

**GaussDB(for MySQL)**

# **User Guide**

**Issue**                45  
**Date**                 2024-02-23



**Copyright © Huawei Technologies Co., Ltd. 2024. All rights reserved.**

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

## **Trademarks and Permissions**



HUAWEI and other Huawei trademarks are trademarks 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 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.

## **Huawei Technologies Co., Ltd.**

Address: Huawei Industrial Base  
Bantian, Longgang  
Shenzhen 518129  
People's Republic of China

Website: <https://www.huawei.com>

Email: [support@huawei.com](mailto:support@huawei.com)

# Security Declaration

## Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

---

# Contents

---

<b>1 Usage Rules.....</b>	<b>1</b>
1.1 Design Rules.....	1
1.1.1 Database Table Usage.....	1
1.1.2 Index Usage.....	2
1.1.3 Database Permissions.....	4
1.2 Development Rules.....	4
1.2.1 SQL Usage.....	4
<b>2 Permissions Management.....</b>	<b>9</b>
2.1 Creating a User and Granting GaussDB(for MySQL) Permissions.....	9
2.2 Creating a GaussDB(for MySQL) Custom Policy.....	10
<b>3 Billing Management.....</b>	<b>13</b>
3.1 Manually Renewing a DB Instance.....	13
3.2 Auto-Renewing a DB Instance.....	14
3.3 Changing the Billing Mode from Yearly/Monthly to Pay-per-Use.....	17
3.4 Changing the Billing Mode from Pay-per-Use to Yearly/Monthly.....	18
3.5 Unsubscribing a Yearly/Monthly Instance.....	20
<b>4 Data Migration.....</b>	<b>23</b>
4.1 Migration Solution Overview.....	23
4.2 Migrating Data to GaussDB(for MySQL) Using DRS.....	25
4.3 Migrating Data to GaussDB(for MySQL) Using mysqldump.....	25
<b>5 Serverless.....</b>	<b>29</b>
5.1 Billing.....	29
5.2 Buying a Serverless DB Instance.....	30
5.3 Changing the Compute Range of a Serverless Instance.....	38
5.4 Changing Serverless Compute Resources.....	39
<b>6 Database Proxy (Read/Write Splitting).....</b>	<b>40</b>
6.1 Introducing Read/Write Splitting.....	40
6.2 Introducing Consistency Levels.....	42
6.3 Create a Proxy Instance.....	43
6.4 Configuring Connection Pools.....	49
6.5 Configuring Transaction Splitting.....	50

6.6 Configuring a Routing Policy.....	52
6.7 Assigning Read Weights.....	53
6.8 Configuring Multi-Statement Processing Mode.....	56
6.9 Changing the Specifications of a Proxy Instance.....	57
6.10 Changing the Number of Proxy Instance Nodes.....	59
6.11 Upgrading the Kernel Version of a Proxy Instance.....	60
6.12 Using a Private Domain Name for a Proxy Instance.....	61
6.13 Changing the IP Address of a Proxy Instance.....	63
6.14 Changing the Port of a Proxy Instance.....	65
6.15 Changing Consistency Level.....	66
6.16 Modifying Proxy Instance Parameters.....	67
6.17 Enabling or Disabling Access Control.....	68
6.18 Rebooting a Proxy Instance.....	69
6.19 Deleting a Proxy Instance.....	70
6.20 Introducing Read Weight Assignment Rules.....	71
6.21 Testing Read/Write Splitting Performance.....	72
<b>7 Instance Lifecycle Management.....</b>	<b>74</b>
7.1 Changing a DB Instance Name.....	74
7.2 Changing a DB Instance Description.....	75
7.3 Deleting a DB Instance.....	76
7.4 Rebooting a DB Instance.....	77
7.5 Rebooting a Node.....	79
7.6 Changing a Node Name.....	81
7.7 Exporting Instance Information.....	82
7.8 Introducing Heterogeneous DR Instances.....	83
7.9 Rebuilding a Deleted Instance from Recycle Bin.....	84
<b>8 Instance Modifications.....</b>	<b>86</b>
8.1 Changing vCPUs and Memory of a DB Instance.....	86
8.2 Changing Storage of a Yearly/Monthly DB Instance.....	88
8.3 Configuring Auto Scaling Policies.....	90
8.4 Changing a Maintenance Window.....	95
8.5 Selecting Instance Displayed Items.....	97
8.6 Upgrading a Minor Version.....	97
8.7 Upgrading the OS of a DB Instance.....	99
<b>9 Read Replicas.....</b>	<b>101</b>
9.1 Introducing Read Replicas.....	101
9.2 Creating a Read Replica.....	102
9.3 Managing a Read Replica.....	103
9.4 Promoting a Read Replica to the Primary Node.....	104
9.5 Deleting a Pay-per-Use Read Replica.....	105
9.6 Unsubscribing a Yearly/Monthly Read Replica.....	107

<b>10 HTAP Analysis.....</b>	<b>109</b>
10.1 What Is HTAP?.....	109
10.2 Constraints.....	113
10.3 Usage Rules.....	113
10.4 HTAP Instance Lifecycle Management.....	114
10.4.1 Creating an HTAP instance.....	114
10.4.2 Rebooting an HTAP Instance.....	116
10.4.3 Deleting an HTAP Instance.....	116
10.5 HTAP Instance Modifications.....	117
10.5.1 Binding and Unbinding an EIP.....	117
10.5.2 Scaling up Storage.....	119
10.5.3 Changing vCPUs and Memory.....	120
10.6 Managing Accounts.....	122
10.7 Viewing Metrics.....	125
10.8 Data Security.....	126
10.8.1 Configuring SSL.....	127
10.9 Data Synchronization.....	128
10.9.1 Introducing DDLs that Cannot Be Synchronized.....	129
10.9.2 Synchronizing Data.....	132
10.9.3 Adding or Deleting Tables in the Blacklist or Whitelist.....	145
10.10 Log Management.....	147
10.10.1 Viewing Slow Query Logs.....	147
10.11 Connecting to an HTAP Instance.....	150
10.11.1 Connecting to an HTAP Instance Through DAS.....	150
10.11.2 Connecting to an HTAP Instance Through JDBC.....	151
10.11.3 Connecting to an HTAP Instance Through HTTPS.....	152
10.12 Developer Guide.....	153
10.12.1 Data Type Conversion.....	153
10.12.2 Syntax.....	154
10.12.3 Database Status.....	155
10.12.4 DB Engines.....	155
10.12.5 Table Engines.....	155
<b>11 Database Management.....</b>	<b>156</b>
11.1 Creating a Database.....	156
11.2 Deleting a Database.....	158
<b>12 Account Management (Non-Administrator).....</b>	<b>160</b>
12.1 Creating a Database Account.....	160
12.2 Resetting a Password for a Database Account.....	161
12.3 Changing Permissions for Database Accounts.....	163
12.4 Deleting a Database Account.....	163
<b>13 Data Security.....</b>	<b>165</b>

13.1 Resetting the Administrator Password.....	165
13.2 Changing a Security Group.....	167
13.3 Configuring Database Account Security.....	167
13.4 Configuring SSL.....	168
13.5 Enabling TDE.....	170
<b>14 Data Backups.....</b>	<b>172</b>
14.1 Backup Principles.....	172
14.2 Backup Types.....	174
14.3 Backup Space and Billing.....	177
14.4 Configuring a Same-Region Backup Policy.....	178
14.5 Configuring a Cross-Region Backup Policy.....	181
14.6 Enabling or Disabling Encrypted Backup.....	184
14.7 Creating a Manual Backup.....	187
14.8 Exporting Backup Information.....	189
14.9 Deleting a Manual Backup.....	190
<b>15 Data Restorations.....</b>	<b>192</b>
15.1 Restoring a DB Instance.....	192
15.2 Restoring Instance Data to a Specific Point in Time.....	192
15.3 Restoring Data to a DB Instance.....	195
15.4 Restoring Table Data to a Specific Point in Time.....	198
<b>16 Connection Management.....</b>	<b>201</b>
16.1 Binding and Unbinding an EIP.....	201
16.2 Changing a Database Port.....	203
16.3 Applying for and Changing a Private Domain Name for a DB Instance.....	204
16.4 Configuring and Changing a Private IP Address.....	205
<b>17 Parameter Template Management.....</b>	<b>207</b>
17.1 Creating a Parameter Template.....	207
17.2 Modifying Parameters of a GaussDB(for MySQL) DB Instance.....	209
17.3 Exporting Parameters.....	211
17.4 Comparing Parameter Templates.....	212
17.5 Viewing Parameter Change History.....	214
17.6 Replicating a Parameter Template.....	215
17.7 Resetting a Parameter Template.....	218
17.8 Applying a Parameter Template.....	219
17.9 Viewing Application Records of a Parameter Template.....	219
17.10 Editing a Parameter Template Description.....	220
17.11 Deleting a Parameter Template.....	221
<b>18 Metrics and Alarms.....</b>	<b>223</b>
18.1 Introducing GaussDB(for MySQL) Metrics.....	223
18.2 Viewing Monitoring Metrics.....	236

18.2.1 Viewing Instance Monitoring Metrics.....	236
18.2.2 Viewing Proxy Monitoring Metrics.....	238
18.3 Configuring Alarm Rules.....	238
18.3.1 Creating Alarm Rules for a DB Instance.....	238
18.3.2 Creating Alarm Rules for a Proxy Instance.....	245
18.4 Event Monitoring.....	248
18.4.1 Introducing Event Monitoring.....	248
18.4.2 Viewing Event Monitoring Data.....	248
18.4.3 Creating Alarm Rules for Event Monitoring.....	249
18.4.4 Events Supported by Event Monitoring.....	251
18.5 Configuring Monitoring by Seconds.....	262
<b>19 Interconnection with CTS.....</b>	<b>266</b>
19.1 Key Operations Supported by CTS.....	266
19.2 Viewing Tracing Events.....	268
<b>20 Log Management.....</b>	<b>271</b>
20.1 Enabling or Disabling Log Reporting.....	271
20.2 Viewing Error Logs.....	274
20.3 Viewing Slow Query Logs.....	276
20.4 Enabling or Disabling SQL Explorer.....	279
<b>21 DBA Assistant.....</b>	<b>280</b>
21.1 Function Overview.....	280
21.2 Dashboard.....	281
21.3 Sessions.....	284
21.4 Performance.....	284
21.5 Slow SQL Analysis.....	285
21.6 Top SQL.....	289
21.7 SQL Insights.....	290
21.8 Concurrency Control.....	292
<b>22 Task Center.....</b>	<b>295</b>
22.1 Viewing a Task.....	295
22.2 Deleting a Task Record.....	298
<b>23 Managing Tags.....</b>	<b>300</b>
<b>24 Managing Quotas.....</b>	<b>303</b>
<b>25 Kernel Functions.....</b>	<b>305</b>
25.1 Parallel Query.....	305
25.1.1 Overview.....	305
25.1.2 Usage.....	306
25.1.2.1 Application Scenarios.....	306
25.1.2.2 Constraints.....	307
25.1.2.3 Parameters and Variables.....	311



25.1.2.4 Enabling or Disabling Parallel Query.....	312
25.1.2.5 Viewing the Statuses of Parallel Query.....	313
25.1.3 Performance Tests.....	314
25.1.3.1 Test Method.....	314
25.1.3.2 Test Results.....	316
<b>25.2 Near Data Processing.....</b>	<b>316</b>
25.2.1 NDP Overview.....	316
25.3 DDL Optimization.....	318
25.3.1 Parallel DDL.....	318
25.3.2 DDL Fast Timeout.....	320
25.4 Statement Outline.....	322
25.5 Idle Transaction Disconnection.....	329
25.5.1 Function.....	329
25.5.2 Parameter Description.....	329
25.5.3 Example.....	331
25.6 LIMIT...OFFSET Pushdown.....	332
25.6.1 Function.....	332
25.6.2 Usage.....	333
25.6.3 Performance Tests.....	334
25.7 Conversion of IN Predicates Into Subqueries.....	335
25.7.1 Function.....	335
25.7.2 Scenarios.....	336
25.7.3 Usage.....	336
25.7.4 Performance Tests.....	338
25.8 Diagnosis on Large Transactions.....	338
<b>A Change History.....</b>	<b>342</b>

# 1 Usage Rules

---

## 1.1 Design Rules

### 1.1.1 Database Table Usage

- All created MySQL tables must use the InnoDB engine.
- The decimal type must be DECIMAL. Do not use FLOAT or DOUBLE.

 **NOTE**

FLOAT and DOUBLE have lower precision than DECIMAL and may cause rounding errors. If a value to be stored is beyond the range of DECIMAL, split the value into INTEGER and DECIMAL parts and store them separately.

- The following reserved words cannot be used: desc, range, match, and delayed. For details, see MySQL official website.
- Every data table can have a primary key. The primary key can be an ordered and unique field related to business or an auto-increment field unrelated to business.
- Each table field must have a default value and NOT NULL. If the field is the numeric type, use 0 as its default value. If the field is the character type (such as VARCHAR), use an empty string ("").

 **NOTE**

No primary key may cause slow execution of the primary database and replication delay.

- You are not advised to use partitioned tables. If necessary, use multiple independent tables.

 **NOTE**

Disadvantages of partitioned tables:

- All partitions will be locked during DDL operations. As a result, operations on the partitions will be blocked.
  - When a partitioned table contains a large amount of data, it is difficult and risky to perform DDL or other O&M operations on the table.
  - Partition tables are seldom used, which may cause unknown risks.
  - When a single server is poor in performance, splitting a partitioned table is expensive.
  - When all partitions are accessed due to improper operations on a partitioned table, severe performance problems may occur.
- Each table contains two DATETIME fields: **create\_time** and **update\_time**.

 **NOTE**

You can obtain the required data from a data warehouse based on these two fields without consulting services.

When an exception occurs in the database, you can use the two fields to determine the time when the data is inserted and updated. In extreme cases, you can determine whether to restore data based on the fields.

- Sharding is recommended only when the number of rows in a single table exceeds 5 million or the size of a single table exceeds 2 GB.
- VARCHAR is a variable-length character data type. The length of VARCHAR cannot exceed 2,048.

If the length of a field exceeds 2,048, define the field type as TEXT or create an independent table and use a primary key to associate the related tables. In this way, the index efficiency of other fields is not affected.

- The length of a single row in a table cannot exceed 1,024 bytes and the number of rows in a single table cannot exceed 5 million.
- The maximum number of fields in a single table is 50.
- If the lengths of strings are almost the same, use the fixed-length character strings.
- On the premise of ensuring data consistency, cross-table redundant fields are allowed to avoid associated queries and improve query performance.

 **NOTE**

Redundant fields must comply with the following rules:

- Fields are not frequently modified.
  - Fields are not larger VARCHAR and TEXT.
- The data types with proper storage size can save database tablespace and index storage space while improving the search speed. LONG TEXT and BLOB are not recommended.

## 1.1.2 Index Usage

- Use the same field type to prevent implicit conversion from causing invalid indexes.
- Create unique indexes on the minimum set of fields or combinations of fields with unique characteristics.

For example, there is a table containing the fields **a**, **b**, **c**, **d**, **e**, and **f**. If the combinations of fields **ab** and **ef** have unique features, you are advised to create unique indexes for **ab** and **ef**, respectively.

 **NOTE**

Even if complete verification control is implemented at the application layer, dirty data is generated as long as there is no unique index according to Murphy's Law.

Before creating a unique index, consider whether it is useful for the query and whether it will affect the INSERT operations.

- Create indexes on fixed-length fields (for example, INT). When creating an index on a VARCHAR field, the index length must be specified. It is not necessary to create an index on the whole field. The index length is determined according to the actual text distinction.

 **NOTE**

The index length and distinction are a pair of contradictions. Generally, for string type data, the distinction of an index with a length of 20 will be more than 90%. The distinction formula is:  $\text{Count}(\text{distinct left}(\text{column name}, \text{index length}))/\text{count} (*)$ . Place the column names with the high distinction in the front.

- If possible, do not use left fuzzy search (for example, **SELECT \* FROM users WHERE u\_name LIKE ' %hk'**) or full fuzzy search on the page to avoid degradation from index scan to full table scan. Solve the problem at the application layer.

 **NOTE**

An index file has the leftmost prefix matching feature of B-tree. If the value on the left is not determined, the index cannot be used.

- Use the overlay index to query and avoid returning to the table. However, the number of fields added to the overlay index cannot be too large, or the write performance will be affected.

 **NOTE**

Types of indexes that can be created include primary key index, unique index, and normal index. Covering index indicates that if you execute EXPLAIN statements, "using index" is displayed in the **Extra** column.

- Optimize the SQL performance as follows: range (minimum requirement), ref (basic requirement), and consts (maximum requirement).
- When creating a composite index, place the column with the highest distinction on the left.
- Ensure that the number of indexes in a single table is at most 5, or does not exceed 20% of the number of table fields.
- Avoid the following misunderstandings when creating indexes:
  - Indexes should be frequently used. An index needs to be created for a query.
  - Indexes should be as few as possible. Indexes consume space and slow down the update and addition speed.
  - Unique indexes cannot be used. Unique features must be resolved at the application layer using the "query first and then insert" method.

## 1.1.3 Database Permissions

- All DDL operations (such as creating tables and modifying table structures) are performed by DBAs through DAS only after being reviewed. Services are launched during off-peak hours.
- Permissions must be managed in a fine-grained manner by separating read permissions from write permissions, and O&M permissions from development permissions.
- DDL operations are recorded in operation logs.

## 1.2 Development Rules

### 1.2.1 SQL Usage

#### Database Design

- Ensure that all characters are stored and represented in utf-8 or utf8mb4 format. Comments must be provided for tables and fields.
- Avoid using large transactions.  
For example, if multiple SELECT and UPDATE statements are executed in a high-frequency transaction, the database concurrency capability is severely affected because resources such as locks held by the transaction can be released only when the transaction is rolled back or committed. In this case, data write consistency must also be considered.

#### Index Design

- Reduce the use of ORDER BY that cannot be used with indexes based on the actual service requirements. The statements such as ORDER BY, GROUP BY, and DISTINCT consume many CPU resources.
- If a complex SQL statement is involved, use the existing index design and add EXPLAIN before the SQL statement. EXPLAIN can help you optimize the index by adding some query restrictions.
- Execute new SELECT, UPDATE, or DELETE statements with EXPLAIN to check the index usage and ensure no **Using filesort** and **Using temporary** are displayed in the **Extra** column. If the number of scanned rows exceeds 1,000, exercise caution when executing these statements. Analyze slow query logs and delete unused slow query statements every day.

 NOTE

EXPLAIN:

- **type:** ALL, index, range, ref, eq\_ref, const, system, NULL (The performance is sorted from poor to good from left to right.)
- **possible\_keys:** indicates the indexes from which MySQL can choose to find the rows in this table. If there is an index on a field, the index is listed but may not be used by the query.
- **key:** indicates the key (index) that MySQL actually decided to use. If key is NULL, MySQL found no index to use for executing the query more efficiently. To force MySQL to use or ignore an index listed in the **possible\_keys** column, use FORCE INDEX, USE INDEX, or IGNORE INDEX in your query.
- **ref:** shows which columns or constants are compared to the index named in the key column to select rows from the table.
- **rows:** indicates the estimated number of rows to be read for required records based on table statistics and index selection.
- **Extra:**
  - **Using temporary:** To resolve the query, MySQL needs to create a temporary table to hold the result. This typically happens if the query contains GROUP BY and ORDER BY clauses that list columns differently.
  - **Using filesort:** MySQL must do an extra pass to find out how to retrieve the rows in sorted order.
  - **Using index:** The column information is retrieved from the table using only information in the index tree without having to do an additional seek to read the actual row. If **Using where** is displayed at the same time, it indicates that desired information needs to be obtained by using the index tree and reading rows of the table.
  - **Using where:** In WHERE clause, **Using where** is displayed when the desired data is obtained without reading all the data in the table or the desired data cannot be obtained by only using indexes. Unless you specifically intend to fetch or examine all rows from the table, you may have something wrong in your query if the **Extra** value is not **Using where** and the table join type is ALL or index.
- If a function is used on a WHERE statement, the index becomes invalid.  
For example, in **WHERE left(name, 5) = 'zhang'**, the left function invalidates the index on **name**.  
You can modify the condition on the service side and delete the function.  
When the returned result set is small, the service side filters the rows that meet the condition.

## Database SQL Query

- Optimize the ORDER BY... LIMIT statements by indexes to improve execution efficiency.
- If statements contain ORDER BY, GROUP BY, or DISTINCT, ensure that the result set filtered by the WHERE condition contains up to 1000 lines. Otherwise, the SQL statements are executed slowly.
- For ORDER BY, GROUP BY, and DISTINCT statements, use indexes to directly retrieve sorted data. For example, use **key(a,b)** in **where a=1 order by b**.
- When using JOIN, use indexes on the same table in the WHERE condition.

Example:

```
select t1.a, t2.b from t1,t2 where t1.a=t2.a and t1.b=123 and t2.c= 4
```

If the **t1.c** and **t2.c** fields are the same, only **b** on **t1** is used.

In this case, if **t2.c=4** in the WHERE condition is changed to **t1.c=4,(b,c)** can be used. This may occur during field redundancy design (anti normal form).

- If deduplication is not required, use UNION ALL, which does not perform sorting operations and is faster than UNION.
- To implement paging query in code, specify that if **count** is set to **0**, the subsequent paging statements are not executed.
- Do not frequently execute COUNT on a table. It takes a long time to perform COUNT on a table with a large amount of data. Generally, the response speed is in seconds. If you need to frequently perform the COUNT operation on a table, introduce a special counting table.
- If only one record is returned, use LIMIT 1. If data is correct and the number of returned records in the result set can be determined, use LIMIT as soon as possible.
- When evaluating the efficiency of DELETE and UPDATE statements, change the statements to SELECT and then run EXPLAIN. A large number of SELECT statements will slow down the database, and write operations will lock tables.
- TRUNCATE TABLE is faster and uses fewer system and log resources than DELETE. If the table to be deleted does not have a trigger and the entire table needs to be deleted, TRUNCATE TABLE is recommended.
  - TRUNCATE TABLE does not write deleted data to log files.
  - A TRUNCATE TABLE statement has the same function as a DELETE statement without a WHERE clause.
  - TRUNCATE TABLE statements cannot be written with other DML statements in the same transaction.
- Do not use negative queries to avoid full table scanning. Negative queries indicate the following negative operators are used: NOT, !=, <>, NOT EXISTS, NOT IN, and NOT LIKE.

If a negative query is used, the index structure cannot be used for binary search. Instead, the entire table needs to be scanned.
- Do not perform JOIN on more than three tables. The data types of the fields to be joined must be the same.
- During multi-table associated query, ensure that the associated fields have indexes. When joining multiple tables, select the table with a smaller result set as the driving table to join other tables. Pay attention to table indexes and SQL performance even if two tables are joined.

## SQL Statement Development

- Splitting single SQL statements.

For example, in the OR condition **f\_phone='10000' or f\_mobile='10000'**, the two fields have their own indexes, but only one of them can be used.

You can split the statement into two SQL statements or use UNION ALL.
- If possible, perform the complex SQL calculation or service logic at the service layer.
- Use a proper paging method to improve paging efficiency. Skipping paging is not recommended for large pages.

- Negative example: **SELECT \* FROM table1 ORDER BY ftime DESC LIMIT 10000,10;**  
It causes a large number of I/O operations because MySQL uses the read-ahead policy.
- Positive example: **SELECT \* FROM table1 WHERE ftime < last\_time ORDER BY ftime DESC LIMIT 10;**  
Recommended pagination mode: Transfer the threshold value of the last page.
- Execute UPDATE statements in transactions based on primary keys or unique keys. Otherwise, a gap lock is generated and the locked data range is expanded. As a result, the system performance deteriorates and a deadlock occurs.
- Do not use foreign keys and cascade operations. The problems of foreign keys can be solved at the application layer.  
Example:  
If **student\_id** is a primary key in the student table, **student\_id** is a foreign key in the score table. If **student\_id** is updated in the student table, **student\_id** in the score table is also updated. This is a cascade update.
  - Foreign keys and cascade updates are suitable for single-node clusters with low concurrency and are not suitable for distributed cluster with high concurrency.
  - Cascade updates may cause strong blocks and foreign keys affect the INSERT operations.
- If possible, do not use IN. If it is required, ensure that the number of set elements after IN should be at most 500.
- To reduce the number of interactions with the database, use batches of SQL statements. For example, **INSERT INTO... VALUES (XX),(XX),(XX)....(XX)**, the number of **XX** should be within 100.
- Do not use stored procedures, which are difficult to debug, extend, and transplant.
- You are not advised to use triggers, event schedulers, and views for service logic. The service logic must be processed at the service layer to avoid logical dependency on the database.
- Do not use implicit type conversion.



 **NOTE**

The conversion rules are as follows:

1. If at least one of the two parameters is NULL, the comparison result is also NULL. However, when <=> is used to compare two NULL values, 1 is returned.
2. If both parameters are character strings, they are compared as character strings.
3. If both parameters are integers, they are compared as integers.
4. When one parameter is a hexadecimal value and the other parameter is a non-digit value, they are compared as binary strings.
5. If one parameter is a TIMESTAMP or DATETIME value and the other parameter is a CONSTANT value, they are compared as TIMESTAMP values.
6. If one parameter is a DECIMAL value and other parameter is a DECIMAL or INTEGER value, they are compared as DECIMAL values. If the other argument is a FLOATING POINT value, they are compared as FLOATING POINT values.
7. In other cases, both parameters are compared as FLOATING POINT values.
8. If one parameter is a string and the other parameter is an INT value, they are compared as FLOATING POINT values (by referring to item 7)

For example, the type of **f\_phone** is varchar. If **f\_phone in (098890)** is used in the WHERE condition, two parameters are compared as FLOATING POINT values. In this case, the index cannot be used, affecting database performance.

