GeminiDB Influx

User Guide

Issue 01

Date 2025-04-09

HUAWEI

HUAWEI CLOUD COMPUTING TECHNOLOGIES CO., LTD.

Trademarks and Permissions


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

Notice

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

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

Contents

1 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 points 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

  The column-oriented storage and dedicated compression algorithm improve the compression ratio of GeminiDB Influx by 5 to 10 times compared with the open-source version.
- Efficient query

  GeminiDB Influx 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 points per day and support multi-node and multi-thread parallel query.

  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 APIs and Versions

GeminiDB Influx instances support the following types: cluster.

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.

1.3 Instance Specifications

Instances of the same type can have different memory specifications. You can select instances of different specifications based on application scenarios.

This section describes the instance specifications supported by GeminiDB Influx. 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. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb.large.4	2	8	100	12,000	250	4	40	1,000	5,000
geminidb.influxdb.xlarge.4	4	16	100	24,000	500	16	40	2,000	20,000
geminidb.influxdb.2xlarge.4	8	32	100	48,000	1,000	64	80	4,000	80,000
geminidb.influxdb.4xlarge.4	16	64	100	96,000	2,000	256	160	8,000	320,000

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. RPs per Cluster	Maximum Fields per Query	Maximum Time Series per Query
geminidb.influxdb.8xlarge.4	32	128	100	192,000	4,000	1,024	320	16,000	1,280,000
geminidb.influxdb.large.8	2	16	100	12,000	750	6	60	1,500	7,500
geminidb.influxdb.xlarge.8	4	32	100	24,000	750	24	60	3,000	30,000
geminidb.influxdb.2xlarge.8	8	64	100	48,000	1,500	96	120	6,000	120,000
geminidb.influxdb.4xlarge.8	16	128	100	96,000	3,000	384	240	12,000	480,000
geminidb.influxdb.8xlarge.8	32	256	100	192,000	6,000	1,536	480	24,000	1,920,000

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, new read and write requests are 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 1-2 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
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 1-3 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.

Status	Description
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.
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.

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 FIELDs.
  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**: The number of FIELDs should not be too large. Each FIELD is calculated independently. Too many FIELDs may cause the fuzzy query to 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 **Connecting to an Instance Using a Load Balancer Address (Recommended)**.
- **Note**: If cold storage is enabled, cold data cannot be written.

2 Getting Started with GeminiDB Influx API

2.1 Getting Started with GeminiDB Influx API

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

Table 2-1 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**

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 2-2 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 **Connection Methods**.

2.2 Buying and Connecting to a Cluster Instance

This section describes how to buy and connect to a GeminiDB Influx cluster 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.

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 customer service.

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

  For details about other connection methods, see **Connection Methods**.

Step 1: Buying a Cluster Instance

For details, see **Buying a GeminiDB Influx Cluster 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 3-1** lists details about the parameters.

Figure 2-1 Billing mode and basic information

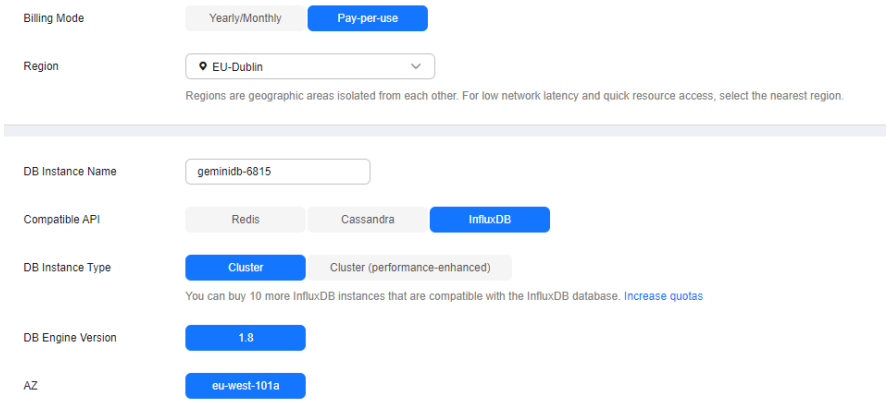

Parameter	Example Value	Description
Billing description	Pay-per-use	Billing mode of an instance • **Yearly/Monthly**: A prepaid billing mode in which you pay for resources before using it. Bills are settled based on the subscription period. The longer the subscription term, the bigger the discount. This mode is a good option for long-term stable services. • **Pay-per-use**: A postpaid billing mode. Pay as you go and just pay for what you use. The DB instance usage is calculated by the second but billed every 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 EU-Dublin.	The region where the tenant is located. It can be changed in the upper left corner. **NOTICE** 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.

Parameter	Example Value	Description
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 (_). If the name contains Chinese characters, the length cannot exceed 64 bytes.
Compatible API	InfluxDB	-
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.
Compatible Version	1.8	• 1.8 • 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.

Figure 2-2 Storage and specifications

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**.
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.
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 2-3 Network and database configurations

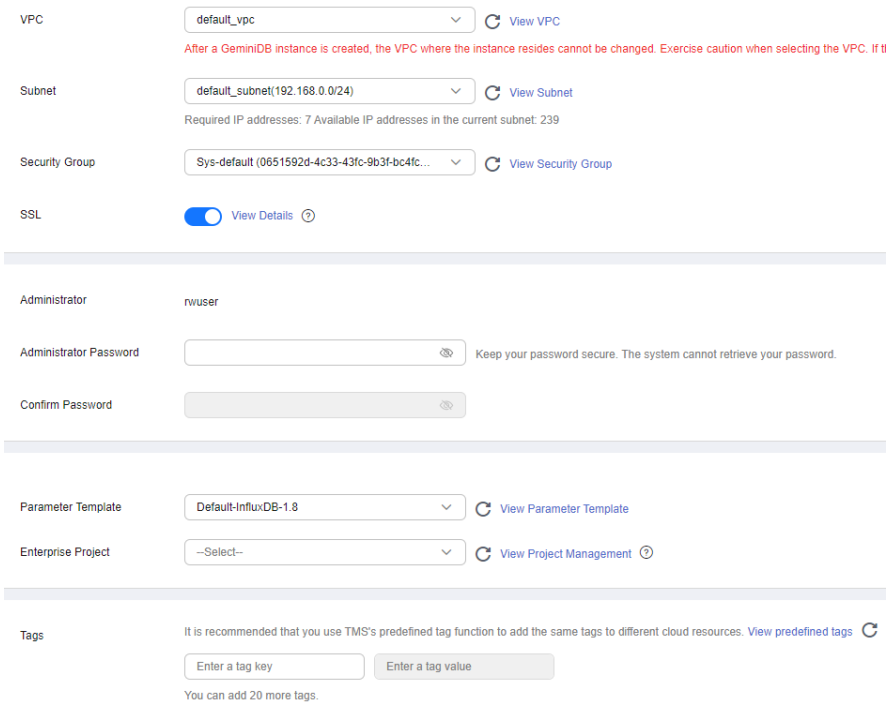

Parameter	Example Value	Description
VPC	default_vpc	Virtual private network where your instances are located. A VPC isolates networks for different services. You can select an existing VPC or create a VPC. **NOTE** • After a GeminiDB Influx instance is created, the VPC where the instance is deployed cannot be changed. • If you want to connect to a GeminiDB Influx instance using an ECS over a private network, the GeminiDB Influx instance and the ECS must be in the same VPC. If they are not, you can 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 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**.
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 creation is complete, the instance status changes to **Available**.

Figure 2-4 Available instance

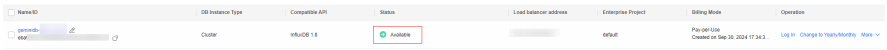

Step 2: Connecting to an Instance Through DAS

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 2-5 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 2-6 Connecting to a GeminiDB Influx Instance

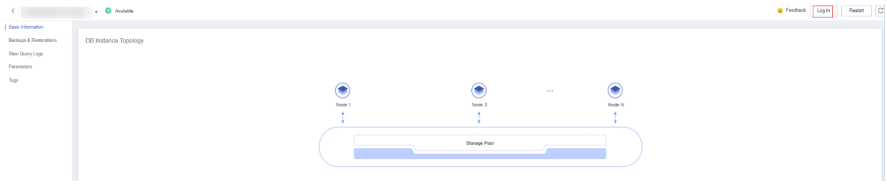

4. Enter a password for logging in to the instance.

   You need to enter the password only when you log in to a GeminiDB Influx instance first time or after you set the password.

Figure 2-7 Logging in to the GeminiDB Influx instance

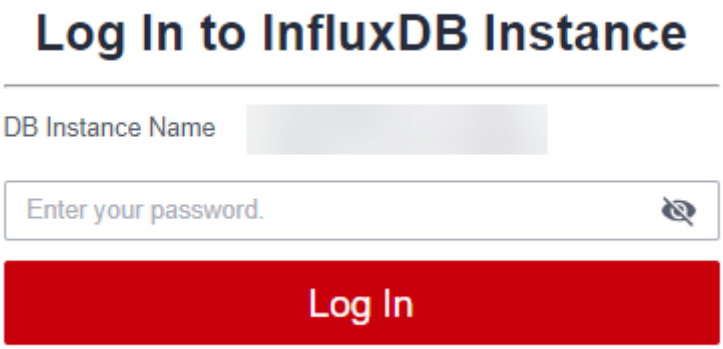

5. Manage relevant databases.

Figure 2-8 Instance homepage

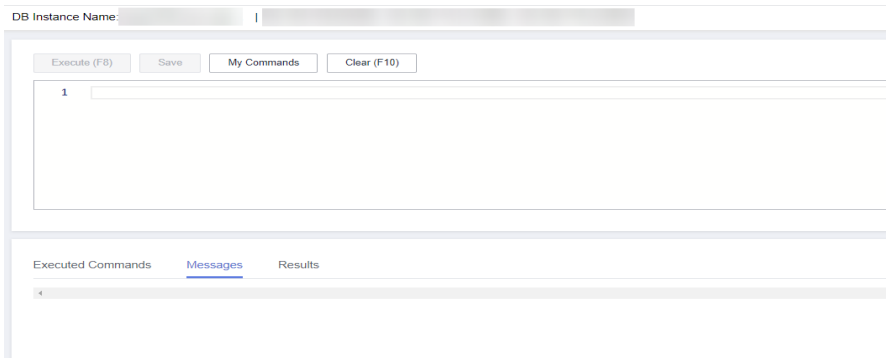

- Save commands to the execution record.

  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 2-9 Executed commands

  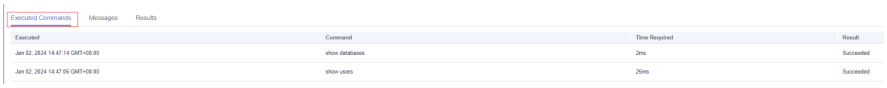

  If this function is disabled, the commands executed subsequently are not displayed. You can click next to **Save Executed SQL Statements** in the upper right corner to disable this function.

- Execute a command.

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

Figure 2-10 Execute a command.

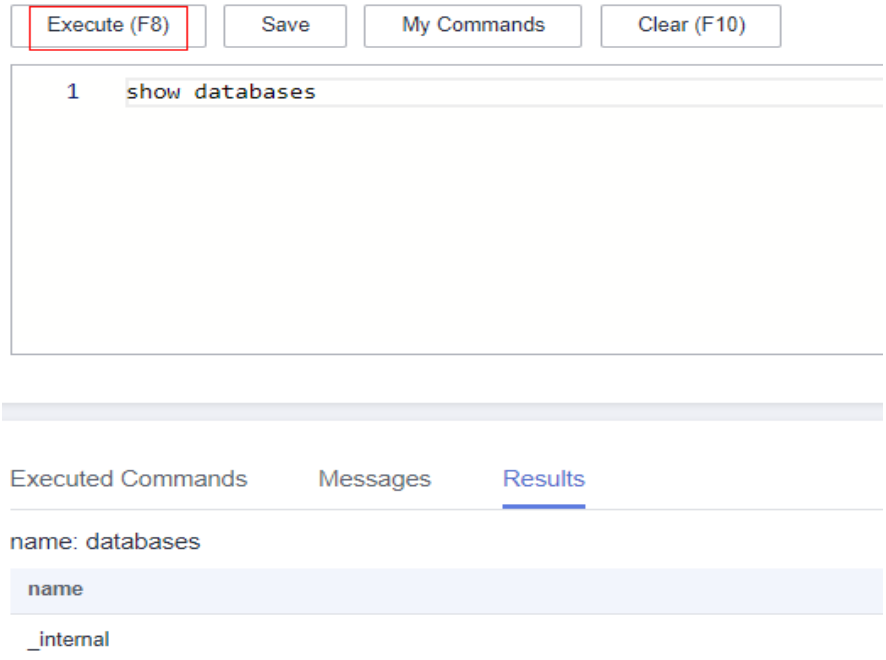

After a command is executed, you can view the execution result on the **Results** page.

- 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 2-11 Save a command.

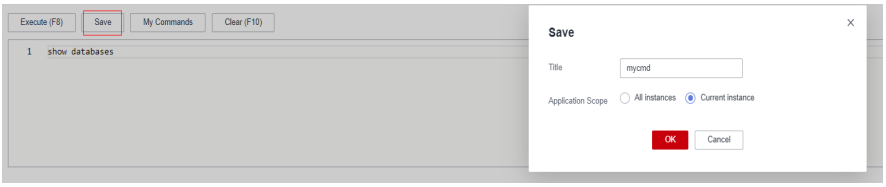

- View my commands.

  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 2-12 Filtering commands

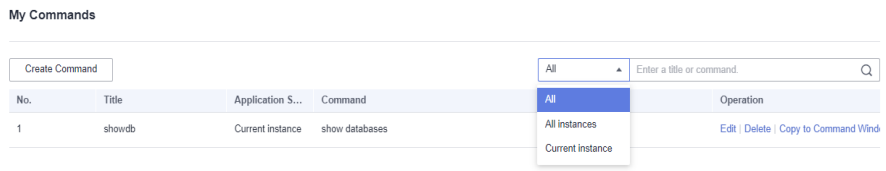

Alternatively, you can enter a command title or statement in the search box to search for the corresponding command.

