

Secure Sockets Layer (SSL) certificates set up encrypted connections between clients and servers, preventing data from being tampered with or stolen during transmission.

To improve data security, GeminiDB Influx provides a default SSL certificate. When creating an instance, you can enable SSL to encrypt connections to the instance.

This section describes how to obtain the default SSL security certificate provided by GeminiDB Influx.

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 whose SSL certificate you want to download and click the instance name.

Step 4 In the **DB Information** area, click in the **SSL** field.

Figure 4-93 Downloading the SSL certificate

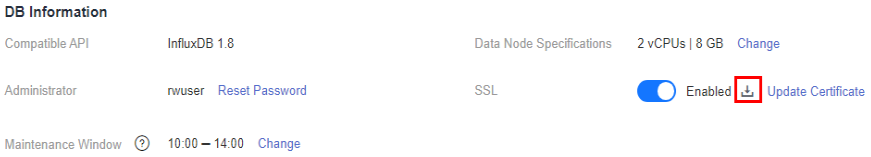

----End

4.10.2 Configuring a CCM Private Certificate

Scenarios

GeminiDB Influx allows you to use the certificate issued by Cloud Certificate Management Service (CCM) to connect to your DB instance. You can select a CCM certificate when you create an instance or update its certificate after the instance is created.

This section describes how to apply for a CCM private certificate to a DB instance in either of the following ways:

1. Select a certificate when you create an instance.
2. Update the certificate after the instance is created.

Precautions

The instance status is **Available**.

Prerequisites

You have created a CCM private certificate. If there are no CCM private certificates, you can apply for a private certificate by referring to **Applying for a Private Certificate** in the *Cloud Certificate Manager User Guide*.

NOTE

- When requesting a private certificate, specify the IP address of the target instance in the **Configure Certificate AltName** area If this parameter is not specified, the database connection will fail.
  - If you select a certificate when creating an instance, you can only add an EIP in the **Configure Certificate AltName** area. Because the instance has not been created, the system cannot assign a private IP address for it.
  - When you update the certificate after an instance is created, you can add private IP addresses or EIPs of all the instance nodes in the **Configure Certificate AltName** area.

Figure 4-94 Creating a CCM private certificate

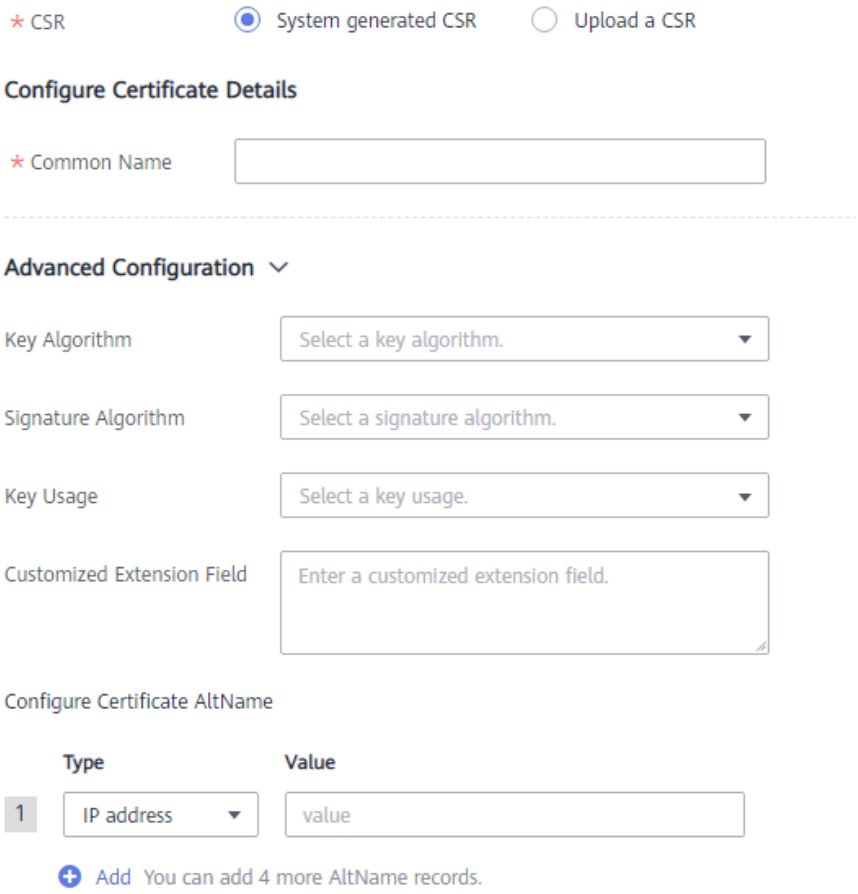

- For details about how to set other parameters, see **Applying for a Private Certificate** in *Cloud Certificate Manager User Guide*.

Scenario 1: Configuring a Private Certificate When Creating an 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 **Buy DB Instance**.

Step 4 On the displayed page, specify required parameters and click **Next**.

- Enable SSL and select an existing CCM private certificate. If there are no certificates available, apply for a certificate by referring to **Prerequisites**.

  Figure 4-95 Selecting a certificate

  

- Configure other parameters by following **Buying a GeminiDB Influx Cluster Instance**.

Step 5 After the instance is created, click its name to go to the **Basic Information** page. In the **DB Information** area, check whether the certificate status is **Available**.

Figure 4-96 Viewing the certificate status

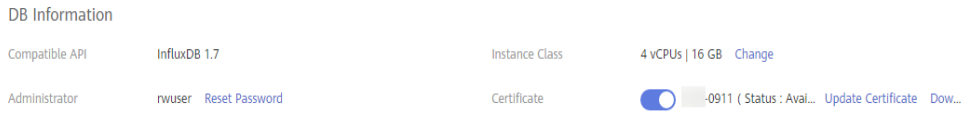

Step 6 Download the certificate.

Click **Download** in the **Certificate** field. On the displayed page, click the **Nginx** tab and click **Download Certificate**.

Figure 4-97 Downloading the certificate

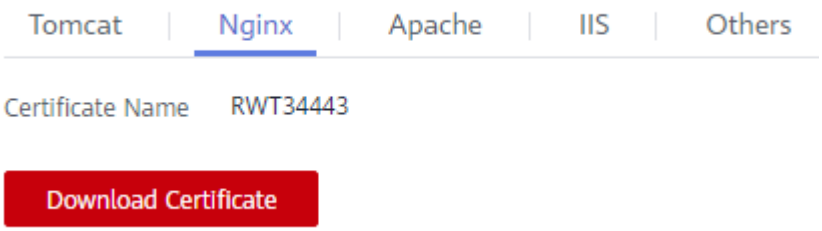

----End

Scenario 2: Updating a Certificate After an Instance Is Created

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 certificate you want to update.

Step 4 In the **DB Information** area, click **Update Certificate** in the **SSL** field.

Figure 4-98 Updating the certificate

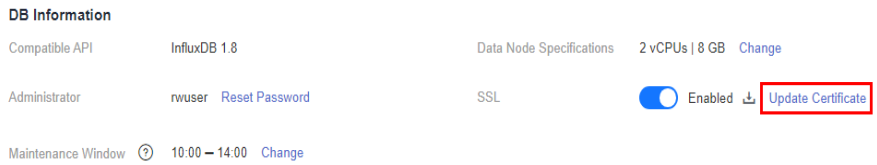

Step 5 In the **Update Certificate** dialog box, select the required certificate and click **OK**.

Figure 4-99 Selecting a certificate

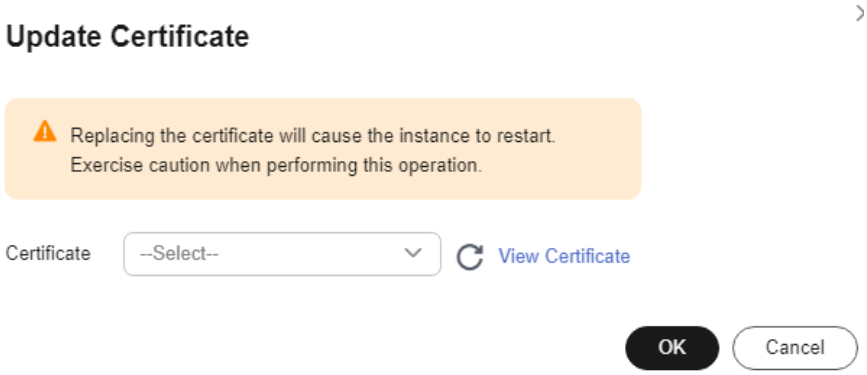

NOTE

- The new certificate takes effect only after the instance is restarted. Perform this operation during off-peak hours to minimize impacts on your services.
- The certificate cannot be changed to the default SSL certificate.

Step 6 After the certificate is updated, check whether the certificate status is **Available** on the **Basic Information** page.

Figure 4-100 Viewing the certificate status

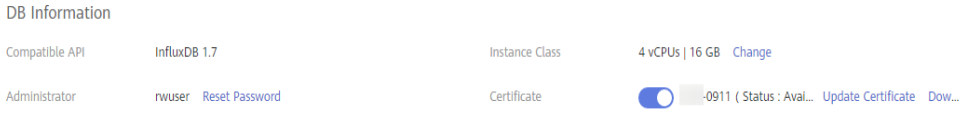

Step 7 Download the certificate.

Click **Download** in the **Certificate** field. On the displayed page, click the **Nginx** tab and click **Download Certificate**.

Figure 4-101 Downloading the certificate

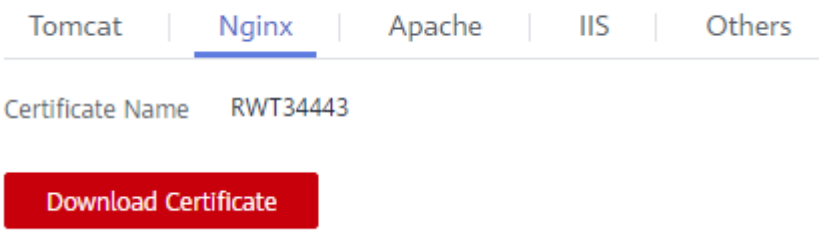

----End

4.11 Data Backup

4.11.1 Overview

You can back up GeminiDB Influx instances to protect your data. 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

This function of instances in a performance-enhanced cluster 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.

Backup Methods

GeminiDB Influx instances support both automatic and manual backups.

- 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 for a length of time you specified.

  Automated backups cannot be manually deleted. You can adjust their retention period by referring to **Modifying an Automated Backup Policy** , and backups that expire 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. You can create a manual backup for your instance at any time to meet service requirements.

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

Table 4-47 Backup methods

Method	Scenario
Automated backup	After you configure a backup policy, the system automatically backs up your database based on the policy. You can also modify the policy based on service requirements.
Manual backup	You can manually create full backups for your instance based on service requirements.

Backup process

As shown in **Figure 4-102**, there are three nodes in the GeminiDB Influx cluster for backing up data. Data snapshots are taken in seconds, and the generated backup files are compressed and stored in OBS, without occupying extra storage space of the GeminiDB Influx instance. The CPU usage may increase 5% to 15% because uploading backups consumes CPU resources.

Figure 4-102 Backup process

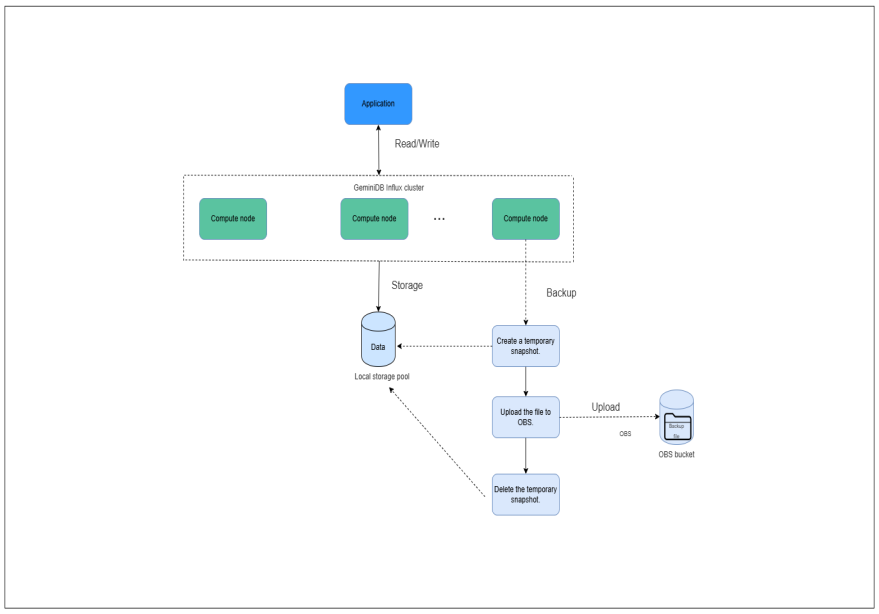

Backup Storage

Backups are stored in OBS buckets to provide disaster recovery and save storage space.

After you purchase an instance, GeminiDB Influx will provide additional backup storage of the same size as what you purchased. For example, if you purchase an instance with 100 GB of storage, you will get another 100 GB of storage free of charge. If the backup data does not exceed 100 GB, it is stored on OBS free of charge. If there is more than 100 GB of data, you will be billed at standard OBS rates.

4.11.2 Managing Automated Backups

GeminiDB Influx creates automated backups to ensure data reliability. If a database or table is deleted, maliciously or accidentally, backups can help recover your data.

NOTE

GeminiDB Influx does not back up **cold storage data**.

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.

Figure 4-103 Enabling the automated backup policy

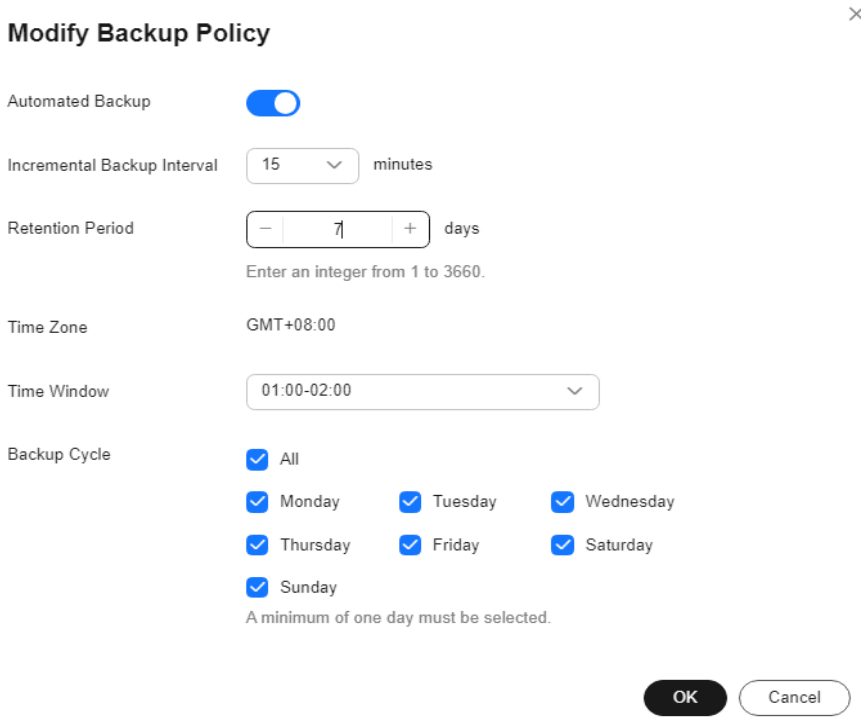

- **Retention Period**: Automated backup files are saved for seven days by default. The retention period ranges from 1 to 3660 days. Full backups are retained till the retention period expires. However, even if the retention period has expired, the most recent backup will be retained.
  - Extending the 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 shorter 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 for, such as 12:00-13:00. The backup time is in GMT format. 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 expires, the most recent backup will be retained, for example,

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.

– The full backup generated on Tuesday will be automatically deleted on Wednesday of the following week. 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.

  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.

- **Incremental Backup Interval**: Incremental backups are generated every 15 minutes. The incremental backup interval can be set to 5 minutes, 10 minutes, or 15 minutes. 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.

- After an instance is created, you can set an automated backup policy. The system will back up data based on the automated backup policy.
- 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 4-104 Modifying the backup policy

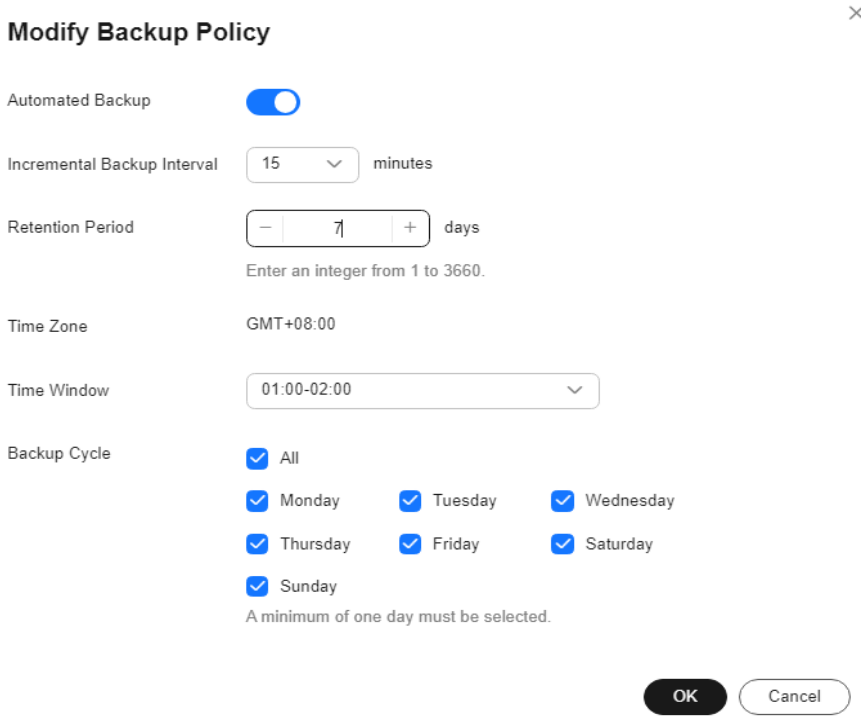

Step 5 Check or manage the generated backups on the **Backups** or **Backups & Restorations** page.