If **f\_user\_id = '1234567'**, the number is directly compared as a character string. For details, see item 2.

- If possible, ensure that the number of SQL statements in a transaction should be as small as possible, no more than 5. Long transactions will lock data for a long time, generate many caches in MySQL, and occupy many connections.
- Do not use NATURAL JOIN.

NATURAL JOIN is used to implicitly join column, which is difficult to understand and may cause problems. The NATURAL JOIN statement cannot be transplanted.

# 2 Permissions Management

---

## 2.1 Creating a User and Granting GaussDB(for MySQL) Permissions

This section describes how to use [IAM](#) for fine-grained permissions control with your GaussDB(for MySQL) 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 GaussDB(for MySQL) resources.
- Grant only the permissions required for users to perform specific tasks.
- Entrust a cloud service account to perform efficient O&M on your GaussDB(for MySQL) resources.

If your account does not require individual IAM users, you can skip this section.

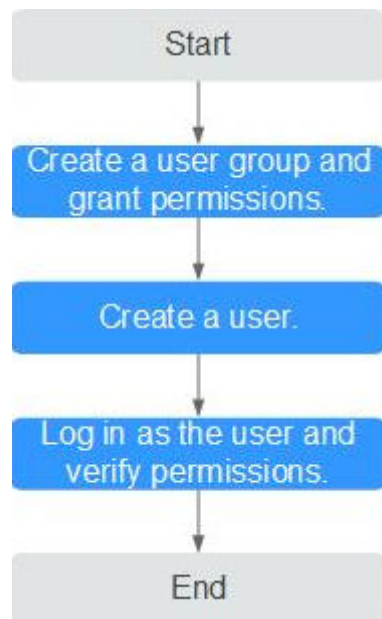
[Figure 2-1](#) describes the procedure for granting permissions.

### Prerequisites

Learn about the permissions (see [system-defined policies](#)) supported by GaussDB(for MySQL) and choose policies or roles according to your requirements. For the permissions of other services, see [System Permissions](#).

## Process Flow

Figure 2-1 Process of granting GaussDB(for MySQL) permissions



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

Create a user group on the IAM console, and attach the **GaussDB(for MySQL) GaussDB FullAccess** policy to the group.

**NOTE**

To use some functions of other services, you need to configure the GaussDB(for MySQL) FullAccess permission and the permission of the corresponding services. For example, when using DAS to connect to a DB instance, you need to configure the GaussDB FullAccess and DAS FullAccess permissions.

2. **Create an IAM user.**

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 GaussDB(for MySQL) console using the created user, and verify that the user only has read permissions for GaussDB(for MySQL).

Under the service list, choose **Databases > GaussDB(for MySQL)**. On the GaussDB(for MySQL) instance list page, click **Buy DB Instance** in the upper right corner. If the DB instance can be purchased, the required permissions policies have already been applied.

## 2.2 Creating a GaussDB(for MySQL) Custom Policy

Custom policies can be created to supplement the system-defined policies of GaussDB(for MySQL).

You can create a custom policy 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: Write policies from scratch or based on an existing policy.

For details, see [Creating a Custom Policy](#). This section provides examples of common GaussDB(for MySQL) custom policies.

## Example Custom Policies

- Example 1: Allowing users to create GaussDB(for MySQL) instances

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

- Example 2: Denying GaussDB(for MySQL) instance deletion

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

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

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

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

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

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "gaussdb:instance:create",
        "gaussdb:instance:modify",
        "gaussdb:instance:delete",
        "vpc:publicips:list",
        "vpc:publicips:update"
      ],
      "Effect": "Allow"
    }
  ]
}
```

```
    }  
  ]  
}
```

- Example 4: Allowing users to manage specified instances and some functions of instances

Assume that your account has multiple instances and you are a database administrator. If you want to allow users to manage specified instances and some functions of instances, you can create the following permission policy.

```
{  
  "Version": "1.1",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "gaussdb:instance:restart",  
        "gaussdb:instance:modify"  
      ],  
      "Resource": [  
        "GAUSSDB:*:*:instance:test*"  
      ]  
    },  
    {  
      "Effect": "Allow",  
      "Action": [  
        "gaussdb:param:list",  
        "gaussdb:tag:list",  
        "gaussdb:backup:list",  
        "gaussdb:instance:create",  
        "gaussdb:instance:list"  
      ]  
    }  
  ]  
}
```

#### NOTE

- Users granted these permissions can view all instances, but can manage only authorized instances. In addition, the database administrator can still use APIs to directly manage these instances. Users granted the permissions can only reboot and modify all instances under this account.
- **test\*** is an example of an instance name for fuzzy match and is mandatory in the permission policy. Otherwise, the authorized users cannot view any instance on the console.
- GaussDB(for MySQL) supports API-level access control. You can use related APIs to perform fine-grained access control on GaussDB(for MySQL). For details, see [API Overview](#).

# 3 Billing Management

---

## 3.1 Manually Renewing a DB Instance

### Scenarios

A yearly/monthly instance can be manually renewed based on service requirements.


### Constraints

- Only yearly/monthly instances can be renewed.
- The statuses of yearly/monthly instances to be renewed must be **Available** or **Abnormal**.

### Procedure

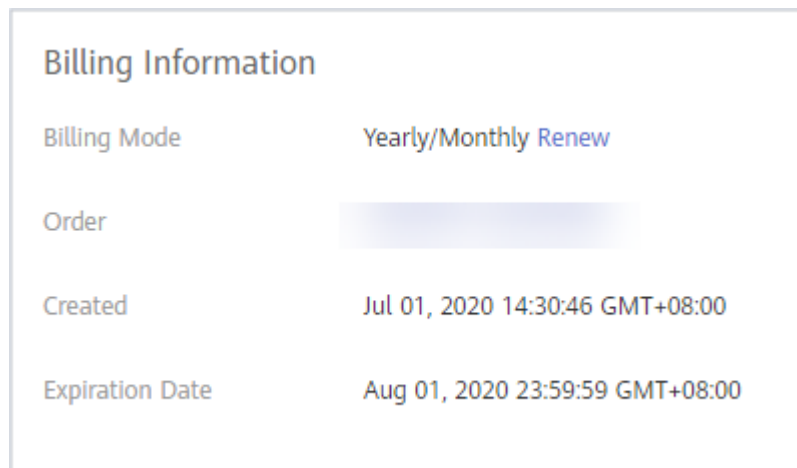
**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance you want to renew and click **Renew** in the **Operation** column.

Alternatively, click the DB instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Renew** on the right of **Billing Mode**.

**Figure 3-1** Renewing an instance

Billing Information	
Billing Mode	Yearly/Monthly Renew
Order	
Created	Jul 01, 2020 14:30:46 GMT+08:00
Expiration Date	Aug 01, 2020 23:59:59 GMT+08:00

**Step 5** On the displayed page, renew the instance.

----End

## 3.2 Auto-Renewing a DB Instance

### Scenarios

Auto-renewal can prevent DB instances from being automatically deleted if you forget to manually renew them.

If auto-renewal is not enabled when you purchase a yearly/monthly DB instance, you can enable it on the **Instances** page and then modify it later as needed.

The auto-renewal rules are as follows:

- The first auto-renewal date is based on when the DB instance expires and the billing cycle.
- The auto-renewal period of a DB instance depends on the subscription term.
  - Monthly subscriptions renew each month by default.
  - Yearly subscriptions renew each year by default.
- After auto-renewal is enabled, you can still renew the DB instance manually if you want to. After a manual renewal is complete, auto-renewal is still valid, and the renewal fee will be deducted from your account seven days before the new expiry date.
- By default, the renewal fee is deducted from your account seven days before the new expiry date. You can change this auto-renewal payment date as required.

For more information about auto-renewal rules, see [Auto-Renewal Rules](#).


### Prerequisites

The yearly/monthly DB instance is not expired.

## Enabling Auto-Renewal During Purchase

You can enable auto-renewal on the DB instance purchase page, as shown in [Figure 3-2](#). For details, see [Step 1: Buy a DB Instance](#).


Figure 3-2 Enabling auto-renewal



## Enabling Auto-Renewal on the Instances Page

**Step 1** [Log in to the management console](#).

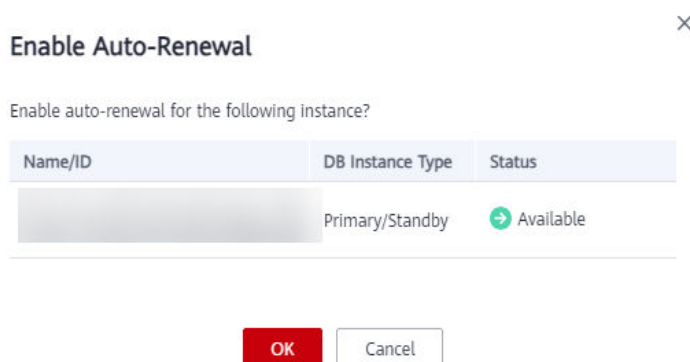
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the yearly/monthly instance and choose **More** > **Enable Auto-Renewal** in the **Operation** column.

**Step 5** In the displayed dialog box, confirm the DB instance information and click **OK**.

Figure 3-3 Enabling auto-renewal



**Step 6** Configure **New Auto-Renew Period** and **Auto-renewals** and click **OK**.

- **New Auto-Renew Period:** Select the period to be renewed.
  - Monthly subscriptions renew each month by default.
  - Yearly subscriptions renew each year by default.
- **Auto-renewals:** Select **Preset Auto-renewals** and then select **Unlimited** or **Custom**.
  - **Unlimited:** By default, the subscription renews every time before it expires.
  - **Custom:** Enter an integer from 1 to 1,000. The auto-renewal is disabled after the DB instance is automatically renewed for the specified number of times.



- After auto-renewal is enabled, you can modify it by referring to [Configuring Auto-Renewal](#).

**Step 7** In the **Billing Information** area on the **Basic Information** page of the instance, check whether **Upon Expiration** is **Automatic renewal**.

**Figure 3-4** Upon expiration




Billing Information			
Billing Mode	Yearly/Monthly Renew Change to Pay-per-use	Order	
Created	Oct 08, 2023 14:53:33 GMT+08:00	Expiration Date	Nov 08, 2023 23:59:59 GMT+08:00
Upon Expiration	Automatic renewal Configure Auto-Renewal		

----End

## Configuring Auto-Renewal

**Step 1** [Log in to the management console](#).

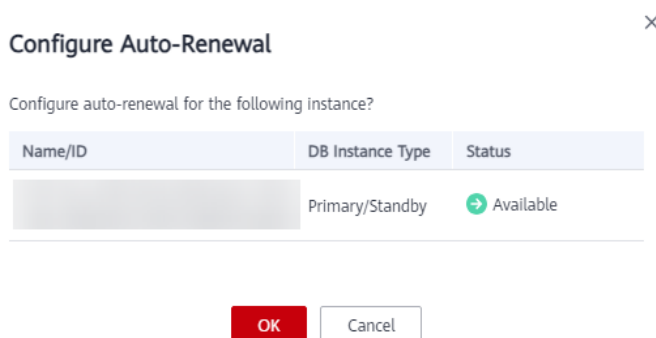
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the yearly/monthly instance and choose **More** > **Configure Auto-Renewal** in the **Operation** column.

**Step 5** In the displayed dialog box, confirm the DB instance information and click **OK**.

**Figure 3-5** Configuring auto-renewal



**Step 6** Configure **Renewal Option**, **Renewal Duration**, **New Auto-Renew Period**, and **Auto-renewals**, and click **OK**.

- If you select **Automatic** for **Renewal Option**, you need to configure **New Auto-Renew Period** and **Auto-renewals**.
- If you select **Manual** for **Renewal Option**, auto-renewal will be disabled and resources will not be renewed automatically before they expire.

- Step 7** In the **Billing Information** area on the **Basic Information** page of the instance, view the **Upon Expiration** field.
- If you select **Automatic** for **Renewal Option**, the field is **Automatic renewal**.
  - If you select **Manual** for **Renewal Option**, the field is **Enter grace period**.
- End

## 3.3 Changing the Billing Mode from Yearly/Monthly to Pay-per-Use

### Scenarios

You can change the billing mode of an instance from yearly/monthly to pay-per-use.



---


#### NOTICE

The pay-per-use billing mode is applied only after a yearly/monthly subscription expires and auto-renew is not in effect.

---

### Changing the Billing Mode from Yearly/Monthly to Pay-per-Use

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, locate the yearly/monthly instance that you want to change to pay-per-use instance and choose **More** > **Change to Pay-per-use** in the **Operation** column.
- Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Change** in the **Billing Mode** field.
- Step 5** On the displayed page, change the billing mode of the instance.
- Step 6** On the **Change to Pay-per-Use After Expiration** page, confirm the instance billing information and click **Change to Pay-per-Use**.
- Step 7** Wait until the billing mode is successfully changed and view the instance on the **Instances** page.

In the upper right corner of the instance list, click  to refresh the list. After the change completes, the instance status will change to **Available** and the billing mode will change to **Pay-per-use**.

----End

## 3.4 Changing the Billing Mode from Pay-per-Use to Yearly/Monthly

### Scenarios

If you want to use GaussDB(for MySQL) for a long time, you can change the billing mode of one or multiple instances from pay-per-use to yearly/monthly for a lower tariff. After the change, you can check whether the operation has taken effect in the order status.


### Constraints

- The billing of the primary node and read replicas for a pay-per-use instance cannot be changed separately to yearly/monthly.
- Pay-per-use instances in any of the following statuses cannot be changed to yearly/monthly instances: frozen, creation failed, changing instance specifications, scaling up, and creating read replicas.

### Changing the Billing Mode of a DB Instance

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the pay-per-use instance that you want to change to the yearly/monthly instance and choose **More** > **Change to Yearly/Monthly** in the **Operation** column. On the displayed page, select your desired storage space and click **Next**.

Alternatively, click the instance name to go to the **Basic Information** page. In the **Billing Information** area, click **Change** next to the **Billing Mode** field. On the displayed page, select your desired storage space and click **Next**.

**Figure 3-6** Changing the billing mode from pay-per-use to yearly/monthly




By default, the new storage space is rounded up to the nearest multiple of ten. For example, if the used storage space of your pay-per-use instance is 91 GB, the minimum storage space you can select when you change the billing mode to yearly/monthly is 100 GB.

**Step 5** Select how many months you want to renew the subscription for. The minimum duration is one month.

- If you do not need to modify your settings, click **Pay** to go to the payment page.
- If you are not sure about the settings, the system will reserve your order. You can choose **Billing Center > Unpaid Orders** in the upper right corner and pay or cancel the order. In addition, the instance status is **Changing to Yearly/Monthly. Payment incomplete. Pay Now.**

**Step 6** Select a payment method and click **Confirm**.

**Step 7** View the results on the **Instances** page.

In the upper right corner of the instance list, click  to refresh the list. After the change completes, the instance status will change to **Available** and the billing mode will change to **Yearly/Monthly**.

----End


## Changing the Billing Mode of Multiple Instance in Batches

### NOTE

Only pay-per-use instances can be changed to yearly/monthly instances.  
The status of pay-per-use instances must be **Available**.

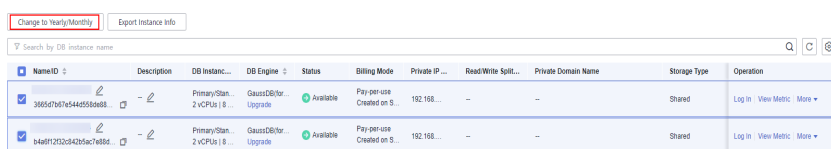
**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, select the desired instances and click **Change to Yearly/Monthly** above the instance list.

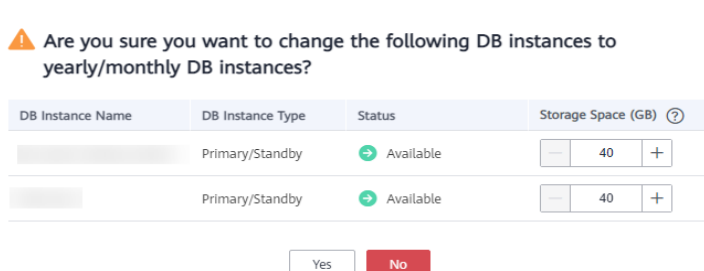
**Figure 3-7** Changing pay-per-use instances to yearly/monthly in batches (1)



NameID	Description	DB Instanc...	DB Engine	Status	Billing Mode	Private IP...	Read/Write Split...	Private Domain Name	Storage Type	Operation
3865767e444550e68		Primary/Stan... 2 vCPU   8...	GaussDB(for... Upgrade	Available	Pay-per-use Created on S...	192.168...	--	--	Shared	Log In   View Metric   More
ba8f72f2c342b5ac7e88d		Primary/Stan... 2 vCPU   8...	GaussDB(for... Upgrade	Available	Pay-per-use Created on S...	192.168...	--	--	Shared	Log In   View Metric   More

**Step 5** In the displayed dialog box, select storage space for the instances and click **Yes**.

**Figure 3-8** Changing pay-per-use instances to yearly/monthly in batches (2)




**Are you sure you want to change the following DB instances to yearly/monthly DB instances?**

DB Instance Name	DB Instance Type	Status	Storage Space (GB)
	Primary/Standby	Available	40
	Primary/Standby	Available	40

- Step 6** Select how many months you want to renew the subscription for. The minimum duration is one month.
- If you do not need to modify your settings, click **Pay** to go to the payment page.
  - If you are not sure about the settings, the system will reserve your order. You can choose **Billing Center > Unpaid Orders** in the upper right corner and pay or cancel the order. In addition, the instance status is **Changing to Yearly/Monthly. Payment incomplete. Pay Now.**

**Step 7** Select a payment method and click **Pay**.

**Step 8** 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 will change to **Yearly/Monthly**.

----End

## 3.5 Unsubscribing a Yearly/Monthly Instance

### Scenarios


To unsubscribe an instance billed on a yearly/monthly basis, you need to unsubscribe the order first. After you unsubscribe the instance order, all resources including read replicas of the instance are also deleted.

### Method 1

Unsubscribe a yearly/monthly instance on the **Instances** page.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance you want to unsubscribe and choose **More > Unsubscribe** in the **Operation** column. 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 unsubscription details, see [Unsubscription Rules](#).

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

**NOTICE**

1. Unsubscribed DB instances will be moved to the recycle bin, but will be permanently deleted after a length of time determined by the recycling policy. Automated backups are deleted, but manual backups are retained and still billed. To delete the manual backups, go to the **Backups** page on the console.
2. If you want to retain data, complete a manual backup before submitting the unsubscription request.

**Step 7** View the unsubscription result. After the instance order is successfully unsubscribed, the instance will be deleted.


----End

## Method 2

Unsubscribe a yearly/monthly instance on the **Billing Center** page.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.


**Step 4** In the upper right corner, click **Billing Center**.

**Step 5** In the navigation pane on the left, choose **Orders** > **Unsubscriptions**.

**Step 6** On the displayed page, select the order to be unsubscribed and click **Unsubscribe** in the **Operation** column.

- You can select **GaussDB(for MySQL)** in the **Service Type** to filter all GaussDB(for MySQL) orders.
- Alternatively, in the search box in the upper left corner of the order list, you can search for target orders by name, order No., or ID.

**Figure 3-9** Searching for orders



**Step 7** On the displayed page, confirm the order to be unsubscribed and select a reason. Then, click **Confirm**.

For unsubscription details, see [Unsubscription Rules](#).

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

---

**NOTICE**

1. Unsubscribed DB instances will be moved to the recycle bin, but will be permanently deleted after a length of time determined by the recycling policy. Automated backups are deleted, but manual backups are retained and still billed. To delete the manual backups, go to the **Backups** page on the console.
  2. If you want to retain data, complete a manual backup before submitting the unsubscription request.
- 

**Step 9** View the unsubscription result. After the instance order is successfully unsubscribed, the instance will be deleted.

----End

## APIs

[Deleting or Unsubscribing from a DB Instance](#)

# 4 Data Migration

---

## 4.1 Migration Solution Overview

You can migrate data from on-premises MySQL databases, on-premises Oracle databases, Huawei Cloud RDS for MySQL databases, MySQL databases built on other clouds to GaussDB(for MySQL) instances, or from a GaussDB(for MySQL) instance to another GaussDB(for MySQL) instance.

Data migration tools include Data Replication Service (DRS) and mysqldump. You are advised to use DRS because it is easy to use and can complete a migration task in minutes. DRS facilitates data transfer between databases, helping you reduce DBA labor costs and hardware costs.

DRS provides real-time migration and synchronization.

- **Real-time migration:** With DRS, you can migrate data from the source to the destination in real time. You just need to create a replication instance to connect to both the source and destination and configure objects to be migrated. DRS will help you compare metrics and data between the source and destination, so you can determine the best time to switch services to the destination while minimizing service downtime.
- **Real-time synchronization:** It refers to the real-time flow of workload data from the source to the destination through a synchronization instance while consistency of data is ensured. It is different from migration. Migration means moving entire data of a database to another. Synchronization refers to the continuous flow of data between different applications.

For more information, see [What Is DRS?](#)



## Migration Solutions

**Table 4-1** Migration solutions

Source Database	Solution	Documentation
RDS for MySQL	Use DRS to migrate table, database, or instance data.	<a href="#">From MySQL to GaussDB(for MySQL)</a>
GaussDB(for MySQL)	Use mysqldump to copy data.	<a href="#">Migrating Data to GaussDB(for MySQL) Using mysqldump</a>
<ul style="list-style-type: none"><li>On-premises MySQL databases</li><li>ECS-hosted MySQL databases</li></ul>	Use DRS to migrate data.	<a href="#">From ECS-hosted MySQL to GaussDB(for MySQL)</a>
MySQL databases on other clouds	Use DRS to migrate data.	<a href="#">From MySQL on Other Clouds to GaussDB(for MySQL)</a>

**Table 4-2** Synchronization solution

Source Database	Solution	Documentation
GaussDB(for MySQL)	Use DRS to synchronize data.	<a href="#">From GaussDB(for MySQL) to GaussDB(for MySQL)</a>
<ul style="list-style-type: none"><li>On-premises MySQL databases</li><li>ECS-hosted MySQL databases</li><li>MySQL databases on other clouds</li><li>RDS for MySQL</li></ul>	Use DRS to synchronize data.	<a href="#">From MySQL to GaussDB(for MySQL)</a>

Source Database	Solution	Documentation
<ul style="list-style-type: none"><li>On-premises Oracle databases</li><li>ECS-hosted Oracle databases</li></ul>	Use DRS to synchronize data.	<a href="#">From Oracle to GaussDB(for MySQL)</a>

## DRS Billing

- Real-time migration tasks are only billed based on a pay-per-use basis. Their configuration and traffic fees are free in the first seven days, lowering your costs for migrating data to the cloud.
- Real-time synchronization and DR tasks are billed based on a pay-per-use or yearly/monthly basis. Real-time migration and synchronization will provide long-term discounts, lowering your costs for data transfers.

For more information, see [Pay-Per-Use](#).

## 4.2 Migrating Data to GaussDB(for MySQL) Using DRS

You can use DRS to smoothly migrate self-managed databases to GaussDB(for MySQL) with no downtime.

For details, see [From MySQL to GaussDB\(for MySQL\) Primary/Standby](#) in the *Data Replication Service User Guide*.

### NOTE

- Currently, only MySQL-> GaussDB(for MySQL) migration and Oracle-> GaussDB(for MySQL) migration are supported.
- For MySQL -> GaussDB(for MySQL) migration, the source database can be GaussDB(for MySQL). You can migrate data from GaussDB(for MySQL) to GaussDB(for MySQL).
- By default, GaussDB(for MySQL) table names are case-insensitive.

## 4.3 Migrating Data to GaussDB(for MySQL) Using mysqldump

---

### NOTICE

Mysqldump is not recommended because it can result in a core dump in parallel backup scenarios. Mysqldump is recommended instead.

---

## Preparing for Data Migration

You can access a GaussDB(for MySQL) instance through a private network or a public network.

1. Prepare an ECS in the same VPC and subnet as the GaussDB(for MySQL) instance or bind an EIP to the GaussDB(for MySQL) instance.
  - To connect to an instance through a private network, an ECS has to be created first.  
[Purchase an ECS](#) and [log in to the ECS](#).
  - To connect to an instance through an EIP, you must:
    - i. Bind the EIP to the instance. For details, see [Binding an EIP](#).
    - ii. Ensure that the local device can access the EIP that has been bound to the instance.
2. Install the MySQL client on the prepared ECS or device.

### NOTE

The MySQL client version must be the same as or later than that installed on the GaussDB(for MySQL) instance. The MySQL database or client provides the mysqldump and mysql tools by default.

## Exporting Data

Before migrating data to GaussDB(for MySQL), data needs to be exported first.

### NOTICE

- The export tool must match the DB engine version.
- Database migration is performed offline. Before the migration, you must stop any applications using the source database.

**Step 1** Log in to the prepared ECS or device that can access the GaussDB(for MySQL) instance.

**Step 2** Use mysqldump to export the metadata into an SQL file.

### NOTICE

MySQL databases are required for GaussDB(for MySQL) management. When exporting metadata, do not specify **--all-database**, or the databases will be unavailable.