Figure 2-13 Searching for a command

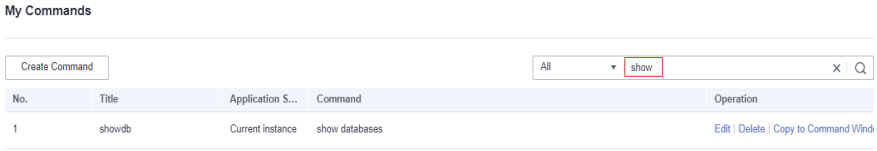

On the **My Commands** page, you can also create, edit, and delete a command or copy it to the command window.

Figure 2-14 Managing a command

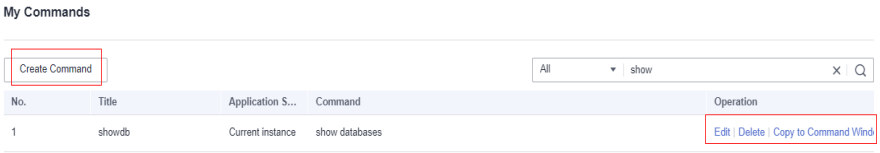

- Clear a command.

  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 Working with GeminiDB Influx API

3.1 Permissions Management

3.1.1 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 3-1**).

Prerequisites

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

Process Flow

Figure 3-1 Process of granting GeminiDB permissions

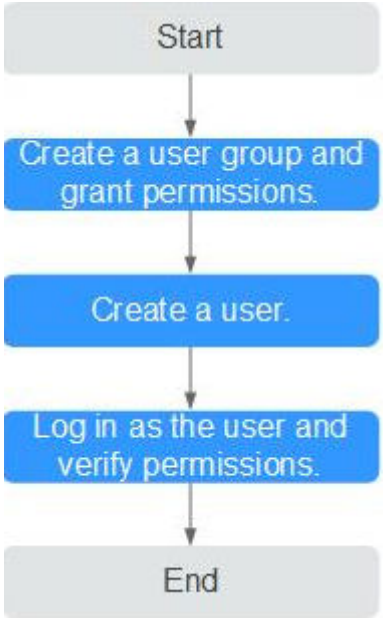

1. **Create a user group and assign permissions** to it.

   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.

3.1.2 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"
            }
      ]
}
```


3.2 Buying a GeminiDB Influx Instance

3.2.1 Buying a GeminiDB Influx Cluster Instance

This section describes how to buy a cluster instance that is compatible with InfluxDB APIs 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 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 instance specifications and click **Next**.

Figure 3-2 Billing mode and basic information

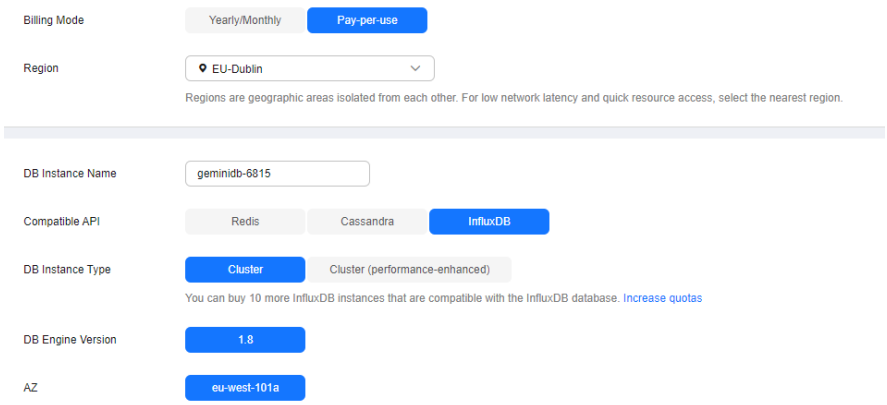

Table 3-1 Billing description

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 Yearly/Monthly 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 **Unsubscribing 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 Pay-per-Use to Yearly/Monthly**.

Table 3-2 Basic information

Parameter	Description
Region	The region where the tenant is located. It can be changed in the upper left corner. **NOTICE** 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 (_). If the name contains Chinese characters, the length cannot exceed 64 bytes. You can change the name of an instance after it is created. For details, see **Changing a GeminiDB Influx Instance Name**.
Compatible API	InfluxDB
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.

Parameter	Description
DB Engine Version	• 1.8 • 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.

Figure 3-3 Specifications and storage

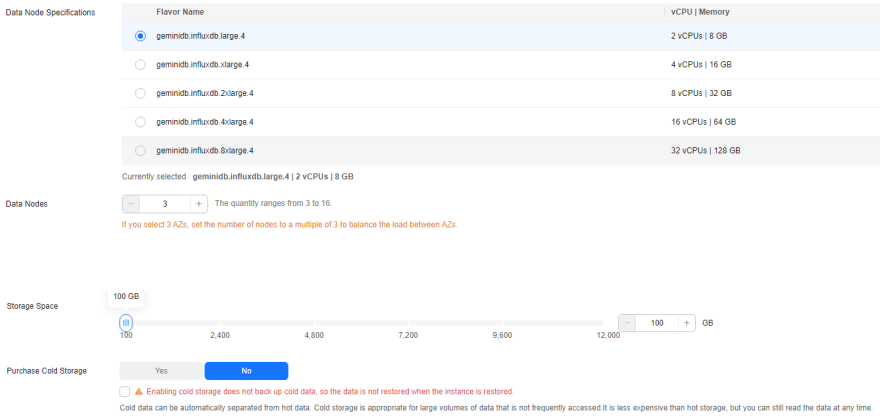

Table 3-3 Specifications and storage

Parameter	Description
Instance 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**.
Nodes	Select the number of 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 create more nodes, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact 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.

Parameter	Description
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 3-4 Network and database configurations

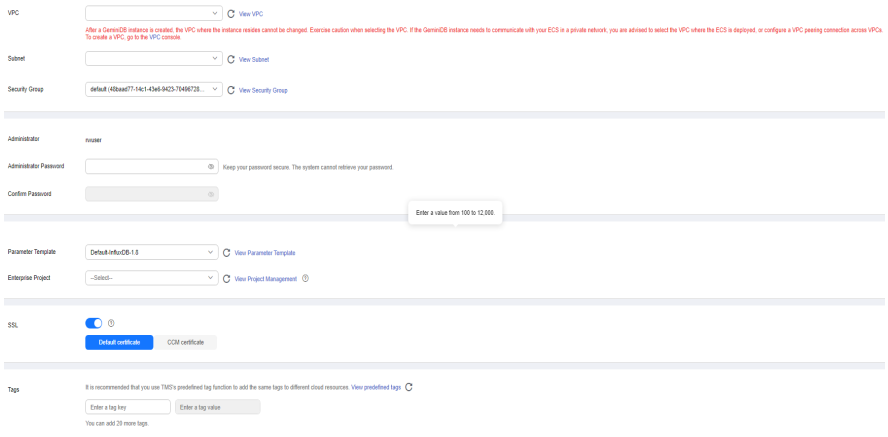

Table 3-4 Network configurations

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. 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. • If you want to connect to a GeminiDB Influx instance using an ECS over a private network, the GeminiDB Influx instance and the ECS must be in the same VPC. If they are not, you can 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 3-5 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	Description
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***.
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. 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 3-6 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. 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 include up to 36 characters, including digits, letters, underscores (_), and hyphens (-). • A tag value is optional if the instance is going to be tagged. The value can be empty. The value can contain up to 43 characters, including digits, letters, underscores (_), periods (.), and hyphens (-). 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 3-7 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-renew	• This option is not selected by default. • If you select this option, the renew cycle is the same as the selected 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 creation is complete, the instance status changes to **Available**.

  You can click  in the upper right corner of the page to refresh the instance status.
- Automated backup 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 3-5 Available instance

----End

3.3 Instance Connection and Management

3.3.1 Connection Methods

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

Figure 3-6 shows the process of connecting to a GeminiDB Influx instance.

Figure 3-6 Connection Methods

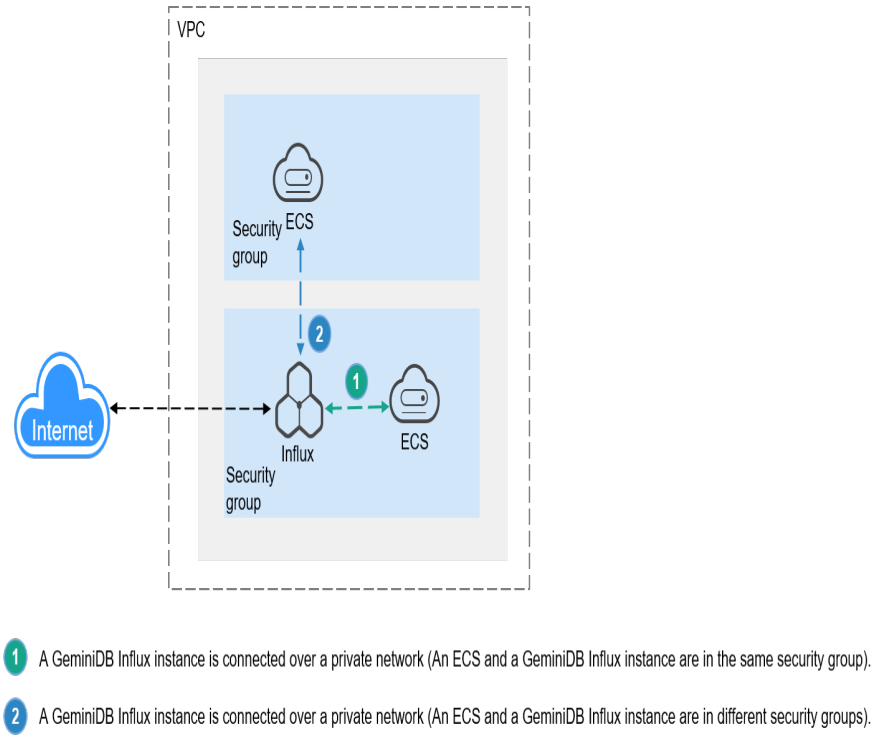

Table 3-8 Connection methods

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.
Private network	Connect to an 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. • .
Program code	Connect to a GeminiDB Influx instance using **Go**, **Java**, or **Python**.	8635	-

3.3.2 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 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

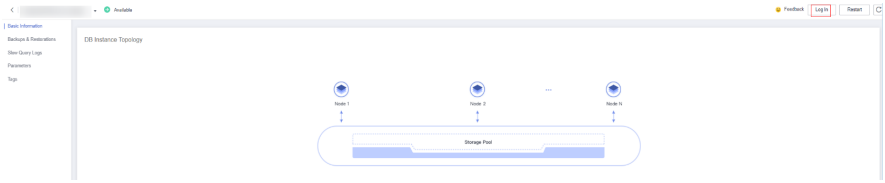

Step 4 Enter the password for logging in to the instance.

You need to enter the password only when you log in to a GeminiDB Influx instance first time or after you set the password.

Figure 3-9 Logging in to the GeminiDB Influx instance

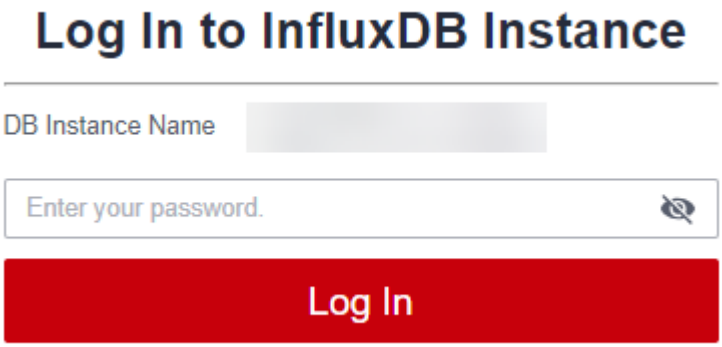

Step 5 Manage relevant databases.

Figure 3-10 Instance homepage

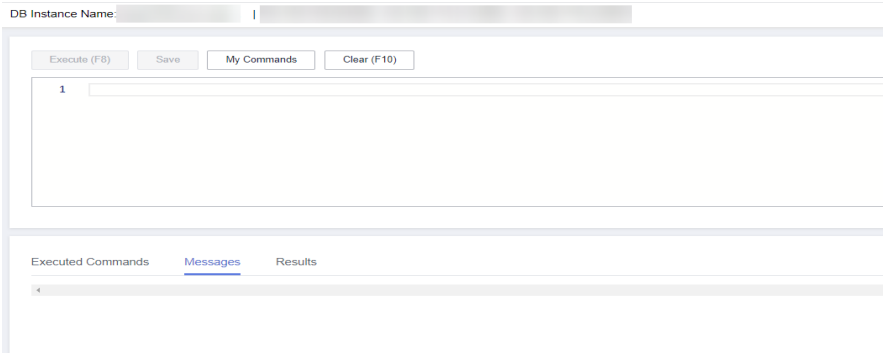

- Save commands to executed commands.

  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 3-11 Executed commands

  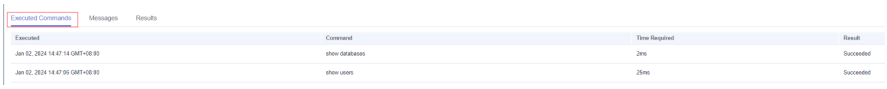

  If this function is disabled, the commands executed subsequently are not displayed any longer. You can click next to **Save Executed SQL Statements** in the upper right corner to disable this function.

- Execute a command.

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

  Figure 3-12 Execute a command.

  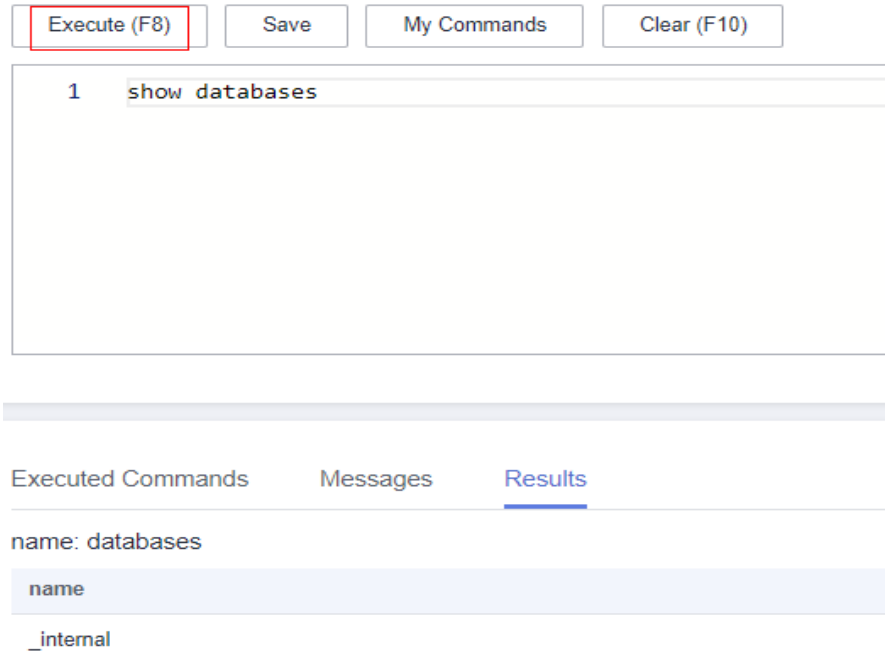

  After a command is executed, you can view the execution result on the **Results** page.

- 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 3-13 Save a command.

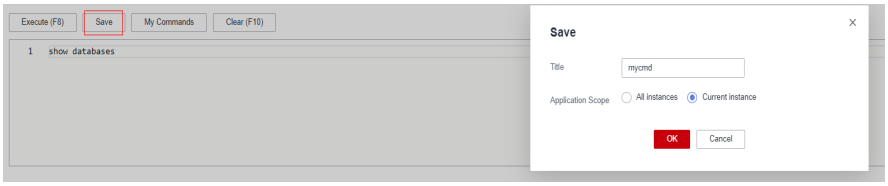

- View my commands.

  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 3-14 Filtering commands

  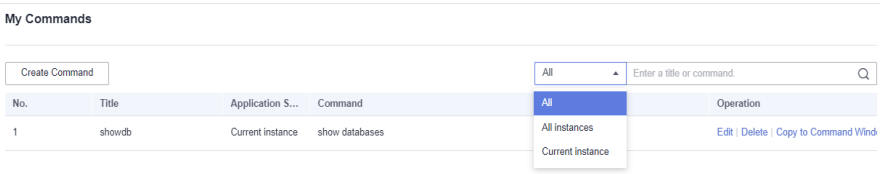

  Alternatively, you can enter a command title or statement in the search box to search for the corresponding command.

  Figure 3-15 Searching for a command

  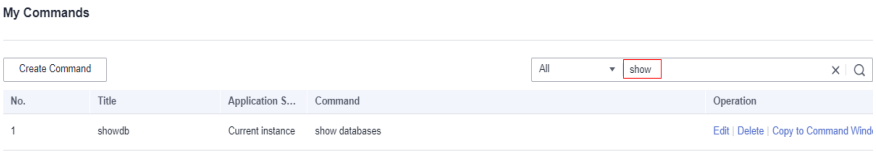

  On the **My Commands** page, you can also create, edit, and delete a command or copy it to the command window.

  Figure 3-16 Managing a command

  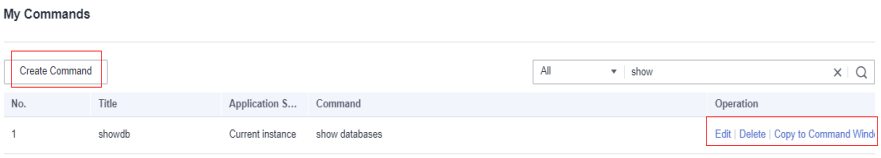

- Clear commands.

  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.

3.3.3 Connecting to a GeminiDB Influx Instance over a Private Network

3.3.3.1 Connecting to an Instance Using a Load Balancer Address (Recommended)

Scenarios

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

Usage Notes

- The DB instances must be in the same VPC and subnet as the ECS.
- The ECS must be allowed by the security group to access DB instances.
  - If the instance is associated with the default security group, you do not need to configure 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.

SSL Connection

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 tool package (the x86 client is used as an example).

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

Step 4 Connect to the DB instance in the directory where the influx tool is located.

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