----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  to disable automatic backup and click **OK**.

Figure 4-105 Disabling backup policies

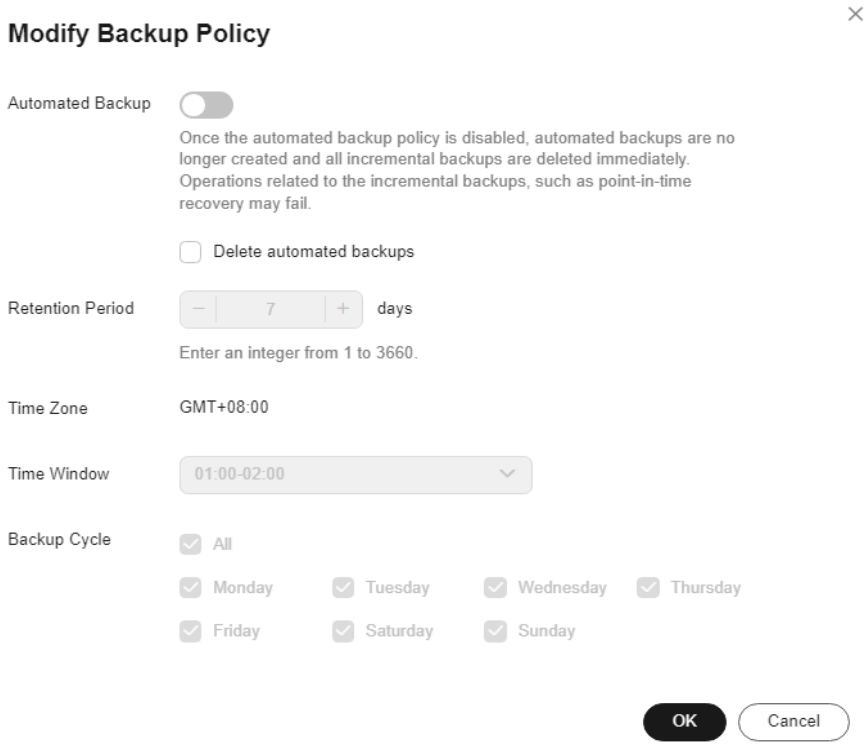

When your disable automated backup, specify whether to delete the automated backups:

- 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.

CAUTION

Deleted backups cannot be restored.

- **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 whose automatic backups you want to delete.
  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**.

4.11.3 Managing Manual Backups

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

Precautions

- Manual backups are full backups.
- GeminiDB Influx does not back up **cold storage data**.
- 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

- Instance with classic storage

  On the **Instances** page, locate the instance you want to back up and choose **More** > **Create Backup** in the **Operation** column.

- Instance with cloud native storage

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

Method 2

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

Method 3

In the navigation pane on the left, choose **Backups**. On the displayed page, click **Create Backup**.

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

Figure 4-106 Creating a manual backup

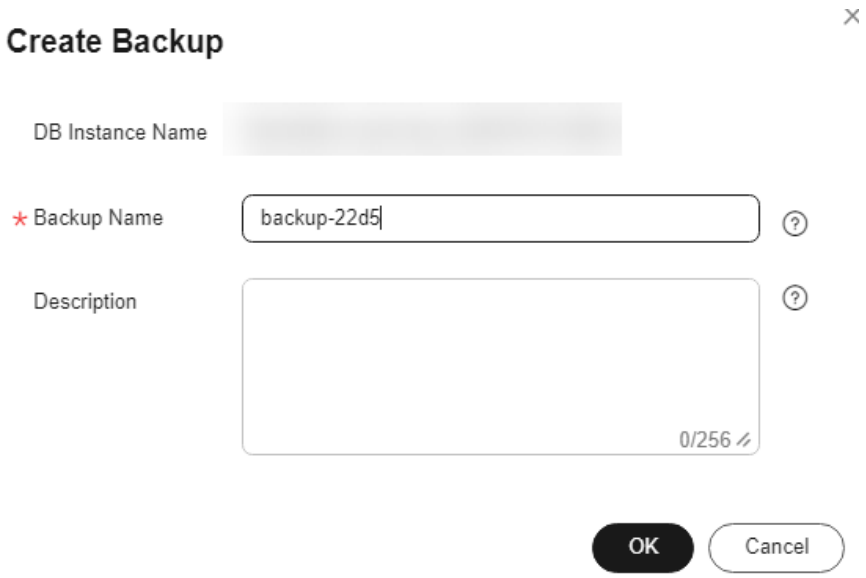

Table 4-48 Parameter description

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 long and start with a letter. It is case-insensitive and contains only letters, digits, hyphens (-), and underscores (_).
Description	Can include a maximum of 256 characters and cannot include 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 changes to **Completed**.

----End

Deleting a Manual Backup

If you no longer need a manual backup, delete it on the **Backups** or **Backups & Restorations** page.

Deleted backups are not displayed in the backup list.

CAUTION

Deleted backups cannot be restored.

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 that you want to delete and click **Delete**.
4. In the displayed dialog box, confirm the backup details and click **Yes**.

4.12 Data Restoration

4.12.1 Restoration Methods

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

Usage Notes

This function of instances in a performance-enhanced cluster 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.

Restoration Methods

Table 4-49 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 Data to a New Instance	You can restore an existing backup to a new instance.

4.12.2 Restoring Data to a New Instance

Scenarios

GeminiDB Influx allows you to use an existing automated or manual backup to restore data to a new instance. The restored instance will have the same data as before.

A full backup will be downloaded from OBS for restoration. The time required depends on the amount of data to be restored.

Precautions

- 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.
- For single-node instances, you can restore an automated backup to a new instance.

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 4-107 Restoration

Method 2

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

Figure 4-108 Restoration

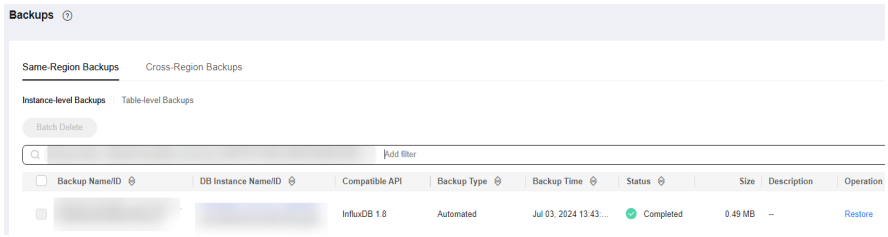

Step 4 In the displayed dialog box, confirm the current instance details and restoration method and click **OK**.

Figure 4-109 Restoring data to a new instance

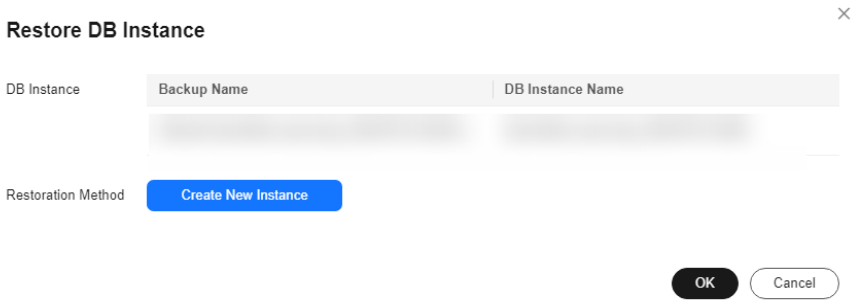

- The default API type and DB engine version are the same as those of the original instance and cannot be changed.
- GeminiDB automatically calculates the minimum storage space required for restoration based on the size of the selected backup file. The storage capacity depends on the instance specifications, and must be an integer.
- You need to set a new administrator password.
- To modify other parameters, see **Buying a GeminiDB Influx Cluster Instance**.

Step 5 View the restoration results.

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

A full backup is triggered after the new DB instance is created.

The new DB instance is independent from the original one.

----End

4.13 Parameter Management

4.13.1 Modifying Parameters of GeminiDB Influx Instances

You can modify parameters in a custom parameter template so that your instance can deliver spectacular performance.

Note that parameter values in default parameter templates cannot be changed.

NOTE

- Exercise caution when modifying parameter values to prevent exceptions.
- Though parameter values in a default template cannot be changed, you can view details about a default parameter template. If a custom parameter template is set incorrectly, the database startup may fail. You can re-configure the custom parameter template according to the configurations of the default parameter template.

Usage Notes

Currently, parameters of GeminiDB Influx instances only on a single node or in a cluster can be modified.

Modifying a Custom Parameter Template and Applying It to an Instance

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 **Parameter Templates**.

Step 4 Click the **Custom Templates** tab, locate the parameter template whose parameters you want to modify, and click its name.

Step 5 On the **Parameters** page, modify parameters. For details about the parameters, see **Modifying Parameters of GeminiDB Influx Instances** or **Table 4-51**.

Figure 4-110 Modifying parameters in the parameter template

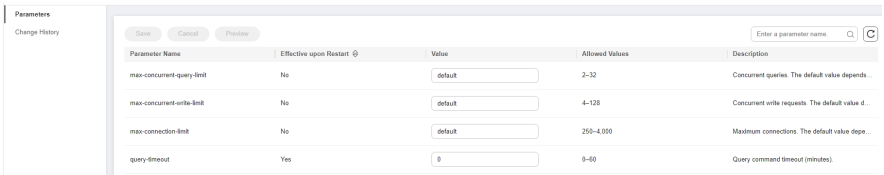

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Table 4-50 Parameters of a GeminiDB Influx cluster instance

Parameter	Effective upon Restart	Default Value	Value Range	Description
max-concurrent-query-limit	Yes	4	4–32	Concurrent queries. If this parameter is set to **default**, the value varies with the CPU specifications.
max-concurrent-write-limit	Yes	16	16–128	Concurrent writes. If this parameter is set to **default**, the value varies with the CPU specifications.
max-connection-limit	Yes	500	500–4,000	Maximum connections. If this parameter is set to **default**, the value varies with the CPU specifications.
query-timeout	Yes	0	0–60	Query command timeout interval in minutes

Table 4-51 Parameters of a single-node GeminiDB Influx instance

Parameter	Effective upon Restart	Default Value	Value Range	Description
max-concurrent-query-limit	No	2	2–16	Concurrent queries. If this parameter is set to **default**, the value varies with the CPU specifications.
max-concurrent-write-limit	No	4	4–64	Concurrent writes. If this parameter is set to **default**, the value varies with the CPU specifications.
max-connection-limit	No	250	250–2,000	Maximum connections. If this parameter is set to **default**, the value varies with the CPU specifications.
query-timeout	Yes	0	0–60	Query command timeout interval in minutes

Figure 4-111 Preview Change

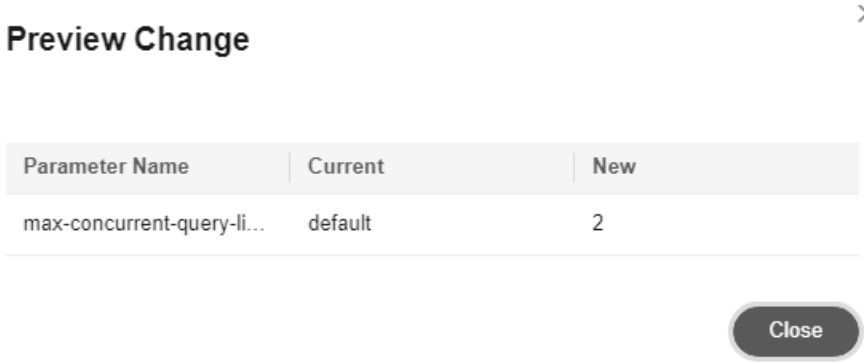

Step 6 After parameters are modified, click **Change History** to view parameter modification details.

For details about how to view parameter modification details, see **Viewing Parameter Change History**.

NOTE

- You need to manually apply the modifications to the current instance. For details, see **Applying a Parameter Template**.
- The change history page displays only modifications in the last seven days.

----End

Modifying Parameters of an Instance

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 **Instances**. In the instance list, locate the instance whose parameters you want to modify and click its name.

Step 4 In the navigation pane on the left, choose **Parameters**. On the displayed page, modify parameters as required.

Figure 4-112 Modifying parameters

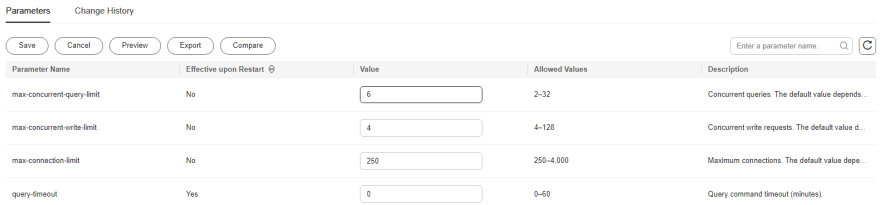

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Step 5 After parameters are modified, click **Change History**.

For details about how to view parameter modification details, see **Viewing Parameter Change History**.

NOTE

The modifications are immediately applied to the current instance.

Check the value in the **Effective upon Restart** column.

- If the value is **Yes** and the instance status on the **Instances** page is **Pending restart**, restart the instance to apply the modifications.
- If the value is **No**, the modifications are applied immediately.

----End

4.13.2 Creating a Parameter Template

You can use database parameter templates to manage DB API configurations. A database parameter template acts as a container for API configuration values that can be applied to one or more DB instances.

Each user can create up to 100 parameter templates. The parameter template quota is shared by all instances in a project.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click **Create Parameter Template**.

Step 5 Select a compatible DB engine version, specify a parameter template name and description, and click **OK**.

Figure 4-113 Creating a parameter template

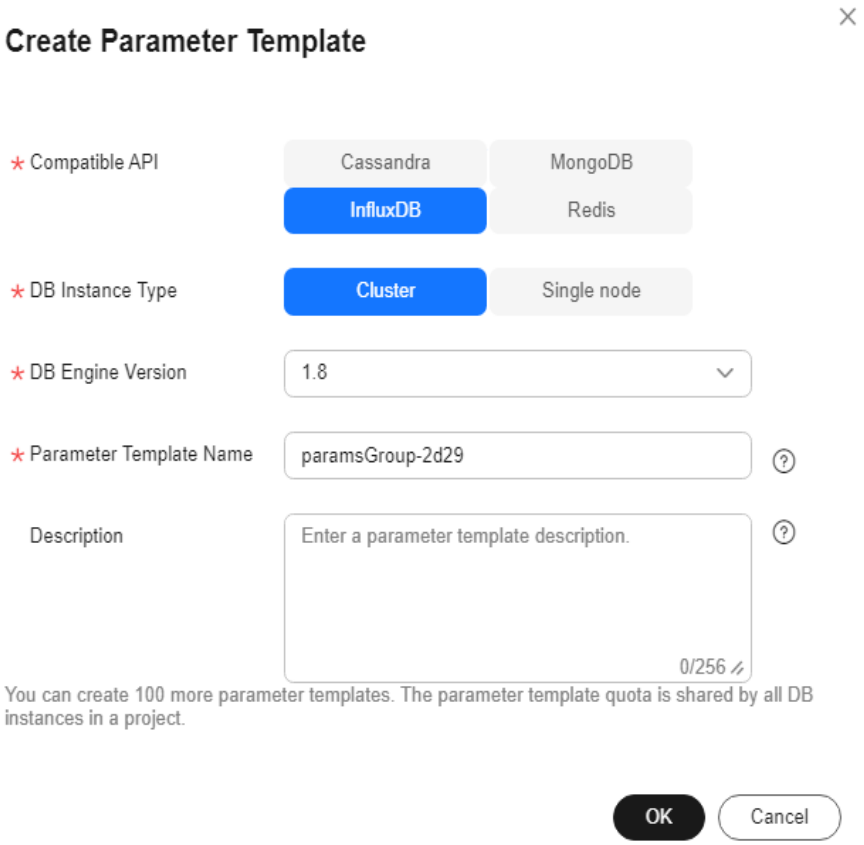

- **Compatible API**: Select the API type that is compatible with your DB engine parameter template.
- **DB Engine Version**: Select a DB engine version, for example, 1.7.
- **Parameter Template Name**: The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description**: The description contains a maximum of 256 characters and cannot include line breaks or the following special characters >!<"&'=

Step 6 On the **Parameter Templates** page, view the created parameter template.

----End

4.13.3 Viewing Parameter Change History

Scenarios

You can view parameter change history of an instance or one of its custom parameter templates based on service requirements.

NOTE

In a newly exported or created parameter template, change history is left blank.

Viewing Change History of a Custom Parameter Template

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 **Parameter Templates**. On the **Custom Templates** page, click the parameter template whose change history you want to view.

Step 4 In the navigation pane on the left, choose **Change History**. Then, view the name, original value, new value, modification status, and modification time of the target parameter.

Figure 4-114 Viewing change history of a customer parameter template

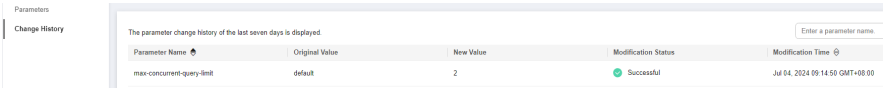

You can apply the parameter template to instances by referring to **Applying a Parameter Template**.

----End

Viewing Parameter Change History of an 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 instance whose parameter change history you want to view and click its name.

Step 4 In the navigation pane on the left, choose **Parameters**. On the **Change History** page, view the name, original value, new value, modification status, and modification time of the target parameter.