```
mysqldump --databases <DB_NAME> --single-transaction --order-by-primary  
--hex-blob --no-data --routines --events --set-gtid-purged=OFF -u <DB_USER>  
-p -h <DB_ADDRESS> -P <DB_PORT> |sed -e 's/DEFINER[ ]*=[ ]*[^\n]*\n/' -e 's/  
DEFINER[ ]*=.*/FUNCTION/FUNCTION/' -e 's/DEFINER[ ]*=.*/PROCEDURE/  
PROCEDURE/' -e 's/DEFINER[ ]*=.*/TRIGGER/TRIGGER/' -e 's/  
DEFINER[ ]*=.*/EVENT/EVENT/' > <BACKUP_FILE>
```

- **DB\_NAME** indicates the name of the database to be migrated.
- **DB\_USER** indicates the database username.
- **DB\_ADDRESS** indicates the database address.
- **DB\_PORT** indicates the database port.
- **BACKUP\_FILE** indicates the name of the file to which the data will be exported.

Enter the database password when prompted.

Example:

```
mysqldump --databases gaussdb --single-transaction --order-by-primary --hex-blob --no-data --routines --events --set-gtid-purged=OFF -u root -p -h 192.xx.xx.xx -P 3306 |sed -e 's/DEFINER[ ]*=[ ]*[^]*\*/\*/' -e 's/DEFINER[ ]*=.*/FUNCTION/FUNCTION/' -e 's/DEFINER[ ]*=.*/PROCEDURE/PROCEDURE/' -e 's/DEFINER[ ]*=.*/TRIGGER/TRIGGER/' -e 's/DEFINER[ ]*=.*/EVENT/EVENT/' > dump-defs.sql
```

Enter password:

After this command is executed, the **dump-defs.sql** file will be generated.

**Step 3** Use `mysqldump` to export the data into an SQL file.

---

**NOTICE**

MySQL databases are required for GaussDB(for MySQL) management. When exporting metadata, do not specify **--all-database**, or the databases will be unavailable.

---

```
mysqldump --databases <DB_NAME> --single-transaction --hex-blob --set-gtid-purged=OFF --no-create-info --skip-triggers -u <DB_USER> -p -h <DB_ADDRESS> -P <DB_PORT> -r <BACKUP_FILE>
```

For details on the parameters in the preceding command, see [Step 2](#).

Enter the database password when prompted.

Example:

```
mysqldump --databases gaussdb --single-transaction --hex-blob --set-gtid-purged=OFF --no-create-info --skip-triggers -u root -p -h 192.xx.xx.xx -P 3306 -r dump-data.sql
```

After this command is executed, the **dump-data.sql** file will be generated.

----End

## Importing Data

You can use a client to connect to the GaussDB(for MySQL) instance through an ECS or a device and then import the exported SQL files into that instance.

**NOTICE**

If the source database calls triggers, stored procedures, functions, or events, you must set **log\_bin\_trust\_function\_creators** to **ON** for the destination database before importing data.

**Step 1** Import metadata into the GaussDB(for MySQL) instance.

```
mysql -f -h <DB_ADDRESS> -P <DB_PORT> -u root -p < <BACKUP_DIR>/dump-  
defs.sql
```

- **DB\_ADDRESS** indicates the IP address of the GaussDB(for MySQL) instance.
- **DB\_PORT** indicates the port of the GaussDB(for MySQL) instance.
- **BACKUP\_DIR** indicates the directory where **dump-defs.sql** will be stored.

Example:

```
mysql -f -h 172.xx.xx.xx -P 3306 -u root -p < dump-defs.sql
```

Enter password:

**Step 2** Import data into the GaussDB(for MySQL) instance.

```
mysql -f -h <DB_ADDRESS> -P <DB_PORT> -u root -p < <BACKUP_DIR>/dump-  
data.sql
```

- **DB\_ADDRESS** indicates the IP address of the GaussDB(for MySQL) instance.
- **DB\_PORT** indicates the port of the GaussDB(for MySQL) instance.
- **BACKUP\_DIR** indicates the directory where **dump-data.sql** will be stored.

Example:

```
mysql -f -h 172.xx.xx.198 -P 3306 -u root -p < dump-data.sql
```

Enter password:

**Step 3** Use the MySQL tool to connect to the instance and view the result.

```
mysql> show databases;
```

In this example, the database named **my\_db** has been imported.