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

2. Connect to a 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 3-9 Description

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the target DB instance. In the **DB Information** area on the **Basic Information** page, you can find the administrator username.
<DB_PWD>	Administrator password
<DB_HOST>	Load balancer IP 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 customer service. Scenario 1: 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. 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 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>	Port for accessing the instance. You can click the name of the instance to go to the **Basic Information** page. In the **Network Information** area, view the port number.

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

Non-SSL Connection

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 tool package (the x86 client is used as an example).

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

Step 4 Connect to the DB instance in the directory where the influx tool is located.

1. Run the following command to go to the InfluxDB directory:
   ```
   cd influxdb-1.8.10-1/usr/bin
   ```
2. Connect to a 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 3-10 Description

Parameter	Description
<DB_USER>	Username of the administrator account. The default value is **rwuser**. On the **Instances** page, click the target DB instance. In the **DB Information** area on the **Basic Information** page, you can find the administrator username.
<DB_PWD>	Administrator password
<DB_HOST>	Load balancer IP 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 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>	Port for accessing the instance. You can click the name of the instance to go to the **Basic Information** page. In the **Network Information** area, view the port number.

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

3.3.3.2 Connecting to an Instance Using a Private IP Address

This section uses the Linux OS as an example to describe how to connect to a GeminiDB Influx instance over a private network.

Precautions

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

    If the security group rules allow the access from the ECS, the ECS can connect to the instance.

    If the security group rule does not allow the access from the ECS, add an inbound rule to the security group.

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.

SSL Connection

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 tool package (the x86 client is used as an example).

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

Step 4 Connect to the DB instance in the directory where the influx tool is located.

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

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

2. Connect to a GeminiDB Influx instance.
   - Use the default certificate for connection.

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

     Example:

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


Table 3-11 Description

Parameter	Description
<DB_HOST>	Specifies the private IP address of the node to be connected. To obtain this IP address, go to the **Instances** page, locate the instance whose node IP addresses you want to view, and click its name. The IP address can be found in the **Private IP Address** column at the **Node Information** area. If the instance you purchased has multiple nodes, select the private IP address of any node.
<DB_PORT>	The port of the DB instance to be connected. The default value is 8635 and cannot be changed. Click the target instance to go to the **Basic Information** page. In the **Network Information** area, you can find the database port.

3. Run the following command for authentication:

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

   Table 3-12 Description

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

Step 5 After the identity verification is successful, run the following command:

show databases

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

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

----End

Non-SSL Connection

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 tool package (the x86 client is used as an example).

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

Step 4 Connect to the DB instance in the directory where the influx tool is located.

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

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

2. Connect to a GeminiDB Influx instance.

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

   Example:

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

   Table 3-13 Description

Parameter	Description
<DB_HOST>	Specifies the private IP address of the node to be connected. To obtain this IP address, go to the **Instances** page, locate the instance whose node IP addresses you want to view, and click its name. The IP address can be found in the **Private IP Address** column at the **Node Information** area. If the instance you purchased has multiple nodes, select the private IP address of any node.
<DB_PORT>	The port of the DB instance to be connected. The default value is 8635 and cannot be changed. Click the target instance to go to the **Basic Information** page. In the **Network Information** area, you can find the database port.

3. Run the following command for authentication:

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

Table 3-14 Description

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

Step 5 After the identity verification is successful, run the following command:

show databases

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

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

----End

3.3.4 Connecting to a GeminiDB Influx Instance 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** 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 tool package (the x86 client is used as an example).

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

Step 4 Connect to the DB instance in the directory where the influx tool is located.

1. Run the following command to go to the InfluxDB directory:
   ```
   cd influxdb-1.8.10-1/usr/bin
   ```
2. Connect to a 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 3-15 Description

Parameter	Description
<DB_HOST>	EIP of the node to be connected To obtain this IP address, go to the **Instances** page and click the target DB instance name. The IP address can be found in the **EIP** field under **Node Information** on the **Basic Information** page. 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**.
<DB_PORT>	The port of the instance to be connected. The default value is 8635 and cannot be changed. Click the instance to go to the **Basic Information** page. In the **Network Information** area, you can find the database port.

3. Run the following command for authentication:

   auth

   Enter the username and password as prompted.

   username:*<DB_USER>*

   password:*<DB_PWD>*

Table 3-16 Description

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

Parameter	Description
<DB_PWD>	Administrator password

Step 5 After the identity verification is successful, run the following command:

show database

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

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

----End

3.3.5 Connecting to a GeminiDB Influx Instance Using Programming Languages

3.3.5.1 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)
    }
}
```