Figure 4-115 Viewing parameter change history of an instance

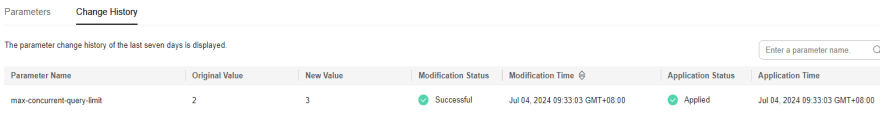

----End

4.13.4 Exporting a Parameter Template

Scenarios

- You can export a parameter template of a DB instance for future use. To learn how to apply the exported parameter template to a DB instance, refer to section **Applying a Parameter Template**.

- You can export the parameter template details (parameter names, values, and descriptions) of a DB instance to a CSV file for review and analysis.

Procedure

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 **Instances**. On the displayed page, locate the instance whose parameters you want to export and click its name.

Step 4 In the navigation pane on the left, choose **Parameters**> **Parameters** and click **Export** above the parameter list.

Figure 4-116 Exporting a parameter template

- **Parameter Template**: You can export parameters of the DB instance to a template for future use.

  In the displayed dialog box, configure required details and click **OK**.

  NOTE

  - **Parameter Template Name**: The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
  - The template description consists of a maximum of 256 characters and cannot include line breaks or the following special characters: >!<"&'=

  After the parameter template is exported, a new template is generated in the list on the **Parameter Templates** page.

- **File**: You can export the parameter template details (parameter names, values, and descriptions) of a DB instance to a CSV file for review and analysis.

  In the displayed dialog box, enter the file name and click **OK**.

NOTE

The file name must start with a letter and consist of 4 to 81 characters. It can contain only letters, digits, hyphens (-), and underscores (_).

----End

4.13.5 Comparing Parameter Templates

Scenarios

This section describes how to compare two parameter templates of the same instance type and compatible API to learn about their configurations.

Comparing Parameter Templates

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 **Parameter Templates**.

Step 4 In the parameter template list, locate the parameter template that you created and click **Compare** in the **Operation** column.

Step 5 In the displayed dialog box, select a parameter template that is of the same instance type and compatible API as the selected template and click **OK**.

Figure 4-117 Comparing two parameter templates

- If their parameters are different, the different parameter names and values are displayed.
- If their parameters are the same, no data is displayed.

----End

Comparing Parameter Templates of a Specific Instance

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 **Instances**.

Step 4 On the **Instances** page, locate the instance whose parameter templates you want to compare and click its name.

Step 5 In the navigation pane on the left, choose **Parameters** and then click **Compare** above the parameter list.

Step 6 In the displayed dialog box, select a parameter template that is of the same instance type as the template of current instance and click **OK**.

Figure 4-118 Comparing the instance parameter template with another parameter template

- If their parameters are different, the different parameter names and values are displayed.
- If their parameters are the same, no data is displayed.

----End

4.13.6 Replicating a Parameter Template

Scenarios

You can replicate a parameter template you have created. When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template. You can also export a parameter template of a DB instance for future use.

Default parameter templates cannot be replicated. You can create parameter templates based on the default ones.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and click **Replicate** in the **Operation** column.

Alternatively, click the target instance on the **Instances** page. On the **Parameters** page, click **Export**.

Step 5 In the displayed dialog box, enter a parameter template name and description and click **OK**.

Figure 4-119 Replicating a parameter template

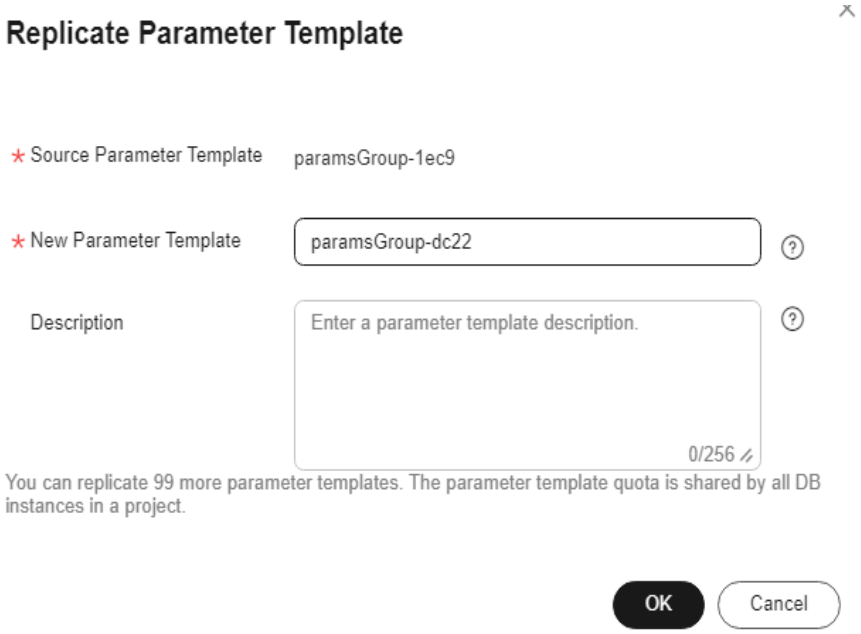

- **New Parameter Template**: The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description**: The description contains a maximum of 256 characters and cannot include line breaks or the following special characters >!<"&'=

After the parameter template is replicated, a new template is generated in the list on the **Parameter Templates** page.

----End

4.13.7 Resetting a Parameter Template

Scenarios

You can reset all parameters in a custom parameter template to their default settings.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and choose **More** > **Reset** in the **Operation** column.

Step 5 Click **Yes** to reset the parameter template.

----End

4.13.8 Applying a Parameter Template

Scenarios

GeminiDB Influx allows you to apply a parameter template. Modifications to parameters in a custom parameter template take effect only after you have applied the template to the target instance.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, perform the following operations based on the template type:

- To apply a default template, click **Default Templates**, locate the template, and in the **Operation** column, click **Apply**.
- To apply a custom template, click **Custom Templates**, locate the target parameter template, and choose **More** > **Apply** in the **Operation** column.

A parameter template can be applied to one or more instances.

Step 5 In the displayed dialog box, select one or more instances that the parameter template will be applied to and click **OK**.

After a parameter template is applied, you can **view its application records**.

----End

4.13.9 Viewing Application Records of a Parameter Template

Scenarios

GeminiDB Influx allows you to view application records of a parameter template.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, perform the following operations based on the template type:

- On the **Default Templates** page, locate the parameter template whose application records you want to view and click **View Application Records** in the **Operation** column.

- On the **Custom Templates** page, locate the target template and choose **More** > **Apply** in the **Operation** column.

You can view the name or ID of the instance that the parameter template applies to, as well as the application status, application time, and causes of any failures that have occurred.

----End

4.13.10 Modifying a Parameter Template Description

Scenarios

You can modify the description of a custom parameter template if needed.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click the **Custom Templates** tab. Locate the target parameter template and click  in the **Description** column.

Step 5 Enter a new description. You can click  to submit or  to cancel the modification.

- After you submit the modification, you can view the new description in the **Description** column.
- The description can include up to 256 characters but cannot contain the following special characters: >!<"&'=

----End

4.13.11 Deleting a Parameter Template

Scenarios

You can delete a custom parameter template that is no longer in use.

Precautions

- Deleted templates cannot be recovered. Exercise caution when performing this operation.
- Default parameter templates cannot be deleted.

Procedure

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 **Parameter Templates**.

Step 4 On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template you want to delete and choose **More** > **Delete** in the **Operation** column.

Step 5 Click **Yes** to delete the parameter template.

----End

4.14 Logs and Audit

4.14.1 Slow Query Logs

You can view slow query logs of GeminiDB Influx databases. Any query that takes longer than an execution time threshold (in milliseconds) will be logged. With slow query logs, you can identify and optimize slow statements.

Usage Notes

- This function is only available for cluster and cluster (performance-enhanced) instances.
- To use this function, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service.

Viewing Log Details

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.

Step 4 In the navigation pane, choose **Slow Query Logs**.

Step 5 On the **Slow Query Logs** page, set search criteria and click **Search** to view log information.

Figure 4-120 Viewing slow query logs

- Select **All nodes** and view slow query logs of all nodes. Alternatively, select a specific node to view its slow query logs.

Figure 4-121 Selecting a node

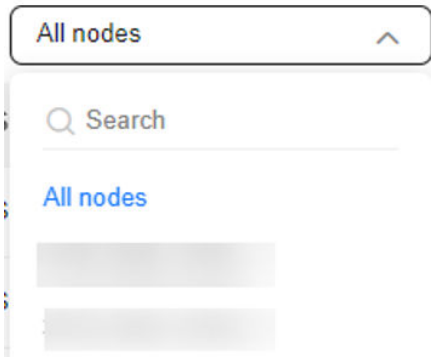

- You can view slow SQL statements of the following types:
  - SELECT
  - DELETE
  - SHOW
  - DROP
  - CREATE
  - ALTER
- After expanding **Advanced Search**, you can filter logs by:
  - Keyword
  - Maximum execution time (ms)
  - Retention policy
  - Database

Figure 4-122 Advanced search

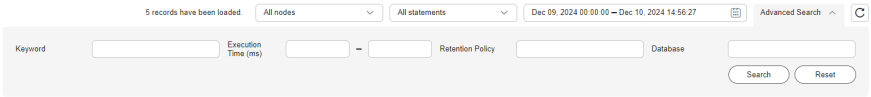

----End

4.14.2 Key Operations Supported by CTS

With CTS, you can record operations on GeminiDB Influx instances for later queries, audit, and backtracking.

Table 4-52 GeminiDB Influx key operations

Operation	Resource Type	Trace Name
Creating an instance	instance	NoSQLCreateInstance
Deleting an instance	instance	NoSQLDeleteInstance
Adding nodes	instance	NoSQLEnlargeInstance
Deleting nodes	instance	NoSQLReduceInstance

Operation	Resource Type	Trace Name
Restarting an instance	instance	NoSQLRestartInstance
Restoring data to a new instance	instance	NoSQLRestoreNewInstance
Scaling up storage space	instance	NoSQLExtendInstanceVolume
Resetting the password of an instance	instance	NoSQLResetPassword
Modifying the name of an instance	instance	NoSQLRenameInstance
Changing specifications	instance	NoSQLResizeInstance
Binding an EIP	instance	NoSQLBindEIP
Unbinding an EIP	instance	NoSQLUnBindEIP
Freezing an instance	instance	NoSQLFreezeInstance
Unfreezing an instance	instance	NoSQLUnfreezeInstance
Creating a backup	backup	NoSQLCreateBackup
Deleting a backup	backup	NoSQLDeleteBackup
Modifying the backup policy of an instance	backup	NoSQLSetBackupPolicy
Adding an instance tag	tag	NoSQLAddTags
Modifying an instance tag	tag	NoSQLModifyInstanceTag
Deleting an instance tag	tag	NoSQLDeleteInstanceTag
Creating a parameter template	parameterGroup	NoSQLCreateConfigurations
Modifying a parameter template	parameterGroup	NoSQLUpdateConfigurations
Modifying instance parameters	parameterGroup	NoSQLUpdateInstanceConfigurations
Replicating a parameter template	parameterGroup	NoSQLCopyConfigurations
Resetting a parameter template	parameterGroup	NoSQLResetConfigurations
Applying a parameter template	parameterGroup	NoSQLApplyConfigurations
Deleting a parameter template	parameterGroup	NoSQLDeleteConfigurations

Operation	Resource Type	Trace Name
Deleting the node that fails to be added	instance	NoSQLDeleteEnlargeFail-Node
Enabling SSL	instance	NoSQLSwitchSSL
Changing the security group of an instance	instance	NoSQLModifySecurityGroup
Exporting parameter template information for an instance	instance	NoSQLSaveConfigurations
Modifying the recycling policy	instance	NoSQLModifyRecyclePolicy

4.14.3 Querying Traces

After CTS is enabled, CTS starts recording operations on cloud resources. The CTS console stores the last seven days of operation records.

This section describes how to query the last seven days of operation records on the CTS console.

Procedure

Step 1 Log in to the management console.

Step 2 Click  in the upper left corner and select a region and project.

Step 3 Click **Service List**. Under **Management & Governance**, click **Cloud Trace Service**.

Step 4 In the navigation pane on the left, click **Trace List**.

Step 5 Specify filter criteria to search for the required traces. The following four filter criteria are available:

- **Trace Source**, **Resource Type**, and **Search By**

  Select filters from the drop-down list.

  When you select **Trace name** for **Search By**, you need to select a specific trace name.

  When you select **Resource ID** for **Search By**, you also need to select or enter a specific resource ID.

  When you select **Resource name** for **Search By**, you also need to select or enter a specific resource name.
- **Operator**: Select a specific operator (a user other than the tenant).
- **Trace Status**: Select **All trace statuses**, **Normal**, **Warning**, or **Incident**.
- **Start Date** and **End Date**: You can specify a time range to query traces.

Step 6 Locate the target trace and click ⌄ to view its details.

Step 7 Click **View Trace** in the **Operation** column. In the displayed dialog box, the trace structure details are displayed.

----End

4.15 Viewing Metrics and Configuring Alarms

4.15.1 Supported Metrics

Description

This section describes GeminiDB Influx metrics reported to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for GeminiDB Influx.

Namespace

SYS.NoSQL

Monitoring Metrics

NOTE

You can view metrics on instance nodes by referring to **Viewing Metrics**.

Table 4-53 GeminiDB Influx API metrics

Metric ID	Metric Name	Description	Value Range	Unit	Number System	Monitored Object	Monitoring Period (Raw Data)
gemini001_cpu_usage	CPU Usage	CPU usage of the monitored system	0–100	%	N/A	GeminiDB Influx instance node	1 minute
gemini002_mem_usage	Memory Usage	Memory usage of the monitored system	0–100	%	N/A	GeminiDB Influx instance node	1 minute
gemini003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0	Bytes/s	1024(IEC)	GeminiDB Influx instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Unit	Number System	Monitored Object	Monitoring Period (Raw Data)
gemini004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second	≥ 0	Bytes/s	1024(IEC)	GeminiDB Influx instance nodes	1 minute
nosql005_disk_usage	Storage Space Usage	Storage usage of the current instance.	0–100	%	N/A	GeminiDB Influx instances	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of the current instance.	≥ 0	GB	1024(IEC)	GeminiDB Influx instances	1 minute
nosql007_disk_used_size	Storage Space Usage	Storage space usage of the current instance.	≥ 0	GB	1024(IEC)	GeminiDB Influx instances	1 minute
influxdb001_series_num	Time Series	Total number of time series	≥ 0	Counts	N/A	GeminiDB Influx instance nodes	1 minute
influxdb002_query_req_ps	Query Requests Per Second	Number of query requests per second	≥ 0	Counts/s	N/A	GeminiDB Influx instance nodes	1 minute
influxdb003_write_req_ps	Write Requests Per Second	Number of write requests per second	≥ 0	Counts/s	N/A	GeminiDB Influx instance nodes	1 minute
influxdb004_write_points_ps	Write Points	Number of write points per second	≥ 0	Counts/s	N/A	GeminiDB Influx instance nodes	1 minute

Metric ID	Metric Name	Description	Value Range	Unit	Number System	Monitored Object	Monitoring Period (Raw Data)
influxdb005_write_concurrency	Concurrent Write Requests	Number of concurrent write requests	≥ 0	Counts	N/A	GeminiDB Influx instance nodes	1 minute
influxdb006_query_concurrency	Concurrent Queries	Number of concurrent query requests	≥ 0	Counts	N/A	GeminiDB Influx instance nodes	1 minute
influxdb010_cold_disk_usage	Cold Storage Space Usage	Cold storage usage of an instance	0–100	%	N/A	GeminiDB Influx instances	1 minute
influxdb011_cold_disk_total_size	Total Cold Storage Space	Total cold storage of an instance	≥ 0	GB	1024(IEC)	GeminiDB Influx instances	1 minute
influxdb012_cold_disk_used_size	Cold Storage Space Usage	Used cold storage of an instance	≥ 0	GB	1024(IEC)	GeminiDB Influx instances	1 minute

Dimensions

Key	Value
influxdb_cluster_id	Cluster ID of the GeminiDB Influx instance
influxdb_node_id	Node ID of the GeminiDB Influx instance

4.15.2 Configuring Alarm Rules

Scenarios

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 management console.

Step 2 Click **Service List**. Under **Management & Governance**, click **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 4-123 Creating an alarm rule

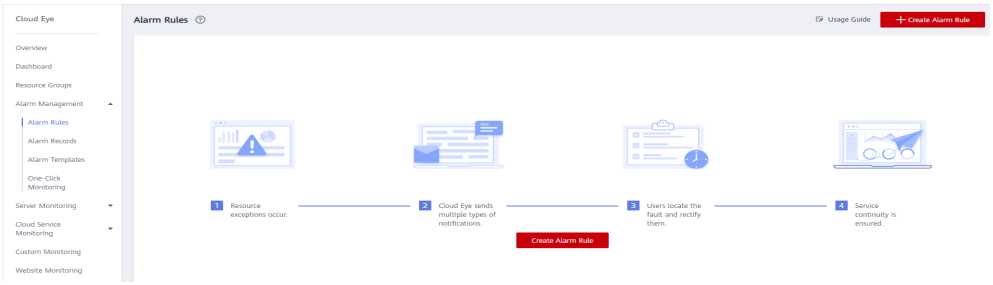

Step 5 Set alarm parameters.

1. Configure basic alarm information.

Figure 4-124 Configuring basic information for an alarm rule

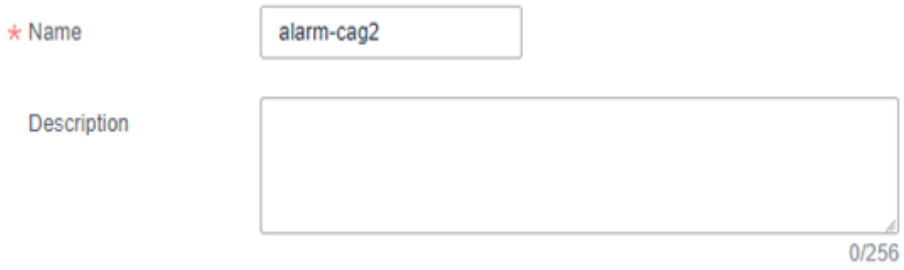

Table 4-54 Basic alarm rule information

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.