```
mysql> show databases;  
+-----+  
| Database      |  
+-----+  
| information_schema |  
| my_db          |  
| mysql          |  
| performance_schema |  
+-----+  
4 rows in set (0.00 sec)
```

----End

# 5 Serverless

## 5.1 Billing

This section describes all billing items of serverless instances.

### Billing Items

Table 5-1 Billing description

Billing Item	Billing Method
DB instance	You are billed for the selected instance specifications. The billing starts immediately after the instance is purchased. For details, see <a href="#">Billing Rules</a> .
Storage space	You do not need to select storage when purchasing a DB instance. Storage will be scaled up dynamically based on the amount of data that needs to be stored, and is billed hourly on a pay-per-use basis.
Backup space	GaussDB provides free backup storage equal to the amount of your used storage space. If the backup space usage exceeds 100% of your provisioned database storage, the additional part will be billed based on the backup pricing.
Public network traffic	GaussDB(for MySQL) instances are accessible from both private and public networks, but the traffic from public networks is billed.
HTAP instance	For billing details, see <a href="#">HTAP Billing</a> .

### Billing Rules

The billing unit is TCU. 1 TCU is approximately equal to the instances with 1 vCPU and 2 GB of memory. The detailed fees are as follows:

**Table 5-2** Billing rules (USD/TCU/hour)

Region	Unit Price (USD/TCU/Hour)
CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	0.063
AP-Singapore	0.19

## 5.2 Buying a Serverless DB Instance

### Scenarios

The capacities of serverless DB instances automatically change based on application requirements.

This section describes how to create a serverless DB instance on the GaussDB(for MySQL) console.

### Constraints

Currently, instances billed at a serverless billing can be purchased in the following regions:

- CN North-Beijing4
- CN East-Shanghai1
- CN South-Guangzhou
- AP-Singapore

### Billing

For details, see [Serverless Billing](#).

### Prerequisites

- You have [created a Huawei ID and enabled Huawei Cloud services](#).
- You can create an IAM user or user group on the IAM console and grant it specific operation permissions, to perform refined management on Huawei Cloud. For details, see [Creating a User and Granting Permissions](#).
- Your account balance is not below zero.

### Procedure

**Step 1** Go to the [Buy DB Instance](#) page.

**Step 2** On the displayed page, select **Serverless** for **Billing Mode**, configure required information, and click **Next**.

**Figure 5-1 Basic information**

The screenshot shows a configuration page for a GaussDB instance. At the top, there are three tabs for Billing Mode: 'Yearly/Monthly', 'Pay-per-use', and 'Serverless'. Below this is a note about serverless features. A 'Region' dropdown menu is set to '0'. A horizontal line separates this from the 'DB Instance Name' field, which has a help icon. Below it is a note about naming conventions for multiple instances. The configuration options are: DB Engine (GaussDB(for MySQL)), DB Engine Version (MySQL 8.0), DB Instance Type (Primary/Standby), Storage Type (Shared), AZ Type (Single-AZ, Multi-AZ), AZ (02, 03), and Time Zone (dropdown menu).

**Table 5-3 Basic information**

Parameter	Description
Region	<p>A region where the DB instance is located. You can change this on the creation page, or go back to the <b>Instances</b> page and change it in the upper left corner.</p> <p><b>NOTICE</b> Products in different regions cannot communicate with each other through a private network. After a DB instance is purchased, the region cannot be changed.</p>
DB Instance Name	<p>A name must start with a letter and consist of 4 to 64 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</p> <ul style="list-style-type: none"> <li>If you create multiple instances at a time, a hyphen (-) followed by a number with four digits will be appended to the instance name, starting with -0001. For example, if you enter <b>instance</b>, the first instance will be named instance-0001, the second instance-0002, and so on.</li> <li>The names for instances created in batches must consist of 4 to 59 characters. Only letters, digits, hyphens (-), and underscores (_) are allowed.</li> </ul>
DB Engine	GaussDB(for MySQL)
DB Engine Version	MySQL 8.0



Parameter	Description
DB Instance Type	Only primary/standby DB instances are supported. <b>Primary/Standby:</b> A primary/standby instance contains one primary node and one read replica. The primary node processes read and write requests, and the read replica processes only read requests. If the primary node becomes unavailable, GaussDB(for MySQL) automatically fails over to the read replica. Primary/standby instances apply to medium- and large-sized enterprises in the Internet, taxation, banking, and insurance sectors.
AZ Type	An AZ is a physical region where resources have their own independent power supply and networks. AZs are physically isolated but interconnected through an internal network. Some regions support both single-AZ and multi-AZ deployment and some only support single-AZ deployment. <ul style="list-style-type: none"> <li>• <b>Single-AZ:</b> The primary node and the read replica are deployed in the same AZ.</li> <li>• <b>Multi-AZ:</b> The primary node and the read replica are deployed in different AZs to ensure high reliability.</li> </ul>
Time Zone	You need to select a time zone for your instance based on the region hosting your instance. The time zone is selected during instance creation and cannot be changed after the instance is created.

**Figure 5-2** Specifications and storage of a serverless instance



**Table 5-4** Specifications and storage

Parameter	Description
Compute Configuration	Currently, only <b>Custom</b> is supported.
Compute Range	1 TCU is approximately equal to the instances with 1 vCPU and 2 GB of memory. Value range: <b>1</b> to <b>32</b> .

Parameter	Description
Nodes	By default, a DB instance contains one primary node and multiple read replicas. When purchasing a serverless DB instance, you can apply for zero or one read replica. Read replicas cannot be created after a DB instance is purchased.
Storage	It contains the system overhead required for inode, reserved block, and database operation. Storage will be scaled up dynamically based on the amount of data that needs to be stored, and is billed hourly on a pay-per-use basis.
Backup Space	GaussDB(for MySQL) provides free backup space equal to the amount of your used storage. After the free backup space is used up, you will be billed for the additional space on a pay-per-use basis.

Figure 5-3 Network

① Relationship among VPCs, subnets, security groups, and DB instances

VPC ①    [View In-use IP Address](#)

After the DB instance is created, the VPC cannot be changed. If you want to create a VPC, go to the VPC console. If you want to create DB instances in batches, the IP addresses are automatically assigned. Available IP addresses: 31762.

Security Group ①  [View Security Group](#)

In a security group, rules that authorize connections to DB instances apply to all DB instances associated with the security group.

▲ The security group rule must allow access from the 100.125.0.0/16 CIDR block.

Ensure that port 3306 of the security group allows traffic from your server IP address to the DB instance.

[Security Group Rules](#) [Add Inbound Rule](#)

---

Database Proxy  Disabled  Enabled ①

Proxy Mode  Read/Write  Read-only ①

Proxy Instance Specifications

---

Administrator root

Administrator Password  Keep your password secure. The system cannot retrieve your password.

Confirm Password

**Table 5-5** Network

Parameter	Description
VPC	<ul style="list-style-type: none"><li>Indicates a dedicated virtual network in which your DB instances are located. It isolates networks for different services. You can select an existing VPC or create a VPC. For details about how to create a VPC, see <a href="#">Creating a VPC</a>. If no VPC is available, GaussDB(for MySQL) allocates a VPC to you by default.</li></ul> <p><b>NOTICE</b> After a DB instance is created, the VPC cannot be changed.</p> <ul style="list-style-type: none"><li>A subnet provides dedicated network resources that are logically isolated from other networks for network security. A private IP address is automatically assigned when you create a DB instance. You can also enter an idle private IP address in the subnet CIDR block.</li></ul> <p><b>NOTE</b> Currently, GaussDB(for MySQL) does not support IPv6.</p>
Security Group	<p>Enhances security by controlling access to GaussDB(for MySQL) from other services. When you select a security group, you must ensure that it allows the client to access instances.</p> <p>If no security group is available or has been created, GaussDB(for MySQL) allocates a security group to you by default.</p> <p><b>NOTE</b> To ensure subsequent database connections and access, click <b>Add Inbound Rule</b> to allow all IP addresses to access your DB instance through port 3306 and over ICMP.</p> <p>For details, see <a href="#">Configuring Security Group Rules</a>.</p>

**Table 5-6** Database proxy

Parameter	Description
Database Proxy	<p>It is enabled by default. After the database proxy is enabled, you can use the IP address of a proxy instance to connect to your DB instance.</p> <p>When the number of nodes is 1, database proxy cannot be enabled.</p> <p><b>NOTE</b> To use this function, submit a service ticket by choosing <a href="#">Service Tickets &gt; Create Service Ticket</a>.</p>

Parameter	Description
Proxy Mode	You can select <b>Read/Write</b> or <b>Read-only</b> as needed. <ul style="list-style-type: none"><li>• <b>Read/Write:</b> All write requests are forwarded only to the primary node, and all read requests are forwarded to the selected nodes based on the read weights.</li><li>• <b>Read-only:</b> The primary node does not process write and read requests, and all read requests are forwarded to the selected read replicas based on read weights.</li></ul>
Proxy Instance Specifications	You can select the proxy instance specifications as needed.

**Table 5-7** Database configuration

Parameter	Description
Administrator	The default login name for the database is <b>root</b> .
Administrator Password	Must consist of 8 to 32 characters and contain at least three of the following: uppercase letters, lowercase letters, digits, and special characters (~!@#%^*_-=+?,)&\$). Enter a strong password and periodically change it to improve security and defend against threats such as brute force cracking attempts. Keep this password secure. If lost, the system cannot retrieve it. After a DB instance is created, you can reset this password. For details, see <a href="#">Resetting the Administrator Password</a> .
Confirm Password	Must be the same as <b>Administrator Password</b> .

**Table 5-8** Parameter template

Parameter	Description
Parameter Template	<p>Contains engine configuration values that can be applied to one or more instances. You can modify the instance parameters as required after the instance is created.</p> <p><b>NOTICE</b></p> <ul style="list-style-type: none"><li>If you use a custom parameter template when creating a DB instance, the following specification-related parameters in the custom template are not applied. Instead, the default values are used. "innodb_buffer_pool_size" "innodb_log_buffer_size" "max_connections" "innodb_buffer_pool_instances" "innodb_page_cleaners" "innodb_parallel_read_threads" "innodb_read_io_threads" "innodb_write_io_threads" "threadpool_size"</li><li>The value of <b>innodb_parallel_select_count</b> is determined by your instance specifications, instead of the parameter value you configured in the parameter template. The default value is <b>OFF</b> for instance with 16 vCPUs or less and <b>ON</b> for instances with more than 16 vCPUs.</li></ul> <p>After a DB instance is created, you can adjust it based on service requirements. For details, see <a href="#">Modifying a Parameter Template</a>.</p>
Table Name	<p>Specifies whether table names are case sensitive. This option cannot be changed later.</p> <ul style="list-style-type: none"><li><b>Case sensitive:</b> Table names are case sensitive.</li><li><b>Case insensitive:</b> Table names are case insensitive and are stored in lowercase letters by default.</li></ul>
Enterprise Project	<p>Only available for enterprise users. If you want to use this function, contact customer service.</p> <p>An enterprise project provides a way to manage cloud resources and enterprise members on a project-by-project basis.</p> <p>You can select an enterprise project from the drop-down list. The default project is <b>default</b>.</p>

**Table 5-9** Tags

Parameter	Description
Tag	This parameter is optional. Adding tags helps you better identify and manage your DB instances. A maximum of 20 tags can be added for each instance. After a DB instance is created, you can view its tag details on the <b>Tags</b> tab. For details, see <a href="#">Managing Tags</a> .

 **NOTE**

The performance of a DB instance depends on its configuration. Hardware configuration items include the instance specifications, storage type, and storage space.

**Step 3** Confirm your specifications.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Submit**.

**Step 4** To view and manage instances, go to the **Instances** page.

- During the creation process, the instance status is **Creating**. After the instance status becomes **Available**, the instance can be used.
- Automated backup is enabled by default during instance creation. After your instance was created, the backup policy cannot be disabled and a full backup will be automatically created.
- After the instance is created, you can confirm the DB instance type on the **Instances** page.
- After the instance is created, you can add a description.
- The default database port is **3306**, but you can change it after instance creation is complete.

For details, see [Changing a Database Port](#).

 **NOTE**

To ensure data and instance security, change the database port immediately after the instance is created.

----End

## APIs



- [Creating a DB Instance](#)
- [Querying DB Instances](#)

## 5.3 Changing the Compute Range of a Serverless Instance

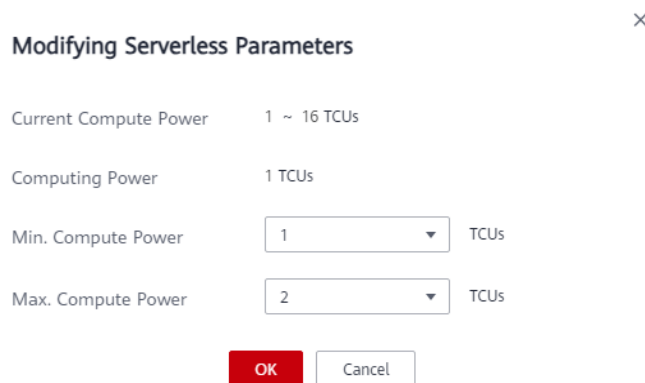
### Scenarios

After [buying a serverless instance](#), you can modify the minimum and maximum compute as needed.

### Procedure

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go the **Basic Information** page.
- Step 5** In the **Serverless Information** area, click **Change** next to **Compute Range**.
- Step 6** In the displayed dialog box, configure **Min. Compute Power** and **Max. Compute Power**, and click **OK**.

**Figure 5-4** Changing the compute range



----End

## 5.4 Changing Serverless Compute Resources

### Scenarios

Cloud Eye monitors CPU usage and memory usage of serverless instances. When certain conditions are met, Cloud Eye automatically scales up or down the serverless compute resources.

### Constraints

When a compute change is triggered, if resources are insufficient, a node with desired specifications is created on a physical machine with sufficient resources. If the resources on the primary node are insufficient, there will be a primary/standby switchover.

### Scaling up Compute

Scale-up is triggered when either of the following conditions is met:

- The CPU usage is greater than 90% for 5 seconds and the interval between two scale-up operations is at least 5 seconds.
- The memory usage is greater than 90% for 5 seconds and the interval between two scale-up operations is at least 5 seconds.

### Scaling down Compute

Scale-down is triggered when either of the following conditions is met:

- The CPU usage is less than 30% for 15 seconds and the interval between two scale-down operations is at least 15 seconds.
- The memory usage is less than 30% for 15 seconds, the CPU usage is less than 70% for 15 seconds, and the interval between two scale-down operations is at least 15 seconds.



# 6 Database Proxy (Read/Write Splitting)

---

## 6.1 Introducing Read/Write Splitting

Read/write splitting enables read and write requests to be automatically routed through a read/write splitting address. You can [create a proxy instance](#) after a GaussDB(for MySQL) instance is created. Thanks to the IP address of a proxy instance, write requests are automatically routed to the primary node and read requests are routed to read replicas and the primary node by user-defined weights. Currently, you can create up to four proxy instances to for complex services that require isolation. You can use the IP address of any proxy instance to connect to your GaussDB(for MySQL) instance based on service requirements.

### Billing

Proxy instances are free of charge.

### Constraints

- If the kernel version of the GaussDB(for MySQL) instance is earlier than 2.0.42.230601, only one proxy instance can be created.
- If the kernel version of the GaussDB(for MySQL) instance is 2.0.42.230601 or later, up to four proxy instances can be created.
- Read/write splitting can be enabled only when at least one read replica is created.
- After read/write splitting is enabled, the port number and read/write private IP address of your GaussDB(for MySQL) instance cannot be changed.
- Read/write splitting does not support the compression protocol.
- If multi-statements are executed, all subsequent requests will be routed to the primary node. To restore the read/write splitting function, you need to disconnect the connection between applications and the read/write splitting address and establish a connection again.
- When the IP address of a proxy instance is used, all transaction requests are routed to the primary node (you can use the transaction splitting feature to route read requests prior to write operations in a transaction to read replicas).

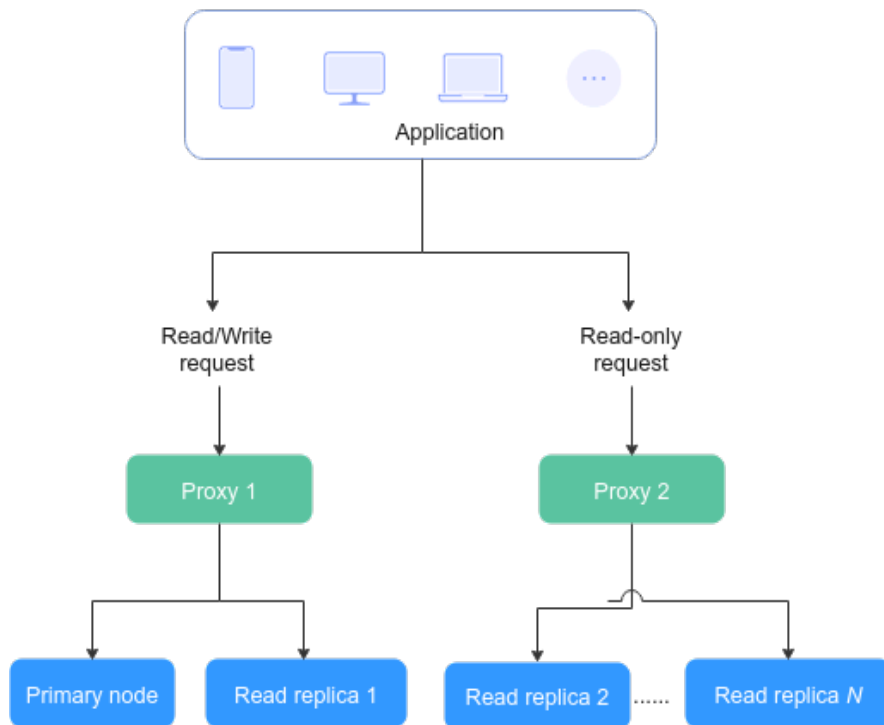
The non-transaction read consistency is not ensured. To ensure read consistency, encapsulate the read requests into a transaction.

- When the IP address of a proxy instance is used, you can run **show processlist** command on the proxy instance or GaussDB(for MySQL) instance. If **show processlist** is executed on a proxy instance, only the services delivered through proxy nodes are displayed.
- When a proxy node is abnormal, running **show processlist** or **kill** on the proxy instance may take a long time, but services are not affected.
- After a proxy node is deleted, workload on the deleted proxy node may be displayed when **show processlist** is executed on the proxy instance.
- When **kill** is executed, error information such as timeout may be displayed occasionally. You can run **show processlist** again to check whether the services are killed successfully.
- If a proxy node is abnormal, there may be frame freezing for 2 seconds when you run **show processlist** on the proxy instance. The result will be returned normally.
- Proxy instances do not support the transaction isolation level READ-UNCOMMITTED.
- To create proxy instances, ensure that the data in a single column of a table cannot exceed 16 MB.
- When proxy instances are used, the total size of statements cannot exceed 100 MB. Otherwise, the proxy instances need to consume too many resources to parse the SQL statements.

## Scenarios

When enabling read/write splitting, you need to associate the nodes (including the primary node and read replica) with proxy instances.

- Different applications can connect to the GaussDB(for MySQL) instance through the IP addresses of different proxy instances. Read requests are routed to the proxy instances that applications connect to. You can also associate nodes with or remove nodes from proxy instances.
- A primary node or read replica can be associated with multiple proxy instances at the same time.
- In the read/write mode, all write requests are routed to the primary node, and read requests are routed to each node based on the read weights or active connections.
- In the read-only mode, only read requests can be routed to read replicas based on the read weights and active connections.
- By default, proxy instances provide overload protection to prevent server OOM (out of memory) due to heavy pressure when users perform operations on large result sets. This function is enabled by default and does not need to be configured separately. If the pressure is caused by the database kernel, you need to configure flow control policy.



## 6.2 Introducing Consistency Levels

GaussDB(for MySQL) provides two consistency levels to meet requirements in different scenarios.

- Eventual consistency (default)
- Session consistency

### Constraints

- To configure consistency level, the kernel version of your GaussDB(for MySQL) instance must be 2.0.28.1 or later.
- To use session consistency, the kernel version of your proxy instance must be 2.7.4.0 or later.

### Eventual Consistency


After a proxy instance is created, requests for SELECT operations are routed to different nodes based on their read weights. Because there is a replication delay between the primary node and each read replica and the replication delay varies for different read replicas, the result returned by each SELECT statement may be different when you repeatedly execute a SELECT statement within a session. In this case, only eventual consistency is ensured.

### Session Consistency

To eliminate data inconsistencies caused by eventual consistency, session consistency is provided. Session consistency ensures the result returned by each SELECT statement in a session is the data that was updated after the last write request.

Proxy instances record the log sequence number (LSN) of each node and session. When data in a session is updated, a proxy instance records the LSN of the primary node as a session LSN. When a read request arrives subsequently, the database proxy compares the session LSN with the LSN of each node and routes the request to a node whose LSN is at least equal to the session LSN. This ensures session consistency.

**Figure 6-1** Principle of session consistency



#### **NOTE**

In session consistency, if there is significant replication delay between the primary node and read replicas and the LSN of each read replica is smaller than the session LSN, requests for SELECT operations will be routed to the primary node. In this case, loads on the primary node are heavy and instance performance suffers.

## 6.3 Create a Proxy Instance

Proxy instances enable read and write requests to be automatically routed through its IP address. Currently, you can create up to four proxy instances to for complex services that require isolation. You can use the corresponding IP address of any proxy instance to connect to your GaussDB(for MySQL) instance based on service requirements.

This section describes how to create a proxy instance.


## Constraints


Proxy instances cannot be created if the GaussDB(for MySQL) kernel version is:

- From 2.0.26.2 to 2.0.28.3
- 2.0.29.1

## Procedure

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Figure 6-2** Going to the Database Proxy page



**Step 6** Click **Create Proxy Instance**.

**Step 7** In the displayed dialog box, configure required parameters by referring to [Table 6-1](#) and click **OK**.

### NOTE

After a proxy instance has been created, you can click **Create Proxy Instance** in the **Database Proxy** page to add a new proxy instance.

**Figure 6-3** Creating a proxy instance

**Create Proxy Instance**

Currently, the database proxy is free. We will notify you of the pricing details later.

Proxy Instance Name: proxy-bb25

Proxy Mode: Read/Write (selected), Read-only

Consistency Level: Eventual consistency (selected), Session consistency

Proxy Instance Specifications: 2 vCPUs | 4 GB (Kunpeng general computin...)

Proxy Instance Nodes: 2

Database Nodes

Available Nodes(2)	
Node Type	Name/ID
<input type="checkbox"/>	Primary
<input type="checkbox"/>	Replica

Selected Nodes(0)			
Node Type	Name/ID	Read Weight	Op...
No nodes selected.			

OK Cancel

**Table 6-1** Parameter description

Parameter	Description
Proxy Instance Name	The name can consist of 4 to 64 characters and must start with a letter. Only letters (case-sensitive), digits, hyphens (-), and underscores (_) are allowed.

Parameter	Description
Proxy Mode	<ul style="list-style-type: none"><li>● <b>Read/Write:</b> All write requests are routed only to the primary node, and all read requests are routed to the selected nodes based on the read weights or active connections. The default read weight of the primary node is <b>100</b>.</li><li>● <b>Read-only:</b> All read requests are routed to the selected read replicas based on the read weights or active connections. The read requests will not be routed to the primary node.</li></ul> <p><b>NOTE</b></p> <ul style="list-style-type: none"><li>- Only read requests are supported. If write requests are forwarded to the selected nodes, an error message is displayed.</li><li>- This mode offloads the pressure of the primary node by routing all read requests to read replicas.</li><li>- DDL, DML, and temporary table operations are not supported in the read-only mode.</li></ul>
Consistency Level	<p>The consistency level can be configured only when the kernel version of your GaussDB(for MySQL) instance is 2.0.28.1 or later.</p> <p>Value:</p> <ul style="list-style-type: none"><li>● <b>Eventual consistency</b> After a proxy instance is created, requests for SELECT operations are routed to different nodes based on their read weights. Because there is a replication delay between the primary node and each read replica and the replication delay varies for different read replicas, the result returned by each SELECT statement may be different when you repeatedly execute a SELECT statement within a session. In this case, only eventual consistency is ensured.  To offload read requests from the primary node to read replicas, you can select eventual consistency.</li><li>● <b>Session consistency</b> To eliminate data inconsistencies caused by eventual consistency, session consistency is provided. Session consistency ensures the result returned by each SELECT statement in a session is the data that was updated after the last write request.  To use session consistency, the kernel version of your proxy instance must be 2.7.4.0 or later.</li></ul> <p>For more information about consistency levels, see <a href="#">Introducing Consistency Levels</a>.</p>

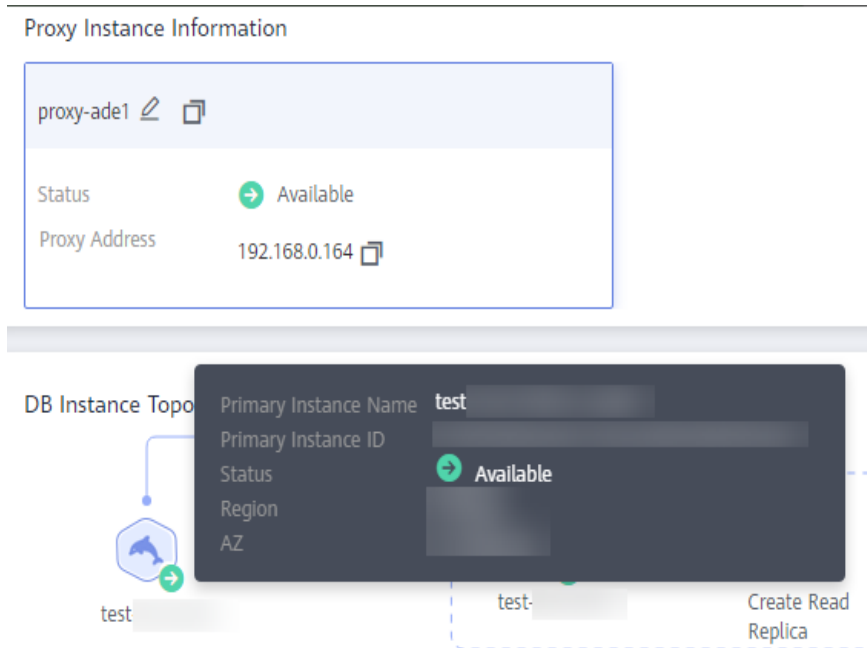
Parameter	Description
Routing Policy	<p>To configure routing policies, submit an application by choosing <a href="#">Service Tickets &gt; Create Service Ticket</a> in the upper right corner of the console.</p> <p>Value:</p> <ul style="list-style-type: none"><li>● <b>Weighted:</b> Read requests are assigned to nodes based on the weights you specify.</li><li>● <b>Load balancing:</b> Read requests are assigned to nodes with fewer active connections. To use load balancing, the kernel version of your proxy instance must be 2.22.07.000 or later.</li></ul>
Proxy Instance Specifications	<p>You can select the proxy instance specifications as needed.</p> <ul style="list-style-type: none"><li>● Kungpeng general computing-plus: 2 vCPUs   4 GB, 4 vCPUs   8 GB, and 8 vCPUs   16 GB</li><li>● General-enhanced: 2 vCPUs   4 GB, 4 vCPUs   8 GB, and 8 vCPUs   16 GB</li></ul>
Subnet	<p>You can specify a subnet when read/write splitting is enabled.</p> <ul style="list-style-type: none"><li>● Submit an application by choosing <a href="#">Service Tickets &gt; Create Service Ticket</a> in the upper right corner of the console.</li><li>● If the subnet where the GaussDB(for MySQL) instance is located is a secondary CIDR block, cross-subnet read/write splitting is not supported. The GaussDB(for MySQL) instance must be in the same subnet as proxy instances.</li></ul>
Proxy Instance Nodes	<p>You can enter an integer from 2 to 32. The default value is 2.</p> <p>Number of recommended proxy instance nodes = (Number of vCPUs of the primary node + Total number of vCPUs of all read replicas)/(4 x Number of vCPUs of the proxy instance), rounded up.</p>



Parameter	Description