3.3.5.2 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);
                });
```

3.3.5.3 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.

3.3.6 Connection Information Management

3.3.6.1 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 3-17**.

Table 3-17 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.

Scenario	Description
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 3-17 Security group

Step 5 Add an inbound rule.

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

Figure 3-18 Inbound rules

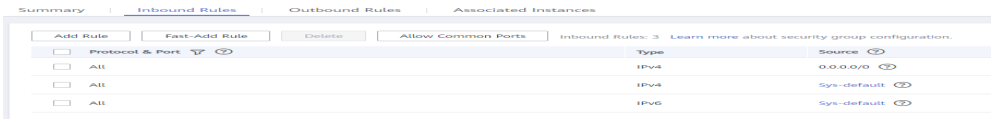

2. Click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Figure 3-19 Adding a rule

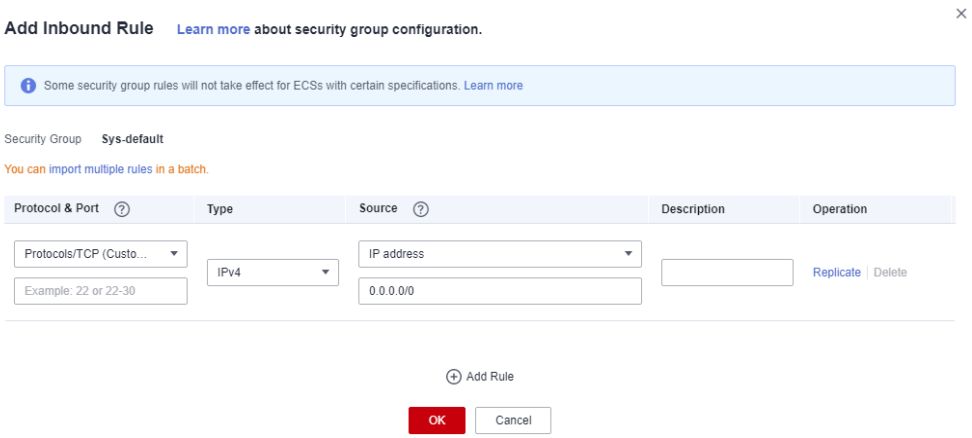

3. In the displayed dialog box, set required parameters.

Table 3-18 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
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

3.3.6.2 Binding an EIP

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 it to allow external access. If later you want to prohibit external access, you can also unbind the EIP from the DB instance.

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 3-20 Binding an EIP

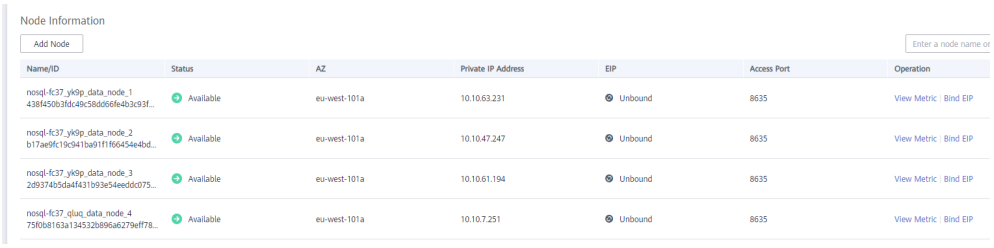

Step 5 In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **Yes**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.

Figure 3-21 Selecting an EIP

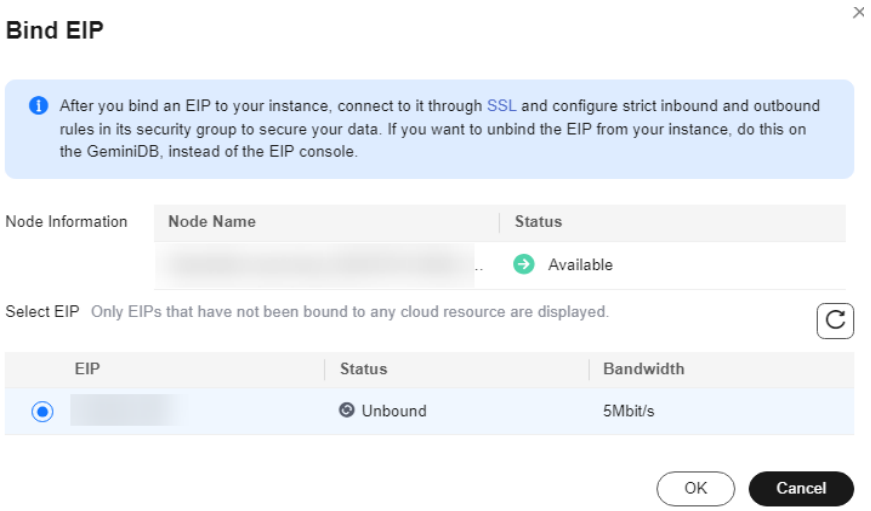

Step 6 In the **EIP** column, view the EIP that is successfully bound.

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 **Instance Management** page, click the instance that you want to unbind an EIP from.

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 3-22 Unbinding an EIP

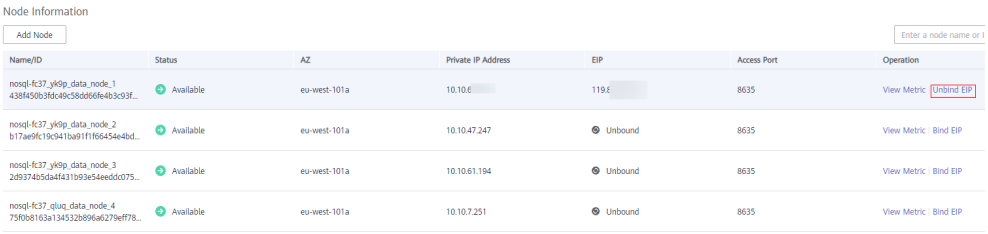

Step 5 In the displayed dialog box, click **Yes**.

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

----End

3.3.6.3 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 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

3.3.6.4 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 3-23 Enabling SSL

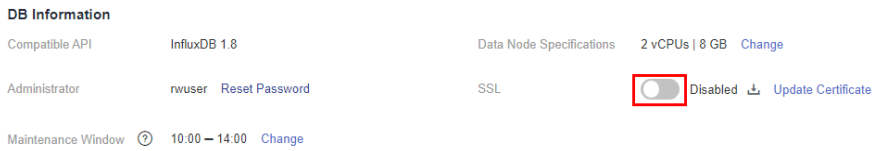

----End

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 3-24 Disabling SSL

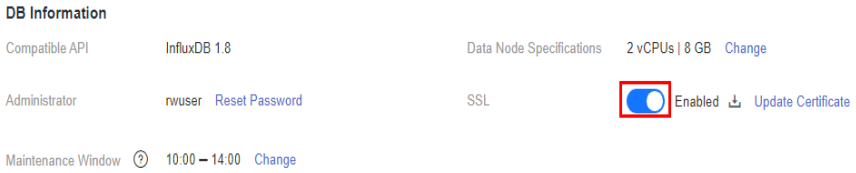

----End

3.4 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.

3.5 Instance Lifecycle Management

3.5.1 Restarting a GeminiDB Influx Instance

Scenarios

You may need to occasionally restart a DB instance to perform 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**.

----End

3.5.2 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 **Unsubscribing 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

3.5.3 Recycling an Instance

Unsubscribed yearly/monthly instances and deleted pay-per-use instances can be moved to the recycle bin, you can restore 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.

Modifying the Recycling Policy

NOTICE

You can modify the retention period, and the changes only apply to the instances deleted after the modification. 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 3-25 Modifying a recycling policy

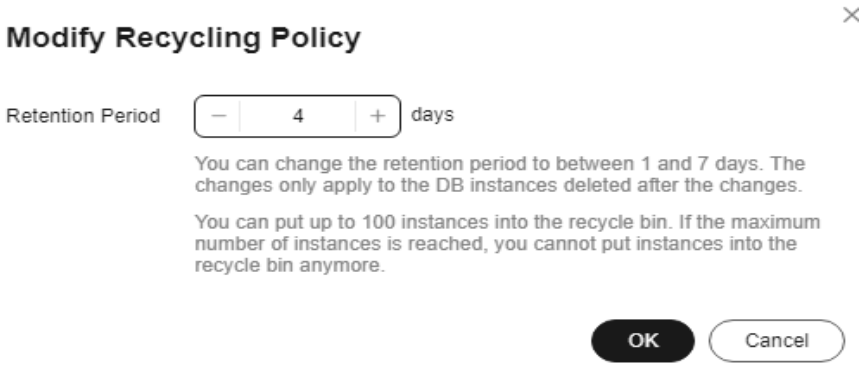

----End

Rebuilding an Instance

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

Step 1 Log in to the Huawei Cloud console.

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

Step 3 On the **Recycling Bin** page, locate the instance that you want to rebuild and click **Rebuild** in the **Operation** column.

Figure 3-26 Rebuilding an instance

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

----End

3.6 Instance Modifications

3.6.1 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

3.6.2 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

3.6.3 Changing CPUs and Memory of an Instance

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** in the specifications field.

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

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.

Go to the **Basic Information** page and in the **DB Information** area, and you can see the new instance 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 **Change Specifications** in the **Operation** column.

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

Step 5 View the change results.

Go to the **Basic Information** page and in the **DB Information** area, and you can see the new instance specifications.

----End

3.6.4 Adding Instance Nodes

Scenarios

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

Usage Notes

- Adding nodes may lead to the decrease of 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, a maximum of 12 nodes are supported. To obtain more nodes, choose **Service Tickets > Create Service Ticket** in the upper right corner of the console and contact customer service.

Method 1

Step 1 In the service list, choose **Databases** > **GeminiDB**.

Step 2 On the **Instances** page, locate the instance that you want to add nodes to and click its name.

Step 3 In the **Node Information** area on the **Basic Information** page, click **Add Node**.

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

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

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 3-27 Adding a node

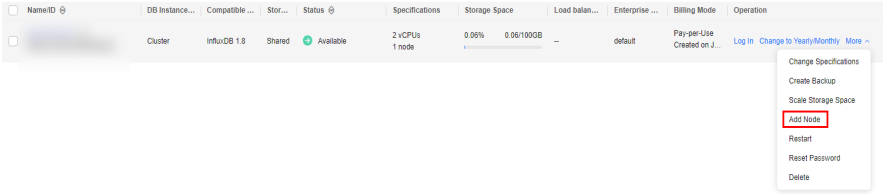

Step 4 Specify **Add Nodes** and click **Next**.

Figure 3-28 Adding a node

Step 5 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

3.6.5 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

- Storage space 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 3-19 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**.

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 3-29 Scaling up storage space

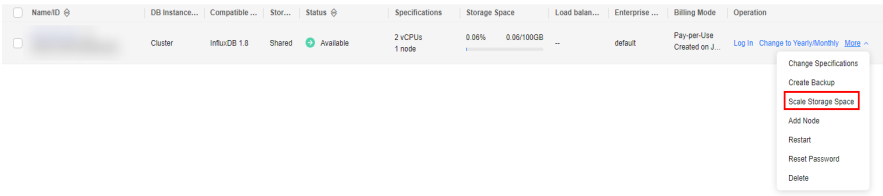

Step 4 On the displayed page, specify new storage and click **Next**.

Figure 3-30 Scaling up storage space

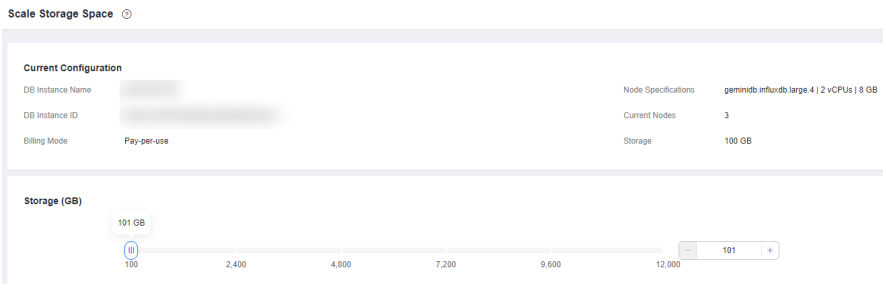

Select at least 1 GB each time, and the value must be an integer.

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

3.7 Database Commands

3.7.1 Supported Commands

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

User Management

Table 3-20 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 3-21 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	√	√

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
precision	√	√
consistency	√	√
history	√	√
settings	√	√
clear	√	√
exit/quit/ctrl+d	√	√

Metadata Management

Table 3-22 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	√	√
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	√	√

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
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	√	√
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 3-23 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 3-24 Commands supported by data points

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
select	√	√

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

Aggregate Functions

Table 3-25 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 3-26 Commands supported by the SELECT function

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
bottom	√	√

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

Conversion Function

Table 3-27 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	√	√

Command	Supported In Read/ Write Mode	Supported In Read-Only Mode
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.

3.8 Cold and Hot Data Separation

3.8.1 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 existing instance is 1.7.4.6 or later. If the kernel version is earlier than 1.7.4.6, choose **Service Tickets > Create Service Ticket** and contact customer service.
- GeminiDB Influx does not back up cold storage data.
- Cold storage cannot be disabled after being enabled.
- For 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 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 3-31 Creating cold storage

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

Figure 3-32 Specifying cold storage

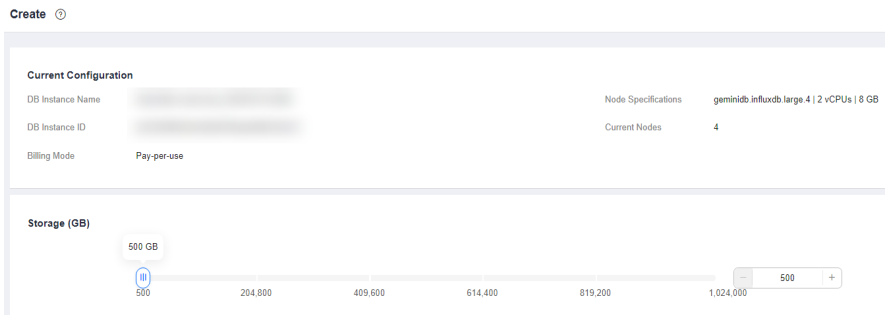

The cold storage is an integer from 500 GB to 1,024,000 GB. You can add a minimum of 1 GB each time you scale up storage space.