Table 4-55 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 **InfluxDB-InfluxDB Nodes**.	-
Monitoring Scope	Monitoring scope the alarm rule applies to. **NOTE** – 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. – To specify **Specific resources**, click **Select Specified Resources**, select one or more resources, and click **OK**.	All Resources
Group	This parameter is mandatory when **Monitoring Scope** is set to **Resource groups**.	-

3. Configure an alarm policy.

Figure 4-125 Configuring the alarm policy

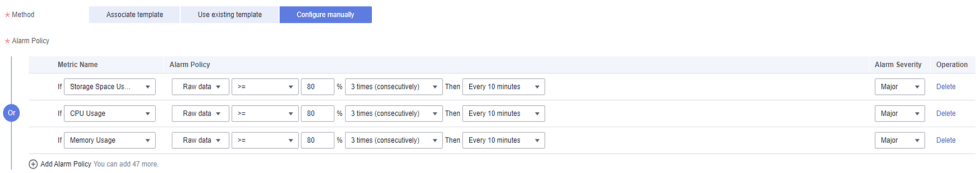

Table 4-56 Parameter description

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

Parameter	Description	Example Value
Template	Select the template to be used. This parameter is available only when you select **Use existing template** for **Method**.	-
Alarm Policy	Policy for triggering an alarm. You can configure the threshold, consecutive periods, alarm interval, and alarm severity based on service requirements. – **Metric Name**: specifies the metric that the alarm rule is created for. The following metrics are recommended: **Storage Space Usage**, which is used to monitor the storage usage of GeminiDB Influx instances. If the storage usage is greater than 80%, scale up the storage in a timely manner by referring to **Manually Scaling Up Storage Space of a GeminiDB Influx Instance**. **CPU Usage** and **Memory Usage**, which are used to monitor the compute resource usage of each GeminiDB Influx instance node. If the CPU usage or memory usage is greater than 80%, you can **add nodes** or **upgrade node specifications** in a timely manner. 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.	Take the CPU usage as an example. The alarm policy configured in **Figure 4-125** indicates that a major alarm notification will be sent to users every 10 minutes if the original CPU usage reaches 80% or above for three consecutive periods.

4. Configure alarm notification information.

Figure 4-126 Configuring alarm notification information

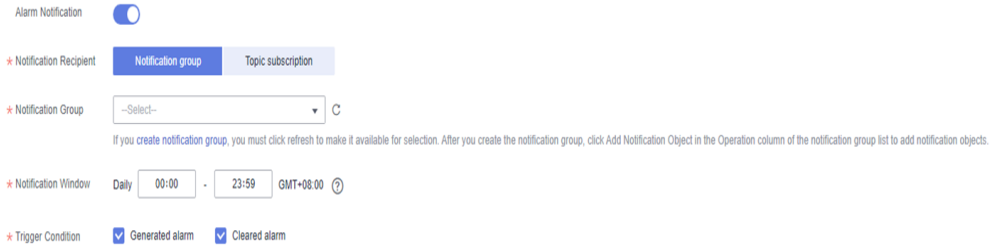

Table 4-57 Parameter description

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 4-127 Advanced settings

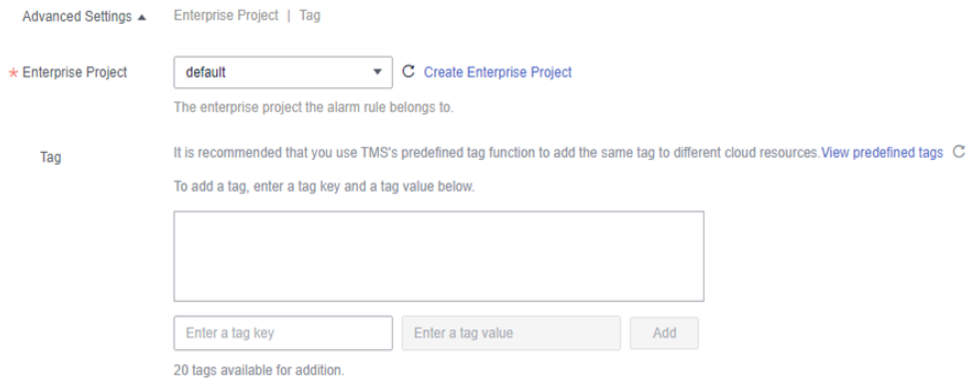

Table 4-58 Parameter description

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

4.15.3 Viewing Metrics

Scenarios

Cloud Eye monitors the GeminiDB Influx 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 4-128 Viewing metrics

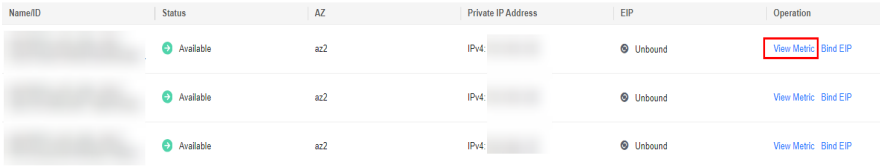

Step 5 In the monitoring area, you can select a duration to view the monitoring data.

The monitoring data generated in the latest 1 hour, 3 hours, 12 hours, 24 hours, or 7 days can be viewed.

To view the monitoring curve in a longer time range, click  to enlarge the graph.

----End

4.15.4 Event Monitoring

4.15.4.1 Introduction to Event Monitoring

Event monitoring provides event data reporting, query, and alarm reporting. 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 Influx resources 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 and allows you to view monitoring details of system events 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.

4.15.4.2 Viewing Event Monitoring Data

Scenarios

Event monitoring provides event data reporting, query, and alarm reporting. 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 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 < 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

4.15.4.3 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. Under **Management & Governance**, click **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 4-59 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.

Parameter	Description
Event Source	Specifies the service the event is generated for. Select GeminiDB.
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 the topics you require are not displayed in the drop-down list, click **Create an SMN topic**.

Table 4-60 Alarm notification parameters

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.
Notification Object	Specifies the 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	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If you set **Validity Period** to **08:00-20:00**, Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification.

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

----End

4.15.4.4 Events Supported by Event Monitoring

Table 4-61 Events Supported by Event Monitoring for GeminiDB

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
NoSQL	Instance creation failure	NoSQLCreateInstanceFailed	Major	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.	Instances fail to be created.
	Specifications change failure	NoSQLResizeInstanceFailed	Major	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.	Services are interrupted.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Node adding failure	NoSQLAddNodesFailed	Major	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	NoSQLDeleteNodesFailed	Major	Releasing underlying resources failed.	Delete the node again.	None
	Storage space scale-up failure	NoSQLScaleUpStorageFailed	Major	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.	Services may be interrupted.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Password resetting failure	NoSQLResetPasswordFailed	Major	Resetting the password times out.	Reset the password again.	None
	Parameter template change failure	NoSQLUpdateInstanceParamGroupFailed	Major	Changing a parameter template times out.	Change the parameter template again.	None
	Backup policy configuration failure	NoSQLSetBackupPolicyFailed	Major	The database connection is abnormal.	Configure the backup policy again.	None
	Manual backup creation failure	NoSQLCreateManualBackupFailed	Major	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 cannot be backed up.
	Automated backup creation failure	NoSQLCreateAutomatedBackupFailed	Major	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 cannot be backed up.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Instance status abnormal	NoSQLFaultyDBInstance	Major	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 database service may be unavailable.
	Instance status recovery	NoSQLDBInstanceRecovered	Major	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	NoSQLFaultyDBNode	Major	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 database service may be unavailable.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Node status recovery	NoSQLDBNodeRecovered	Major	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 switchover or failover	NoSQLPrimaryStandbySwitched	Major	This event is reported when a primary/secondary switchover or a failover is triggered.	No further action is required.	None
	Occurrence of hotspot partitioning keys	HotKeyOccurs	Major	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 service request success rate is affected, and the cluster performance and stability deteriorates.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	BigKey occurrence	BigKeyOccurs	Major	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 partition, cluster stability deteriorates.
	Insufficient storage space	NoSQLRiskyDataDiskUsage	Major	The storage space is insufficient.	Scale up storage space. For details, see section "Scaling Up Storage Space" in the user guide of GeminiDB.	The instance is set to read-only and data cannot be written to the instance.
	Data disk expanded and being writable	NoSQLDataDiskUsageRecovered	Major	The data disk has been expanded and becomes writable.	No further action is required.	None

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Index creation failure	NoSQLCreateIndexFailed	Major	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 off-peak hours. Create indexes in the background. Select indexes as required.	The index fails to be created or is incomplete. Delete the index and create a new one.
	Write speed decrease	NoSQLStallingOccurs	Major	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 success rate of service requests is affected.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Data write stopped	NoSQLStoppingOccurs	Major	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 success rate of service requests is affected.
	Database restart failure	NoSQLRestartDBFailed	Major	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 instance status may be abnormal.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Restoration to new instance failure	NoSQLRestoreToNewInstanceFailed	Major	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 cannot be restored to a new instance.
	Restoration to existing instance failure	NoSQLRestoreToExistInstanceFailed	Major	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 current instance may be unavailable.
	Backup file deletion failure	NoSQLDeleteBackupFailed	Major	The backup files fail to be deleted from OBS.	Delete the backup files again.	None

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Failure to display slow query logs in plaintext	NoSQLSwitchSlowlogPlainTextFailed	Major	The DB 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	NoSQLBindEipFailed	Major	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 instance cannot be accessed from a public network.
	EIP unbinding failure	NoSQLUnbindEipFailed	Major	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 Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Parameter modification failure	NoSQLModifyParameterFailed	Major	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 application failure	NoSQLApplyParameterGroupFailed	Major	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 Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Enabling or disabling SSL failure	NoSQLSwitchSSLFailed	Major	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 connection mode cannot be changed.
	Too much data in a single row	LargeRowOccurs	Major	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 records in a single row, cluster stability will deteriorate as the data volume increases.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Schedule for deleting a KMS key	planDeleteKmsKey	Major	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 automatically deleted after it expires. Deleting the key will affect the instance services.
	Too many tombstones	TooManyQueryTombstones	Major	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	TooLargeCollectionColumn	Major	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 collection column will fail.

4.16 Managing Tags

Scenarios

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.

Adding tags to GeminiDB Influx instance helps you better identify and manage them. An instance can be tagged when or after it is created.

After a DB 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 naming rules of tag keys and tag values, see **Table 4-62**.
- A maximum of 20 tags can be added for each instance.
- The tag name must comply with the naming rules described in **Table 4-62**.

Table 4-62 Naming rules

Parameter	Requirement	Example Value
Tag key	• Cannot be left blank. • Must be unique for each instance. • 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: -_@.:/+=	Organization
Tag value	• Can be left blank. • Can contain a maximum of 255 characters. • Only letters, digits, spaces, and the following special characters are allowed: -_@.:/+=	nosql_01

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, locate the instance you want to add tags to and click its name.

Step 4 In the navigation pane on the left, choose **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 the tag 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, locate the instance whose tags you want to edit and click its name.

Step 4 In the navigation pane on the left, choose **Tags**.

Step 5 On the **Tags** page, locate the tag to be edited 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 the tag 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, locate the instance whose tags you want to delete and click its name.

Step 4 In the navigation pane on the left, choose **Tags**.

Step 5 On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

Step 6 Verify that the tag is no longer displayed on the **Tags** page.

----End

Search 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 4-129 Selecting tags

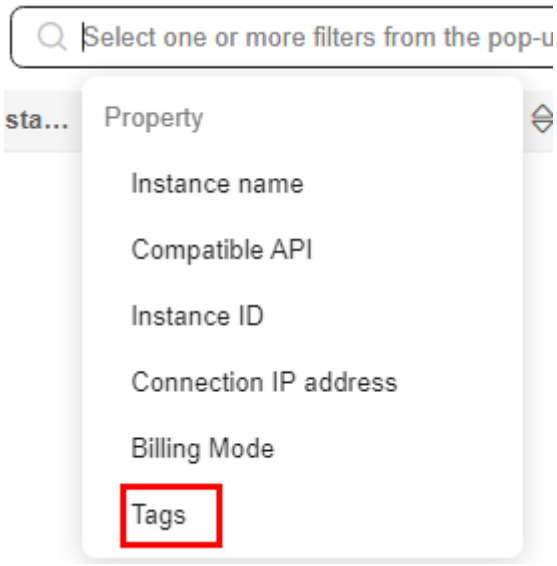

Step 4 Select the tag to be queried and click **OK** to query information about instances associated with the tag.

Figure 4-130 Searching by tag

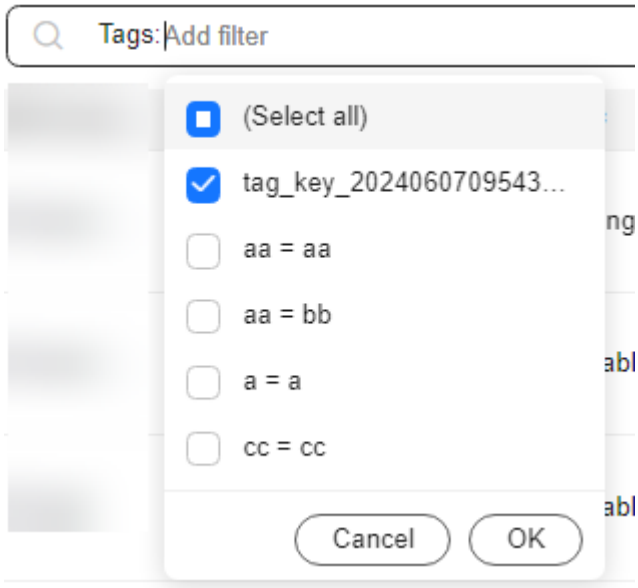

----End

4.17 Viewing 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.

Viewing Quotas

Step 1 Log in to the Huawei Cloud console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 Click in the upper left corner and select a region and project.

Step 4 In the upper right corner, choose **Resources** > **My Quotas**.

The **Quotas** page is displayed.

Figure 4-131 My quotas

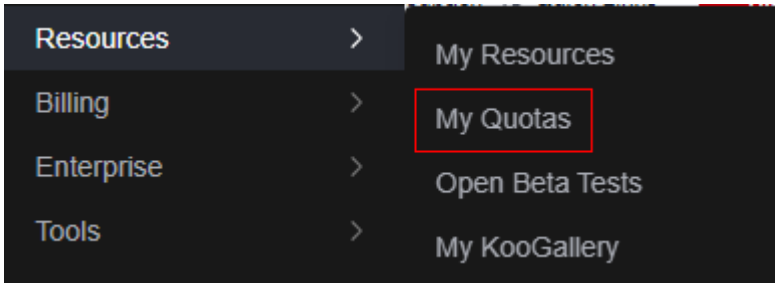

Step 5 Check the used and total quotas of each type of GeminiDB instance resources.

----End

Increasing Quotas

Step 1 Log in to the Huawei Cloud console.

Step 2 In the service list, choose **Databases** > **GeminiDB**.

Step 3 Click in 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 4-132 Increasing quotas

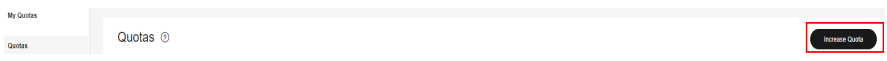

Step 6 On the **Create Service Ticket** page, configure parameters.

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 Buying and Connecting to a GeminiDB Influx Instance

This section describes how to buy a GeminiDB Influx instance and uses a Linux ECS as an example to describe how to connect to the instance over a private network.

- **Buying a GeminiDB Influx Instance**
- **Buying an ECS**
- **Connecting to the GeminiDB Influx Instance**

Buying a GeminiDB Influx 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 **Buy DB Instance**.

Step 4 Click **Buy DB Instance**, select a billing mode, and configure instance parameters. Then, click **Next** and complete subsequent operations.