Database Nodes	<p>Select the nodes that need to be associated with the database proxy to process read requests.</p> <ul style="list-style-type: none"> <li>If <b>Routing Policy</b> is <b>Load balancing</b>, read requests are assigned to nodes with fewer active connections. You do not need to configure read weights for nodes.</li> <li>If <b>Routing Policy</b> is <b>Weighted</b>, read requests are assigned to nodes based on the weights you specify. You need to set the read weights for the primary node and read replicas. For example, read weights assigned to one primary node and two read replicas are 100, 200, and 200, respectively. In the read/write mode, the primary node and two read replicas process read requests in the ratio of 1:2:2. The primary node processes 20% of read requests, and each read replica processes 40% of read requests. Write requests are automatically routed to the primary node.</li> </ul> <p>In the read-only mode, the read weight of the primary node does not take effect, and the two read replicas process 50% of read requests, respectively.</p>

**Step 8** Click the name of the GaussDB(for MySQL) instance to go to the **Basic Information** page. Click the created proxy instance and view the nodes associated with the proxy instance in the **DB Instance Topology** area. You can move the pointer to the node name to view node details.

**Figure 6-4** Information about the nodes associated with a proxy instance



----End

## APIs

- [Enabling Database Proxy](#)
- [Querying Database Proxy Instances](#)
- [Querying Database Proxy Specifications](#)
- [Disabling Database Proxy](#)

## 6.4 Configuring Connection Pools

### Scenarios

A session-level connection pool helps reduce the database load caused by frequent establishment of short connections.

Connection Pool is disabled by default. You can enable a session-level connection pool.

### Constraints

- To configure a connection pool, submit an application by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the console.
- The kernel version of proxy instances must be 2.22.07.000 or later.

### How a Session-Level Connection Pool Works


A session-level connection pool is suitable for short connections.

When your client disconnects from your database, the system checks whether the connection is idle. If it is, the system places the connection in the connection pool of a proxy instance and retains the connection for a short period of time. When your client re-initiates a connection, any available connection in the connection pool is used, reducing the overhead of establishing a new connection to the database. If no connections are available in the connection pool, a new connection will be established.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.


**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** Click the name of a proxy instance.

**Step 7** On the **Basic Information** page, click **Change** next to **Connection Pool**.

**Step 8** Set **Connection Pool** to **Session level** and click **OK**.

**Figure 6-5** Configuring a connection pool



----End

## APIs

- [Changing the Connection Pool Type of a Database Proxy Instance](#)
- [Querying Database Proxy Instances](#)
- [Querying Database Proxy Specifications](#)

## 6.5 Configuring Transaction Splitting

### Scenarios

In most cases, a proxy instance sends all requests in transactions to the primary node to ensure transaction correctness. However, in some frameworks, all requests are encapsulated into transactions that are not automatically committed using **set autocommit=0**. This causes heavy loads on the primary node.

### Constraints

- The kernel version of the proxy instances must be 2.3.9.5 or later.
- Transaction splitting is only available for instances whose transaction isolation level is READ UNCOMMITTED or READ COMMITTED. The default isolation level is REPEATABLE READ.
- To enable transaction splitting, the proxy mode must be set to read/write.

### Function

Proxy instances support transaction splitting. With transaction splitting is enabled, GaussDB(for MySQL) can route the read requests prior to write operations in a transaction to read replicas, reducing the pressure on the primary node.

Transaction splitting is disabled by default.


After transaction splitting is enabled and **autocommit** is set to **0**, GaussDB(for MySQL) starts a transaction only for write requests. Before the transaction starts, read requests are routed to read replicas through load balancers.


## Precautions

After transaction splitting is enabled, the transaction isolation level can only be changed to READ-UNCOMMITTED or READ-COMMITTED. To change the isolation level to a higher level, disable the function.

## Procedure

**Step 1** Log in to the management console.


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** Click the name of a proxy instance.

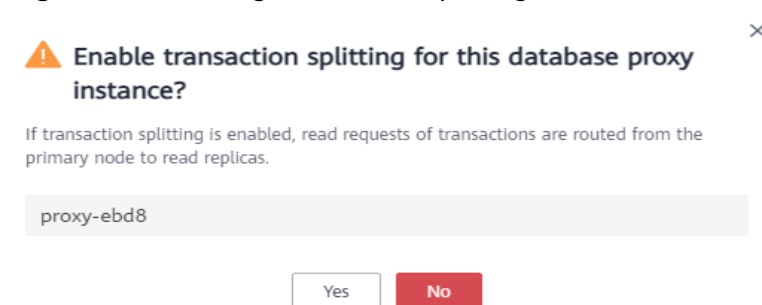
**Step 7** On the **Basic Information** page, click  next to **Transaction Splitting**.

**Figure 6-6** Configuring transaction splitting



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

**Figure 6-7** Enabling transaction splitting



----End

### NOTE

Transaction splitting takes effect only for new connections established after this function is enabled or disabled.

## APIs

- [Enabling or Disabling Transaction Splitting of a Database Proxy Instance](#)
- [Querying Database Proxy Instances](#)
- [Querying Database Proxy Specifications](#)
- [Disabling Database Proxy](#)

## 6.6 Configuring a Routing Policy

Proxy instances support weighted and load balancing routing policies.

To configure a routing policy, you need to:


- Create a proxy instance. For details, see [Create a Proxy Instance](#).
- Select a routing policy by referring to this section.


### Constraints

- To configure routing policies, submit an application by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the console.
- To use the load balancing policy, the kernel versions of proxy instances must be 2.22.07.000 or later. To upgrade a kernel version, see [Upgrading the Kernel Version of a Proxy Instance](#).

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** On the **Database Proxy** page, click the name of a proxy instance.

**Step 7** On the **Basic Information** page, click **Change** next to **Routing Policy**.

**Step 8** In the displayed dialog box, select a routing policy.

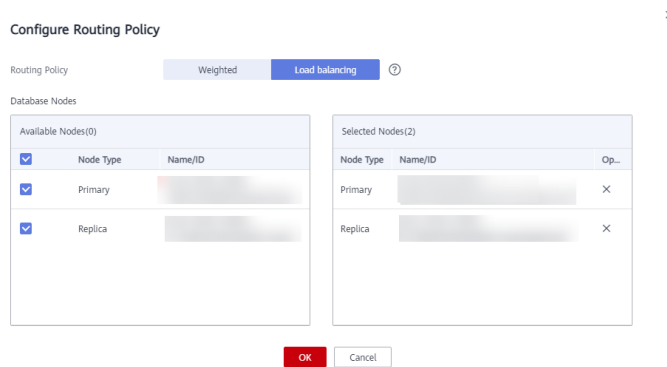
- **Weighted:** Read requests are assigned to nodes based on the weights you specify.
- **Load balancing:** Read requests are assigned to nodes with fewer active connections. In load balancing policy, you do not need to configure the weights of nodes.

**NOTE**

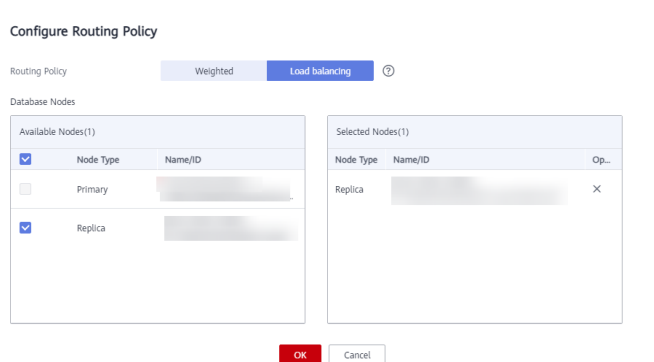
The proxy mode of a proxy instance affects read requests assigned to different nodes.

- Read-only mode: All read requests are assigned to the selected read replica based on the routing policy and weights you specify, but not to the primary node.
- Read/write mode: All read requests are assigned to the selected nodes (including primary nodes and read replicas) based on the routing policy and weights you specify.

**Figure 6-8** Changing the routing policy of a proxy instance in read/write mode



**Figure 6-9** Changing the routing policy of a proxy instance in read-only mode



----End

## 6.7 Assigning Read Weights

After read/write splitting is enabled, you can assign read weights as required. You can also adjust the number of nodes and their read weights when configuring a routing policy.

### Description


- After read/write splitting is enabled, you can assign read weights for the primary node and read replicas.
- The default read weight of the primary node is 0. The higher read weight the primary node is assigned, the more read requests it can process.

- When the read weights of all nodes are 0, services are not affected. In this case, the primary node processes all read and write requests by default.
- The weight of a read replica ranges from 0 to 1000.
- Newly created read replicas are automatically associated with proxy instances and their read weights are 0 by default. For details about the weight assignment rules, see [Introducing Read Weight Assignment Rules](#).

## Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

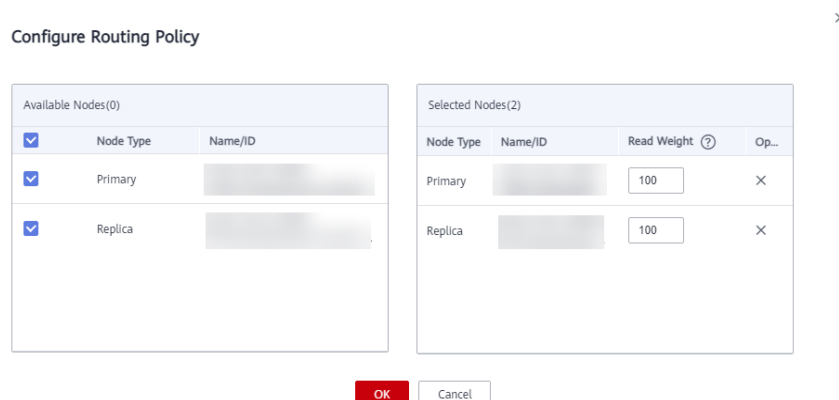
**Step 5** Click the name of a proxy instance.

**Step 6** On the **Basic Information** page, click **Change** next to **Routing Policy**.

**Step 7** In the displayed dialog box, select nodes for which you want to assign weights on the left and assign the weights on the right.

- Different applications can connect to the GaussDB(for MySQL) instance through different proxy addresses. Read and write requests are forwarded to associated nodes. You can also add nodes to or remove nodes from proxy instances.
- In the read/write mode, all write requests are routed to the primary node, and read requests are routed to each node based on the read weights.
- In the read-only mode, only read requests can be routed to read replicas based on the read weights.
- In the load balancing policy, read requests are routed only to the selected nodes. The weights of nodes cannot be modified.
- In the weighted policy, you can set weights for selected nodes.

**Figure 6-10** Setting weights of nodes (in the weighted policy)

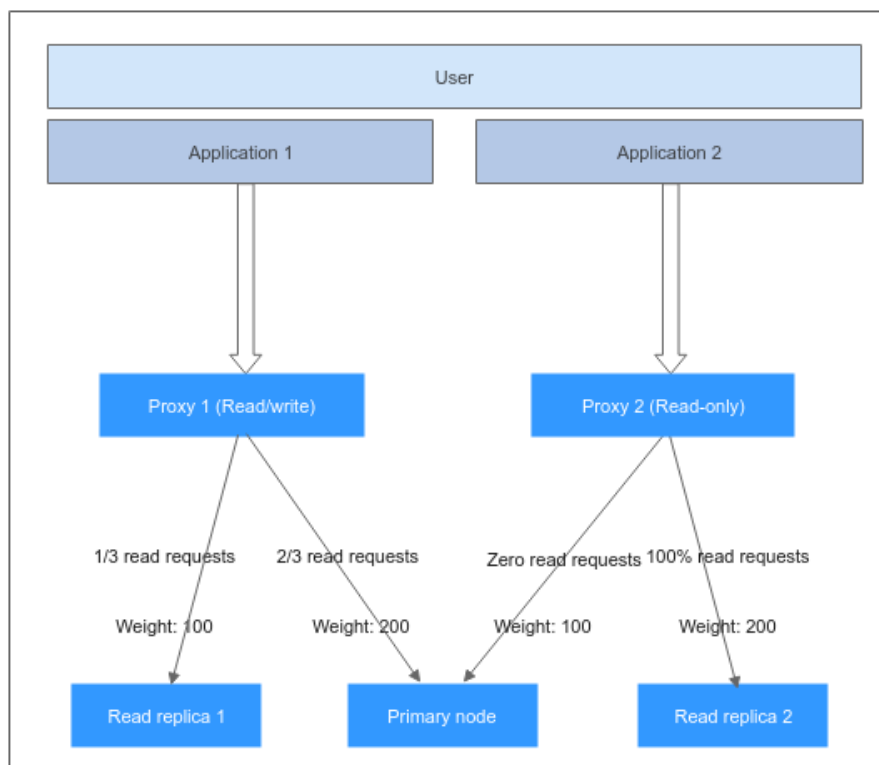


Example:

For example, a DB instance contains one primary node and two read replicas, and two proxy instances have been created.

- Proxy instance 1 is in the read/write mode. The primary node and read replica 1 are connected to proxy instance 1 and assigned with a read weight of 200 and 100, respectively. They process read requests in the ratio of 2:1, that is, the primary node processes 2/3 read requests and read replica 1 processes 1/3 read requests. Write requests are automatically routed to the primary node.
- Proxy instance 2 is in read-only mode. The primary node and read replica 2 are connected to proxy instance 2 and assigned with a read weight of 100 and 200, respectively. In this case, the weight of the primary node does not take effect, and read replica 2 processes all read requests.

**Figure 6-11** Ratio of read requests processed by each node in multiple proxy instances



**NOTE**

- When there are multiple proxy instances, newly created read replicas are automatically associated with proxy instances and their read weights are 0 by default.
- After a read replica is deleted, its weight is automatically removed while the weights of other read replicas remain unchanged.

----End

## APIs

### Assigning Read Weights



## 6.8 Configuring Multi-Statement Processing Mode



### Prerequisites

To configure multi-statement processing mode on the console, submit an application by choosing [Service Tickets > Create Service Ticket](#).

### Scenarios

When you enable [multi-statement execution](#) for a proxy instance, you can set multi-statement processing mode to **Strict** (default value), **Loose**, or **Parse**. For details, see [Multi-Statement Processing Mode](#).

### Procedure

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Database Proxy**, select a proxy instance and click its name.
- Step 6** In the navigation pane on the left, choose **Parameter Modification**, change the value of **multiStatementType** to **Strict**, **Loose**, and **Parse**. For details, see [Multi-Statement Processing Mode](#).
- Step 7** Click **Save** to save your change. In the displayed dialog box, click **Yes**.

----End

### Multi-Statement Processing Mode

The changed multi-statement processing mode applies to your proxy instance immediately. You do not need to reboot the proxy instance. If a read/write splitting connection fails due to a multi-statement execution, changing the multi-statement processing mode will not restore the connection. You will need to reconnect the connection manually.

- **Strict** (default mode)  
If a request containing multiple statements is routed to the primary node, the subsequent requests are all routed to the primary node. Read/write splitting can be restored only after you disconnect the current connection and reconnect it.  
Your proxy instances will not parse these statements, so the performance is better. It is suitable for short connections.

- **Loose**  
If a request containing multiple statements is routed to the primary node, the subsequent requests of the current connection can still be routed to the primary node or read replicas.  
Your proxy instances will not parse these statements, so the performance is better. It is good for when multiple statements contain only DML SQL statements and do not contain operations like setting session variables, creating temporary tables, creating stored procedures, or executing uncommitted transactions.
- **Parse**  
If a request containing multiple statements is routed to the primary node, your database proxy instance parses these statements and determines whether to restore read/write splitting for subsequent requests of the current connection based on the operations in the SQL statements. For details about operations in SQL statements, see [Parse Description](#).  
Parsing statements affects the proxy instance performance. The degree of the impact depends on the length and complexity of statements. It is recommended that the statements be less than 100 MB.

## Parse Description

If multi-statements contain the operations, listed here, all subsequent requests are routed to the primary node. To restore read/write splitting, you need to disconnect the connection and then re-establish it.

- Creating temporary tables
- Creating stored procedures
- Executing uncommitted transactions (for example, **begin** is executed but **commit** or **rollback** is not executed)
- Executing complex or special syntax. These statements will not be parsed.

## 6.9 Changing the Specifications of a Proxy Instance


### Constraints

- You can change the proxy instance specifications only when the statuses of the DB instance, primary node, read replicas, and proxy instance are **Available**.
- A proxy instance cannot be deleted when its CPU and memory specifications are being changed.

### Procedure


**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Database Proxy**.
- Step 6** On the **Database Proxy** page, locate the desired proxy instance and click **Change Specifications** in the **Operation** column.

**Figure 6-12** Going to the proxy instance list



You can also click the proxy instance name. In the **Proxy Instance Information** area, click **Change** next to the **Specifications** field.

**Figure 6-13** Changing proxy instance specifications (1)

### Proxy Instance Information

Proxy Instance Name  ✎ 📄

DB engine version 2.23.12.000 Upgrade

Consistency Level ? Eventual consistency ✎

Specifications 2 vCPUs | 4 GB Change

- Step 7** In the displayed dialog box, select new specifications and click **OK**. You can reduce or expand the specifications as required.

**Figure 6-14** Changing proxy instance specifications (2)

### Change Specifications

Proxy Instance ID

Current Proxy Instance Specifications 2 vCPUs | 4 GB

New Proxy Instance Specifications

Proxy Instance Nodes 2

OK Cancel

**Step 8** View the new specifications on the **Database Proxy** page.

----End

## APIs

- [Changing Specifications of a Database Proxy Instance](#)
- [Querying Database Proxy Instances](#)
- [Querying Database Proxy Specifications](#)

# 6.10 Changing the Number of Proxy Instance Nodes

## Scenarios

You can change the number of proxy nodes as required.

## Prerequisites

- Read/write splitting has been enabled.
- The primary node and proxy instance are all available.


## Constraints

There can be 2 to 32 proxy nodes.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

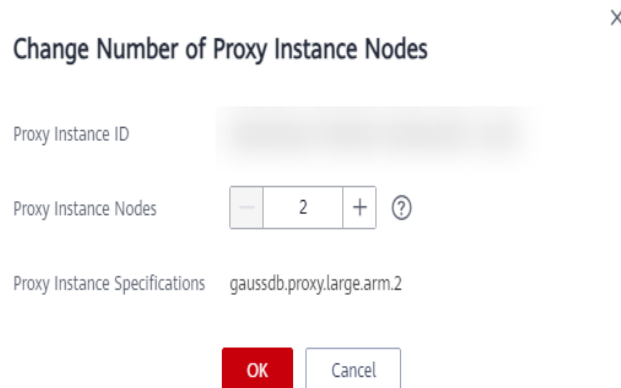
**Step 5** In the navigation pane on the left, choose **Database Proxy**. Click the proxy instance name.

**Step 6** In the **Proxy Instance Information** area, click **Change** next to the **Proxy Instance Nodes** field.

### NOTE

Number of recommended proxy instance nodes = (Number of vCPUs of the primary node + Total number of vCPUs of all read replicas)/(4 x Number of vCPUs of the proxy instance), rounded up.

**Step 7** In the displayed dialog box, set the number of proxy instance nodes and click **OK**.

**Figure 6-15** Changing the number of proxy nodes

----End

## APIs

- [Adding Database Proxy Nodes](#)
- [Querying Database Proxy Instances](#)

## 6.11 Upgrading the Kernel Version of a Proxy Instance

### Scenarios

You can manually upgrade your database proxy instance to the latest kernel version to improve performance, add new functions, and fix problems.

A kernel version can be upgraded in either of the following ways:


- Upon submission: The system **upgrades the proxy instance version** immediately after you submit the upgrade request.
- In maintenance window: The system upgrades the proxy instance version during a maintenance window you specify. For details about how to change the maintenance window, see [Changing a Maintenance Window](#).


### Precautions

Intermittent disconnections occur during an upgrade. The time required to complete the upgrade depends on how many proxy instances there are. Perform the upgrade during off-peak hours.

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Database Proxy**.
- Step 6** Click the name of a proxy instance to go to the **Basic Information** page.
- Step 7** In the **Proxy Instance Information** area, click **Upgrade** in the **DB Engine Version** field.
- Step 8** In the displayed dialog box, select a scheduled time and click **OK**.
- Upon submission: The system upgrades the proxy instance to the latest version immediately after you submit the request. You can view the task progress in **Task Center > Instant Tasks**.
  - In maintenance window: The system upgrades the proxy instance to the latest version during a maintenance window. You can view the task progress in **Task Center > Scheduled Tasks**.
- End



## 6.12 Using a Private Domain Name for a Proxy Instance

You can use a private network domain name to connect to a proxy instance.

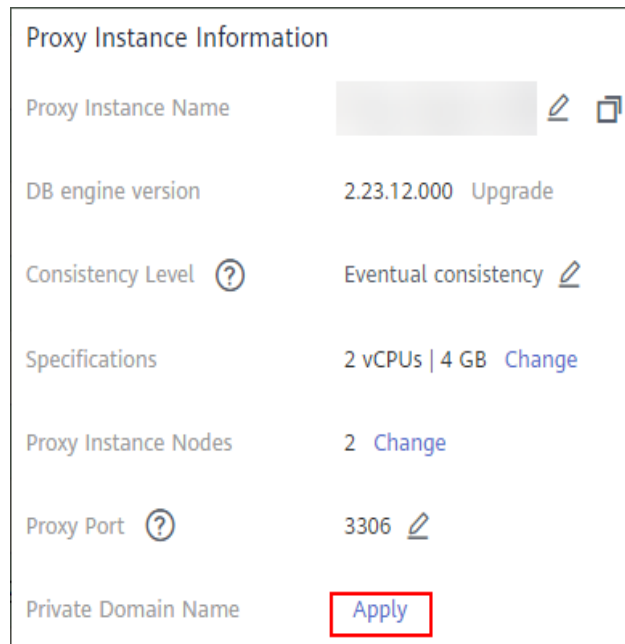
### Constraints

To apply for a private domain name, submit an application by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the console.

### Applying for a Private Domain Name for a Proxy Instance


- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Database Proxy**.
- Step 6** Click the name of a proxy instance.
- Step 7** In the **Proxy Instance Information** area on the **Basic Information** page, click **Apply** in the **Private Domain Name** field.

**Figure 6-16** Applying for a private domain name (1)



**Step 8** Click **OK**.

**Figure 6-17** Applying for a private domain name (2)




**Step 9** In the **Private Domain Name** field, view the generated private domain name.

----End

## Changing a Private Domain Name for a Proxy Instance

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** In the **Proxy Instance Information** area on the **Basic Information** page, click **Change** in the **Private Domain Name** field.

**Step 7** In the displayed dialog box, enter a new domain name and click **OK**.

 **NOTE**


- Only the prefix of a private domain name can be modified.
- The prefix of a private domain name contains 8 to 63 characters, and can include only lowercase letters and digits.
- The new private domain name must be different from existing ones.

----End

## Deleting a Private Domain Name for a Proxy Instance

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** In the **Proxy Instance Information** area on the **Basic Information** page, click **Delete** in the **Private Domain Name** field.

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

----End

## 6.13 Changing the IP Address of a Proxy Instance

### Scenarios

You can change the IP address of a proxy instance.

### Precautions



Changing a proxy address will interrupt database connections and services. Perform the operation during off-peak hours or when services are stopped.

### Constraints

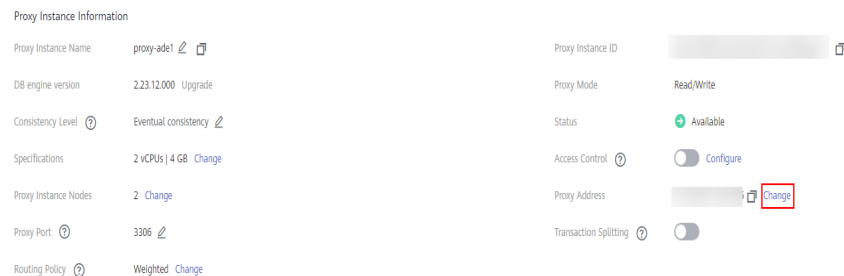
- The new IP address is not in use and must be in the same subnet as the GaussDB(for MySQL) instance.
- To change a proxy address, you need to contact customer service to apply for permissions required.



## Procedure

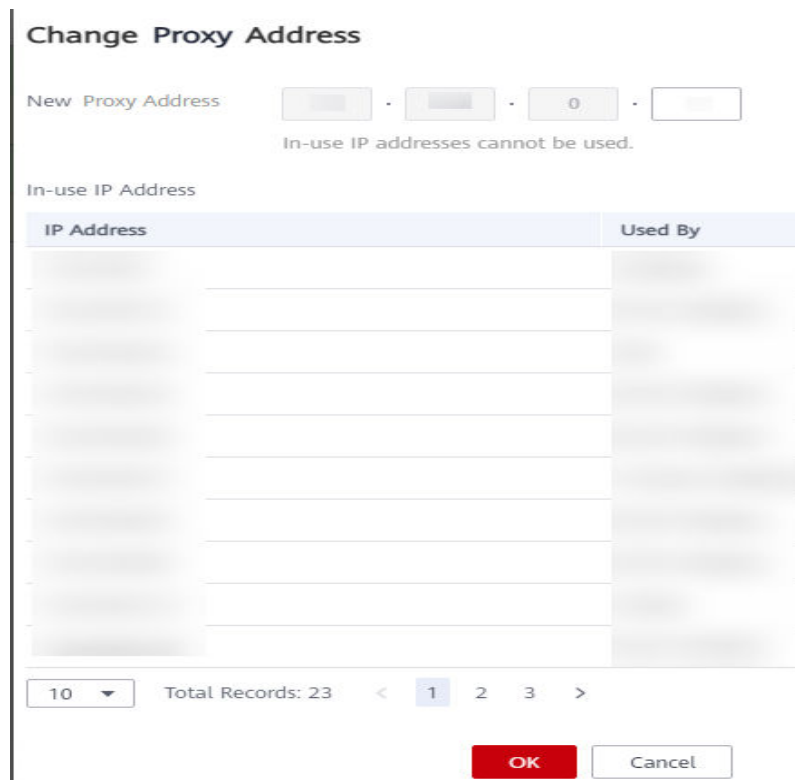
- Step 1** Log in to the management console.
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, click **Database Proxy**.
- Step 6** Click the desired proxy instance name. In the **Proxy Instance Information** area, click **Change** next to the **Proxy Address** field.

**Figure 6-18** Changing the IP address of a proxy instance (1)



- Step 7** In the displayed dialog box, enter a new IP address and click **OK**.  
In-use IP addresses cannot be used.

**Figure 6-19** Changing the IP address of a proxy instance (2)



----End

## 6.14 Changing the Port of a Proxy Instance

### Scenarios

After a proxy instance is created, you can change its port as needed.


### Precautions

- Changing a proxy port will interrupt the database connection. You are advised to change the port number during off-peak hours.
- Only the port of the current proxy instance will be changed.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.



**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**.

**Step 6** Click the name of a proxy instance.

**Step 7** On the **Basic Information** page, click  next to **Proxy Port**.

Proxy port range: 1025 to 65534 (except for 1033, 5342, 5343, 5344, 5345, 12017, 20000, 20201, 20202, 33062, and 33071, which are reserved by the system)

- To submit the change, click .
  - In the displayed dialog box, click **Yes** to confirm the change.
  - In the displayed dialog box, click **No** to cancel the change.
- To cancel the change, click .

----End

## 6.15 Changing Consistency Level

### Scenarios


After a proxy instance is created, you can change its consistency level.


### Constraints

- To configure consistency level, the kernel version of your GaussDB(for MySQL) instance must be 2.0.28.1 or later.
- To use session consistency, the kernel version of your proxy instance must be 2.7.4.0 or later.

### Procedure


**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, click **Database Proxy**.

**Step 6** Click the desired proxy instance name. In the **Proxy Instance Information** area, click  next to the **Consistency Level** field.

**Figure 6-20** Changing consistency level



**Step 7** Select a consistency level and click .

---

**NOTICE**

After the consistency level is changed, you need to manually reboot the proxy instance or reconnect your application to the proxy instance on the management console.

For details about how to reboot a proxy instance, see [Rebooting a Proxy Instance](#).

---

----End

## APIs

[Changing Session Consistency of a Database Proxy Instance](#)

# 6.16 Modifying Proxy Instance Parameters

## Prerequisites

To modify the parameters of a proxy instance, submit an application by choosing [Service Tickets > Create Service Ticket](#).


## Scenarios

You can change parameter for your database proxy instances.

## Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy**, select a proxy instance and click its name.

**Step 6** In the navigation pane on the left, choose **Parameter Modification**. On the displayed page, change parameters if needed.

You can save, cancel, or preview your changes.

- To save your changes, click **Save**.
- To cancel your changes, click **Cancel**.

- To preview your changes, click **Preview**.

----End

## 6.17 Enabling or Disabling Access Control


If load balancing is enabled for a proxy instance, the security group associated with the proxy instance does not apply. You need to use access control to limit access from specific IP addresses.


### NOTE

If access control is not displayed on the management console, the security group associated with the proxy instance is used.

### Enabling Access Control

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

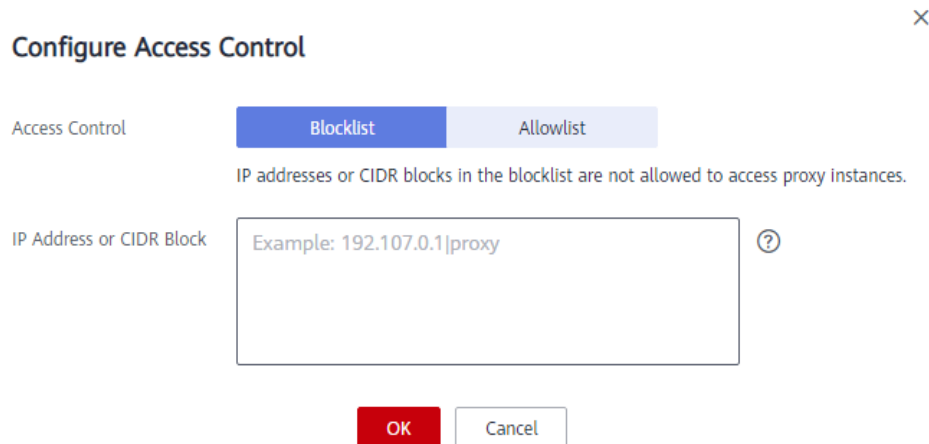
**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy** and click the name of a proxy instance to go to the **Basic Information** page.

**Step 6** Click  next to **Access Control**.

**Step 7** Click **Configure**. In the displayed dialog box, configure required parameters.


- **Access Control:** The blacklist and allowlist cannot be configured at the same time. If you switch between lists, your previously entered settings will be lost. IP addresses or CIDR blocks in the blacklist are not allowed to access the proxy instance.
- **IP Address or CIDR Block:** Enter valid IP addresses or CIDR blocks that meet the following requirements:
  - Each line contains an IP address or a CIDR block and ends with a line break.
  - Each IP address or CIDR block can include a description separated by a vertical bar symbol (|), for example, 192.168.10.10|GaussDBforMySQL01. The description can include up to 50 characters but cannot contain angle brackets (<>).
  - Up to 300 IP addresses or CIDR blocks can be added.


**Figure 6-21** Configuring access control

----End

## Disabling Access Control

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Database Proxy** and click the name of a proxy instance to go to the **Basic Information** page.

**Step 6** Click  next to **Access Control**.

**Step 7** In the displayed dialog box, click **Yes** to disable access control.

----End

## 6.18 Rebooting a Proxy Instance

### Scenarios


You can reboot a proxy instance as needed.


### Constraints

- You have obtained the required permissions from customer service.
- If the proxy instance status is **Abnormal**, the reboot may fail.
- To shorten the time required, reduce database activities during the reboot to reduce rollback of transit transactions.

## Procedure

**Step 1** Log in to the management console.


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

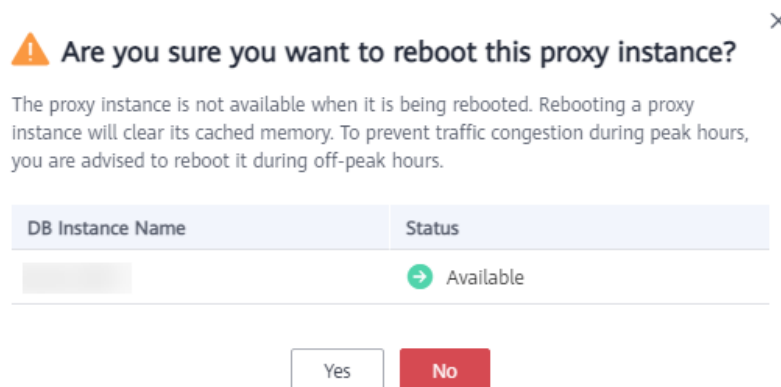
**Step 5** In the navigation pane on the left, choose **Database Proxy**, locate the target proxy instance, and choose **More** > **Reboot** in the **Operation** column.

**Figure 6-22** Rebooting a proxy instance



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

**Figure 6-23** Confirming information



### NOTE



Reboot a proxy instance interrupts the database connection. You are advised to reboot it during off-peak hours.

----End


## 6.19 Deleting a Proxy Instance

You can delete a proxy instance as required.

## Procedure

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Database Proxy**.
- Step 6** On the **Database Proxy** page, locate the desired proxy instance and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

**Figure 6-24** Deleting a proxy instance



### NOTE

If the proxy instance is deleted, read/write splitting is disabled and services using the proxy address are interrupted. You need to switch your applications to the instance address.

----End

## APIs

- [Enabling Database Proxy](#)
- [Disabling Database Proxy](#)

## 6.20 Introducing Read Weight Assignment Rules

Read weights are assigned to read replicas by their specifications.

### Weight Assignment Rules

The system automatically assigns read weights to read replicas based on their specifications, as shown in the following table.

### NOTE

The default weight of a read replica is (Number of vCPUs) x 50.  
If the number of vCPUs is 4, the weight is 200 (4 x 50 = 200).



**Table 6-2** Weights assigned to read replicas

vCPUs	Memory (GB)	Weight
4	16	200
8	32	400
16	64	800
32	128	1,000
48	192	1,000
60	256	1,000

## Specifying Whether a SQL Statement Is Sent to the Primary Node or Read Replica By Adding a Hint

Hints supported by read/write splitting are as follows:

**/\*FORCE\_MASTER\*/**: A SQL statement is executed on a primary node.

**/\*FORCE\_SLAVE\*/**: A SQL statement is executed on read replicas.

### NOTE

- In addition to the weight assignment rules of read/write splitting, hints serve as a complementary SQL syntax to specify whether a SQL statement is executed on a primary node or read replica.
- Hints are only used as routing suggestions. In non-read-only SQL and non-transaction scenarios, SQL statements cannot be routed to read replicas.
- If you want to connect to an instance using the MySQL CLI and Hints, add the **-c** option.

## 6.21 Testing Read/Write Splitting Performance

After a proxy instance is created, you can connect your GaussDB(for MySQL) instance through a proxy address. You can use internal SQL commands to verify the read/write splitting performance.

### Procedure

**Step 1** [Log in to an ECS](#).

**Step 2** Connect to a DB instance through a proxy address.