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

3.8.2 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. GeminiDB Influx provides cold and hot data separation and uses low-cost media to store cold data. It can help you greatly reduce storage costs in 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 3-33 Diagram

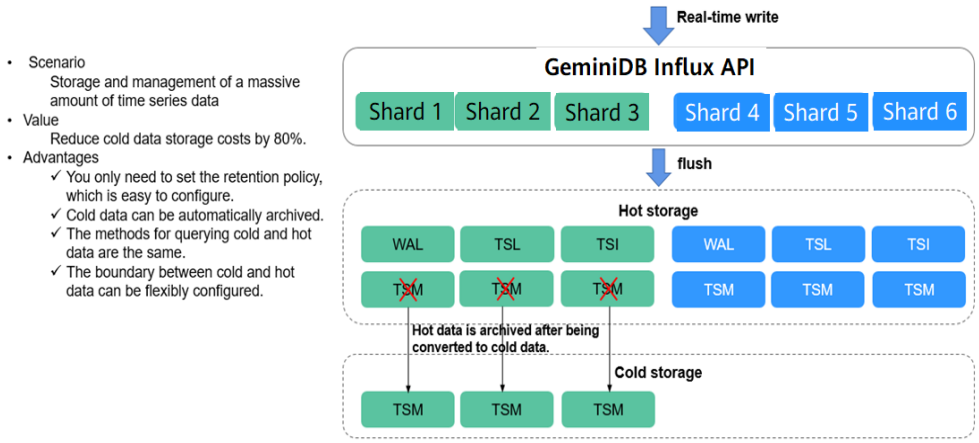

Basic Usage

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.

3.8.3 Scaling Up Cold Storage

Scenarios

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

Usage Notes

- Cold storage scaling does not interrupt your services. After the scaling is complete, you do not need to restart your instance.
- Cold storage can only be scaled up and cannot be scaled down.

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 3-34 Scaling up cold storage of cluster and single-node instances

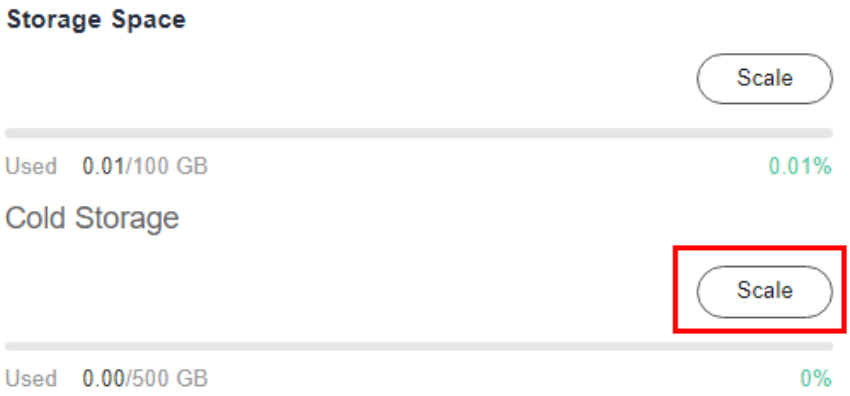

Step 5 On the displayed page, specify desired cold storage space and click **Next**.

Figure 3-35 Scaling up cold storage of cluster and single-node instances

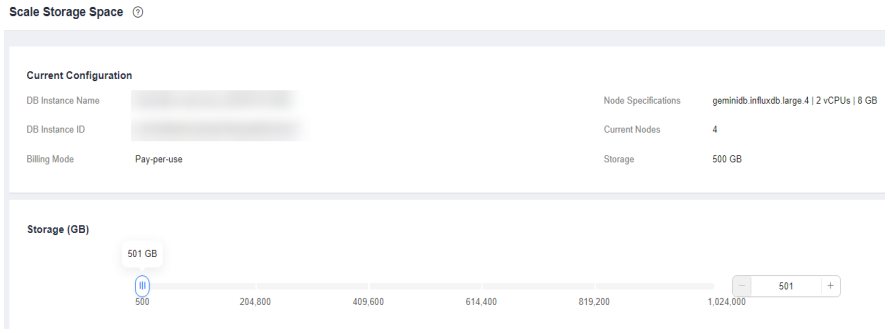

Select at least 1 GB each time, and the value must be an integer.

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

3.9 Data Backup

3.9.1 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.

Backup Methods

GeminiDB Influx instances support both automatic and manual backups.

- Automated backup

  You can click **Modify 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 configure in the backup policy and will store the data for a length of time you specify.

  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 3-28 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 3-36**, 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 3-36 Backup process

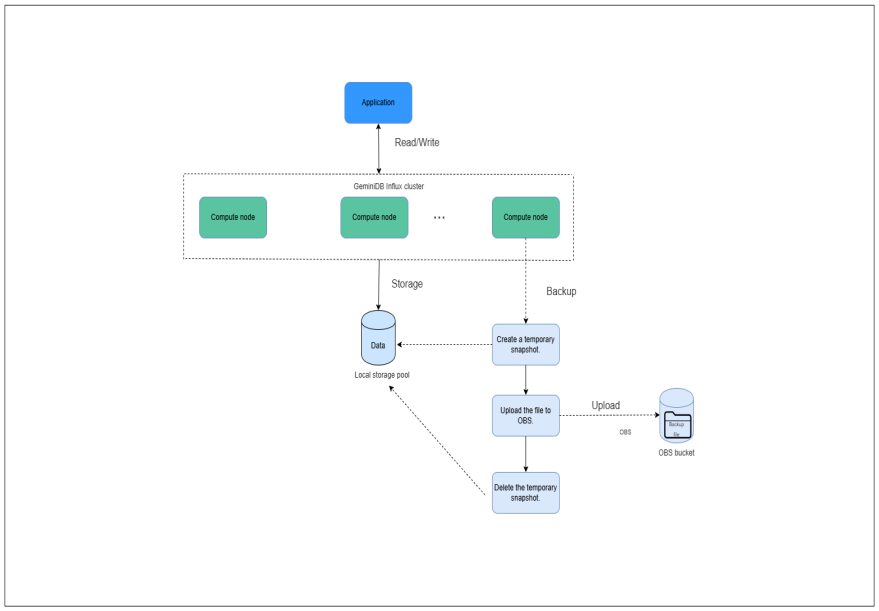

Backup Storage

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.

3.9.2 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 3-37 Enabling automated backup

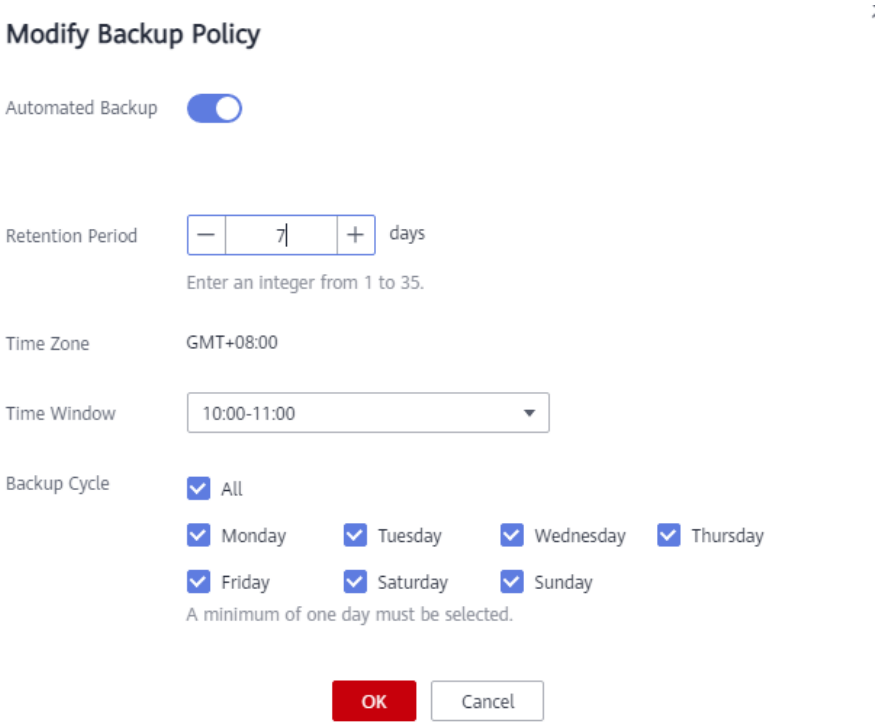

- **Retention Period**: Automated backup files are saved for seven days by default. The retention period ranges from 1 to 3660 days. Full backups are retained till the retention period expires. However, even if the retention period has expired, the most recent backup will be retained.
  - Extending the 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 less than seven days, the system automatically backs up data daily.
  - The system checks existing automated backups and deletes any backups that exceed the backup retention period you configured.
  - **Time Window**: A one-hour period the backup will be scheduled for, such as 10:00-11:00. 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.
- A full backup generated on Tuesday will be automatically deleted on the following Wednesday. The reasons are as follows:

  The backup generated on Tuesday will expire on Thursday, but as it is the last backup, so it will be retained until a new backup expires. The next backup will be generated on the following Monday and will expire on the following Wednesday. So the full backup generated on Tuesday will not be automatically deleted until the following Wednesday.

- **Backup Cycle**: All options are selected by default.
  - **All**: Each day of the week is selected. The system automatically backs up data every day.
  - You can select one or more days in a week. The system automatically backs up data at the specified time.

  NOTE

  A full backup starts within one hour of the time you specify. The amount of time required for the backup depends on the amount of data to be backed up. The more data has to be backed up, the longer it will take.
- After the DB instance is created, you can modify the automated backup policy as needed. You can change the time window after the DB instance is created. The system backs up data based on an automated backup policy you configure.
- If **Automated Backup** is disabled, any automated backups in progress stop immediately.

Modifying an Automated Backup Policy

Step 1 Log in to the Huawei Cloud console.

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

Step 3 On the **Instances** page, click the instance you want to back up.

Step 4 Choose **Backups & Restorations** in the navigation pane one the left, and click **Modify Backup Policy**. In the displayed dialog box, configure the backup policy. Click **OK**.

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

Figure 3-38 Modifying the backup policy

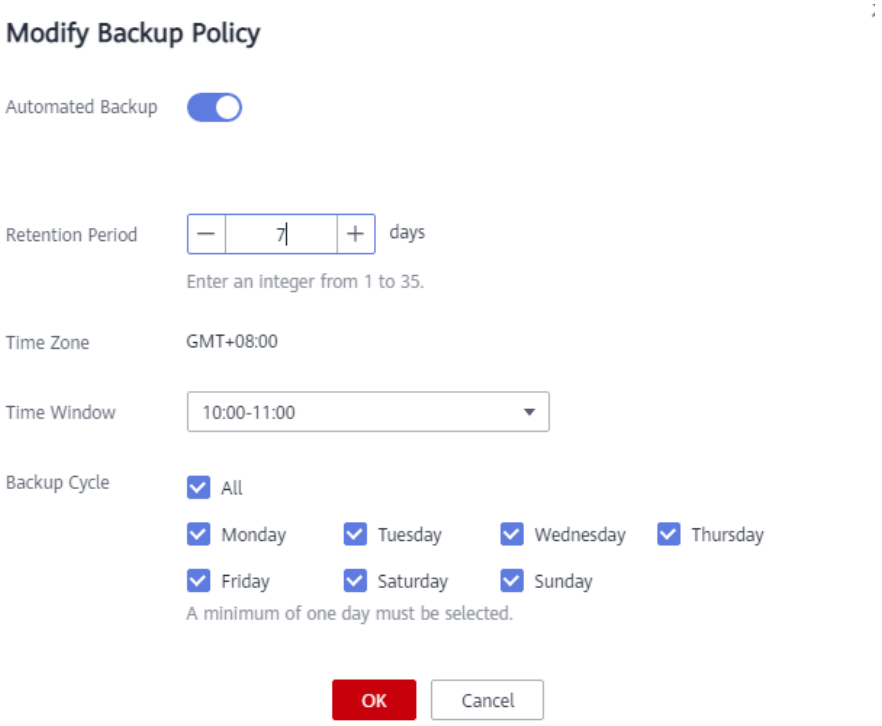

Step 5 Check or manage the generated backups on the **Backups** or **Backups & Restorations** page.

----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 3-39 Disabling automated backup

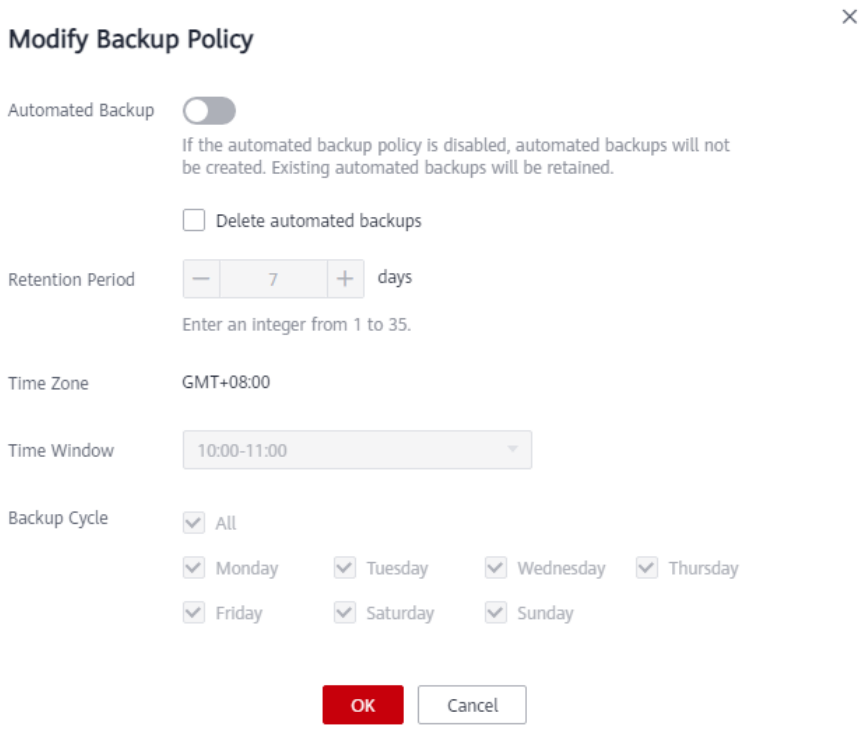

When your disable automated backup, specify whether to delete the automated backups:

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

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

----End

Deleting an Automated Backup

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

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

NOTICE

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

- **Method 1**
  a. Log in to the Huawei Cloud console.
  b. In the service list, choose **Databases** > **GeminiDB**.
  c. On the **Instances** page, click the instance 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**.

3.9.3 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

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 3-40 Creating a manual backup

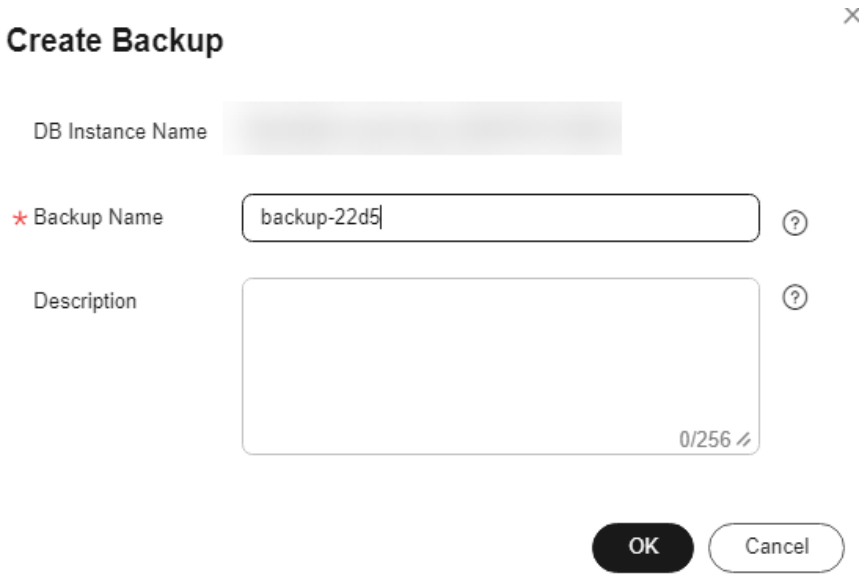

Table 3-29 Parameter description

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

Deleted backups are not displayed in the backup list.

NOTICE

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

Method 1

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

Method 2

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

3.10 Data Restoration

3.10.1 Restoration Methods

GeminiDB Influx supports multiple forms of data restoration. You can select one based on service requirements.

Table 3-30 Restoration methods

Method	Scenario
Restoring Data to a New Instance	You can restore an existing backup file to a new instance.

3.10.2 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.

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 3-41 Restoration

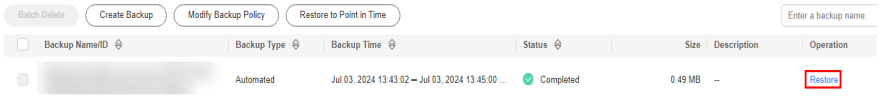

Method 2

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

Figure 3-42 Restoration

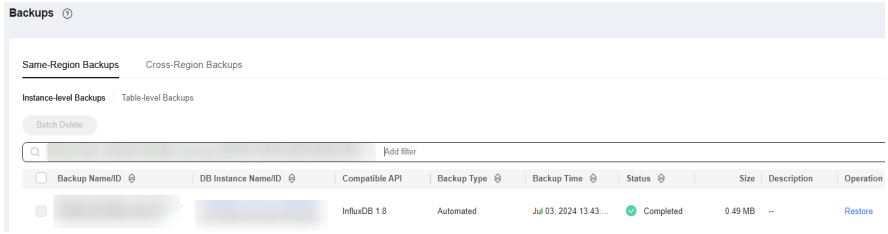

Step 4 In the displayed dialog box, confirm the current instance details and restoration method and click **OK**.

Figure 3-43 Restoring data to a new instance

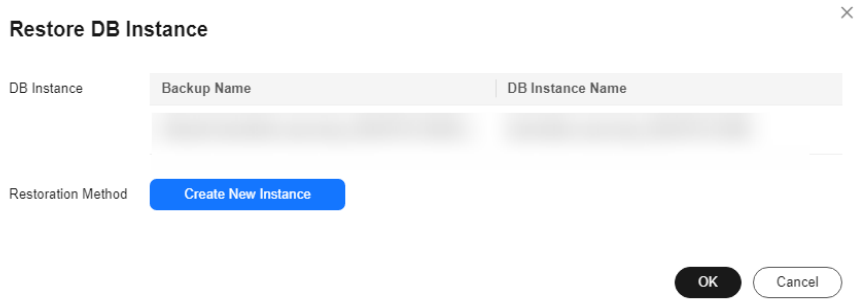

- The default API type and DB engine version are the same as those of the original instance and cannot be changed.
- GeminiDB automatically calculates the minimum storage space required for restoration based on the size of the selected backup file. The storage capacity depends on the instance specifications, and must be an integer.
- You need to set a new administrator password.
- To modify other parameters, see .

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

3.11 Parameter Management

3.11.1 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 3-32**.

Figure 3-44 Modifying parameters in the parameter template

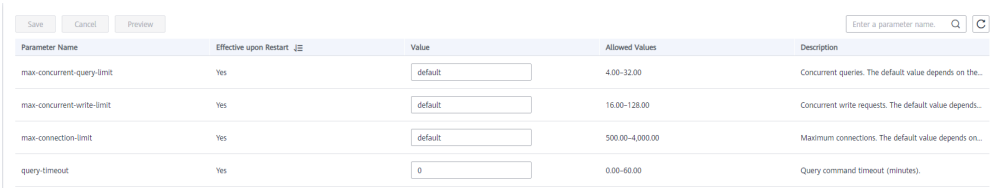

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Table 3-31 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 3-32 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 3-45 Preview Change

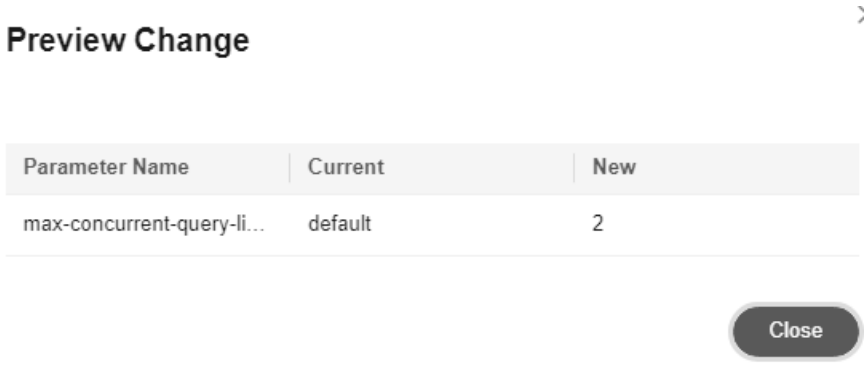

Step 6 After parameters are modified, click **Change History** to view parameter modification details.

For details about how to view parameter modification details, see **Viewing Parameter Change History**.

NOTICE

- The modifications take effect only after you apply the parameter template to instances. For details, see **Applying a Parameter Template**.
- The change history page displays only the modifications of 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 3-46 Modifying parameters of the instance