Figure 5-1 Basic information

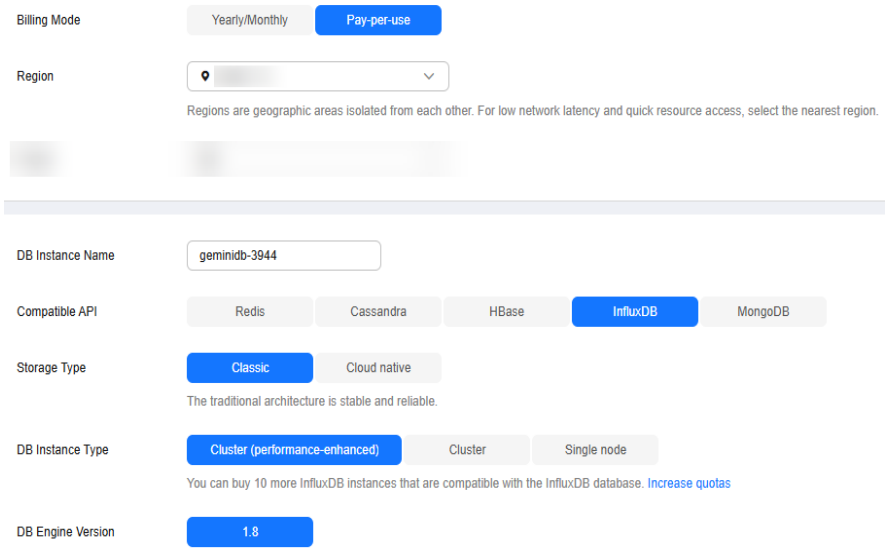

Figure 5-2 Selecting specifications

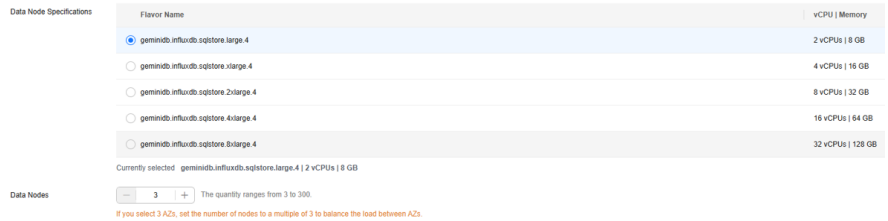

Figure 5-3 Network settings

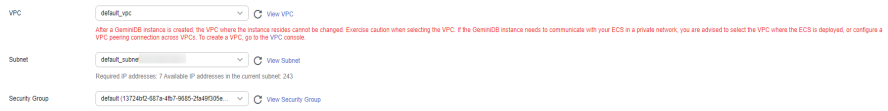

Figure 5-4 Setting a password

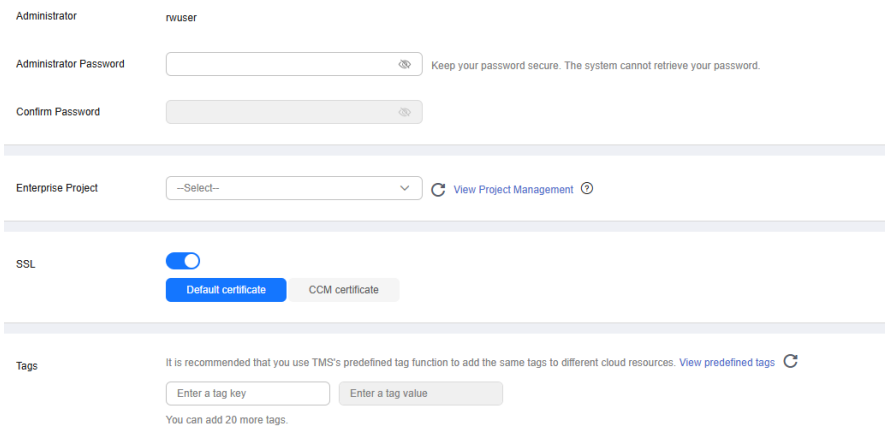

Step 5 View the purchased GeminiDB Influx instance.

Figure 5-5 Available instance

----End

Buying an ECS

Step 1 **Log in to the Huawei Cloud console.**

Step 2 In the service list, choose **Compute** > **Elastic Cloud Server**. On the Elastic Cloud Server console, click **Buy ECS**.

Step 3 Configure basic settings and click **Next: Configure Network**. Make sure that the ECS is in the same region, AZ, VPC, and security group as the GeminiDB Influx instance you created.

Figure 5-6 Basic settings

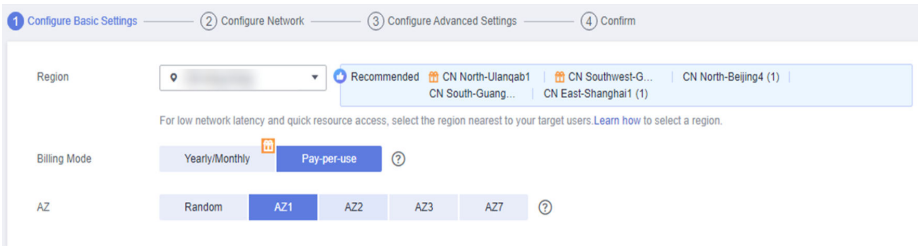

Figure 5-7 Selecting specifications

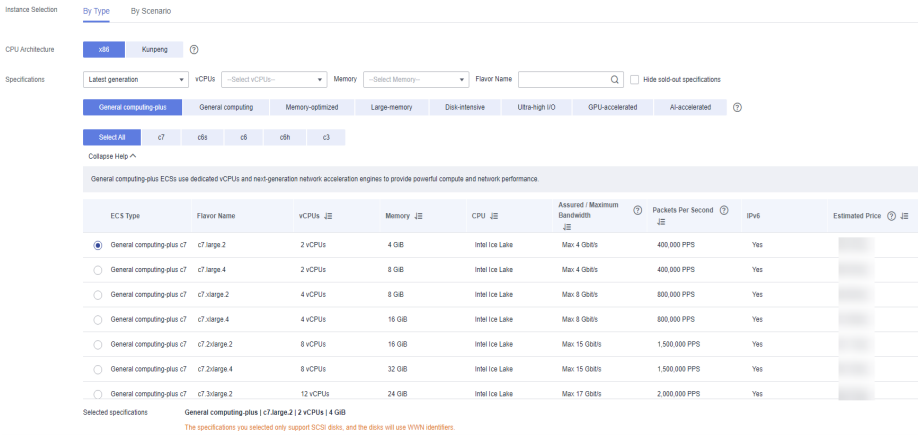

Figure 5-8 Selecting an image

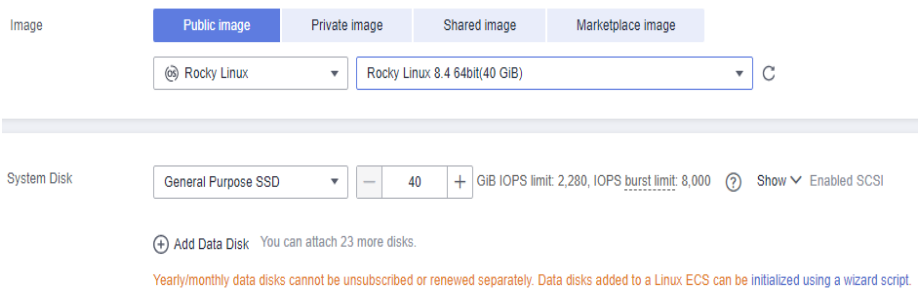

Step 4 Configure the network and click **Next: Configure Advanced Settings**. Make sure that the ECS is in the same VPC and security group as the GeminiDB Influx instance.

- If security group rules allow access from the ECS, you can connect to the instance using the ECS.
- If the security group rules do not allow access from the ECS, add an inbound rule to the security group.

Figure 5-9 Network settings

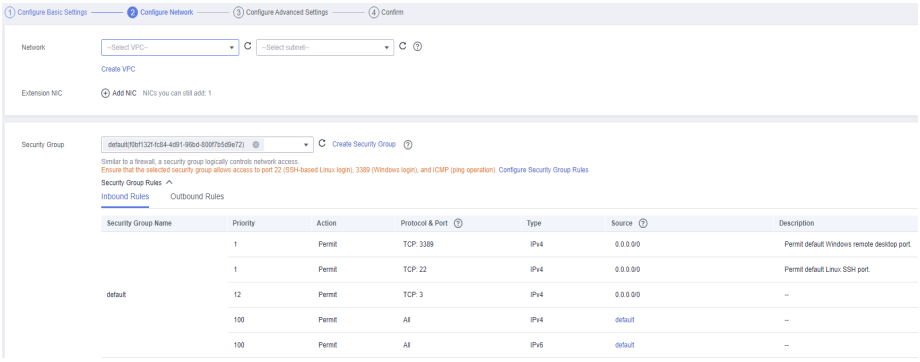

Figure 5-10 Selecting an EIP

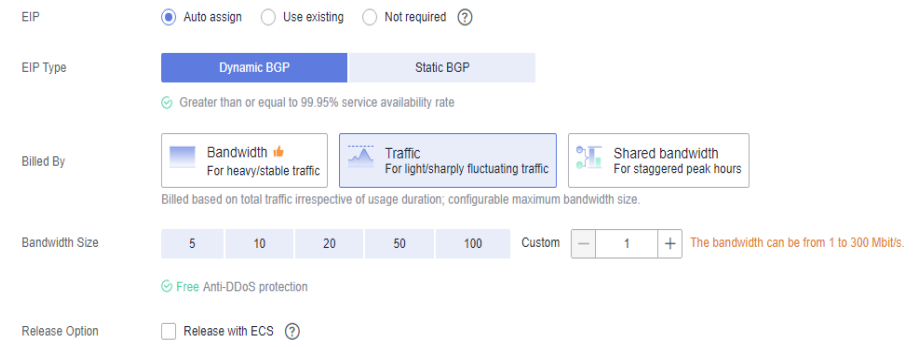

Step 5 Configure a password for the ECS and click **Next: Confirm**.

Figure 5-11 Advanced settings

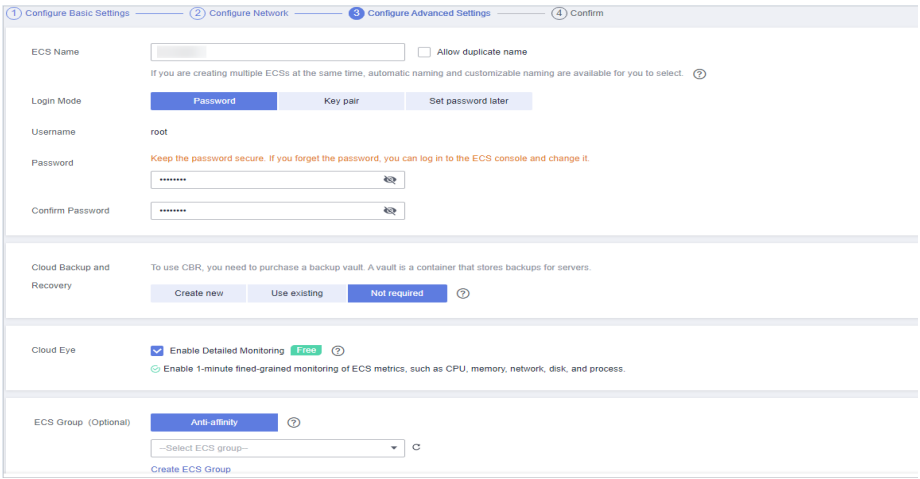

Step 6 Confirm the configurations and click **Submit**.

Figure 5-12 Confirming the configurations

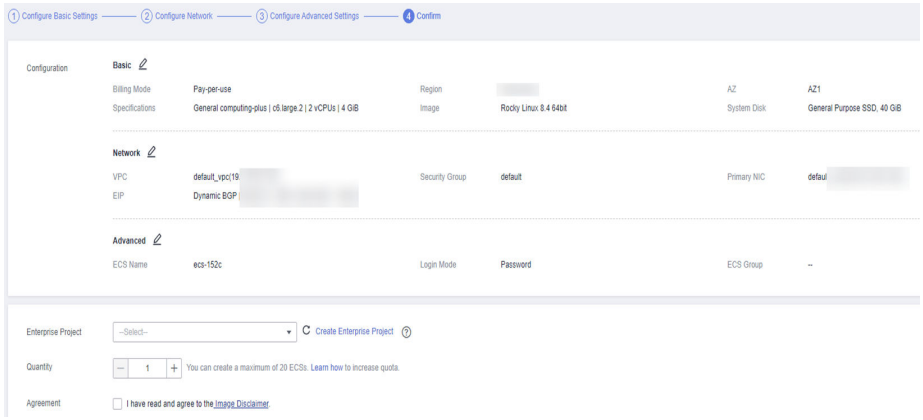

Step 7 View the purchased ECS.

----End

Connecting to the GeminiDB Influx Instance

Step 1 On the ECS console, log in to the ECS using the remote login option.

Figure 5-13 Remote login

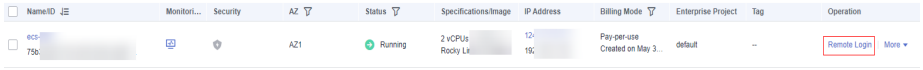

Step 2 Enter the username and password of the ECS.

Figure 5-14 Entering the username and password

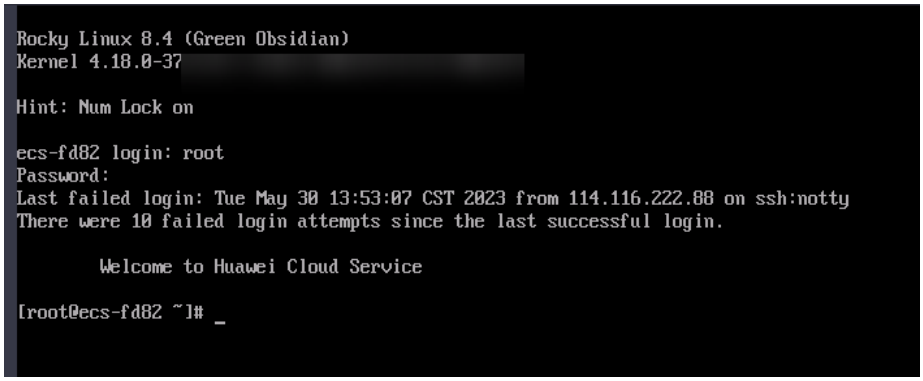

Step 3 Obtain the x86 or Arm InfluxDB client.

Download the **x86** or **Arm** InfluxDB client and upload the InfluxDB client installation package to the ECS.

Step 4 Decompress the client tool package (the x86 client is used as an example).

```
tar -xzf influxdb-1.8.10_linux_amd64.tar.gz
```

Step 5 Connect to your instance in the directory where the InfluxDB client is located.

1. Run the following command to go to the InfluxDB directory:
   cd influxdb-1.8.10-1/usr/bin
2. Connect to the GeminiDB Influx instance.
   ./influx -ssl -unsafeSsl -username '<*DB_USER*>**' -password '**<*DB_PWD*>**' -host** <*DB_HOST*> **-port** <*DB_PORT*>

   Example:

   ./influx -ssl -unsafeSsl -username 'rwuser' -password '<*DB_PWD*>**' -host 192.xx.xx.xx -port 8635**

Table 5-1 Required description

Parameter	Description
<*DB_USER*>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, locate the instance and click its name. In the **DB Information** area on the **Basic Information** page, you can find the administrator username.
<*DB_PWD*>	Administrator password
<*DB_HOST*>	Load balancer address of the instance to be connected. Connecting to an instance using a load balancer address 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. Scenario 1: If you have obtained a load balancer address before creating an instance, you can view that the load balancer address is selected by default on the instance creation page. After the instance is created, click its name to go to the **Basic Information** page and obtain the load balancer IP address in the **Network Information** area. Scenario 2: To use a load balancer address after the instance is created, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact the customer service. Then you can click the instance name to view the load balancer IP address in the **Network Information** area on the **Basic Information** page.
<*DB_PORT*>	Port for accessing the instance. Click the name of the instance to go to the **Basic Information** page. In the **Network Information** area, you can find the database port.

Step 6 If information similar to the following is displayed, the connection was successful.


```
Connected to https://host:port version x.x.x
InfluxDB shell version: 1.8.10
>
```

----End

5.2 Comparison Between GeminiDB Influx and Self-Managed InfluxDB Instances

This section describes differences between GeminiDB Influx and self-managed InfluxDB instances.

Feature Comparison

Table 5-2 Comparison between GeminiDB Influx and self-managed InfluxDB instances

Item	Self-Managed InfluxDB Instance	GeminiDB Influx
Cloud native	Not supported	Supported
Cluster	Not supported	Supported
Tiered storage of hot and cold data	Not supported	Two types of storage media ensure high performance and low costs.
System security	Database vulnerabilities are automatically fixed.	You do not need to give serious attention to database vulnerabilities.
DR	High availability is not provided.	Instances can be deployed across three AZs, ensuring 99.95% service availability.
Backup	Users back up data by themselves.	Data is automatically backed up.
O&M difficulty	It is difficult for users to maintain hardware and software by themselves.	Users perform basic management operations on the GUI. Instructional documents and 24/7 technical support are provided.

5.3 GeminiDB Time Series IoV Solution

Application Scenarios

Fueled by immense popularity of intelligent new energy vehicles, time series data generated in real time experienced tremendous growth. There were urgent demands for vehicle enterprises and owners to query the real-time status of vehicles, but the traditional HBase-based vehicle monitoring platform cannot meet the requirements.

Solution Overview

The GeminiDB time series IoV solution is designed for real-time queries of vehicle data through the dedicated Influx API, which parses, sorts, merges, analyzes, and writes millions of time series data of vehicles in real time. This solution supports high compression ratio and separation of cold and hot data, effectively reducing costs.

- Figure 5-15 shows the technical architecture of HBase.
  - Vehicle data is collected and then reported to Kafka.
  - Real-time data is written to HBase through Flink. All monitoring data is written to HBase as a character string.
  - HBase supports point queries.
  - Offline data is parsed by Flink and written to Hive using Spark for analysis.

Figure 5-15 HBase technical architecture

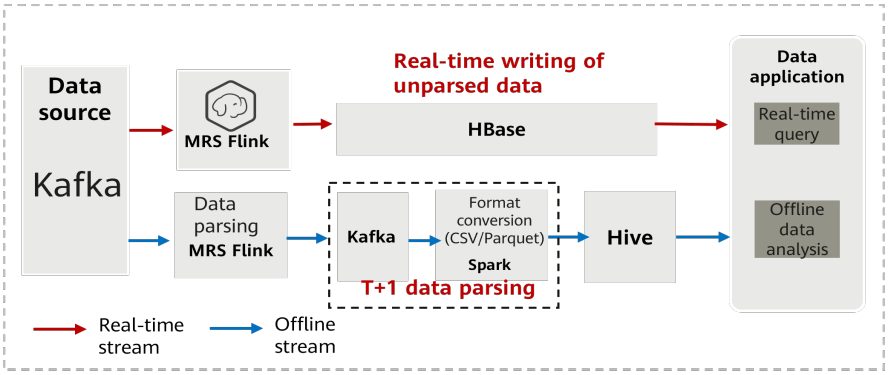

- Figure 5-16 shows the technical architecture of GeminiDB Influx API.
  - Vehicle data is collected and then reported to Kafka.
  - Real-time data is parsed by Flink and written to GeminiDB Influx API.
  - GeminiDB Influx API supports real-time queries and data analysis.
  - GeminiDB Influx API can convert data into Parquet format, which can be directly read and analyzed by Hive. For details, see **Converting Data into a Parquet file and Exporting the Data to OBS**.