```
mysql -h <hostIP> -P <port> -u <userName> -p <password>
```

**Table 6-3** Parameter description

Parameter	Description
<hostIP>	Proxy address.

Parameter	Description
<port>	Database port. By default, the value is <b>3306</b> .
<userName>	Username of the GaussDB(for MySQL) database administrator account. The default username is <b>root</b> .
<password>	Password

**Step 3** Run the following command to view the instance that executes the SQL command:

Run **show last route;** to view the routing result of the previous SQL statement.

**Figure 6-25** Query result

```
Copyright (c) 2000, 2019, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.

mysql> select 1;
+----+
| 1 |
+----+
1 row in set (0.08 sec)

mysql> show last route;
+-----+
| LAST ROUTE |
+-----+
| 192.168.120.92 |
+-----+
1 row in set (0.05 sec)

mysql>
```

Do not use **show last route** for service code or multi-statement execution.

----End

# 7 Instance Lifecycle Management

---

## 7.1 Changing a DB Instance Name


### Scenarios


You can change the name of a GaussDB(for MySQL) instance.


### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a DB instance and click  in the **Name/ID** column to edit the DB instance name.

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Instance Name** field in the **DB Instance Information** area, click  to edit the instance name.

- The instance name must start with a letter and consist of 4 to 64 characters. Only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (\_) are allowed.
- When changing the instance name, you can determine whether to select **Change node names synchronously** as required. If this option is selected, the names of the corresponding nodes are changed when the instance name is changed. If this option is not selected, only the instance name is changed, and the corresponding node names are not changed.
- If you want to submit the change, click **OK**. If you want to cancel the change, click **Cancel**.

**Step 5** View that the instance name has been changed. It takes less than 1 minute to change a DB instance name.

----End

## APIs

- [Changing a DB Instance Name](#)
- [Querying DB Instances](#)
- [Querying Details of a DB Instance](#)
- [Querying Details of DB Instances in Batches](#)


## 7.2 Changing a DB Instance Description


### Scenarios


After a GaussDB(for MySQL) instance is created, you can add a description for it.

### Procedure


**Step 1** [Log in to the management console.](#)



**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance and click  in the **Description** column to edit the instance description.

- If you want to submit the change, click **OK**.
- If you want to cancel the change, click **Cancel**.

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Instance Information** area, click  in the **Description** field to edit the instance description.

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

#### NOTE

The instance description can contain up to 64 characters. Only letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed.

**Step 5** View the results on the **Basic Information** page.

Alternatively, view the results on the **Instances** page.

----End

## APIs

- [Changing a DB Instance Description](#)
- [Querying DB Instances](#)
- [Querying Details of a DB Instance](#)
- [Querying Details of DB Instances in Batches](#)

## 7.3 Deleting a DB Instance

### Scenarios

You can manually delete a DB instance billed on a pay-per-use or serverless basis on the GaussDB(for MySQL) **Instances** page.

---


#### NOTICE

- Instances cannot be deleted when operations are being performed on them.
  - If you delete a DB instance billed on a pay-per-use basis, its automated backups are also deleted and you are no longer billed for them. Manual backups are still retained and will incur additional costs.
  - If you delete a DB instance billed on a pay-per-use or serverless basis, the read replicas associated with it are also deleted.
  - You can rebuild deleted instances from the recycle bin. For details, see [Rebuilding a Deleted Instance from Recycle Bin](#).
- 

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance you want to delete and click **More** > **Delete** in the **Operation** column.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

**Step 6** In the displayed dialog box, enter **DELETE** and click **OK**. Refresh the **Instances** page later to check that the deletion is successful.

----End

## APIs

- [Deleting a DB Instance](#)
- [Deleting a Read Replica](#)
- [Querying DB Instances](#)

# 7.4 Rebooting a DB Instance

## Scenarios

You may need to reboot a GaussDB(for MySQL) instance for maintenance reasons. For example, after changing some parameters, you must reboot the instance for the modifications to take effect.

## Constraints

- If the DB instance status is **Abnormal**, the reboot may fail.
- To shorten the time required, you need to reduce database activities during the reboot to reduce a rollback of transit transactions.

---


### NOTICE

- Rebooting a DB instance will interrupt services. During this period, the instance status is **Rebooting**.
  - Rebooting DB instances will cause instance unavailability. To prevent traffic congestion during peak hours, reboot instances during off-peak hours.
- 

## Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance you want to reboot and choose **More > Reboot** in the **Operation** column.

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

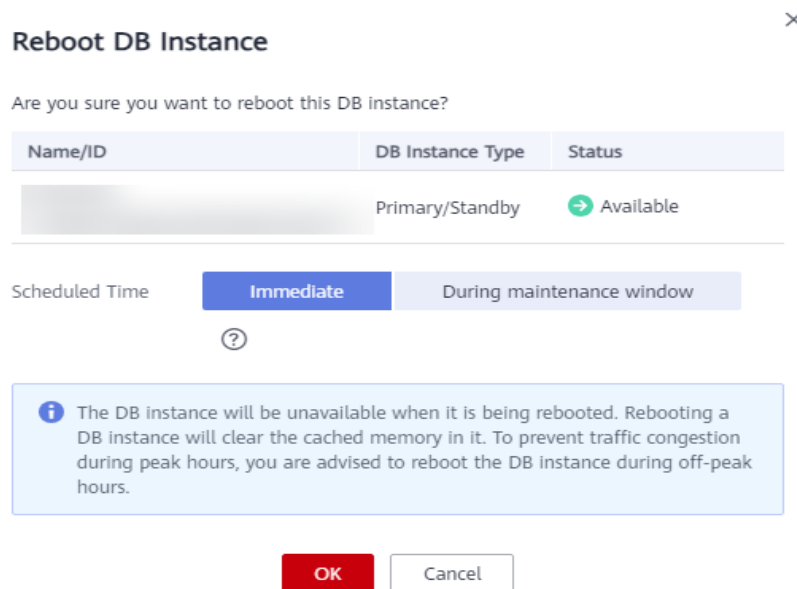
The read replicas are also rebooted.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

**Step 6** In the displayed dialog box, set **Scheduled Time** to **Immediate** or **During maintenance window** as needed.

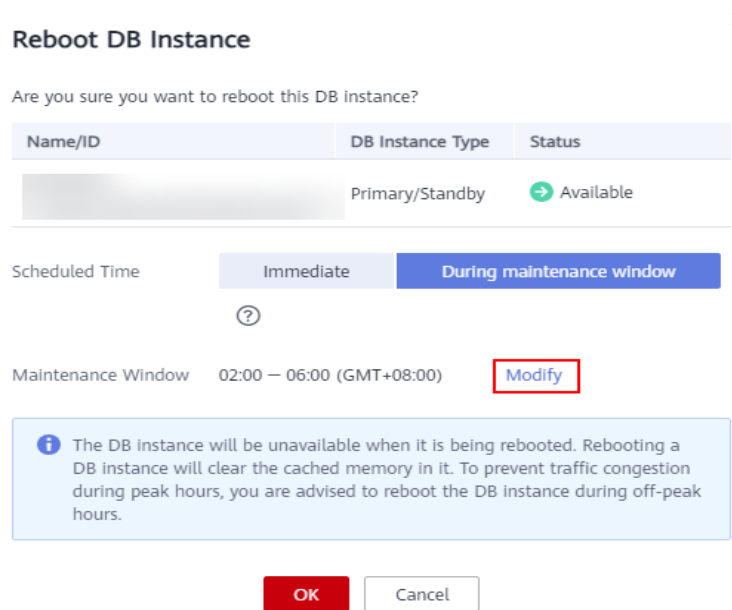
**Figure 7-1** Rebooting an instance immediately



1. To reboot the instance immediately, set **Scheduled Time** to **Immediate** and click **Yes**.
2. To reboot the instance during the maintenance window, set **Scheduled Time** to **During maintenance window** and click **Yes**. You can also change the maintenance window by clicking **Change**.

Changing the maintenance window will not affect the timing that has already been scheduled.

A reboot task configured during a current maintenance window will not be executed until the next maintenance window.

**Figure 7-2** Rebooting a DB instance during maintenance window

**Step 7** View the task execution progress on the **Task Center** page. If its status is **Available**, it has been rebooted.

----End

## APIs

- [Rebooting a DB Instance](#)
- [Rebooting a Node](#)

## 7.5 Rebooting a Node

### Scenarios

You can reboot a node of your DB instance on the management console to reconnect the database.

### Constraints

- The nodes in the **Abnormal** state can be rebooted.
- The time required for rebooting a node depends on the crash recovery process of the DB engine. To shorten the time required, reduce database activities during the reboot to reduce rollback of transit transactions.




**NOTICE**

- Rebooting a node will cause service interruption. During this period, the node status is **Rebooting node**.
- The node is not available when it is being rebooted. To prevent traffic congestion during peak hours, you are advised to reboot the node during off-peak hours.
- If a parameter of your DB instance is modified, you need to first reboot the DB instance for the modification to take effect, and then reboot a node of the DB instance.

**Procedure**

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

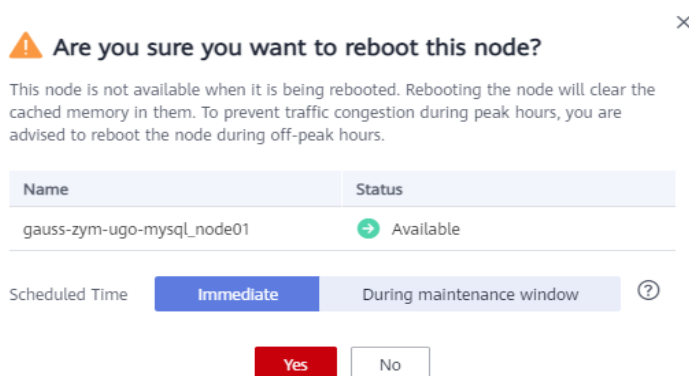
**Step 5** In the **Node List** area, locate a node that you want to reboot and choose **Reboot** in the **Operation** column.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

**Step 7** In the displayed dialog box, set **Scheduled Time** to **Immediate** or **During maintenance window**.

**Figure 7-3** Rebooting a node

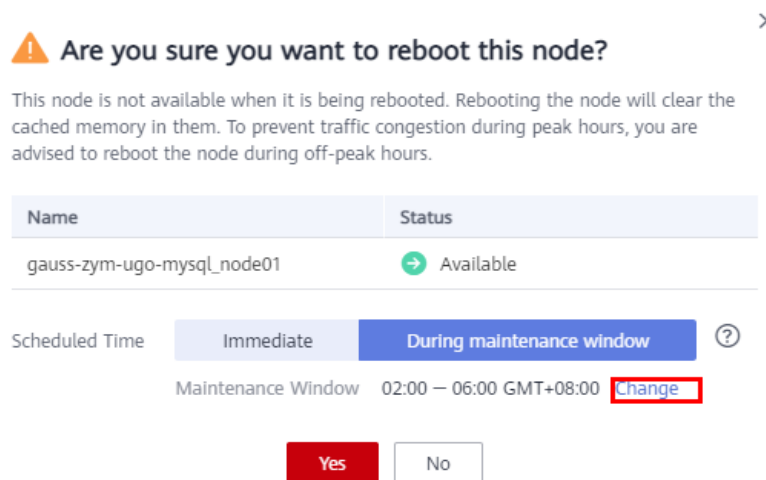


1. To reboot the instance immediately, set **Scheduled Time** to **Immediate** and click **Yes**.
2. To reboot the instance during the maintenance window, set **Scheduled Time** to **During maintenance window** and click **Yes**. You can also change the maintenance window by clicking **Change**.

Changing the maintenance window will not affect the timing that has already been scheduled.

A reboot task configured during a current maintenance window will not be executed until the next maintenance window.

**Figure 7-4** Rebooting a node during maintenance window



**Step 8** View the task execution progress on the **Task Center** page. If the node status is **Available**, the reboot is complete.

----End

## APIs

- [Rebooting a Node](#)
- [Rebooting a DB Instance](#)

## 7.6 Changing a Node Name


### Scenarios

GaussDB(for MySQL) allows you to change the node name for easy identification.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.


- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the **Node List** area on the **Basic Information** page, select one or more nodes, click **Change Node Name**.
- Click **OK** to save the modifications.
  - Click **Cancel** to cancel the modifications.

**Figure 7-5** Change node names

Node List

[Change Node Name](#)

Name/ID	Node Type	Status	AZ	Private IP Address for Read	Follower Priority	Operation	
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Primary	Available	az1	192.168.0.75	1 <input type="text"/>	<a href="#">View Metric</a> <a href="#">Reboot</a>
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Replica	Available	az4	192.168.0.229	1 <input type="text"/>	<a href="#">View Metric</a> <a href="#">Promote to Primary</a> <a href="#">Reboot</a>

You can also click  next to a node name, enter the new node name, and click **OK**.

 **NOTE**

- The node name must start with a letter and consist of 4 to 128 characters. Only letters (case-sensitive), digits, hyphens (-), and underscores (\_) are allowed.
- The node name must be unique.

- Step 6** View that the node names have been changed.



----End

## 7.7 Exporting Instance Information

### Scenarios

You can export information about all instances for review and analysis.

### Exporting Information About All Instances

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click **Export Instance Info** above the instance list. In the displayed dialog box, select the items to be exported and click **OK**.
- Step 5** After the export task is complete, a .csv file is generated locally.

----End

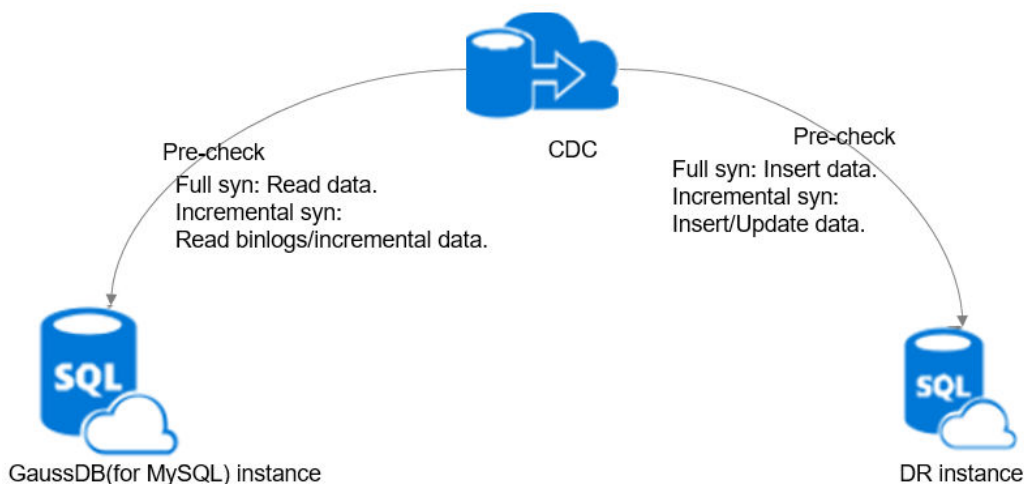
## 7.8 Introducing Heterogeneous DR Instances

GaussDB(for MySQL) supports heterogeneous DR instances to ensure high availability. If your DB instance is unavailable due to unknown community bugs, misoperations or AZ faults, the heterogeneous DR instance can take over the services.

The DR instance creation, maintenance, and switchover are all completed by Huawei Cloud and you do not need to do anything.

The free trial of heterogeneous DR instances has ended on January 1, 2024.

**Figure 7-6** Heterogeneous DR instance



### Constraints

Heterogeneous DR instances cannot be created in the scenarios listed in the following table.

**Table 7-1** Constraints

Scenario	Reason
Binlog is not enabled.	Binlogs are used to synchronize data between the GaussDB(for MySQL) instance and DR instance.
<b>binlog_format</b> is not set to <b>ROW</b> .	The DR link may be disconnected.
There are triggers.	Data may be inconsistent between the GaussDB(for MySQL) instance and the DR instance.

Scenario	Reason
There are events.	Data may be inconsistent between the GaussDB(for MySQL) instance and the DR instance.
Database names, table names, field names, or indexes contain special characters.	The DR link may be disconnected.
The data volume exceeds 4 TB.	The capacity of the DR instance has an upper limit.
The database access is restricted by rules in the security group.	The DR instance cannot connect to the GaussDB(for MySQL) instance to synchronize data.
The GaussDB(for MySQL) instance is frozen.	The DR instance cannot be created.
There are no sufficient IP addresses in the VPC subnet.	The DR instance and GaussDB(for MySQL) instance need to use IP addresses in the same subnet.
There are tables without primary keys or indexes and a large number of operations.	The DR link may be disconnected.
Cascade operations are performed on tables with foreign keys.	Data may be inconsistent between the GaussDB(for MySQL) instance and the DR instance.

## 7.9 Rebuilding a Deleted Instance from Recycle Bin

Unsubscribed yearly/monthly instances and deleted pay-per-use instances in the recycle bin can still be restored.

The recycle bin is enabled by default and cannot be disabled.


### Modifying the Recycling Policy

#### NOTICE

The new recycling policy takes effect only for instances that are put into the recycle bin after the modification. For instances that already exist in the recycle bin before the modification, the original recycling policy takes effect.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Recycle Bin** page, click **Modify Recycling Policy**. In the displayed dialog box, set the retention period for the deleted DB instances (value range: 1 to 7 days).
- Step 5** Click **OK**.

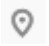

**Figure 7-7** Modifying the recycling policy



----End

## Rebuilding a DB Instance

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

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Recycle Bin** page, locate the instance you want to rebuild and click **Rebuild** in the **Operation** column.
- Step 5** On the **Rebuild DB Instance** page, configure required information and submit the task. For details, see [Restoring Data to a DB Instance](#).

----End

# 8 Instance Modifications

---

## 8.1 Changing vCPUs and Memory of a DB Instance

### Scenarios

You can change the vCPUs and memory of a yearly/monthly or pay-per-use DB instance if needed. If the status of a DB instance changes from **Changing instance specifications** to **Available**, the change was successful.

You can scale up or down your instance specifications.

### Constraints

- A DB instance cannot be deleted when its specifications are being changed.
- The vCPUs and memory can be changed only at the instance level. It means that the specifications of the primary node or read replicas cannot be changed separately for a given instance.
- Instance specifications can only be changed from the general-purpose edition to the dedicated edition.
- You can change the specifications of yearly/monthly or pay-per-use DB instances immediately or during a maintenance window. Serverless DB instances do not support specification changes.
- If you want to change instance specifications during a maintenance window, you can cancel the task before it starts. Once started, the task cannot be canceled.


### NOTICE

- Changing instance specifications will cause a primary/standby switchover. To prevent service interruptions, change the instance specifications during off-peak hours.
- The time required for modifying specifications depends on factors such as the number of nodes, database load, and number of database tables.
- Changing instance specifications will change the private IP addresses for read of the primary node and read replicas. The connection addresses in your application need to be changed to prevent your services from being affected. You are advised to use the private IP address of a DB instance to connect your application.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** Change the instance specifications in either of the following ways:

- On the **Instances** page, locate the instance for which you want to change specifications and choose **More > Change Instance Specifications** in the **Operation** column.
- Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Instance Information** area, click **Change** next to the **Instance Specifications** field.

**Step 5** On the displayed page, select new specifications as required and the scheduled time, and click **Next**.

Choose either of the following scheduled time:

- **Upon submission:** The instance specifications will be changed immediately after the task is submitted.
- **In maintenance window:** The instance specifications will be changed during the maintenance window you specify.

**Step 6** On the displayed page, confirm the instance specifications.

- If you need to modify your settings, click **Previous** to go back to the page where you specify details.
- For pay-per-use instances, click **Submit**.  
To view the cost incurred by the instance specifications change, choose **Billing Center > Billing Dashboard** in the upper right corner.
- For yearly/monthly instances:
  - Scaling down the instance specifications: click **Submit**.



The refund is automatically returned to your account. You can click **Billing Center** in the upper right corner and then choose **Orders > My Orders** in the navigation pane on the left to view the details.

- Scaling up the instance specifications: click **Submit**. The scaling starts only after the payment is successful.

#### Step 7 View the results.

Changing the instance specifications takes 5–15 minutes. During this period, the status of the instance on the **Instances** page is **Changing instance specifications**. After a few minutes, you can click the instance name to view the new instance specifications on the displayed **Basic Information** page.

---

#### NOTICE

- After the instance specifications of GaussDB(for MySQL) 8.0 are changed, the system will change the values of the following parameters accordingly: **innodb\_buffer\_pool\_size**, **innodb\_log\_buffer\_size**, **max\_connections**, **innodb\_buffer\_pool\_instances**, **innodb\_page\_cleaners**, **innodb\_parallel\_read\_threads**, **innodb\_read\_io\_threads**, **innodb\_write\_io\_threads**, and **threadpool\_size**.
- The default value of **innodb\_parallel\_select\_count** is **OFF** for instance with 16 vCPUs or less and **ON** for instances with more than 16 vCPUs.  
If you have modified value of the parameter, the parameter value remains unchanged after the specifications are changed, or the default value is used.

---

----End

## APIs

- [Changing DB Instance Specifications](#)
- [Promoting a Read Replica to Primary](#)
- [Querying Database Specifications](#)
- [Querying Details of a DB Instance](#)

## 8.2 Changing Storage of a Yearly/Monthly DB Instance

### Scenarios

If the original storage space of your yearly/monthly instance is insufficient or redundant as your services change, you can scale up or down the storage.

### Constraints

- The storage of pay-per-use instances grows as needed, so you cannot manually scale up their storage. The storage of pay-per-use instances is not limited.
- When you purchase a yearly/monthly instance, you need to select storage for it as needed. If your purchased storage cannot meet service requirements, the


system will automatically scale up the storage as needed and you will be billed on a pay-per-use basis for additional storage. If services requirements decrease later, the system preferentially scales down the storage that was automatically scaled up.


For example, you purchased 10 GB of storage when purchasing a DB instance. Later, as services increased, GaussDB(for MySQL) automatically scaled the storage to 18 GB as needed and you would be billed on a pay-per-use basis for the additional 8 GB of storage. Then, you manually scaled up the storage to 20 GB, which could meet service requirements. The 8 GB of storage that was scaled up by the system will be scaled down. You would only pay for the 20 GB of storage at yearly/monthly rates.

- The system changes storage of your DB instance as your services change, but you can change storage only by a multiple of 10 GB.
- The storage that is automatically scaled up can only be scaled down by the system, cannot by users.
- During a storage change, services including backup service are not interrupted.
- You can change storage of a DB instance numerous times.
- If the storage of a DB instance is being changed, you cannot reboot or delete the DB instance.

## Procedure

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the DB instance and choose **More** > **Change Storage** in the **Operation** column.

Alternatively, click the instance name to go to the **Basic Information** page. In the **Storage/Backup Space** area, click **Change Storage**.

**Figure 8-1** Scaling storage



**Step 5** Select the new storage space and click **Next**.

Storage space can be scaled up to 128,000 GB only by a multiple of 10 GB. Price after scaling is displayed in the lower left corner of the page.

Storage space can be scaled down to 40 GB only by a multiple of 10 GB. Refund price is displayed in the lower left corner of the page.

 **NOTE**

To reduce the storage of a DB instance to 10 GB, contact customer service.

**Step 6** Confirm your settings.

- If you need to modify your settings, click **Previous**.
- If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.

**Step 7** View the new storage.

On the **Instances** page, you can click the instance name to go to the **Basic Information** page and view the new storage.

----End

## APIs

- [Scaling up Storage of a Yearly/Monthly DB Instance](#)
- [Querying Database Specifications](#)
- [Querying Details of a DB Instance](#)

## 8.3 Configuring Auto Scaling Policies

### Scenarios

You can configure auto scaling policies for your pay-per-use and yearly/monthly DB instances on the **Basic Information** page. When configuring auto scaling policies, you can enable or disable **Auto Scale-up** or **Auto Scale-down**. The scaling type includes changing instance specifications or the number of read replicas.

### Constraints

- This function is only available for pay-per-use and yearly/monthly DB instances.
- To set **Scaling Type** to **Number of read replicas** for a yearly/monthly instance, choose [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console. The read replicas that are automatically added or deleted will be billed based on a pay-per-use basis.
- To configure auto scaling policies, you must have the iam:agencies:listAgencies permission. If you do not have this permission, [create a custom policy](#).
- Changing DB instance specifications will briefly interrupt services.

- If you want to set **Scaling Type** to **Number of read replicas**, there must be only one proxy instance. For details, see [Create a Proxy Instance](#).
- The system will delete or add read replicas. To prevent your services from being affected, you are advised not to use an IP address for read to connect to your applications.
- The pricing standard for auto scaling is the same as that for manual scaling. For details, see [Billing](#).

## Billing

- **Pay-per-Use Instances**  
The instance specifications and number of read replicas can be automatically changed.  
Pricing is listed on a per-hour basis, but bills are calculated down to the second. The old order automatically becomes invalid.  
To view the cost incurred by auto scaling, choose **Billing Center > Billing Dashboard** in the upper right corner of the management console.
- **Yearly/Monthly Instances**  
The instance specifications and number of read replicas can be automatically changed.  
You will be billed for the new specifications. For details, see [GaussDB\(for MySQL\) Pricing Details](#).  
If the new specifications are less than the specifications that you purchased, the refund is automatically returned to your account. You can click **Billing Center** in the upper right corner of the management console and then choose **Orders > My Orders** in the navigation pane on the left to view the details.


**Table 8-1** Pricing description for yearly/monthly instances

Billing Item	Description
Specifications after scale-up	<p>You need to pay the following fee:</p> $\text{Price of new specifications} \times \text{Remaining duration} \times \text{Number of nodes} - \text{Price of old specifications} \times \text{Remaining duration} \times \text{Number of nodes}$ <p>Note: Remaining duration = Number of remaining days in a calendar month/Total number of days in the calendar month</p> <p>Example:</p> <p>A customer placed and paid a monthly order for a GaussDB(for MySQL) instance on April 1, 2023. The instance contains 2 nodes and its specifications are 2 vCPUs and 8 GB of memory. The total subscription period would be 30 days and the instance would expire on April 30, 2023. On April 18, 2023, the instance specifications were automatically expanded to 4 vCPUs and 16 GB memory. The remaining duration is 0.4 (12/30).</p> <p>The monthly price of the new specifications is \$290 USD and that of the old specifications is \$145 USD, so you need to pay for \$116 USD (<math>290 \times 0.4 \times 2 - 145 \times 0.4 \times 2</math>).</p>
Specifications after scale-down	For pricing details, see <a href="#">Unsubscriptions</a> .
Added read replicas	New read replicas are billed based on the actual usage duration.
Deleted read replicas	Deleted read replicas are no longer billed.

## Modifying Auto Scaling Policies


**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

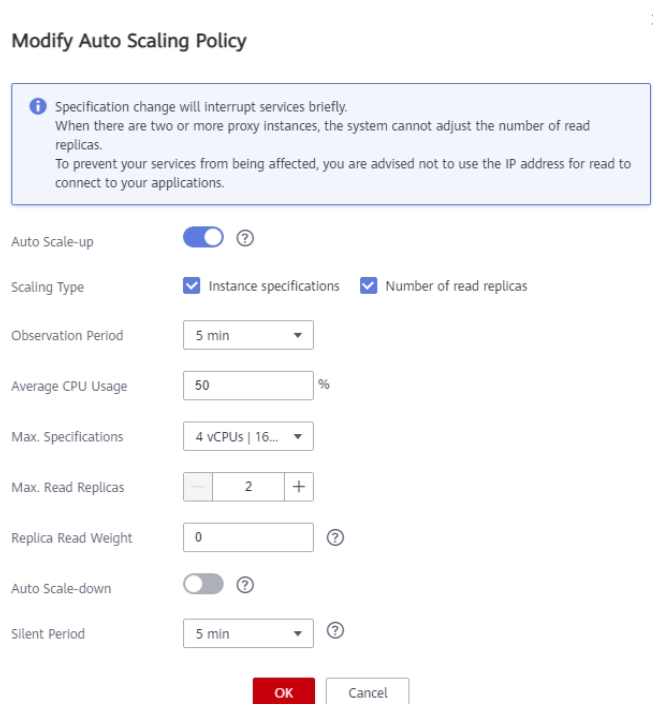
**Step 4** Click the instance name to go to the **Basic Information** page. In the **DB Instance Information** area, click **Modify** next to the **Auto Scaling** field.

**Figure 8-2** Modifying auto scaling policies



**Step 5** In the displayed dialog box, configure the required parameters.

**Figure 8-3** Modifying the auto scaling policy



**Table 8-2** Parameter configuration

Parameter	Parameter description
Auto Scale-up	You can enable or disable it as needed.


Parameter	Parameter description