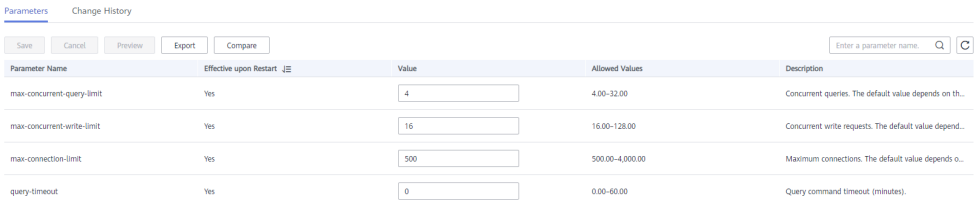

- To save the modifications, click **Save**.
- To cancel the modifications, click **Cancel**.
- To preview the modifications, click **Preview**.

Step 5 After parameters are modified, click **Change History**.

For details about how to view parameter modification details, see **Viewing Parameter Change History**.

NOTICE

After you modify instance parameters, the modifications immediately take effect for the 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 for the modifications to take effect.
- If the value is **No**, the modifications take effect immediately.

----End

3.11.2 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 3-47 Creating a parameter template

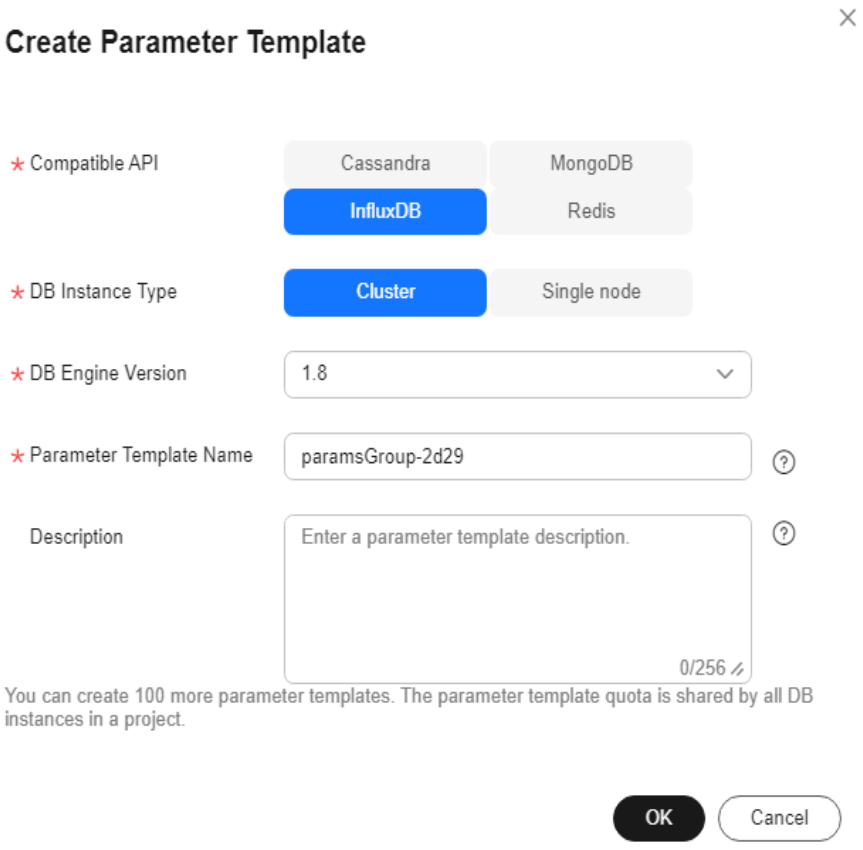

Figure 3-48 Creating a parameter template

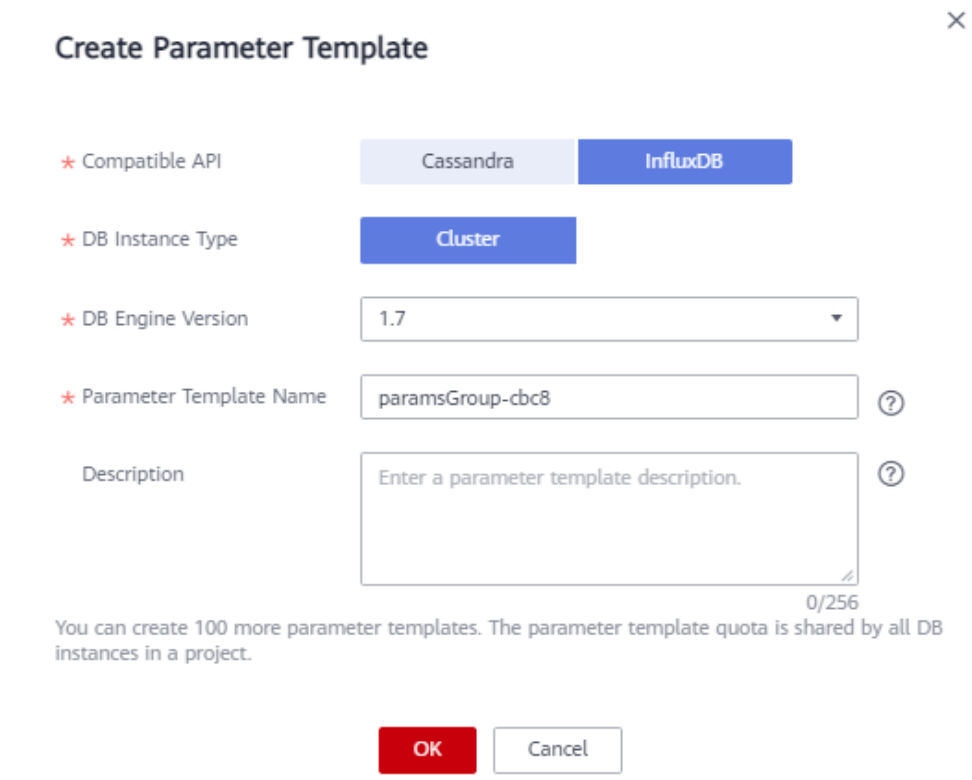

- **Compatible API**: Select the API type that is compatible with your DB engine parameter template.
- **DB Engine Version**: Select a DB engine version, for example, 1.7.
- **Parameter Template Name**: The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase 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

3.11.3 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 3-49 Viewing change history of a customer parameter template

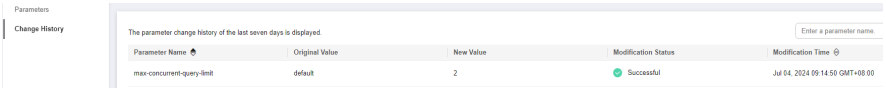

You can apply the parameter template to instances by referring to **Applying a Parameter Template**.

----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 3-50 Viewing parameter change history of an instance

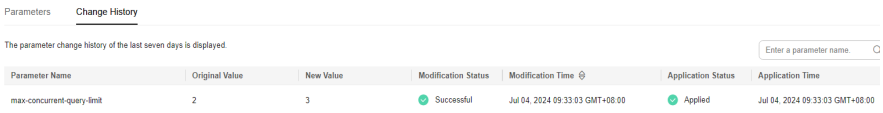

----End

3.11.4 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 3-51 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

3.11.5 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 3-52 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 3-53 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

3.11.6 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 3-54 Replicating a parameter template

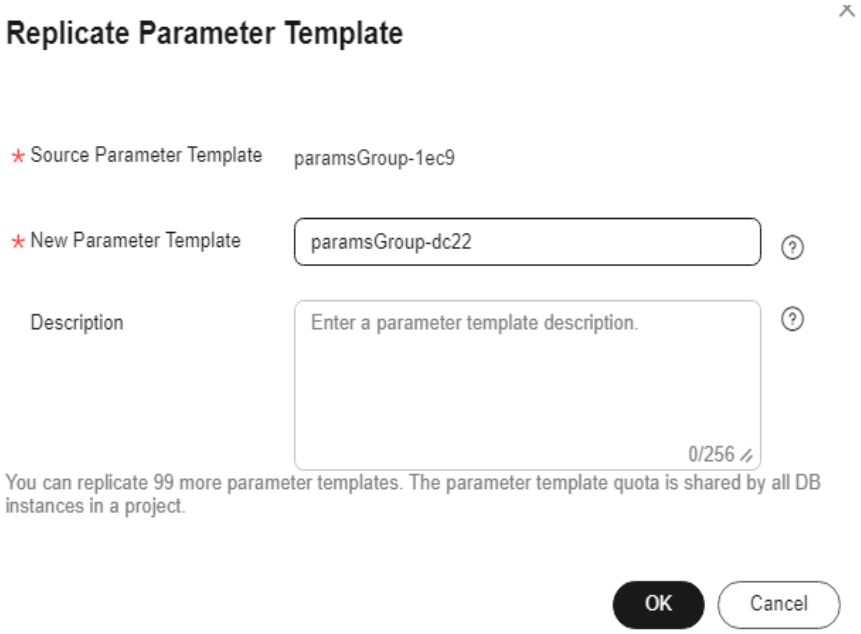

- **New Parameter Template**: The template name can be up to 64 characters long. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (_), and periods (.).
- **Description**: The description contains a maximum of 256 characters and cannot include line breaks or the following special characters >!<"&'=

After the parameter template is replicated, a new template is generated in the list on the **Parameter Templates** page.

----End

3.11.7 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

3.11.8 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

3.11.9 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

3.11.10 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

3.11.11 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

3.12 CTS

3.12.1 Key Operations Supported by CTS

With CTS, you can record operations on GeminiDB Influx instances for later queries, audit, and backtracking.

Table 3-33 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
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

Operation	Resource Type	Trace Name
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
Deleting the node that fails to be added	instance	NoSQLDeleteEnlargeFailNode
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

3.12.2 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

3.13 Viewing Metrics and Configuring Alarms

3.13.1 GeminiDB Influx 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 3-34 GeminiDB Influx metrics

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
gemini001_cpu_usage	CPU Usage	CPU usage of the monitored system Unit: Percent	0–100	GeminiDB Influx instance node	1 minute
gemini002_mem_usage	Memory Usage	Memory usage of the monitored system Unit: Percent	0–100	GeminiDB Influx instance node	1 minute
gemini003_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second Unit: kbit/s	≥ 0	GeminiDB Influx instance nodes	1 minute
gemini004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second Unit: kbit/s	≥ 0	GeminiDB Influx instance nodes	1 minute
nosql005_disk_usage	Storage Space Usage	Storage space usage of the monitored object. Unit: Percent	0–100	GeminiDB Influx instances	1 minute
nosql006_disk_total_size	Total Storage Space	Total storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Influx instances	1 minute
nosql007_disk_used_size	Used Storage Space	Used storage space of the monitored object. Unit: GB	≥ 0	GeminiDB Influx instances	1 minute

Metric ID	Metric Name	Description	Value Range	Monitored Object	Monitoring Period (Raw Data)
influxdb001_series_num	Time Series	Total number of time series Unit: count	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb002_query_req_ps	Query Requests Per Second	Number of query requests per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb003_write_req_ps	Write Requests Per Second	Number of write requests per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb004_write_points_ps	Write Points	Number of write points per second Unit: count/s	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb005_write_concurrency	Concurrent Write Requests	Number of concurrent write requests Unit: count	≥ 0	GeminiDB Influx instance nodes	1 minute
influxdb006_query_concurrency	Concurrent Queries	Number of concurrent query requests Unit: count	≥ 0	GeminiDB Influx instance nodes	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

3.13.2 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 & Deployment**, 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 3-55 Creating an alarm rule

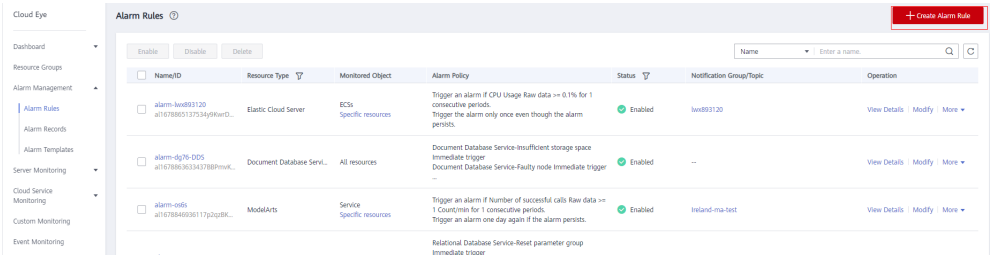

Step 5 Set alarm parameters.

1. Configure basic alarm information.

Figure 3-56 Configuring basic information for an alarm rule

Table 3-35 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

Parameter	Description	Example Value
Description	(Optional) Alarm rule description.	-

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

Table 3-36 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 **Resource groups** and any resource in the group meets the alarm policy, an alarm notification will be sent. – After you select **Specific resources**, select one or more resources and click » to add them to the box on the right.	Specified Resources
Group	This parameter is mandatory when **Monitoring Scope** is set to **Resource groups**.	-

3. Configure an alarm policy.

Figure 3-57 Configuring the alarm policy

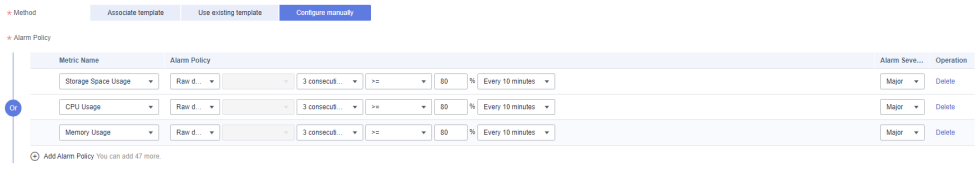

Table 3-37 Parameter description

Parameter	Description	Example Value
Method	Select **Associate template**, **Use existing template**, or **Configure manually**. **NOTE** If you set **Monitoring Scope** to **Specific resources**, you can set **Method** to **Use existing template**.	Configure manually
Template	Select the template to be used. This parameter is available only when you 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 **GeminiDB Influx 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 3-57** 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 3-58 Configuring alarm notification information

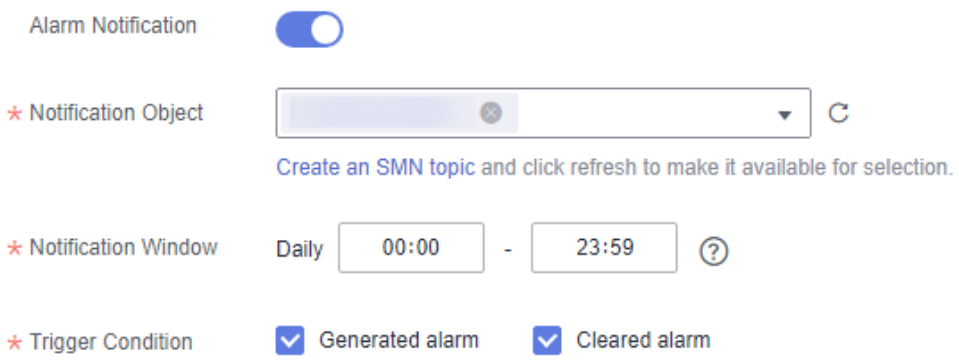

Table 3-38 Parameter description

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 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.	-

Parameter	Description	Example Value
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.	-
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 3-59 Advanced settings

Table 3-39 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.	default

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.

----End

3.13.3 Viewing Metrics

Scenarios

Cloud Eye monitors the 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 3-60 Viewing metrics

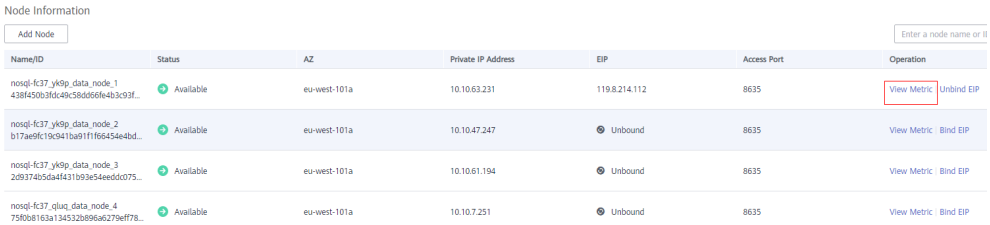

Step 5 In the monitoring area, you can select a duration to view the monitoring data.

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

3.14 Billing Management

3.14.1 Renewing Instances

This section describes how to renew your yearly/monthly GeminiDB Influx instances.

Precautions