Figure 5-16 Technical architecture of GeminiDB Influx API

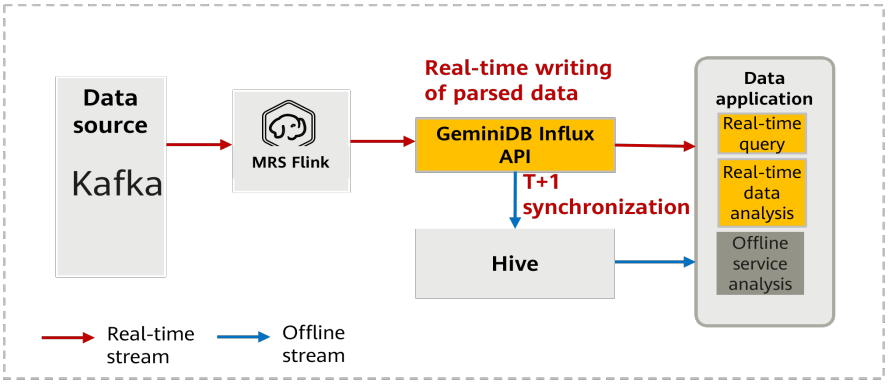

Advantages

- Parsing and writing massive volumes of data in real time; simplifying application development

HBase: Thousands of monitoring metrics reported by vehicles are written into HBase as character strings. When an application reads a metric, it needs to read and parse all character strings. This process is complex and inefficient.

GeminiDB Influx API: Thousands of monitoring metrics reported by vehicles are directly written into GeminiDB Influx instances as thousands of columns. Metrics can be directly queried and without being parsed again.

- Automatically sorting and combining data; simplifying the intermediate process

Multi-dimensional metric data reported at the same time point by vehicles is processed by different components under different network delays, so the data cannot be reported and written at a time in sequence.

HBase: Applications need to use Spark to combine and sort HBase data, which is complex and cannot meet real-time query requirements.

GeminiDB Influx API: When time series data is written, it is automatically merged and sorted. Applications can directly access GeminiDB Influx instances to obtain the result.

- Real-time analysis

Traditional HBase: Raw data consists of unresolved character strings and cannot be used directly. It must be read and resolved, or it cannot meet requirements for efficiently analyzing and processing large volumes. Typically, data analysis involves synchronizing data to platforms like Hive, which lacks real-time capabilities.

GeminiDB Influx API: Data can be queried and analyzed based on metrics at a time. Just one database is enough for effective real-time query and analysis.

- High compression ratio

HBase: The compression algorithm can be set only by column family. Only the GZIP, Snappy, LZO, and LZ4 algorithms are supported.

GeminiDB Influx API: Different compression algorithms are used for data types of each column. Multiple compression algorithms, such as Simple8b, Delta, Delta-Of-

Delta, RLE, ZigZag, ZSTD, Snappy, and bit-packing, are supported. The compression ratio is 10 times that of HBase.

- Separation of hot and cold data

Users can configure hot and cold data policies to automatically dump data to cold storage without changing applications, which effectively reducing the overall cost.

5.4 GeminiDB's Functions for Efficient Data Analysis

Application Scenarios

In IoV and IoT scenarios, a large amount of data needs to be collected frequently. Service personnel have been seeking a better way to analyze data faster. Traditionally, queried data is analyzed at the service layer. The design of a service's functionality and behavior must meet stringent criteria and a large number of service resources will be consumed, so the traditional solution does not fit the need of analyzing massive amount of data. GeminiDB Influx API is a time series database that can efficiently process massive data write and analysis requests. It provides various advanced data analysis functions, making data analysis efficient and convenient. This section uses histogram functions to demonstrate GeminiDB's advantages in advanced data analysis.

Use Cases

In statistics, a histogram is a valuable tool for gaining insights into data characteristics. It is widely used in many scenarios:

- Network monitoring: Histograms clearly show abnormal data distributions, facilitating network self-diagnosis and recovery.
- IoT data analysis: Histograms reveal data distribution characteristics, facilitating identification of time series data.

GeminiDB Influx API supports two types of histograms: equi-height and bounded histograms. Each bucket in an equi-height histogram contains roughly the same number of values. Values in each bucket have both upper and lower limits in a bounded histogram. Custom bounded histograms can be created by users to identify key distribution characteristics of data. Integers, floating-point data, strings, and Boolean values can meet data analysis requirements in all sectors.

Using Histograms for GeminiDB Influx API

Example

In the example, **mst** is the table name. There are four fields of different data types and two tags. The raw data is as follows:

```
> select * from mst
name: mst
time                address   age alive country  height name
----                --------   --- ----- --------  ------ ----
1629129600000000000 shenzhen  12  true  china    70     azhu
1629129601000000000 shanghai  20  false american 80     alan
1629129602000000000 beijin    3   true  germany  90     alang
1629129603000000000 guangzhou 30  false japan    121    ahui
1629129604000000000 chengdu   35  true  canada   138    aqiu
```



```
1629129605000000000 wuhan     48        china    149    agang
1629129606000000000           52  true  american 153    agan
1629129607000000000 anhui     28  false germany         alin
1629129608000000000 xian          true  japan    179    ali
1629129609000000000 hangzhou  60  false canada   180
1629129610000000000 nanjin    102 true           191    ahuang
1629129611000000000 zhengzhou 123 false china    203    ayin
```

Equi-height histogram

Query syntax:

```
SELECT HISTOGRAM([ * | <field_key> | /<regular_expression>/ ], <N>) [WINTO_clause] FROM_clause
[WHERE_clause] [GROUP_BY_clause] [ORDER_BY_clause] [LIMIT_clause] [OFFSET_clause]  [SLIMIT_clause]
[SOFFSET_clause]
HISTOGRAM(field_key, N) generates a histogram visualization of the data distribution for field_key, with N
representing the desired number of bins in the histogram.
HISTOGRAM(/regular_expression/, N) calculates values of the field that matches the regular expression in
each bin.
HISTOGRAM(*, N) calculates values of integer and floating-point fields in each bin.
```

Example

1. Query an equi-height histogram in which *field_key* is specified as **age** and *N* as **5**.

```
> select histogram(age, 5) from mst where time >= 1629129600000000000 and time <=
1629129611000000000
name: mst
time histogram          value
---- ---------          -----
0    20                 3
0    30                 2
0    48                 2
0    60                 2
0    9223372036854775807 2
```

2. Query an equi-height histogram in which *field_key* matches regular expression **/hei/** (heights only) and *N* is specified as **5**.

```
> select histogram(/hei/, 5) from mst where time >= 1629129600000000000 and time <=
1629129611000000000
name: mst
time histogram_height        value
---- ----------------        -----
0    90                     3
0    138                     2
0    153                     2
0    180                     2
0    1.7976931348623157e+308  2
```

3. Query an equi-height histogram with addresses of the string type and an equi-height histogram with the **alive** field of the Boolean type.

```
> select histogram(address, 5) from mst where time >= 1629129600000000000 and time <=
1629129611000000000
name: mst
time                histogram value
----                --------- -----
1629129600000000000 chengdu   3
1629129600000000000 hangzhou  2
1629129600000000000 shanghai  2
1629129600000000000 wuhan     2
1629129600000000000 zhengzhou 2

> select histogram(alive, 3) from mst where time >= 1629129600000000000 and time <=
1629129611000000000
name: mst
time                histogram  value
----                ---------  -----
```



```
1629129600000000000 false     5
1629129600000000000 true      6
```

Bounded histogram

Query syntax:

```
SELECT HISTOGRAM([ * | <field_key> | /<regular_expression>/ ], 'specifyBins', boundary1, boundary2, ..., boundaryN) [WINTO_clause] FROM_clause [WHERE_clause] [GROUP_BY_clause] [ORDER_BY_clause] [LIMIT_clause] [OFFSET_clause] [SLIMIT_clause] [SOFFSET_clause]

HISTOGRAM (field_key, 'specifyBins', boundary1, boundary2,..., boundaryN) calculates values of a specified metric field in a specified bin. specifyBins indicates a flag and boundaryN the specified boundary value. Integers, floating-point data, strings, and Boolean values are supported.
HISTOGRAM(/regular_expression/, 'specifyBins',  boundary1, boundary2, ..., boundaryN) calculates values of the field that matches the regular expression in each bin.

HISTOGRAM (*,'specifyBins', boundary1, boundary2,..., boundaryN) calculates values of integer and floating-point fields in the specified bin.
```

Example

1. Query a bounded histogram in which *field_key* is specified as **age** with the following bins: [0, 10), [10, 20), [20, 30), [30, 40), and [40, 50)

```
> select histogram(age, 'specifyBins', 10, 20, 30, 40, 50) from mst
name: mst
time histogram value
---- --------- -----
0    10        1
0    20        2
0    30        2
0    40        1
0    50        1
```

2. Query a bounded histogram in which *field_key* matches regular expression **/eight/** (heights only) in the following bins [0, 160), [160, 170), [170, 180), [180, 190), [190, 200):

```
> select histogram(/eight/, 'specifyBins', 160.0, 170.0, 180.0, 190.0, 200.0) from mst
name: mst
time histogram value
---- --------- -----
0    160       7
0    170       0
0    180       2
0    190       0
0    200       1
```

3. Query a bounded histogram with addresses of the string type and a bounded histogram with the **alive** field of the Boolean type.

```
> select histogram(address, 'specifyBins', 'anhui', 'beijin', 'chengdu') from mst
name: mst
time histogram value
---- --------- -----
0    anhui     1
0    beijin    1
0    chengdu   1
# Boolean type
> select histogram(alive, 'specifyBins', false, true) from mst
name: mst
time histogram value
---- --------- -----
0    false     5
0    true      6
```

In addition to histogram functions, GeminiDB supports multiple other types of advanced analysis functions. For details, see **Supported Commands**.

Advantages

GeminiDB's advanced functions can be used to quickly analyze data and directly generate results. Databases handle the massive amount of data, eliminating the need for processing logic at the service layer. Using such functions will simplify your service design and reduce unnecessary resource consumption.

5.5 Multi-Level Downsampling

Application Scenarios

In DevOps or IoT scenarios, users focus on statistical metrics (such as maximum, minimum, and average values) instead of historical data details.

A conventional approach involves storing all granular data, querying it, and performing calculations on demand. If historical data is retained only for feature calculation purposes, the conventional approach presents several drawbacks: the longer granular data is retained, the higher the storage costs incurred.

Introduction to Multi-Level Downsampling

Multi-level downsampling policies are available to data across time ranges. For instance, data from the last week to month is aggregated every 15 minutes (regardless of aggregation methods), while historical data in the last month is aggregated every hour.

Users may be sensitive to recent data but have little requirements for long-term data, so different multi-level downsampling policies need to be used. Multi-level downsampling not only meets query requirements for high-value data, but also ensures storage efficiency.

For example, raw data in the last 7 days is directly imported to a database. Data from the last 7 to 30 days is downsampled every 15 minutes and then stored in a database. Data from the last 30 days to 12 months is downsampled every hour and then stored in a database. As shown in **Figure 5-17**, assume that today is December 31, 2022. The dark blue area on the left displays data within the last 7 days. The gray blue area in the middle displays data from the last 7 to 30 days. The light blue area on the right displays data from the last 30 days to 12 months. After a period of time, data in the dark blue area is aggregated into the gray blue area every 15 minutes, and data in the gray blue area is aggregated into the light blue area every hour. There are four 15-minute segments in an hour, so aggregation is available as a convenience to avoid extra implementation overhead.

Figure 5-17 Multi-level downsampling grid

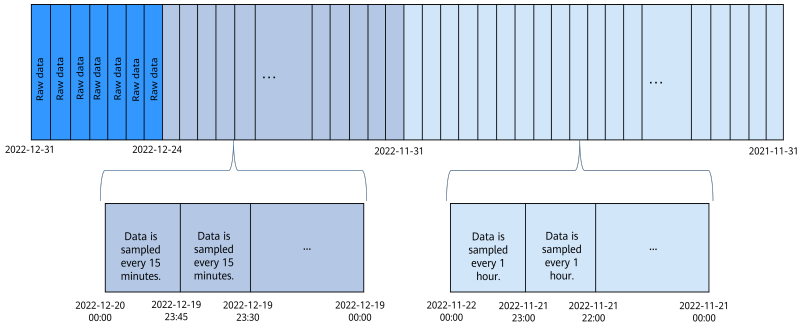

NOTE

Raw data from the last 7 to 30 days as well as 30 days to 12 months will not be retained after downsampling.

Creating a Downsampling Task

Syntax:

```
Create DownSample [on <rp_name>| on <dbname>.<rp_name>| ]((dataType(aggregators)...)) With
Duration <timeDuration> SampleInterval(time Durations) TimeInterval(time Durations)
```

Table 5-3 Parameters

Duration	SampleInterval	TimeInterval
Data retention period after downsampling	Next-level downsampling time	Sampling interval

Aggregation method definition:

```
dataType(aggfunctions...)
```

Example of an aggregation method:

```
integer(first,sum,count,last,min,max)
integer(min,max),float(sum)
```

For example:

Create a retention rule named rp1. Data is retained for seven days, and one shard is created every day.

```
create retention policy rp1 on mydb duration 7d replication 1 shard duration 1d
```

To create a downsampling task based on rp1, set the retention period of the sampled data to seven days. Details about data of the last day are stored. Data of the last one to two days is sampled every minute. Data generated two days ago is downsampled every 3 minutes.

```
Create DownSample on rp1 (float(sum,last),integer(max,min)) With Duration 7d sampleinterval(1d,2d)
timeinterval(1m,3m)
```


NOTE

- The second values of **sampleinterval** and **timeinterval** must be integer multiples of their first values.

  Valid example:

  - **sampleinterval(1d,2d)**
  - **timeinterval(1m,3m)**

  Invalid example:

  - **sampleinterval(2d,3d)**
  - **timeinterval(5m,6m)**
- The number of values of **sampleinterval** and **timeinterval** must be consistent.

  Valid example:

  - **sampleinterval(1d,2d)** and **timeinterval(1m,3m)**

  Invalid example:

  - **sampleinterval(1d,2d)** and **timeinterval(3m)**
  - **sampleinterval(1d)** and **timeinterval(1m,3m)**
- **Duration** controls the retention period of downsampled data. The value specified for rp1 is updated synchronously and can be the same as that specified for rp. Only **first**, **last**, **sum**, **max**, **min**, **mean**, and **count** are allowed.
- **float(sum,last)** indicates that FLOAT data is downsampled using the sum() and last() functions.
- The mechanism of **integer(max,min)** is similar to that of **float(sum,last)**.

Showing Downsampling Tasks

Show all downsampling tasks of the default database.

```
SHOW DOWNSAMPLES
```

Show all downsampling tasks of a specified database.

```
SHOW DOWNSAMPLES ON <database name>
```

For example:

```
> show downsamples on ExampleDatabase
+--------+-------------------------+----------+----------------+--------------+
| rpName |     field_operator      | duration | sampleInterval | timeInterval |
+--------+-------------------------+----------+----------------+--------------+
| rp1    | float{max},integer{min} | 24h0m0s  | 1h0m0s,2h0m0s  | 1m0s,5m0s    |
+--------+-------------------------+----------+----------------+--------------+
```

Deleting Downsampling Tasks

Delete all downsampling tasks of a specified database.

```
Drop DownSamples
Drop DownSamples on db0
```

Delete all downsampling tasks of a specified RP.

```
Drop DownSample on rp1
Drop DownSample on db0.rp1
```


5.6 Suggestions on Alarm Rules of GeminiDB Influx Instance Metrics

After setting alarm rules on the Cloud Eye console, for example, specifying monitored objects and notification policies, you can stay ahead of your instance status. For details, see **Configuring Alarm Rules**.

This section describes recommended alarm rules of GeminiDB Influx instances.

Table 5-4 Suggestions on alarm rules of GeminiDB Influx instance metrics

Metric ID	Metric Name	Dimension	Threshold (Raw Value) in Best Practices	Alarm Severity in Best Practices	Alarm Handling Suggestion
gemini001_cpu_usage	CPU Usage	Node	> 90% for 3 consecutive periods	Major	Upgrade instance specifications. For details, see **Changing vCPUs and Memory**.
gemini002_mem_usage	Memory Usage	Node	> 80% for 3 consecutive periods	Major	Upgrade instance specifications. For details, see **Changing vCPUs and Memory**.
nosql005_disk_usage	Storage Space Usage	Instance	> 80% for 3 consecutive periods	Major	Evaluate how much storage needs to be added based on data growth. For details, see **Manually Scaling Up Storage Space of a GeminiDB Influx Instance**.
influxdb005_write_concurrency	Concurrent Write Requests	Node	≥ CPU cores x 4 for 3 consecutive periods	Major	Check whether the service traffic increases sharply and whether the database is normal. For details, see **Viewing Metrics**.
influxdb006_query_concurrency	Concurrent Queries	Node	≥ CPU cores for 3 consecutive periods	Major	Check whether the service traffic increases sharply and whether the database is normal. For details, see **Viewing Metrics**.

5.7 How Do I Improve Write Efficiency of GeminiDB Influx Instances?

This section uses influxdb-java as an example to describe how to set the write policy to improve write efficiency of GeminiDB Influx instances.

Core strategy: Write data concurrently in batches. More than 256 concurrent connections in more than 400 batches are recommended.

1. SDK Lifecycle Management

- **Create**: In one process, you only need to create a single global instance using the InfluxDB client.
- **Use**: Call the write or query method without close() after each operation.
- **Destroy**: Call close() only once when the process is shut down to release resources.

2. Submitting Data Points In Batches

- **Advantage**: Compared with submission of a single data point, batch submission can significantly minimize network overhead and greatly improve the overall throughput.
- **Default SDK configuration**: The default batch size of influxdb-java in asynchronous submission mode is 1,000 data points/batch.