Scaling Type	<ul style="list-style-type: none"><li>Instance specifications</li><li>Number of read replicas</li></ul> <b>NOTE</b> <ul style="list-style-type: none"><li>You can select one or more scaling types.</li><li>The read replicas that are automatically added or deleted will be billed based on a pay-per-use basis.</li><li>If you deselect <b>Number of read replicas</b> for <b>Scaling Type</b>, pay-per-use nodes created in the current instance will be automatically deleted. Exercise caution when performing this operation.</li><li>The account balance must be sufficient, or scaling up the specifications or adding read replicas may fail.</li><li>After <b>Auto Scale-up</b> is enabled, read replicas that are automatically added cannot be promoted to primary.</li></ul>
Observation Period	<ul style="list-style-type: none"><li>Once auto scale-up is enabled, if the system observes any increases in the average CPU usage over the preset value, it upgrades the specifications or adds read replicas based on the read and write traffic. The system then enters a silent period.</li><li>The minimum observation period is 5 minutes.</li></ul>
Average CPU Usage	Indicates threshold for triggering an auto scale-up. Allowed range: 50%–100%
Max. Specifications	Indicates the maximum specifications after the final auto scale-up. The specifications can only be scaled up gradually and the system enters the silent period after each scale-up.
Max. Read Replicas	Only one read replica can be added at a time.
Replica Read Weight	If you have enabled read/write splitting, the new read replicas are automatically associated with the proxy instance.
Auto Scale-down	You can enable or disable it as needed. <b>NOTE</b> Once auto scale-down is enabled, if the system observes an average CPU usage of 99% drops below 30% within the observation period, it gradually restores the original configuration. The system enters a silent period after each scale-down.
Silent Period	The silent period is the minimum interval between two changes (triggered automatically or manually), where no more changes can happen.


**Step 6** Click **OK**.

----End

## Viewing Change History

**Step 1** [Log in to the management console.](#)

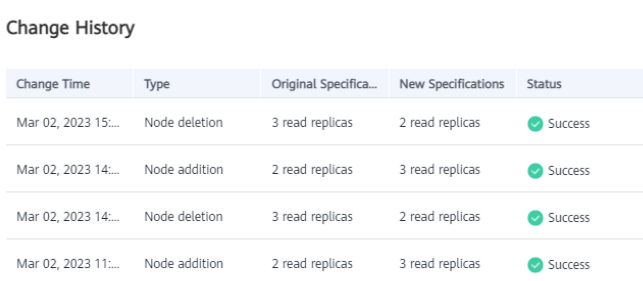
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** Click the instance name to go to the **Basic Information** page. In the **DB Instance Information** area, click **View Change History** next to the **Auto Scaling** field.

**Step 5** In the displayed dialog box, view the change time, change type, status, original specifications, and new specifications.

**Figure 8-4** Viewing change history



Change Time	Type	Original Specifica...	New Specifications	Status
Mar 02, 2023 15:...	Node deletion	3 read replicas	2 read replicas	Success
Mar 02, 2023 14:...	Node addition	2 read replicas	3 read replicas	Success
Mar 02, 2023 14:...	Node deletion	3 read replicas	2 read replicas	Success
Mar 02, 2023 11:...	Node addition	2 read replicas	3 read replicas	Success

----End

## 8.4 Changing a Maintenance Window

### Scenarios

The maintenance window is 02:00–06:00 by default, but you can change it if needed. To prevent service interruption, set the maintenance window to off-peak hours.


### Precautions

Before maintenance is performed, GaussDB(for MySQL) will send SMS messages and emails to the contact person that has been set in the Huawei ID.

### Procedure

**Step 1** [Log in to the management console.](#)

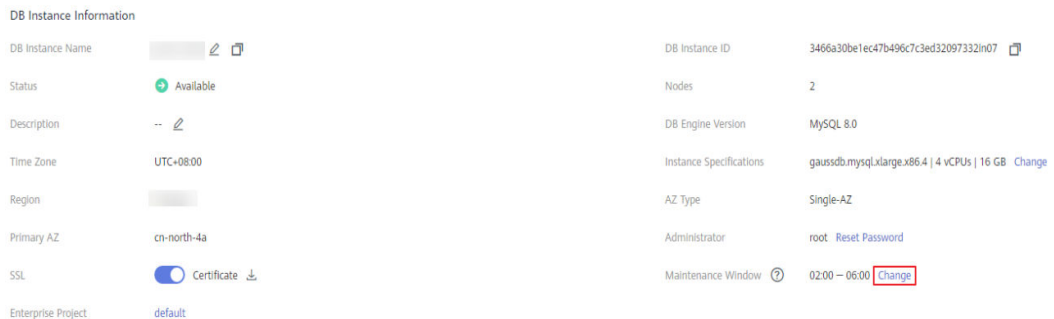
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.



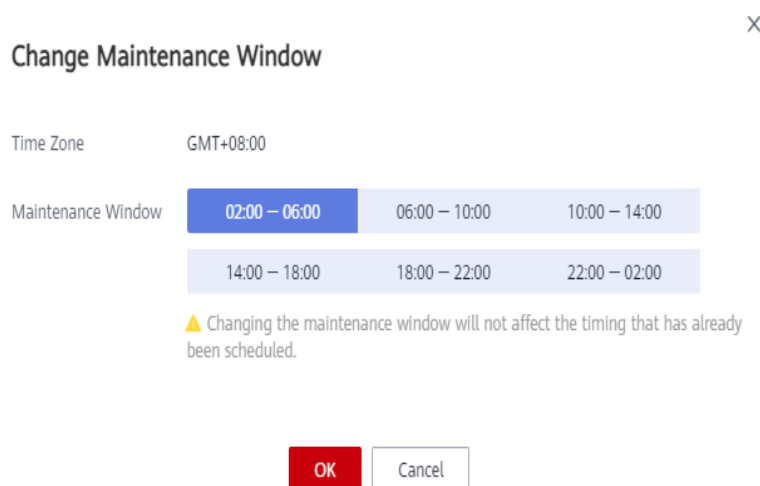
**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page. In the **DB Instance Information** area, click **Change** in the **Maintenance Window** field.

**Figure 8-5** Changing a maintenance window (1)



**Step 5** In the displayed dialog box, select a maintenance window and click **OK**.

**Figure 8-6** Changing a maintenance window (2)



**NOTE**

Changing the maintenance window will not affect the timing that has already been scheduled.

----End

**APIs**

**Changing a Maintenance Window**

## 8.5 Selecting Instance Displayed Items


### Scenarios


You can customize instance information items displayed on the **Instances** page based on your requirements.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click  to edit items displayed in the instance list.

- The following items are displayed by default: instance name/ID, instance type, description, DB engine, status, enterprise project, billing mode, private IP address, and operation.  
These default displayed items cannot be hidden.
- You can also select other items, including the creation time, database port, and storage type.

----End

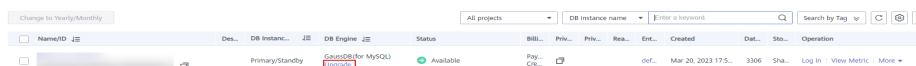
## 8.6 Upgrading a Minor Version

### Scenarios

GaussDB(for MySQL) supports manual minor version upgrades, which can improve performance, add new functions, and fix bugs.

When a new minor version is released on Huawei Cloud, you will see **Upgrade** in the **DB Engine** column on the **Instances** page.

**Figure 8-7** Upgrading a minor version



Name/ID	DB Instance	DB Engine	Status	Billing Mode	Private IP	Private IP	Release	End	Created	DB Size	Storage	Operation
	Primary/Standby	GaussDB(for MySQL) <small>Upgrade</small>	Available	Pay-As-You-Go			def...	Mar 20, 2023 17:5...	3306	Sha...	Log In   View Metric   More	

For details about the minor versions, see [Kernel Version Release History](#).

### Precautions


- When any new minor version is released for addressing issues and vulnerabilities from the open source community, upgrade the minor version of your instance **immediately** or **during the maintenance window**.

- The upgrade will cause the instance to reboot and briefly interrupt services. To limit the impact of the upgrade, perform the upgrade during off-peak hours, or ensure that your applications support automatic reconnection.
- If a DB instance contains a large number of table partitions (more than 1 million), it may take more than 2 hours to reboot the instance.
- If you want to upgrade the minor version of your DB instance from 8.0.18 to 8.0.22 and there are more than 1,000 partitions, the upgrade may fail. Contact Huawei Cloud engineers to check the version compatibility before the upgrade.
- If the primary node and read replicas of a DB instance are deployed in the same AZ, a minor version upgrade will trigger a failover. If they are in different AZs, a minor version upgrade will trigger two failovers. A failover means that the system fails over to a read replica in case the primary node is unavailable.
- When you upgrade a minor version of a DB instance, minor versions of read replicas (if any) will also be upgraded automatically. Minor versions of read replicas cannot be upgraded separately. A minor version upgrade cannot be rolled back after the upgrade is complete.
- DDL operations, such as creating events, dropping events, and altering events, are not allowed during a minor version upgrade.
- If the replication delay between primary node and read replicas is longer than 300 seconds, the minor version cannot be upgraded.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.


**Step 5** In the **DB Instance Information** area, click **Upgrade** next to the **DB Engine Version** field.

Alternatively, go to the **Instances** page, click **Upgrade** next to the **DB Engine Version** column.

**Figure 8-8** Upgrading a minor version

Name/ID	Description	DB Instanc...	DB Engine	Status	Billing Mode	Private IP...	Private Do...	Enterprise ...	Storage Type	Operation
		Primary/Stan... 2 vCPUs   8 ...	GaussDB(for... <b>Upgrade</b>		Pay-per-use Created on ...	192.168...	-	default	Shared	Log In   View Metric   More

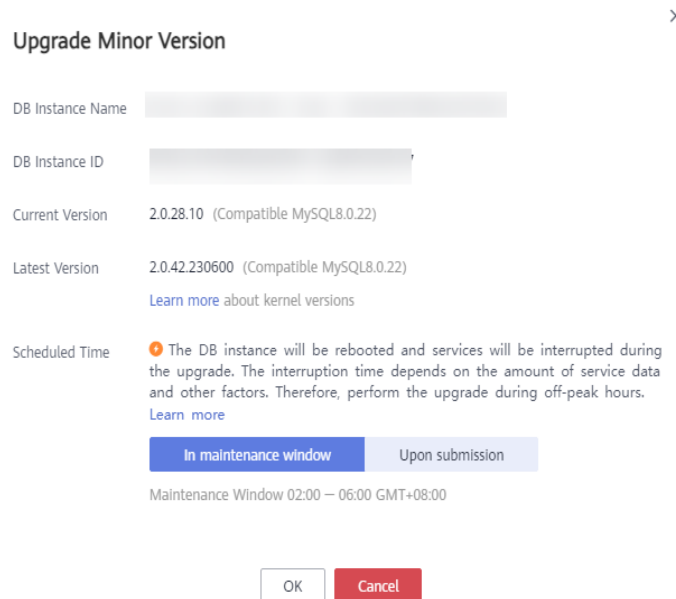
**Figure 8-9** Upgrading a minor version



**Step 6** In the displayed dialog box, select a scheduled time and click **OK**.

- Upon submission: The system upgrades the minor version immediately after your submission of the upgrade request. After the operation is complete, in the **Task Center** page, click **Instant Tasks** and view the information about the upgrade task.
- In maintenance window: The system will upgrade the minor version **during a maintenance window you specify**. After the operation is complete, in the **Task Center** page, click **Scheduled Tasks** and view the information about the upgrade task.

**Figure 8-10** Selecting a scheduled time



-----End

## APIs

### Upgrading the Kernel Version of a DB Instance

## 8.7 Upgrading the OS of a DB Instance

To improve database performance and security, the OS of your GaussDB(for MySQL) instance needs to be updated timely.

Every time you upgrade the kernel version of your instance, GaussDB(for MySQL) determines whether to update the OS and selects the right cold patch to upgrade the OS if necessary.

Updating the OS does not change the DB instance version or other information.

In addition, GaussDB(for MySQL) installs hot patches as required to fix major OS vulnerabilities within the maintenance window you specified.

# 9 Read Replicas

---

## 9.1 Introducing Read Replicas

### Scenarios

A GaussDB(for MySQL) instance contains read replicas in addition to a primary node.

In read-intensive scenarios, a primary node may be unable to handle the read pressure and service performance may be affected. To offload read pressure from the primary node, you can create one or more read replicas. These read replicas can process a large number of read requests and increase application throughput. To do this, connection addresses need to be scheduled separately for the primary node and each read replica on your applications so that all read requests can be sent to read replicas and write requests to the primary node.

### Billing Standards

Read replicas are billed as well. The billing mode is the same as that of the primary node.

### Functions

- Specifications of read replicas are the same as those of the primary node.
- You do not need to maintain accounts and databases for read replicas. They are synchronized from the primary node.
- The system can monitor the performance of read replicas.

### Constraints

- There are up to 15 read replicas for a yearly/monthly or pay-per-use instance, and one read replica for a serverless instance.
- Read replicas do not support restoration from backups.
- Data cannot be migrated to read replicas.
- You cannot create or delete databases on read replicas.

- You cannot create database accounts for read replicas.
- There may be a latency between the read replicas and the primary node. The latency of the full-text index is significant due to its special mechanism. For latency-sensitive application workloads, you are advised to send queries to the primary node.

## 9.2 Creating a Read Replica

### Scenarios

Read replicas of a DB instance are used to enhance instance capabilities and reduce the read pressure on the primary node.

You can add read replicas to an existing yearly/monthly or pay-per-use instance.

- If you select single-AZ deployment, read replicas are deployed in the same AZ as the primary node.
- If you select multi-AZ deployment, read replicas are evenly deployed in different AZs to ensure high reliability.


### Constraints

- There are up to 15 read replicas for a yearly/monthly or pay-per-use instance.
- You cannot add read replicas for serverless instances.


### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance that you want to create read replicas for and choose **More > Create Read Replica** in the **Operation** column.

Alternatively, click the instance name to go to the **Basic Information** page. In the **DB Instance Topology** area, click  to create read replicas.

**Step 5** On the displayed page, configure related parameters.

**Table 9-1** Parameter description

Parameter	Description
Failover Priority	Failover priority ranges from 1 for the first priority to 16 for the last priority. This priority determines the order in which read replicas are promoted when recovering from a primary node failure. Read replicas with the same priority have a same probability of being promoted to the new primary node. You can configure a failover priority for up to 9 read replicas, and the default priority for the remaining read replicas is <b>-1</b> , indicating these read replicas cannot be promoted to primary. You can change the failover priority of a read replica.
Quantity	A DB instance can contain up to 15 read replicas.

**Step 6** For a yearly/monthly instance, click **Buy Now** and select a payment mode.

**Step 7** For a pay-per-use instance, click **Next**.

**Step 8** Check the read replica settings.

- If you need to modify the settings, click **Previous**.
- If you do not need to modify the settings, click **Submit**.

**Step 9** View the new read replica information in the **Node List** area of the **Basic Information** page. You can also promote a read replica to primary or delete a read replica.

----End

## APIs

- [Creating a Read Replica](#)
- [Deleting or Unsubscribing from a Read Replica](#)


## 9.3 Managing a Read Replica

You can manage read replicas after they are created, such as [promoting a read replica to the primary node](#) and [deleting a read replica](#).

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** Click the instance name to go to the **Basic Information** page.



**Step 5** In the **Node List** area, view and manage all read replicas associated with the instance.

**Figure 9-1** Viewing and managing read replicas

Node List

On the Database Proxy page, you can enable the database proxy to access your DB instance through the read/write splitting address. In this case, you do not need to reconfigure the connection address in your application if the private network addresses of the primary node and read replicas are changed due to the instance specification change.

Name/ID	Node Type	Status	AZ	Private IP Address for Read	Fallover Priority	Operation
	Primary	Available	az3	192.168.0.90	1	View Metric
	Replica	Available	az3	192.168.0.107	1	View Metric   Promote to Primary

**NOTE**

The private IP addresses for read of the primary node and read replicas cannot be changed.

-----End


## 9.4 Promoting a Read Replica to the Primary Node

A GaussDB(for MySQL) instance consists of a primary node and multiple read replicas. In addition to **automatic failover** scenarios, you can perform a **manual switchover** to promote a read replica to the new primary node.

### Manual Switchover

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Node List** area, locate the read replica to be promoted and click **Promote to Primary** in the **Operation** column.

**Figure 9-2** Promoting a read replica to the new primary node

Node List

On the Database Proxy page, you can enable the database proxy to access your DB instance through the read/write splitting address. In this case, you do not need to reconfigure the connection address in your application if the private network addresses of the primary node and read replicas are changed due to the instance specification change.

Name/ID	Node Type	Status	AZ	Private IP Address for Read	Fallover Priority	Operation
	Primary	Available	az3	192.168.0.90	1	View Metric
	Replica	Available	az3	192.168.0.107	1	View Metric   Promote to Primary

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

- During a manual switchover, there may be a brief disconnection lasting about 30 seconds. Ensure that your applications support automatic reconnection.
- During a manual switchover, the DB instance status is **Promoting to primary** and this process takes several seconds or minutes.

- After a switchover is complete, the node types of the original primary node and read replica have been exchanged, and the read replica status changes to **Available**.

---

**NOTICE**

- A read replica whose failover priority is **-1** cannot be promoted to the primary node.
  - Services may be intermittently interrupted for several seconds or minutes when a read replica is promoted to the primary node.
  - Promoting a read replica to primary will switch over the private IP addresses for read of the primary node and read replica. To ensure the services are not interrupted, connect to your DB instance using the private IP address from the **Network Information** area in the **Basic Information** page or the proxy address from the **Database Proxy** page. For details about the differences between the two addresses, see [Description of Each IP Address](#).
- 

----End

## Automatic Failover

GaussDB(for MySQL) uses a high availability active-active architecture that automatically fails over to a read replica automatically selected by the system.

Each read replica has a failover priority that determines which read replica is promoted if the primary node fails.

- Priorities range from 1 for the highest priority to 16 for the lowest priority.
- If two or more read replicas share the same priority, they have a same probability of being promoted to the new primary node.

GaussDB(for MySQL) selects a read replica and promotes it to the new primary node as follows:

1. Read replicas available for promotion are identified.
2. One or more read replicas with the highest priority are identified.
3. One of the read replicas with the highest priority is selected and promoted. If the promotion fails due to network faults or abnormal replication status, GaussDB(for MySQL) attempts to promote another read replica by priority and repeats the process until a read replica is successfully promoted.

## 9.5 Deleting a Pay-per-Use Read Replica

### Scenarios

You can manually delete read replicas billed on a pay-per-use basis on the **Instances** page.

#### NOTICE


Deleted read replicas cannot be restored. Exercise caution when performing this operation.


## Constraints

- You can delete a read replica only when a DB instance has two or more read replicas.
- If another operation is being performed on a DB instance, the read replicas of the instance cannot be manually deleted.
- For multi-AZ deployment, make sure that the primary node and remaining read replicas are located in different AZs after a read replica is deleted.  
If a primary node and a read replica are deployed in AZ1 and the other read replica is deployed in AZ2, the read replica in AZ2 cannot be deleted.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Node List** area, locate the read replica to be deleted and choose **More** > **Delete** in the **Operation** column.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

**Step 7** In the displayed dialog box, click **Yes**. Refresh the **Instances** page later to confirm that the deletion has completed.

----End

## APIs

- [Creating a Read Replica](#)
- [Deleting or Unsubscribing from a Read Replica](#)

## 9.6 Unsubscribing a Yearly/Monthly Read Replica

### Scenarios

You can unsubscribe a read replica of a yearly/monthly instance.


### Constraints

- You can unsubscribe a read replica only when the DB instance has two or more read replicas.
- Only isolated read replicas can be unsubscribed.
- If a read replica of a DB instance is being isolated, you cannot perform the following operations for the instance:
  - Creating read replicas
  - Scaling up storage space
  - Changing instance specifications
  - Rebooting the instance
  - Resetting the password
  - Upgrading the patch
  - Changing the private IP address
  - Changing the database port
  - Enabling or disabling SSL
  - Binding an EIP
  - Operations related to proxy instances
- The following operations cannot be performed on other read replicas of the instance:
  - Changing a failover priority
  - Promoting a read replica to primary
  - Isolating a read replica

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Node List** area, locate a read replica and choose **More > Isolate** in the **Operation** column.

**NOTICE**

- When a read replica is isolated, you can only unsubscribe or release it.
- When the workloads are heavy, you can release the isolated read replica if necessary.

**Step 6** After read replica status changes to **Isolated**, choose **More > Unsubscribe** in the **Operation** column.

**Figure 9-3** Unsubscribing a read replica

Name/ID	Node Type	Status	AZ	Private IP Address for Read	Fallover Priority	Operation
	Primary	Available	az1		1	View Metric
	Replica	Available	az1		2	View Metric   Promote to P   <b>Isolate</b>   <b>Unsubscribe</b>
	Replica	Isolated	az1		3	View Metric   Promote to Primary   <b>More</b>

**NOTE**

- It takes about 1 minute to isolate a read replica.
- When a read replica is isolated, read operations and database synchronization cannot be performed.
- To avoid being billed for an isolated read replica, unsubscribe it in a timely manner.

**Step 7** On the displayed page, confirm the order to be unsubscribed and select a reason. Then, click **Confirm**.

For unsubscription details, see [Unsubscription Rules](#).

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

**NOTICE**

After an unsubscription request is submitted, resources and data will be deleted and cannot be retrieved.

**Step 9** View the unsubscription result. After the order is successfully unsubscribed, the unsubscribed read replica of the instance will be deleted.

----End

**APIs**

- [Creating a Read Replica](#)
- [Deleting or Unsubscribing from a Read Replica](#)

# 10 HTAP Analysis

---

## 10.1 What Is HTAP?

Hybrid Transaction and Analytical Process (HTAP) instances are based on open-source ClickHouse. They use column-based storage engine and Single Instruction Multiple Data (SIMD) for parallel compute, improving query performance in massive data analysis, especially for large and wide tables.

HTAP instances free you independently maintaining data extraction and synchronization links, reduce data management costs, and provide simple and efficient real-time data analysis capabilities.

### Overview

An HTAP instance can be used as a standby database of a GaussDB(for MySQL) instance and provides high-performance data analysis capabilities. Data is synchronized to the HTAP instance in real time. You can perform online transaction processing and online data analysis on your GaussDB(for MySQL) DB instance.

### Supported Regions

HTAP instances are only available in the following regions:

- CN North-Beijing4
- CN East-Shanghai1
- CN South-Guangzhou
- AP-Singapore

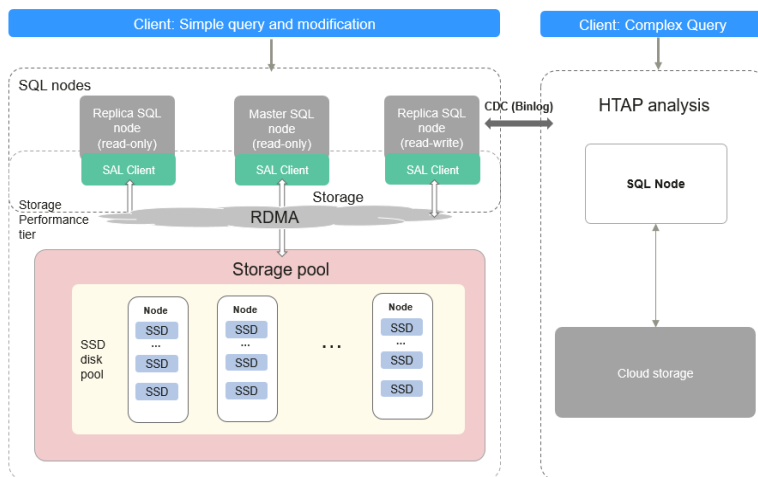
### Architecture

HTAP instances are deployed on ECSs and use extreme SSDs or ultra-high I/O disks.

You can [enable binlog of your GaussDB\(for MySQL\) instance](#) to synchronize data and operations to HTAP instances. The operations include inserting, deleting, modifying, and querying tables and changing table structures. After data is

synchronized to an HTAP instance, you can access the HTAP instance through its private IP address and EIP for data analysis.

**Figure 10-1** Architecture



## Features

- Multi-Version Concurrency Control (MVCC) and transaction-level read consistency  
You can select required isolation levels among four isolation levels by configuring parameters in [data synchronization task creation](#).
  - READ\_UNCOMMITTED: Read operations are not committed, and transaction consistency cannot be ensured.
  - READ\_COMMITTED: To ensure read consistency, read data is committed last.
  - QUERY\_SNAPSHOT: Snapshot query can avoid data deduplication and merging, providing high query performance and ensuring read consistency.
  - QUERY\_RAW: All raw data is returned, including data of different versions that have been deleted and updated.
- Quick deduplication  
Based on snapshots, data is quickly deduplicated to improve query performance.
- Data compression for storage  
In HTAP instances, data is compressed for storage by default, which greatly reduces storage costs under any given set of conditions.
- Parallel data synchronization  
In the initial full data synchronization phase, data is automatically sliced based on data statistics, and parallel processing improves synchronization performance. You can set the number of concurrent threads when creating a database for synchronization.
- Table definition rewriting  
When creating a synchronization task, you can modify tables to further improve the analysis and query performance. The modification operations

include ORDER BY, PARTITION BY, SAMPLE BY, PRIMARY KEY, TTL and COLUMNS.

- Table filtering based on a blacklist and a whitelist  
When creating a synchronization task, you can select required tables or excluded tables based on a blacklist and a whitelist.
- Binlog  
When a database has multiple tasks for data synchronization, one binlog is used to reduce network resource consumption.
- Higher stability for data replication  
Most GaussDB(for MySQL) DDLs are supported for synchronization. The character set of the source database can be automatically converted to the UTF-8 character set of the destination database.
- Various data types  
All data types of GaussDB(for MySQL) are supported. For details, see [Data Type Conversion](#).
- Aggregation of multiple data sources  
Data in multiple GaussDB(for MySQL) databases can be synchronized to the same HTAP instance.
- Enhanced security  
User account information is encrypted for storage.

## Billing

**Table 10-1** Billing items

Billing Item	Description
HTAP Instance	Yearly/monthly or pay-per-use
Storage space	Pay-per-use. If you select the storage space when purchasing an HTAP instance, the storage will be billed by the hour.
Public network traffic	GaussDB(for MySQL) instances are accessible from both private and public networks, but only the traffic from public networks is billed.

**Table 10-2** Specifications billing description for pay-per-use HTAP instances

Specification	Region	Price (USD/Hour)	
		Single	Primary/Standby
4 vCPUs   16 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	0.37	0.74



	AP-Singapore	0.544	1.088
8 vCPUs   32 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	0.75	1.50
	AP-Singapore	1.088	2.176
16 vCPUs   64 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	1.49	2.98
	AP-Singapore	2.176	4.352
32 vCPUs   128 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	2.98	5.96
	AP-Singapore	4.352	8.704
64 vCPUs   256 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	5.96	11.92
	AP-Singapore	8.704	17.408
88 vCPUs   352 GB	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	8.19	16.18
	AP-Singapore	11.968	23.936

**Table 10-3** Storage billing for pay-per-use HTAP instances

Storage	Region	Price (USD/GB/Hour)	