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

Renewing a Single Yearly/Monthly Instance

Step 1 Log in to the Huawei Cloud console.

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

Step 3 On the **Instances** page, locate the target instance and click **Renew** in the **Operation** column.

Figure 3-61 Renewal button

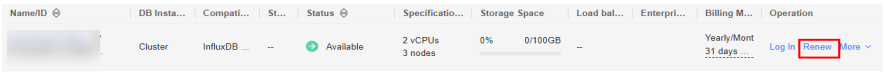

Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** next to the **Billing Mode** field.

Figure 3-62 Renewal button

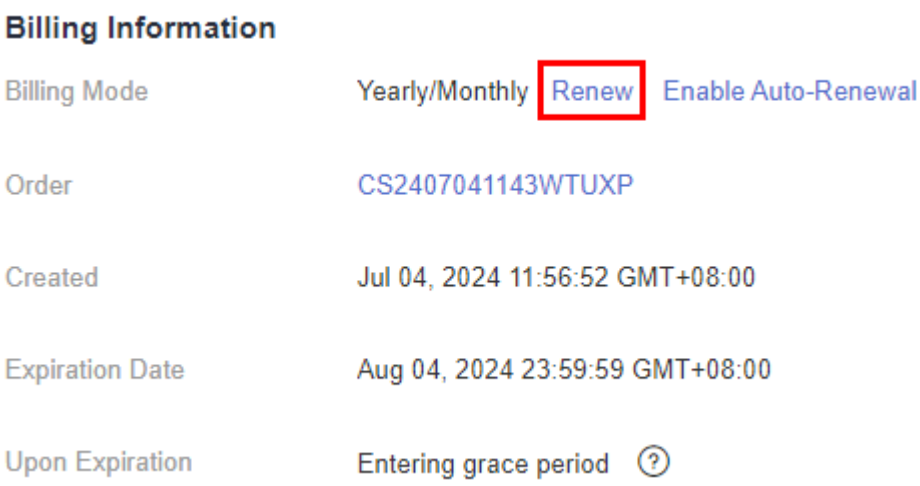

Step 4 On the displayed page, renew the instance.

----End

Renewing Multiple Yearly/Monthly Instances in Batches

Step 1 Log in to the Huawei Cloud console.

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

Step 3 On the **Instances** page, select the instances that you want to renew and click **Renew** above the instance list.

Figure 3-63 Batch renewing

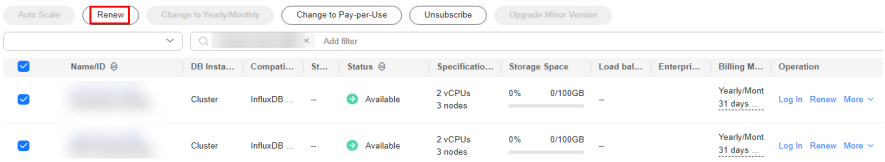

Step 4 In the displayed dialog box, click **Yes**.

----End

3.14.2 Changing Pay-per-Use to Yearly/Monthly

This section describes how to change the billing mode of a GeminiDB Influx instance from pay-per-use to yearly/monthly. If you want to use a pay-per-use instance for a long time, change its billing mode to yearly/monthly to reduce costs.

Precautions

- Only when the status of a pay-per-use instance is **Available**, its billing mode can be changed to yearly/monthly.

Changing the Billing Mode of a Single Pay-per-Use Instance to Yearly/Monthly

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**.

Figure 3-64 Change to Yearly/Monthly

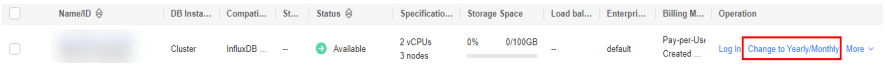

Step 4 On the displayed page, select the renewal duration in month. The minimum duration is one month.

Confirm the settings and click **Pay Now**.

Step 5 Select a payment method and click **Pay**.

Step 6 View the 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

Batch Changing Pay-per-Use to Yearly/Monthly

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 whose billing mode you want to change and click **Change to Yearly/Monthly** above the instance list. In displayed dialog box, click **Yes**.

Figure 3-65 Change to Yearly/Monthly

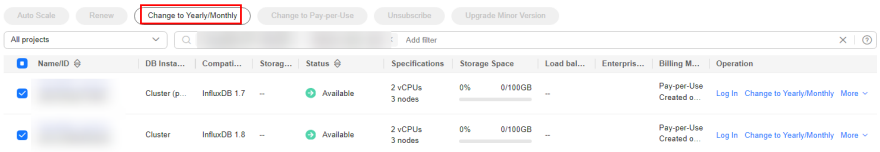

Step 4 On the displayed page, select the renewal duration in month. The minimum duration is one month.

Confirm the settings and click **Pay Now**.

Step 5 Select a payment method and click **Pay**.

Step 6 View the 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 becomes to **Yearly/Monthly**.

----End

3.14.3 Changing Yearly/Monthly to Pay-per-Use

You can change the billing mode of a GeminiDB Influx instance from yearly/monthly to pay-per-use and then pay only for the actual usage of your resources.

Precautions

- The billing mode of a yearly/monthly instance can only be changed to pay-per-use when the instance is in the **Available** status.

Changing the Billing Mode of a Single Yearly/Monthly Instance to Pay-per-Use

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 billing mode you want to change and click **Change to Pay-per-Use** in the **Operation** column.

Figure 3-66 Change to Pay-per-Use

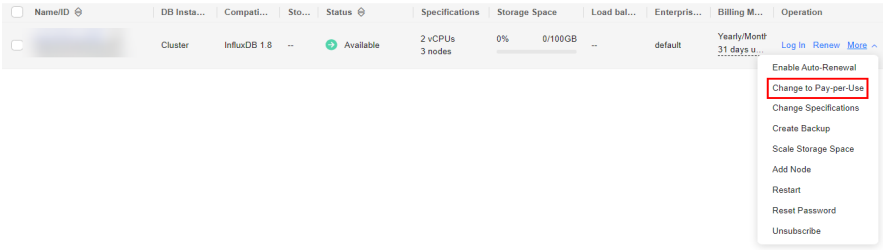

Step 4 On the displayed page, confirm the instance information and click **Change to Pay-per-Use**. The billing mode will change to pay-per-use after the instance expires.

NOTICE

Auto renewal will be disabled after the billing mode of your instances change to pay-per-use. Exercise caution when performing this operation.

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 instance expires.

Step 6 To cancel the change, choose **Billing** > **Renewal** to enter the Billing Center. On the **Renewals** page, locate the instance and click **More** > **Cancel Change to Pay-per-Use**.

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

----End

Batch Changing Yearly/Monthly to Pay-per-Use

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 whose billing mode you want to change and click **Change to Pay-per-Use** above the instance list.

Figure 3-67 Batch changing yearly/monthly to pay-per-use

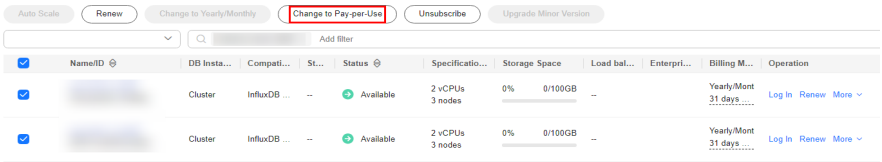

Step 4 In the displayed dialog box, click **Yes**.

Step 5 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.

NOTICE

Auto renewal will be disabled after the billing mode of your instances change to pay-per-use. Exercise caution when performing this operation.

Step 6 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 instance expires.

Step 7 To cancel the change, choose **Billing** > **Renewal** to enter the Billing Center. On the **Renewals** page, locate the instance and click **More** > **Cancel Change to Pay-per-Use**.

Step 8 In the displayed dialog box, click **Yes**.

----End

3.14.4 Unsubscribing from a Yearly/Monthly Instance

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

Precautions

- Unsubscribed operations cannot be undone. Exercise caution when performing this operation. To retain data, create a manual backup before unsubscription. For details, see **Creating a Manual Backup**.
- After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved. Ensure that the manual backup is complete before submitting the unsubscription request.

Unsubscribing from a Single Yearly/Monthly Instance

Step 1 Log in to the Huawei Cloud console.

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

Step 3 On the **Instances** page, locate the instance you want to unsubscribe from and click **Unsubscribe** or choose **More** > **Unsubscribe** in the **Operation** column.

Figure 3-68 Unsubscribing from a yearly/monthly instance

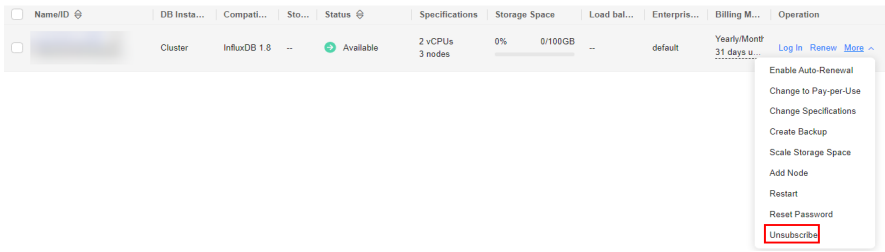

Step 4 In the displayed dialog box, click **Yes**.

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

For details, see **Unsubscription Rules**.

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

NOTICE

1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
2. If you want to retain data, complete a manual backup before submitting the unsubscription request.

Step 7 View the results. After the instance order is successfully unsubscribed, 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 3-69 Batch unsubscribing from yearly/monthly instances

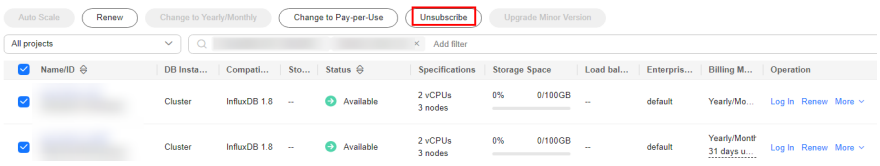

Step 4 In the displayed dialog box, click **Yes**.

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

For details, see **Unsubscription Rules**.

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

NOTICE

1. After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.
2. If you want to retain data, complete a manual backup before submitting the unsubscription request.

Step 7 View the results. After the instance order is successfully unsubscribed, the instance is no longer displayed in the instance list on the **Instances** page.

----End

3.15 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 3-40**.
- A maximum of 20 tags can be added for each instance.
- The tag name must comply with the naming rules described in **Table 3-40**.

Table 3-40 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. • Can only consist of digits, letters, underscores (_), and hyphens (-).	Organization
Tag value	• Can be left blank. • Can contain a maximum of 255 characters. • Can only consist of digits, letters, underscores (_), periods (.), and hyphens (-).	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, click **Search by Tag** in the upper right corner of the instance list.

Figure 3-70 Search by Tag

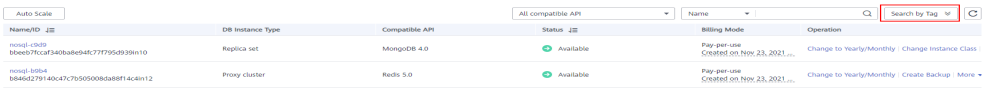

Step 4 Enter a tag key or value and click **Search** to query the instance associated with the tag.

Figure 3-71 Searching by tag key

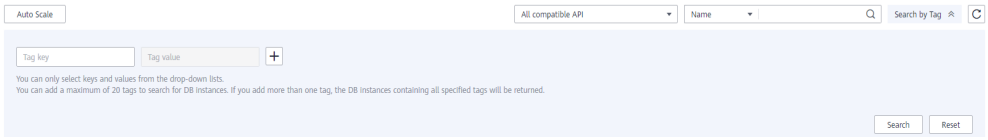

----End

4 FAQs

4.1 Product Consulting

4.1.1 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.

4.1.2 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.

4.1.3 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.

4.1.4 How Much Data Can a GeminiDB Influx Instance Hold?

For details, see **Instance Specifications**.

4.1.5 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?**.

4.1.6 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;
```