3. Optimizing Asynchronous Writes

- **Enable asynchronous writes**: Enable **enableBatch** and call the write method.
- After **enableBatch** is enabled, one client can asynchronously write data to GeminiDB using only one thread. If the write load is too heavy for a single client, you can initialize multiple clients to write data.

```
InfluxDB influxDB = InfluxDBFactory.connect(serverURL, "username", "password", client);
BatchOptions batchOptions = BatchOptions.DEFAULTS.
//Number of points submitted in each batch
    .actions(1000)
    //Length of the write error buffer queue
    .bufferLimit(20000)
    //Asynchronous submission interval
    .flushDuration(100)
    //Handles a write exception.
    .exceptionHandler((points, throwable) -> {});
influxDB.enableBatch(batchOptions);
influxDB.write(point)
```

4. Optimizing Synchronous Writes

- **Enable synchronous writes**: Enable **disableBatch** and call the write method.
- The performance depends on the number of data points submitted each time. You are advised to submit data points in batches each time.
- To improve the throughput, you can maintain a queue outside the SDK. An external program can control the queue to obtain data in batches and call write() to submit the request.


```
InfluxDB influxDB = InfluxDBFactory.connect(serverURL, "username", "password", client);
influxDB.disableBatch();
influxDB.write(point)
```

5.8 Basic Syntax Examples of GeminiDB Influx Instances

This section describes the basic syntax of GeminiDB Influx instances.

- Database syntax
  - Create a database.

    create_database_stmt = "CREATE DATABASE" db_name

    [WITH

    [retention_policy_duration]

    [retention_policy_replication]

    [retention_policy_shard_group_duration]

    [retention_policy_name]

    ] .

    NOTE

    The commands in square brackets ([]) are optional.

    Example:

    - Create a database named **mydb**.

      CREATE DATABASE "mydb"

    - Create a database named **mydb** using the specified retention policy **myrp**. Set the retention period to 1 day, the number of copies to 1, and the storage duration of **shardGroup** to 30 minutes.

      CREATE DATABASE "mydb" WITH DURATION 1d REPLICATION 1 SHARD DURATION 30m NAME "myrp"

    - Create a database named **mydb** using the default retention policy **myrp**.

      CREATE DATABASE "mydb" WITH NAME "myrp"

  - Query databases.

    SHOW DATABASES

  - Switch to another database.

    USE db_name

  - Delete a database.

    DROP DATABASE "db_name"


```
InfluxDB shell version: 1.7.4
> create database demo
> show databases
name: databases
name
----
_internal
mydb
demo
> use mydb
Using database mydb
>
```

- RETENTION POLICY
  - Create a retention policy and ensure that the policy name does not contain periods (,), colons (:), semicolons (;), or dots (.).

    create_retention_policy_stmt = "CREATE RETENTION POLICY" policy_name on_clause

    retention_policy_duration

    retention_policy_replication

    [retention_policy_shard_group_duration]

    ["DEFAULT"] .

    NOTE

    The commands in square brackets ([]) are optional.

    Example:

    - Create a data retention policy.

      CREATE RETENTION POLICY "10m_events" ON "somedb" DURATION 60m REPLICATION 2

    - Create a data retention policy and set it as the default one.

      CREATE RETENTION POLICY "10m_events" ON "somedb" DURATION 60m REPLICATION 2 DEFAULT

    - Create a data retention policy and specify the storage duration of **shardGroup**.

      CREATE RETENTION POLICY "10m_events" ON "somedb" DURATION 60m REPLICATION 2 SHARD DURATION 30m

  - View a retention policy.

    show retention policies on <database name>

```
> CREATE RETENTION POLICY "rp_1_day" ON "mydb" DURATION 1d REPLICATION 2 SHARD DURATION 1h
> show retention policies
name       duration warmDuration shardGroupDuration replicaN default
----       -------- ------------ ------------------ -------- -------
autogen    0s       0s           168h0m0s           1        true
rp_1_hours 1h0m0s   0s           1h0m0s             1        false
rp_1_day   24h0m0s  0s           1h0m0s             1        false
>
```


NOTE

If you specify both parameters **retention_policy_duration** and **retention_policy_shard_group_duration**, ensure that the former parameter has a larger value than the latter.

- Delete a retention policy.

  DROP RETENTION POLICY policy_name ON db_name

- Modify a retention policy.

  Alter_retention_policy_stmt = "ALTER RETENTION POLICY" policy_name on_clause

  retention_policy_option

  [retention_policy_option]

  [retention_policy_option]

  [retention_policy_option] .

  NOTE

  The commands in square brackets ([]) are optional.

  Example:

  - Modify the default retention policy.

    ALTER RETENTION POLICY "1h_cpu" ON "mydb" DEFAULT

  - Modify the retention period and number of copies.

    ALTER RETENTION POLICY "policy1" ON "somedb" DURATION 1h REPLICATION 4

- Add data.

  insert into <retention policy> **measurement,tagKey=tagValue fieldKey=fieldValue timestamp**

  NOTE

  When data is inserted, the system creates a measurement as required.

  - Use the default retention policy.

    insert demo,name=LiSi math=99,english=90,language=95

    Add data. Set **measurement** to **demo**, **tag** to **name**, and **field** to **math**, **english**, and **language**.

  - Use the specified retention policy.

    insert into rp_1_hours demo,name=ZhangSan math=99,english=90,language=95

- Query data.

  - Query data from the default retention policy.

    select * from demo where time < xxx and time > xxx

  - Query data from the specified retention policy.

    select * from rp_1_hours.demo where time < xxx and time > xxx

    NOTE

    Specify a time range in the query statement.

- Modify data.

  When you modify data using INSERT, if all tags and timestamps are the same, the existing data will be overwritten.

```
> select * from demo
name: demo
time                           english language math name
----                           ------- -------- ---- ----
2019-07-26T13:55:27.925320596Z 90      86       95   LiLei
2022-07-14T08:14:54.593459723Z 90      95       99   LiSi
2022-07-14T09:07:48.520893767Z 70      86       77   ZhangSan
>
> insert demo,name=LiLei math=90,english=91,language=88 1564149327925320596
> select * from demo
name: demo
time                           english language math name
----                           ------- -------- ---- ----
2019-07-26T13:55:27.925320596Z 91      88       90   LiLei
2022-07-14T08:14:54.593459723Z 90      95       99   LiSi
2022-07-14T09:07:48.520893767Z 70      86       77   ZhangSan
>
```

- Delete data.

  You can create a retention policy to automatically delete data.

- **HELP** command
  - Run **HELP** to view all supported commands.

    Figure 5-18 Viewing all supported commands

```
InfluxDB shell version: 1.7.4
> help
Usage:
        connect <host:port>   connects to another node specified by host:port
        auth                  prompts for username and password
        pretty                toggles pretty print for the json format
        chunked               turns on chunked responses from server
        chunk size <size>     sets the size of the chunked responses.  Set to 0 to reset to the default chunked size
        use <db_name>         sets current database
        format <format>       specifies the format of the server responses: json, csv, or column
        precision <format>    specifies the format of the timestamp: rfc3339, h, m, s, ms, u or ns
        consistency <level>   sets write consistency level: any, one, quorum, or all
        history               displays command history
        settings              outputs the current settings for the shell
        clear                 clears settings such as database or retention policy.  run 'clear' for help
        exit/quit/ctrl+d      quits the influx shell

        show databases        show database names
        show series           show series information
        show measurements     show measurement information
        show tag keys         show tag key information
        show field keys       show field key information

        A full list of influxql commands can be found at:
        https://docs.influxdata.com/influxdb/latest/query_language/spec/
>
```

  - Run **HELP** *<COMMAND>* to query the usage of a command.

    Example: **HELP DESC**

6 Performance White Paper

6.1 Performance Test Methods

This section describes performance testing of GeminiDB Influx instances, including the test environment, procedure, and results.

Test Environment

- Region: CN-Hong Kong
- AZ: AZ1
- Elastic Cloud Server (ECS): m6.2xlarge.8 with 8 vCPUs, 64 GB of memory, and CentOS 7.6 64-bit image
- Nodes per instance: 3
- Instance specifications: 4 vCPUs | 16 GB, 8 vCPUs | 32 GB, 16 vCPUs | 64 GB, and 32 vCPUs | 128 GB

Test Tool

Time Series Benchmark Suite (TSBS) of the open-source community is used.

Test Metrics

- Write performance test: Data points per second
- Query performance test: Latency and OPS (operations per second)

Test Procedure

Step 1 Run the following command to generate the data to be written:

tsbs_generate_data --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:00Z" --log-interval="10s" --format="influx" | gzip > /tmp/influx-data.gz

NOTE

--scale indicates the number of time series to be generated.

--log-interval indicates the interval for collecting data.

Step 2 Run the following command to test write performance and obtain the required data:

NUM_WORKERS=${numWorkers} BATCH_SIZE=${batchSize} DATABASE_HOST=${influxIP} DATABASE_PORT=${influxPORT} BULK_DATA_DIR=/tmp scripts/load_influx.sh

Step 3 Run the following commands to generate query statements:

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=20 --query-type="high-cpu-all" --format="influx" | gzip > /tmp/influx-20queries-high-cpu-all-12h-frequency.gz

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=1000000 --query-type="single-groupby-1-8-1" --format="influx" | gzip > /tmp/influx-1000000queries-single-groupby-1-8-1-12h-frequency.gz

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=500 --query-type="double-groupby-1" --format="influx" | gzip > /tmp/influx-500queries-double-groupby-1-12h-frequency.gz

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=50 --query-type="double-groupby-all" --format="influx" | gzip > /tmp/influx-50queries-double-groupby-all-12h-frequency.gz

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=200 --query-type="lastpoint" --format="influx" | gzip > /tmp/influx-200queries-lastpoint-12h-frequency.gz

tsbs_generate_queries --use-case="devops" --seed=123 --scale=10000 --timestamp-start="2016-01-01T00:00:00Z" --timestamp-end="2016-01-01T12:00:01Z" --queries=500 --query-type="groupby-orderby-limit" --format="influx" | gzip > /tmp/influx-500queries-groupby-orderby-limit-12h-frequency.gz

NOTE

Ensure that values of fields **--use-case**, **-seed**, **--scale**, and **--timestamp-start** must be the same as those values set when data is generated in **Step 1**.

--timestamp-end indicates the second after data generation ends.

--queries indicates the number of generated queries.

--queries-type indicates the type of generated queries. For details, see **Table 6-1**.

Step 4 Run the following commands to query performance data:

cat /tmp/influx-20queries-high-cpu-all-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 10 --urls=(http|https)://${influxIP}:${influxPORT}

cat /tmp/influx-1000000queries-single-groupby-1-8-1-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 10000 --urls=(http|https)://${influxIP}:${influxPORT}

cat /tmp/influx-500queries-double-groupby-1-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 50 --urls=(http|https)://${influxIP}:${influxPORT}

cat /tmp/influx-50queries-double-groupby-all-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 10 --urls=(http|https)://${influxIP}:${influxPORT}

cat /tmp/influx-200queries-lastpoint-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 10 --urls=(http|https)://${influxIP}:${influxPORT}

cat /tmp/influx-500queries-groupby-orderby-limit-12h-frequency.gz | gunzip | tsbs_run_queries_influx --workers=${numWorkers} --print-interval 50 --urls=(http|https)://${influxIP}:${influxPORT}

----End

Test Models

Table 6-1 Test models involved

Test Model	Description	Example Statement
load	100% insertion.	-
high-cpu-all	Queries all the readings where one metric is above a threshold across all hosts for a period of time.	SELECT * from cpu where usage_user > 90.0 and time >= '2020-11-01T05:24:55Z' and time < '2020-11-01T17:24:55Z'
single-groupby-1-8-1	Queries the maximum value of one metric for 8 hosts for a period of time.	SELECT max(usage_user) from cpu where (hostname = 'host_61885' or hostname = 'host_51710' or hostname = 'host_9380' or hostname = 'host_46446' or hostname = 'host_67623' or hostname = 'host_54344' or hostname = 'host_82215' or hostname = 'host_7458') and time >= '2020-11-01T19:38:15Z' and time < '2020-11-01T20:38:15Z' group by time(1m)

Test Model	Description	Example Statement
single-groupby-1-1-1	Queries the maximum value of one metric for 1 host for a period of time.	SELECT max(usage_user) from cpu where (hostname = 'host_6334') and time >= '2016-01-01T03:03:21Z' and time < '2016-01-01T04:03:21Z' group by time(1m)
cpu-max-all-1	Queries the maximum value of all metrics for 1 host for a period of time.	SELECT max(usage_user),max(usage_system),max(usage_idle),max(usage_nice),max(usage_iowait),max(usage_irq),max(usage_softirq),max(usage_steal),max(usage_guest),max(usage_guest_nice) from cpu where (hostname = 'host_1166') and time >= '2016-01-01T00:23:32Z' and time < '2016-01-01T08:23:32Z' group by time(1h)

6.2 Performance Test Data

Write Performance Testing

Table 6-2 Data used for testing write performance of cluster instances

Instance Specifications	Concurrent Requests	Write Performance (unit: rows/sec)
4 vCPUs | 16 GB	20	123648.75
8 vCPUs | 32 GB	40	221034.80
16 vCPUs | 64 GB	80	348762.25
32 vCPUs | 128 GB	160	496511.06

Table 6-3 Data used for testing write performance of single-node instances

Instance Specifications	Concurrent Requests	Write Performance (unit: rows/sec)
4 vCPUs | 8 GB	10	50113
8 vCPUs | 16 GB	10	108781
16 vCPUs | 32 GB	20	158744

Query Performance Testing

The test data depends on service models and instance specifications.

NOTE

Table 6-1 describes the service models involved in this test.

- **The test instance of the cluster type, with specifications of 4 vCPUs and 16 GB and concurrent requests of 20. All metrics in the following table are average values calculated from 1 million executions.**

Table 6-4 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	710.81	28.12	9.35	714.11
single-groupby-1-8-1	1308.56	13.74	2.56	148.96
single-groupby-1-1-1	6393.67	3.10	1.43	45.02
cpu-max-all-1	850.51	23.49	6.16	715.23

- **The test instance of the cluster type, with specifications of 8 vCPUs and 32 GB and concurrent requests of 40. All metrics in the following table are average values calculated from 1 million executions.**

Table 6-5 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	1236.46	32.34	9.70	412.86
single-groupby-1-8-1	2663.19	12.47	2.58	222.84
single-groupby-1-1-1	9696.13	4.03	1.56	141.06
cpu-max-all-1	1406.48	28.42	8.97	444.16

- **The test instance of the cluster type, with specifications of 16 vCPUs and 64 GB and concurrent requests of 80. All metrics in the following table are average values calculated from 1 million executions.**

  Table 6-6 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	2107.83	37.94	11.33	252.74
single-groupby-1-8-1	4707.25	15.40	3.29	225.18
single-groupby-1-1-1	17658.59	4.44	1.80	51.16
cpu-max-all-1	2262.40	35.35	12.80	247.85

- **The test instance of the cluster type, with specifications of 32 vCPUs and 128 GB and concurrent requests of 160. All metrics in the following table are average values calculated from 1 million executions.**

  Table 6-7 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	3468.89	46.10	19.14	290.61
single-groupby-1-8-1	5107.15	13.84	3.58	118.97
single-groupby-1-1-1	23023.11	6.72	1.80	74.45
cpu-max-all-1	3715.62	43.04	14.24	186.80

- **The test instance of the single-node type, with specifications of 4 vCPUs and 8 GB and concurrent requests of 10. All metrics in the following table are average values calculated from 1 million executions.**

Table 6-8 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	423.07	14.17	5.24	693.53
single-groupby-1-8-1	1278.77	4.68	2.01	822.53
single-groupby-1-1-1	3138.4	1.9	1.1	424.77
cpu-max-all-1	357.93	16.75	8.51	992.06

- **The test instance of the single-node type, with specifications of 8 vCPUs and 16 GB and concurrent requests of 20. All metrics in the following table are average values calculated from 1 million executions.**

Table 6-9 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	1031.77	15.49	5.37	614.3
single-groupby-1-8-1	3082.18	5.18	2.12	154.53
single-groupby-1-1-1	7604.41	2.1	0.96	31.93
cpu-max-all-1	856.75	18.66	7.76	573.18

- **The test instance of the single-node type, with specifications of 16 vCPUs and 32 GB and concurrent requests of 20. All metrics in the following table are average values calculated from 1 million executions.**

Table 6-10 Test data

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
high-cpu-all	1646.46	12.13	4.4	409.82

Test Model	OPS (unit: queries/sec)	Average Latency (unit: ms)	Minimum Latency (unit: ms)	Maximum Latency (unit: ms)
single-groupby-1-8-1	3909.19	5.11	2.15	122.95
single-groupby-1-1-1	10340.02	1.93	1.02	146.8
cpu-max-all-1	1181.3	16.92	7.79	175.29

7 FAQs

7.1 Product Consulting

7.1.1 What Do I Need to Note When Using GeminiDB Influx 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. The backup files stored in OBS and the system containers used by GeminiDB Influx API are invisible to you. They are visible only in the GeminiDB Influx API management system.
3. Precautions after purchasing instances:

   After purchasing 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 insufficient for the DB instances.
   b. The instance has performance problems and whether optimization is required.

7.1.2 What Does the Availability of GeminiDB Influx Instances Mean?

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 a DB instance after you buy the instance. The total service duration refers to the total running time of the instance.

7.1.3 Can GeminiDB Influx API Convert Multiple Columns to Multiple Rows?

GeminiDB Influx API does not support the function for converting multiple columns into multiple rows.

7.1.4 How Much Data Can a GeminiDB Influx Instance Hold?

For details, see **Instance Specifications**.

7.1.5 Can I Access GeminiDB Influx Instances Using Grafana?

Yes. You can access GeminiDB Influx Instances using Grafana. For details, see **How Do I Connect to a GeminiDB Influx Instance Using Grafana?**.

7.1.6 How Do I Use GeminiDB Influx Hints?

GeminiDB Influx API supports hints, improving query performance. Hints can be used only when you need to specify a value for each tag in a query statement. To use hints, add **/*+ full_series */** before an SQL statement.

For example:

A common query statement is as follows:

```
select value from cpu where server_id=1;
```

If a hint is used, the corresponding syntax is:

```
select /*+ full_series */ value from cpu where server_id=1;
```

7.1.7 What Do I Do If Error "select *" query without time range is not allowed Is Reported?

When you execute a query statement like SELECT* and give no constraints on the time range, error "select *" query without time range is not allowed will be reported. To resolve this problem, you need to rectify the query statement and specify time range constraints.

Example:

- `select * from measurement where time > '2023-01-19T12:00:00Z' and time <= '2023-01-19T13:00:00Z'`
- `select * from measurement where time = '2023-01-19T12:30:00Z'`

7.1.8 What Do I Do If the Error Message "ERR: Max-select-series Limit Exceeded" Is Displayed?

If the timeline involved in the result returned by a query statement exceeds the limit, the error "max-select-series limit exceeded" is triggered. Two solutions are available:

1. Optimize the query statement and add timeline constraints. Tag restriction information is added to the WHERE statement to narrow down the tag query scope and ensure that the timeline restriction is not exceeded.

2. Scale up the instance specifications. The number of timelines allowed for query is related to the instance specifications. The larger the instance specifications, the larger the number of timelines allowed.

The **limit** keyword cannot reduce the timelines involved in the query. Therefore, the error cannot be rectified by using the keyword.

7.1.9 What Do I Do If "delete is forbidden" Is Reported?

When a logical deletion command, such as **delete/drop measurement**, is executed, error message "delete is forbidden" is displayed.

Executing a logical deletion command is inefficient, and the system may be suspended. Set a retention period so that data can be automatically deleted.

7.1.10 What Should I Do If "THE TOTAL NUMBER OF DBs EXCEEDS THE LIMIT 16" Is Displayed?

The number of databases that can be created varies depending on the specifications. If the number exceeds the threshold, this error message is displayed.

You are advised to specify an appropriate value. For details about the relationship between the number of databases and specifications, see **Instance Specifications**.

7.1.11 What Should I Do If "THE TOTAL NUMBER OF RPs EXCEEDS THE LIMIT 16" Is Displayed?

The number of RPs that can be created varies depending on the specifications. If the number exceeds the threshold, this error message is displayed.

You are advised to specify an appropriate value. For details about the relationship between the number of RPs and specifications, see **Instance Specifications**.

7.2 Billing

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

Yearly/Monthly is a prepaid billing mode in which resources are billed based on the service duration. This cost-effective mode is ideal when the duration of resource usage is predictable. It is recommended for long-term users.

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 Payments?

You can change the billing mode from yearly/monthly to pay-per-use or vice versa.

- If you want to change the billing mode from yearly/monthly to pay-per-use, see **Changing a Yearly/Monthly Instance to Pay-per-Use**.
- If you want to change the billing mode from pay-per-use 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*.
   - The ECS to be created must be in the same VPC with the GeminiDB Influx instance to which it connects.
   - 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 User Guide*.

7.3.2 Can I Change the VPC of a GeminiDB Influx Instance?

Once a GeminiDB Influx instance is created, the VPC where the instance resides cannot be changed.

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 Data to a New Instance**.

7.3.3 How Do I Connect to a GeminiDB Influx Instance Locally?

You can connect to a GeminiDB Influx instance using a private network, public network, or program code. For details, see **Connecting to a GeminiDB Influx Instance**.

7.3.4 How Do I Connect to a GeminiDB Influx Instance Using Grafana?

Grafana is a cross-platform open source analytics and interactive visualization web application. It provides charts, graphs, and alerts for the web when connected to supported data sources.

This section describes how to connect to a GeminiDB Influx instance using Grafana.

Procedure

Step 1 Start Grafana on the server and access **http://IP:3000** using a browser.

NOTE

The **IP** field can be an elastic IP address of a cloud server or the IP address of an on-premises server.

Step 2 Log in to the Grafana homepage.

Figure 7-1 Logging in to the Grafana homepage

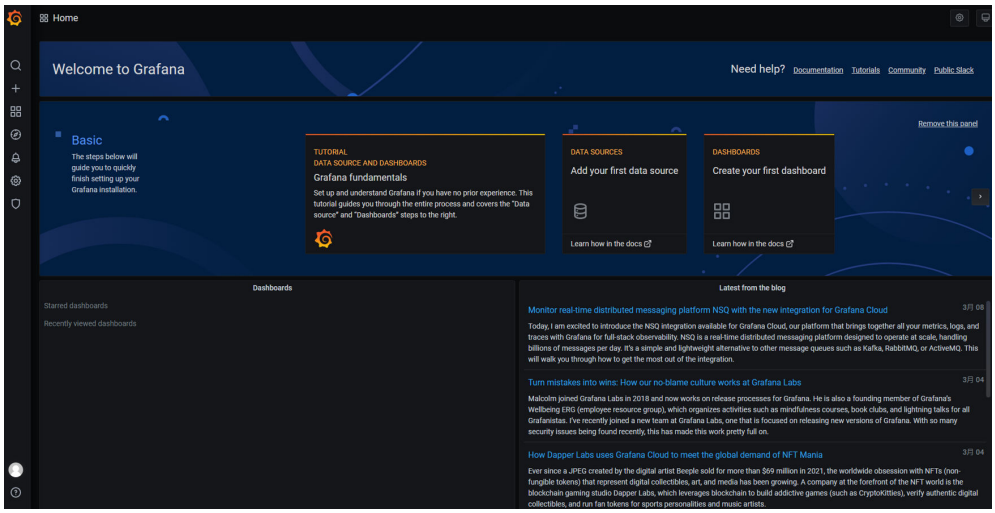

Step 3 Create a data source.

Figure 7-2 Creating a data source

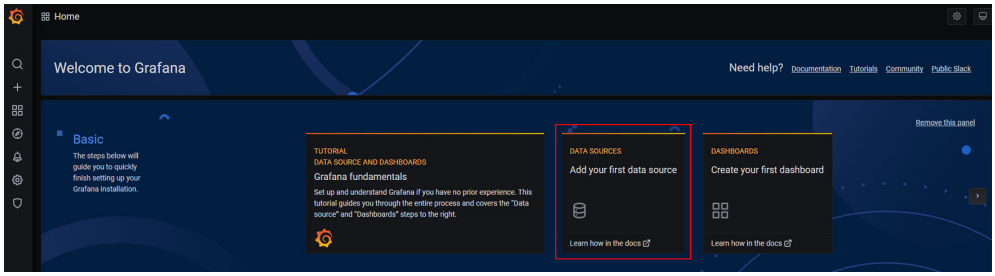

Step 4 Select **InfluxDB**.

Figure 7-3 Selecting **InfluxDB**

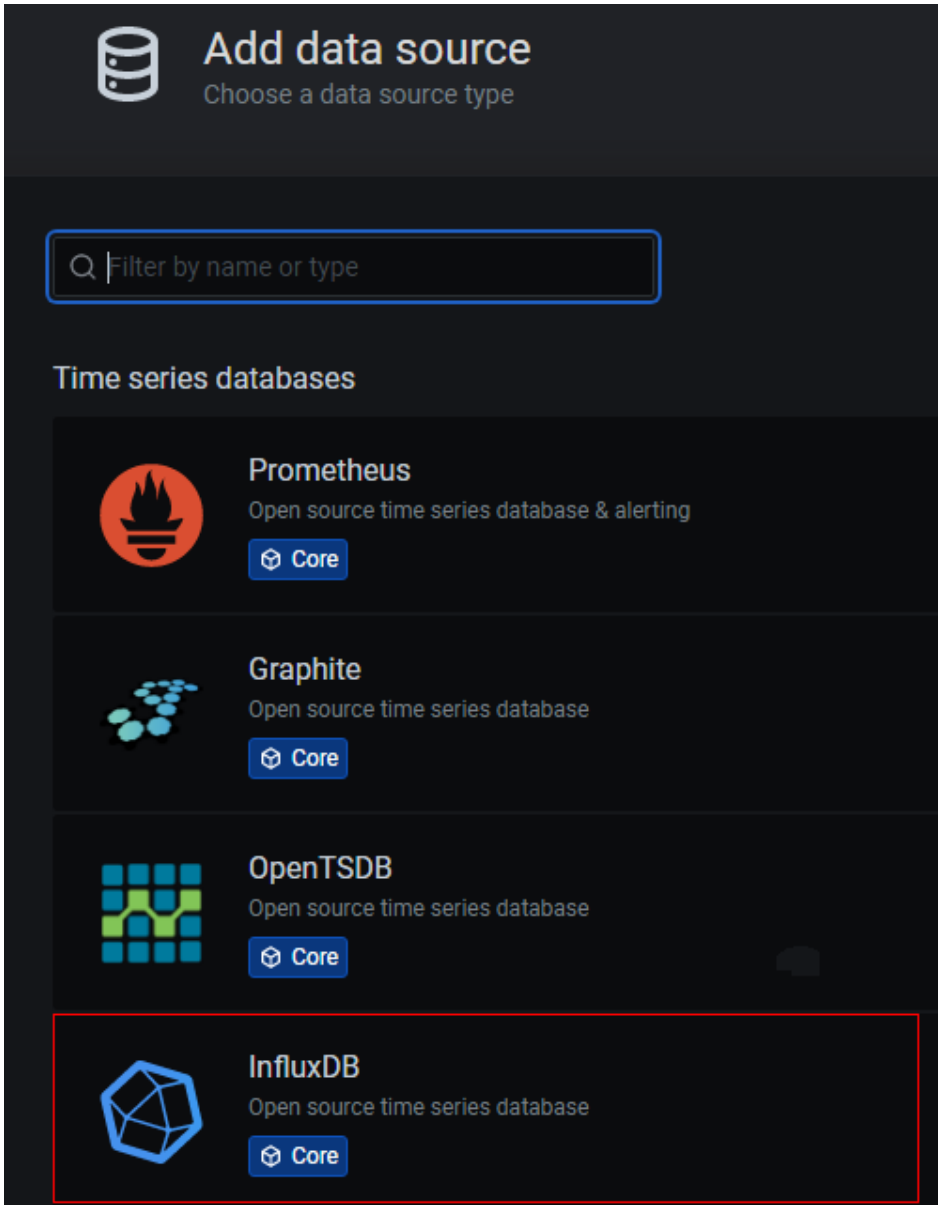

Step 5 Configure the required parameters.

Figure 7-4 Configuring parameters

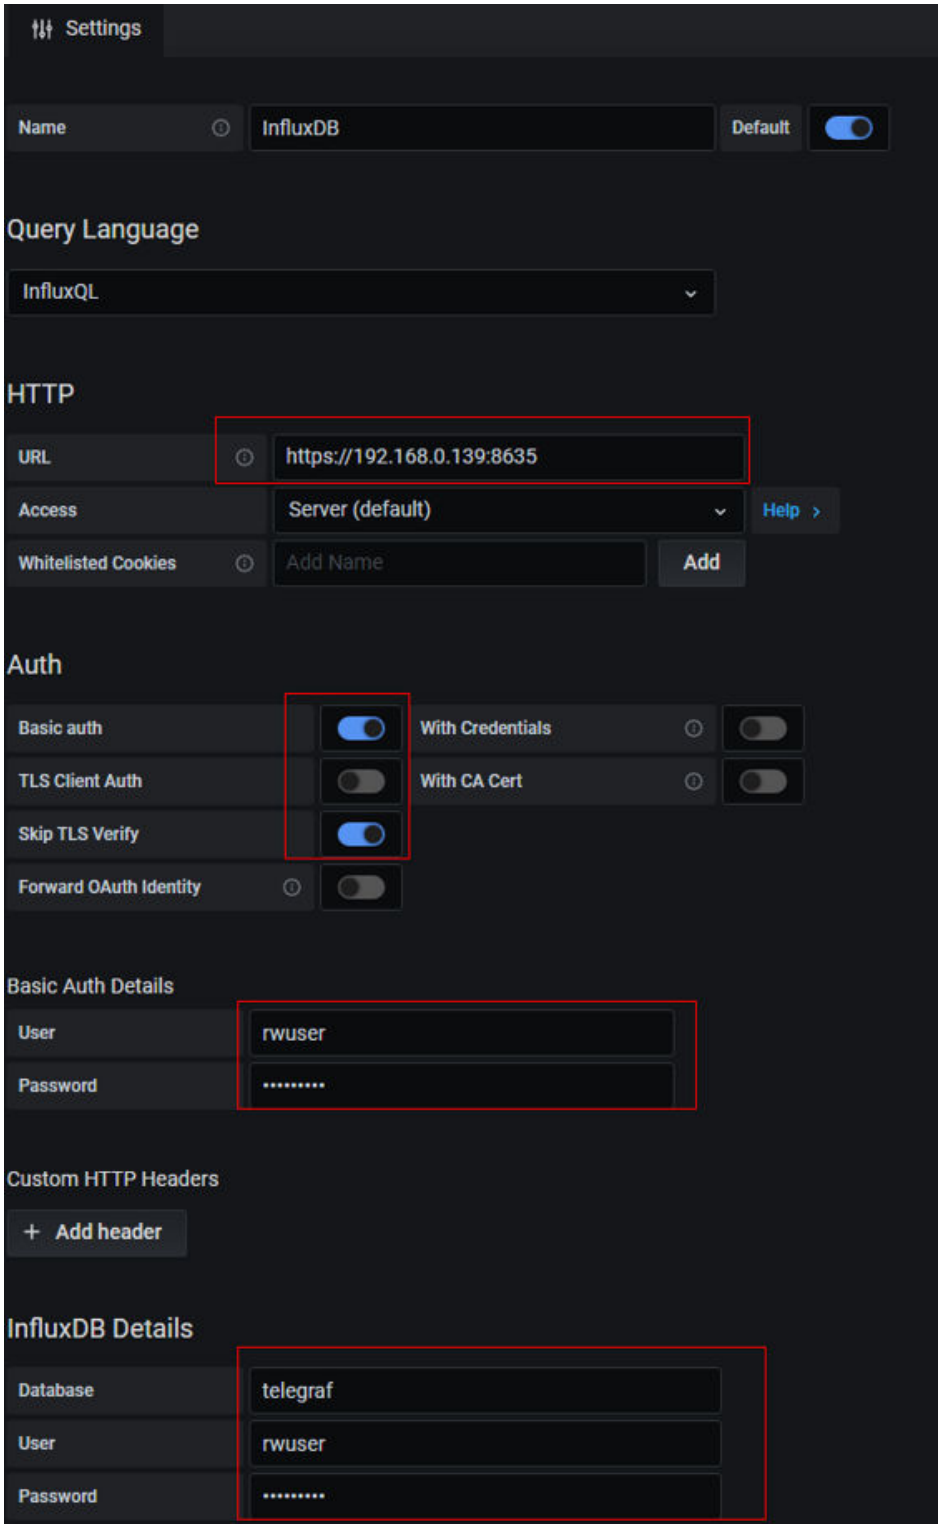

Table 7-1 Parameter description

Parameter	Description
URL	URL format: https://<IP>:8635 The **IP** field indicates the private IP address of the database instance.
Auth	Open **Basic auth** and **skip TSL Verify**.
Basic Auth Details	• **User**: Username, for example, **rwuser** • **Password**: The password you set when you buy a GeminiDB Influx instance
InfluxDB Details	• **Database**: Name of the created database, for example telegraf • **User**: **rwuser** • **Password**: The password you set when you buy a GeminiDB Influx instance

Step 6 Click **Save**.

Step 7 Create a dashboard based on service requirements.

----End

Related Issues

If you fail to connect to a GeminiDB Influx instance using Grafana, the causes may be as follows:

- Network connection is abnormal.
- The URL address is incorrect. When you enter a URL, make sure to type colons (:) and **https** correctly.
- SSL authentication failed. Note to select **skip ssl verify**.

7.4 Backup and Restoration

7.4.1 How Long Can a GeminiDB Influx Instance Backup Be Saved?

Automated backup data is kept based on the backup retention period you specified. There is no limit for the manual backup retention period. You can delete manual backups as needed.

7.5 Regions and AZs

7.5.1 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**.

7.5.2 Can I Change the Region of a GeminiDB Influx Instance?

No. After an instance is created, its region cannot be changed.

7.6 Instance Freezing, Release, Deletion, and Unsubscription

Why Are My GeminiDB Influx Instances 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 the subscriptions or top up your account after the grace period expires, your instances will enter a retention period and become unavailable. 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 **Service Suspension and Resource Release**.

Why Are My GeminiDB Influx Instances Frozen?

Your instances may be frozen for a variety of reasons. The most common reason is that you are in arrears.

Can I Still Back Up Data If My Instances Are Frozen?

No. If your instances are frozen because your account is in arrears, go to top up your account to unfreeze your instances and then back up instance data.

How Do I Unfreeze My Instances?

If your instances are frozen because your account is in arrears, you can unfreeze them by renewing them or topping up your account. Frozen GeminiDB Influx instances can be renewed, released, or deleted. Expired yearly/monthly GeminiDB Influx instances cannot be unsubscribed from, while those that have not expired can be unsubscribed from.

What Impacts Does Instance Freezing, Unfreezing or Release Have on My Services?

- After an instance is frozen:
  - It cannot be accessed, and your services will be interrupted. For example, if a GeminiDB Influx 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.
- Releasing an instance means deleting it. Before the deletion, GeminiDB Influx API determines whether to **move the instance to the recycle bin** based on the recycling policy you specified.

How Do I Renew My Instances?

After a yearly/monthly GeminiDB Influx instance expires, you can renew it on the **Renewal Management** page. For details, see **Renewal Management**.

Can My Instances Be Recovered After They Are Released or Unsubscribed From?

If your instance is moved to the recycle bin after being deleted, you can recover it from the recycle bin by referring to **GeminiDB Influx Instance Recycle Bin**. If the recycling policy is not enabled, you cannot recover 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 Influx 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?**.