4.1.7 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'`

4.2 Billing

4.2.1 What Are the Differences Between Yearly/Monthly and Pay-per-use Billing Mode?

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 post payment mode, so you can start or stop an instance at any time. Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration.

4.2.2 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 Yearly/Monthly to Pay-per-Use**.
- If you want to change the billing mode from pay-per-use to yearly/monthly, see **Changing Pay-per-Use to Yearly/Monthly**.

4.3 Database Connection

4.3.1 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*.

4.3.2 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**.

4.3.3 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 **Connection Methods**.

4.3.4 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 4-1 Logging in to the Grafana homepage

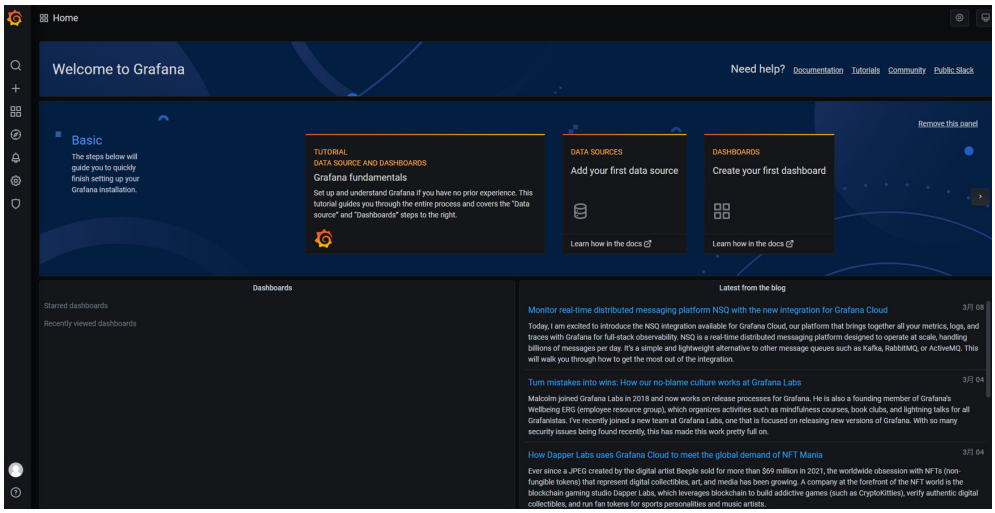

Step 3 Create a data source.

Figure 4-2 Creating a data source

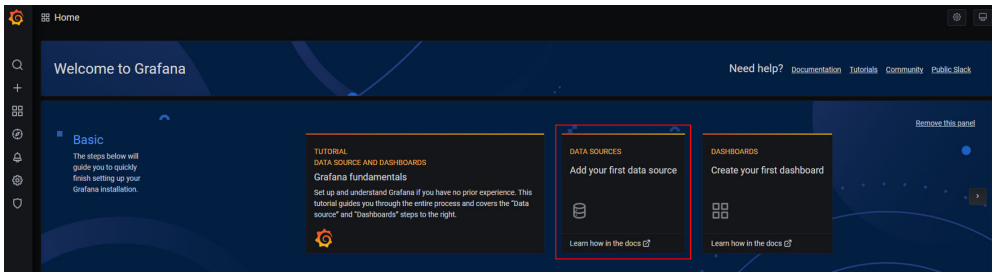

Step 4 Select **InfluxDB**.

Figure 4-3 Selecting **InfluxDB**

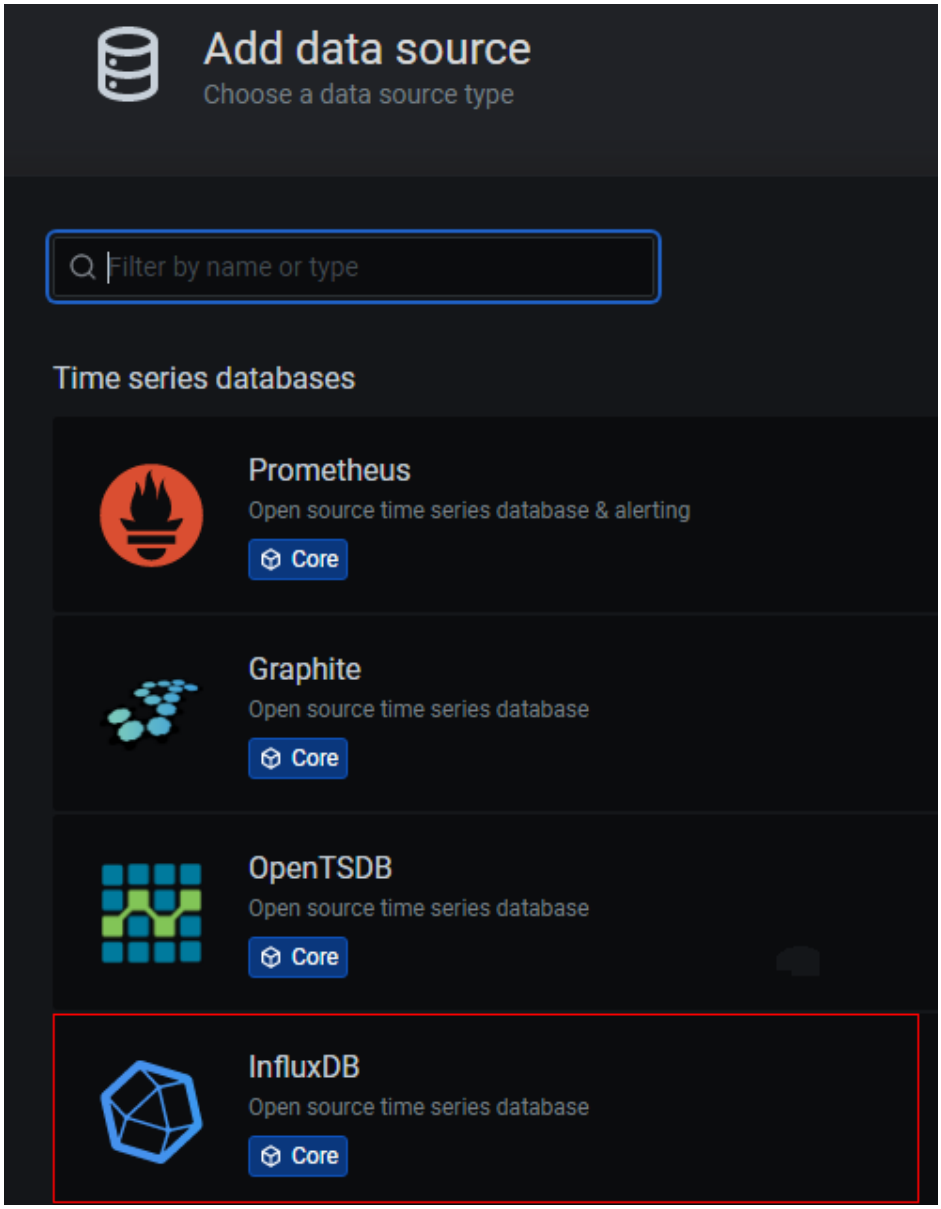

Step 5 Configure the required parameters.

Figure 4-4 Configuring parameters

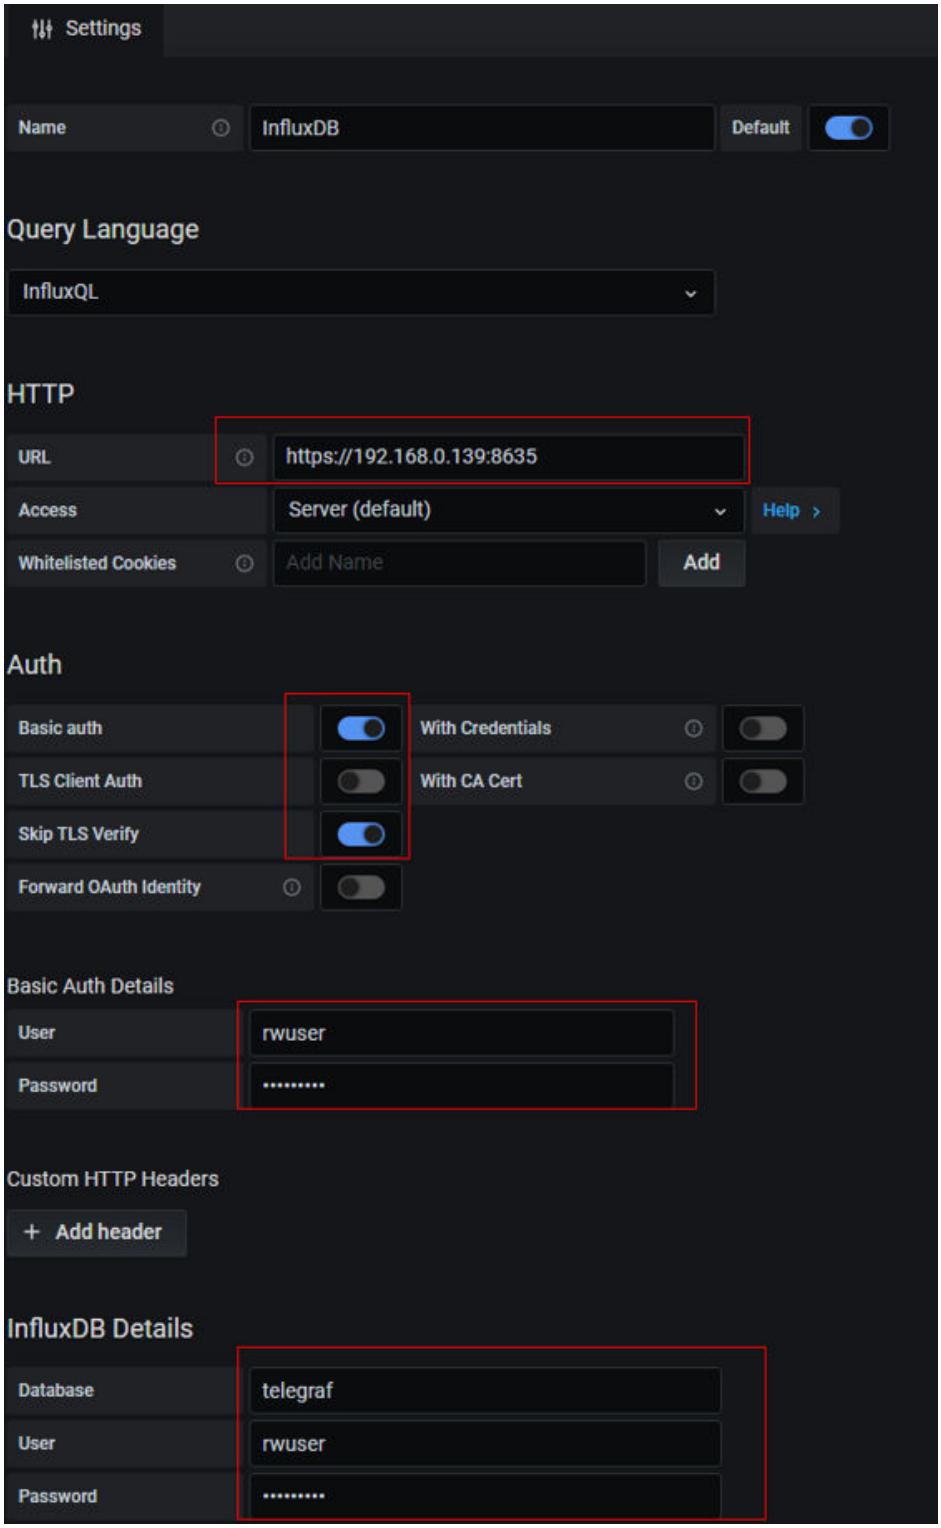

Table 4-1 Parameter description

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**.

4.4 Backup and Restoration

4.4.1 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.

4.5 Regions and AZs

4.5.1 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**.

4.5.2 Can I Change the Region of a GeminiDB Influx Instance?

No. After an instance is created, its region cannot be changed.

4.6 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 instance expires, you can renew it on the **Renewals** 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 **Recycling an Instance**. 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 **Unsubscribing from a Yearly/Monthly Instance**.