		Single	Primary/Standby
Ultra-high I/O	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	0.00022	0.00044
	AP-Singapore	0.00028	0.00056

Extreme SSD	CN North-Beijing 4, CN East-Shanghai 1, and CN South-Guangzhou	0.00065	0.0013
	AP-Singapore	0.001	0.002

## 10.2 Constraints

To improve the stability and security of HTAP instances, note the following constraints:

1. SQL statements are case-sensitive.
2. A primary key must be defined for a table during synchronization, or synchronization is skipped.
3. You need to select proper partition keys. If the number of partitions is greater than 2 but does not exceed 20, the RANGE partitions created in GaussDB(for MySQL) DB instance will be automatically synchronized to HTAP instances. Other partitions will not be synchronized. You can re-specify partition keys during synchronization task creation.
4. You need to properly define sorting keys and sort the query results based on the query frequency and left matching. When there are different sorting requirements, you can create materialized views.
5. Some DDL statements executed on GaussDB(for MySQL) DB instances cannot be synchronized to HTAP instances, which may cause synchronization failures or data inconsistencies. For details, see [Introducing DDLs that Cannot Be Synchronized](#).
6. HTAP instances use snapshots to accelerate query. Currently, projection-based query is not supported.
7. You can only insert base tables for materialized views. You cannot modify and delete data in materialized views. Materialized views cannot be created in synchronized databases. They can be created only in self-created databases.
8. Databases created during data synchronization cannot be modified. You can only modify the databases that you created on the HTAP instances.
9. The names of the databases and tables to be synchronized cannot contain Chinese characters.

## 10.3 Usage Rules

HTAP instances deliver good query performance. A single query uses more than half of the vCPUs. It is recommended that the QPS be less than 100.

To achieve good performance, you can use the following methods to optimize queries:

1. Simplify SQL statements by reducing invalid calculations, deleting unused fields, and avoiding SELECT.

2. Instead of querying all columns, delete those that are unnecessary.
3. Use the original data types to define field types (for example, do not define the TIME type as String). NULLABLE cannot be used to specify non-empty fields.
4. Manage data by partition. Partitions must be differentiated and are often used for queries. Only synchronize RANGE partitions created in GaussDB(for MySQL) to the HTAP instances. Do not synchronize other partitions. You can re-specify a partition key when creating a synchronization task. Too many partitions may affect performance.
5. Sort data based on the query frequency and remove unnecessary sort keys. Sort keys are critical to high performance, so you are advised to use up to 5 sort keys.
6. Add a hop index when the query condition is not in the sort key.
7. Instead of associating multiple tables to ensure data redundancy, use large wide tables. Use a small table as the right table when joining tables. You need to replace JOIN with IN.
8. Create data dictionary tables. If a table is frequently used and contains less than 500,000 records, you can use it as a data dictionary table. The dictionary tables are loaded to memory, reducing the number of JOIN operations and improving query performance.
9. Create materialized views for frequently used time-consuming queries to improve performance. It ensures that only the base table can be inserted and updated when creating a materialized view through multi-table association. The right table cannot be updated.



## 10.4 HTAP Instance Lifecycle Management

### 10.4.1 Creating an HTAP instance

#### Scenarios

You can create an HTAP instance on the management console.

#### Procedure

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **HTAP Analysis**. On the displayed page, click **Create HTAP Instance**.
- Step 6** On the **Create HTAP Instance** page, configure parameters as needed.

**Table 10-4** Parameter description

Parameter	Description
DB Instance Information	GaussDB(for MySQL) instance information is displayed, including the instance name, region, DB engine, primary AZ, time zone, billing mode, VPC, subnet, and security group.
Billing Mode	Pay-per-use
DB Instance Type	Single or primary/standby. <ul style="list-style-type: none"><li>● <b>Single:</b> There is only one node for an HTAP instance.</li><li>● <b>Primary/Standby:</b> There are two nodes that provide cross-AZ HA. Only the primary node provides services. If the primary node fails and cannot provide services, the standby node will take over services.</li></ul>
Storage Type	Extreme SSD and ultra-high I/O. <ul style="list-style-type: none"><li>● <b>Ultra-high I/O:</b> uses multi-disk striping to balance I/O loads among multiple disks, improving read and write bandwidth. The maximum throughput is 1.7 GB/s.</li><li>● <b>Extreme SSD:</b> uses 25GE network and RDMA to provide you with up to 1 million random read/write performance per disk and low latency per channel.</li></ul>
AZ	AZ where the HTAP instance is located.
Instance Specifications	Only general-enhanced is available.
Nodes	Currently, only single-node HTAP instances without read replicas can be created. A maximum of 10 HTAP instances can be created for each GaussDB(for MySQL).
Storage Space (GB)	Storage space that can be used by the HTAP instance. <ul style="list-style-type: none"><li>● If the storage type is extreme SSD, the default storage space is 40 GB, but it can be scaled up to 32 TB.</li><li>● When the storage type is ultra-high I/O, the default storage space is 50 GB, but it can be scaled up to 32 TB.</li></ul>
Time Zone	You need to select a time zone for your instance based on the region it is hosted in. The time zone is selected during instance creation and cannot be changed later.
Administrator	The default username is <b>root</b> .
Administrator Password	Must consist of 8 to 32 characters and contain at least three of the following types of characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%^*_-=+?,). Enter a strong password and periodically change it to improve security. Periodic password changes protect your account from brute force attacks and other risks.
Confirm Password	Must be the same as <b>Administrator Password</b> .

**Step 7** After configuration, click **Next**.

**Step 8** Confirm the configuration and click **Submit**.

**Step 9** On the HTAP instance list page, view and manage the HTAP instance.

----End

## 10.4.2 Rebooting an HTAP Instance

### Scenarios


You may need to reboot an HTAP instance for maintenance reasons.


### Constraints

- Only available HTAP instances can be rebooted.
- It takes about 1 to 2 minutes to reboot an HTAP instance. During the reboot, the instance is unavailable. Rebooting an HTAP instance will clear its cache. To prevent traffic congestion during peak hours, you are advised to reboot it during off-peak hours.

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Locate an HTAP instance and choose **More > Reboot** in the **Operation** column.

**Step 7** Click **Yes**. It takes about 1 to 2 minutes to reboot an instance.

----End

## 10.4.3 Deleting an HTAP Instance

### Scenarios

You can manually delete pay-per-use HTAP instances as required to release resources.


---


**NOTICE**

- HTAP instances cannot be deleted when operations are being performed on them.
  - Deleted HTAP instances cannot be recovered. Exercise caution when performing this operation.
- 

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Locate an HTAP instance and choose **More** > **Delete** in the **Operation** column.

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

----End

## 10.5 HTAP Instance Modifications

### 10.5.1 Binding and Unbinding an EIP

#### Scenarios

By default, an HTAP instance is not publicly accessible (not bound with an EIP) after being created. You can bind an EIP to the HTAP instance for public accessibility and can unbind the EIP from the HTAP instance if needed.

---

**NOTICE**

To ensure that the database can be accessed, the security group used by the database must allow access to the database port. For example, if the database port is **3306**, ensure that the security group allows access to the port **3306**.


---


#### Constraints

If an HTAP instance already has an EIP bound, that EIP needs to be unbound before a new one can be configured.

## Binding an EIP

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the **Network Information** area, click **Bind** in the **Public IP Address (EIP)** field.

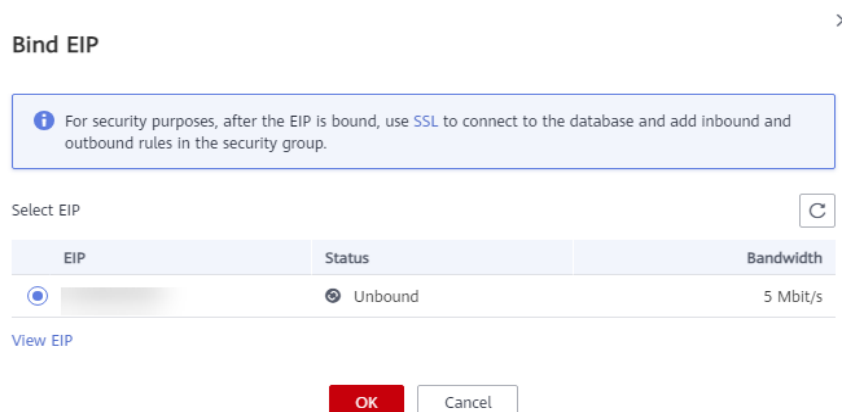
**Figure 10-2** Binding an EIP



**Step 8** In the displayed dialog box, select an EIP and click **OK**.

If no EIPs are available, click **View EIP** to create an EIP on the network console.

**Figure 10-3** Selecting an EIP



**Step 9** In the **Public IP Address (EIP)** field of the **Network Information** area, view the EIP that was bound.

To unbind the EIP from the instance, see [Unbinding an EIP](#).

----End

## Unbinding an EIP

- Step 1** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **HTAP Analysis**. Locate an HTAP instance that an EIP was bound to and click its name to go to the **Basic Information** page.
- Step 2** In the **Network Information** area, click **Unbind** in the **Public IP Address (EIP)** field.

**Figure 10-4** Unbinding an EIP



Network Information	Public IP Address (EIP)
VPC	Subnet
Security Group	TCP Port
HTTP Port	MySQL Protocol Port
Private IP Address	Public IP Address (EIP)

- Step 3** In the displayed dialog box, click **Yes**.
- Step 4** If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account and cloud product. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

- Step 5** To bind an EIP to the HTAP instance again, see [Binding an EIP](#).

----End

## 10.5.2 Scaling up Storage



### Scenarios

You can scale up the storage of an HTAP instance to meet service requirements.

### Constraints

- HTAP instances cannot be rebooted or deleted when the storage is being scaled up.
- The storage space can be scale down.

### Procedure

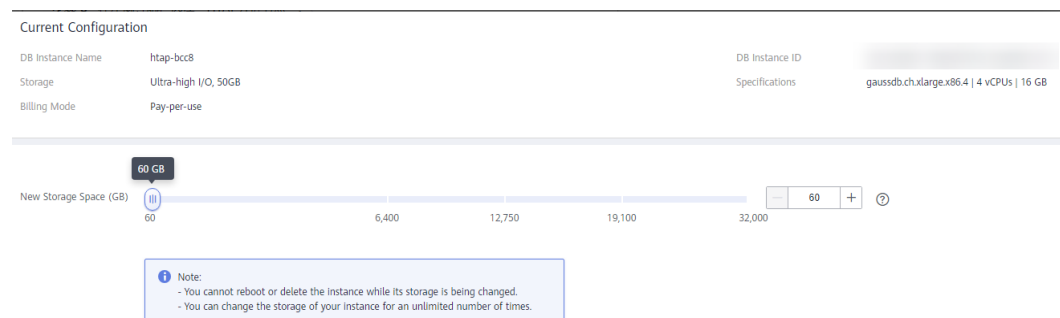
- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.



- Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **HTAP Analysis**.
- Step 6** Locate an HTAP instance and choose **More > Change Storage** in the **Operation** column.
- Step 7** Select the amount of storage you need. The maximum value is 32,000 GB.

The storage will be used for the system overhead required for inodes, reserved blocks, and database operation. The storage must be a multiple of 10.

**Figure 10-5** Scaling up storage



- Step 8** Click **Next**.
- Step 9** Confirm the new configuration, including the specifications, billing mode, and price.
- Step 10** Click **Submit**.

#### NOTE

Click **Go to Instance List** after the request is submitted and the storage of the HTAP instance is being scaling up.

If the **Change Storage** button in the **Operation** column is unavailable, the storage is currently being scaled. When the button becomes available, that means the scaling is complete.

----End

## 10.5.3 Changing vCPUs and Memory

### Scenarios

You can change the CPU and memory specifications of an HTAP instance if needed.

### Constraints


- An HTAP instance cannot be deleted when its vCPUs and memory are being changed.
- You can scale the specifications of your HTAP instance up or down.

- After you change vCPUs and memory, your HTAP instance will be rebooted and the cache will be cleared. You are advised to reboot the instance during off-peak hours.

## Procedure

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

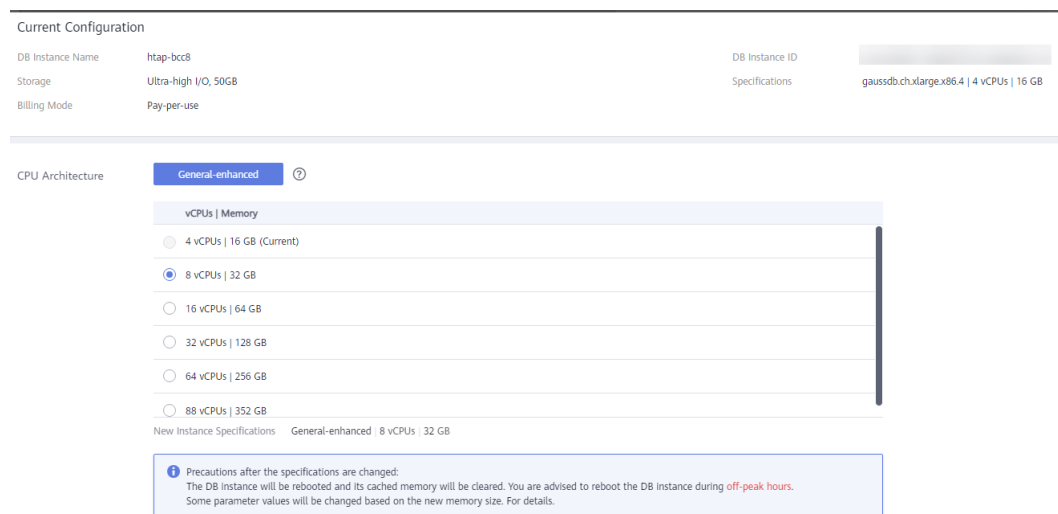
**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Locate an HTAP instance and choose **More** > **Change Instance Specifications** in the **Operation** column.

**Step 7** Reduce or increase the specifications as required and click **Next**.

**Figure 10-6** Changing vCPUs and memory



**Step 8** On the displayed page, confirm the specifications.

- If you need to modify your settings, click **Previous** to go back to the page where you specify details.
- If you do not need to modify your settings, click **Submit**.

**Step 9** View the new specifications.

On the HTAP instance list page, the instance status becomes **Changing instance specifications**.

----End

## 10.6 Managing Accounts

You can create and delete instance accounts, reset account passwords, and modify account permissions.

### System Accounts

To provide O&M services, the system automatically creates system accounts when you create HTAP instances. These system accounts are unavailable to you.

---

#### NOTICE


Deleting, renaming, and changing passwords or permissions for these accounts will cause the instance to run abnormally. Exercise caution when performing these operations.

- 
- **rdsAdmin**: a management account with superuser permissions, which is used to query and modify instance information, rectify faults, migrate data, and restore data.
  - **rdsMetric**: a metric monitoring account, which is used by watchdog to collect database status data.

### Creating an Account

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Accounts**. On the displayed page, click **Create Account**.

**Figure 10-7** Creating an account

**Create Account**

Username  ?

Authorization Scope  All databases  Specific databases

DML Permissions  Read and configure

Password

Confirm Password

**Table 10-5** Parameter description

Parameter	Example Value
Username	Contains 1 to 32 characters. It must start with a lowercase letter and end with a lowercase letter or digit. Only lowercase letters, digits, and underscores (_) are allowed.
Authorization Scope	<ul style="list-style-type: none"> <li>• All databases</li> <li>• Specific databases                             <ul style="list-style-type: none"> <li>- <b>Databases Not Authorized:</b> When creating an account, do not select any database in this area. The created account cannot perform operations on any database. To grant the permission of a database to the account, see <a href="#">Modifying Account Permissions</a>.</li> <li>- <b>Databases Authorized:</b> The databases selected in the <b>Databases Not Authorized</b> area are displayed.</li> </ul> </li> </ul>
DML Permissions	Includes the read-only permission and read and configure permission.

Parameter	Example Value
Password	<ul style="list-style-type: none"><li>• Contains 8 to 32 characters.</li><li>• Contains at least three of the following types of characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%^*_-=+?,.).</li><li>• Cannot be the username or the username backwards.</li></ul>
Confirm Password	Must be the same as the new password.

**Step 8** Click **OK**.

**Step 9** In the account list, view the account information including the username, authorized database, and DML permission.

 **NOTE**


You can [reset account passwords](#), [change account permissions](#), or [delete accounts](#).

----End

## Resetting Account Passwords

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Accounts**. On the displayed page, locate an account and click **Reset Password** in the **Operation** column.

**Step 8** In the displayed dialog box, enter a new password, confirm the password, and click **OK**.


----End


## Modifying Account Permissions

 **NOTE**

If you do not delete a database on the HTAP console but delete a database in other ways, permissions granted specifically for the database are not automatically deleted. They must be deleted manually.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Accounts**. On the displayed page, locate an account and choose **More > Change Permission** in the **Operation** column.


**Step 8** In the displayed dialog box, modify permissions as required and click **OK**.

----End

## Deleting an Account

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Accounts**. On the displayed page, locate an account and choose **More > Delete** in the **Operation** column.

**Step 8** In the displayed dialog box, confirm the information and click **OK**.

----End

## 10.7 Viewing Metrics

### Scenarios

The Cloud Eye service monitors HTAP instance statuses. You can obtain the monitoring metrics on the management console.

### Prerequisites

- The HTAP instance is running properly.

Cloud Eye does not display the metrics of faulty or deleted HTAP instances. When the status of HTAP instances becomes **Available**, you can view their metrics.

#### NOTE


If an HTAP has been faulty for 24 hours, Cloud Eye considers it to no longer exist and deletes it from the monitoring object list. You need to manually clear the alarm rules created for such instances.

- The HTAP instance has been properly running for at least 10 minutes.  
The monitoring data and graphics are available for a new HTAP instance after the instance runs for a period of time.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.


**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Locate an HTAP instance and click **View Metric** in the **Operation** column to go to the Cloud Eye console and view the monitoring metrics of the instance.

**Figure 10-8** Viewing metrics



#### NOTE

Cloud Eye can display performance metrics from the last 1 hour, last 3 hours, last 12 hours, last 24 hours or last 7 days. If you enable auto refresh, monitoring data is automatically refreshed every 60s.

-----End

## 10.8 Data Security

## 10.8.1 Configuring SSL

Secure Socket Layer (SSL) is an encryption-based Internet security protocol for establishing secure links between a server and a client. It provides privacy, authentication, and integrity to Internet communications. SSL:

- Authenticates users and servers, ensuring that data is sent to the correct clients and servers.
- Encrypts data to prevent data theft.
- Ensures data integrity during transmission.

SSL is enabled by default for HTAP instances. Enabling SSL increases the network connection response time and CPU usage, and you are advised to evaluate the impact on service performance before enabling SSL.

You can use a client to connect to an HTAP instance through a non-SSL or SSL connection.

- If SSL is enabled for the instance, you can connect to the instance using SSL, which is more secure.
- If SSL is disabled, you can only connect to the instance using a non-SSL connection.

---


### NOTICE

- Enabling or disabling SSL will reboot the instance immediately. During the reboot, the instance is unavailable. Rebooting an instance will clear its cache. To prevent traffic congestion during peak hours, you are advised to reboot it during off-peak hours.
  - After SSL encryption is enabled, the client can access the HTAP instance only through HTTPS port 8443, TCP port 9440, and MySQL port 3306. After SSL encryption is disabled, the client can access the HTAP instance only through HTTP port 8123, TCP port 9000, and MySQL port 3306.
- 

## Enabling SSL

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.


**Step 5** In the navigation pane on the left, choose **HTAP Analysis**. Locate an HTAP instance and click its name to go to the **Basic Information** page.

**Step 6** In the **DB Instance Information** area, click  in the **SSL** field.

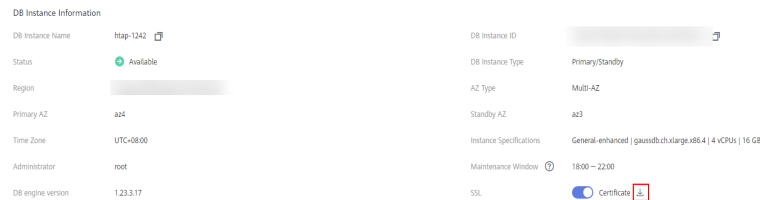



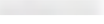



**Step 7** In the displayed dialog box, click **Yes**.

**Step 8** On the **Basic Information** page, view the results.

- To download the SSL certificate, click .

**Figure 10-9** Downloading the SSL certificate



DB Instance Information	
DB Instance Name	htap-1242 
Status	<span style="color: green;">●</span> Available
Region	
Primary AZ	az4
Time Zone	UTC+08:00
Administrator	root
DB engine version	1.23.3.17
DB Instance ID	
DB Instance Type	Primary/Standby
AZ Type	Multi-AZ
Standby AZ	az3
Instance Specifications	General-enhanced   gaussdb.ch.large.v86.4   4 vCPUs   16 GB
Maintenance Window	 18:00 - 22:00
SSL	<input checked="" type="checkbox"/> Certificate 

- To connect the HTAP instance using HTTPS, see [Connecting to an HTAP Instance Through HTTPS](#).

----End

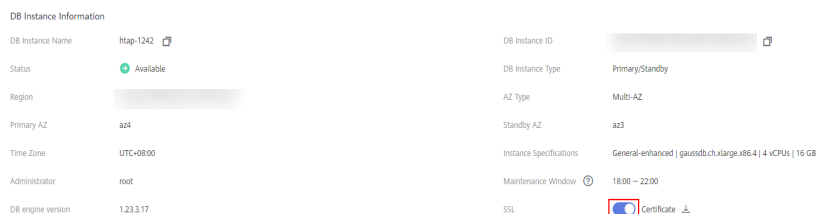
## Disabling SSL






**Step 1** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 2** In the navigation pane on the left, choose **HTAP Analysis**. Locate an HTAP instance and click its name to go to the **Basic Information** page.

**Step 3** In the **DB Instance Information** area, click  in the **SSL** field.

**Figure 10-10** Disabling SSL



DB Instance Information	
DB Instance Name	htap-1242 
Status	<span style="color: green;">●</span> Available
Region	
Primary AZ	az4
Time Zone	UTC+08:00
Administrator	root
DB engine version	1.23.3.17
DB Instance ID	
DB Instance Type	Primary/Standby
AZ Type	Multi-AZ
Standby AZ	az3
Instance Specifications	General-enhanced   gaussdb.ch.large.v86.4   4 vCPUs   16 GB
Maintenance Window	 18:00 - 22:00
SSL	<input type="checkbox"/> Certificate 

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

**Step 5** On the **Basic Information** page, view the results.

To connect the HTAP instance using a non-SSL connection, see [Connecting to an HTAP Instance Through JDBC](#).

----End

## 10.9 Data Synchronization

## 10.9.1 Introducing DDLs that Cannot Be Synchronized

- After a database synchronization task, only tables and data can be synchronized. Operations for databases, tablespaces, indexes, foreign keys, functions, stored procedures, triggers, views, partitions (DELETE operations), primary keys (INSERT/DELETE/ALTER operations), transactions, users, roles, and permissions, events cannot be synchronized.

Index- and partitioned table- related operations in [Table 10-6](#) and [Table 10-7](#) cannot be synchronized, but they do not affect data query and analysis on HTAP instance

- Table creation statements cannot contain CHECK or table options.
- During data synchronization, operations in [Table 10-8](#) may cause data inconsistency between the HTAP and GaussDB(for MySQL) instances. You are not advised to use these operations.

**Table 10-6** Index-related operations that do not support synchronization

DDL Name	SQL Example
Adding an index	ALTER TABLE tbl_name ADD INDEX name;
Renaming an index	ALTER TABLE tbl_name RENAME INDEX old_index_name TO new_index_name;
Deleting an index	DROP INDEX name ON table;
Adding a full-text index	CREATE FULLTEXT INDEX name ON table(column);
Adding a spatial index	ALTER TABLE geom ADD SPATIAL INDEX(g);
Modifying the type of an index	ALTER TABLE tbl_name DROP INDEX i1, ADD INDEX i1(key_part,...) USING BTREE;
Adding an index constraint	ALTER TABLE tbl_name ADD CONSTRAINT UNIQUE USING BTREE (column); ALTER TABLE tbl_name ADD CONSTRAINT UNIQUE USING HASH(column);

**Table 10-7** Partitioned table-related operations that do not support synchronization

DDL Name	SQL Example
Analyzing a table partition	ALTER TABLE {db}.tp ANALYZE PARTITION p0;
Checking a table partition	ALTER TABLE {db}.tp CHECK PARTITION p0;
Optimizing table partitions	ALTER TABLE {db}.tp OPTIMIZE PARTITION p0;

DDL Name	SQL Example
Re-building a table partition	ALTER TABLE {db}.tp REBUILD PARTITION p0;
Repairing a table partition	ALTER TABLE {db}.tp REPAIR PARTITION p0;
Creating a database	CREATE DATABASE ddl_test_2;
Modify the row format	ALTER TABLE tbl_name ROW_FORMAT = row_format;
Setting persistent table statistics	ALTER TABLE tbl_name STATS_PERSISTENT=0, STATS_SAMPLE_PAGES=20,STATS_AUTO_RECALC=1, ALGORITHM=INPLACE, LOCK=NONE;
Setting table character set	ALTER TABLE tbl_name CHARACTER SET = charset_name;
Converting table character set	ALTER TABLE tbl_name CONVERT TO CHARACTER SET charset_name;
Optimizing a table	OPTIMIZE TABLE tbl_name;
Rebuilding a table using the FORCE option	ALTER TABLE tbl_name FORCE;
Rebuilding a table without data	ALTER TABLE tbl_name ENGINE=InnoDB;
Renaming a tablespace	ALTER TABLESPACE tablespace_name RENAME TO new_tablespace_name;
Adding a table partition	ALTER TABLE {db}.tp ADD PARTITION (PARTITION p3 VALUES LESS THAN (2006) );
Specifying a sorting rule	ALTER TABLE tbl_name DEFAULT CHARACTER SET = utf8 COLLATE = utf8_general_ci;

**Table 10-8** DDL operations that cause data inconsistency

DDL Name	SQL Example
Deleting a primary key	ALTER TABLE tbl_name DROP PRIMARY KEY;
Adding a primary key	ALTER TABLE {db}.t1 ADD PRIMARY KEY (id);
Adding a primary key and deleting a primary key	ALTER TABLE tbl_name DROP PRIMARY KEY, ADD PRIMARY KEY (column);
Setting a primary key to NULL	ALTER TABLE tbl_name MODIFY COLUMN key_column_name data_type;

DDL Name	SQL Example
Changing the type of a primary key	ALTER TABLE tbl_name MODIFY COLUMN key_column_name data_type not null;
Adding a foreign key	ALTER TABLE tbl1 ADD CONSTRAINT fk_name FOREIGN KEY index (col1)REFERENCES tbl2(col2) referential_actions;
Deleting a foreign key	ALTER TABLE tbl DROP FOREIGN KEY fk_name;
Adding a column <b>NOTE</b> Common columns can be added. If columns contain the following default values, they cannot be added. <ul style="list-style-type: none"> <li>• Double quotation marks (")</li> <li>• Functions, character strings, and identifier that do not exist in HTAP instances</li> </ul>	ALTER TABLE tbl_name ADD COLUMN column_name column_definition c VARCHAR(10) DEFAULT (CONCAT('1', '2'));
Setting the default value of a column <b>NOTE</b> If columns contain the following default values, you cannot reset default values for the columns. <ul style="list-style-type: none"> <li>• Double quotation marks (")</li> <li>• Functions, character strings, and identifier that do not exist in HTAP instances</li> </ul>	ALTER TABLE tbl_name ALTER COLUMN col SET DEFAULT literal;
Changing NULL in tables to NOT NULL	ALTER TABLE tbl_name MODIFY COLUMN column_name data_type NOT NULL;
Changing the column name and type at the same time	ALTER TABLE t1 CHANGE b b1 VARCHAR(100);
Creating a table with a primary key	ALTER TABLE t1 ADD COLUMN (c2 INT GENERATED ALWAYS AS (c1 + 1)STORED);
Adding a STORED derived column	ALTER TABLE {db}.t1 ADD COLUMN (st2 INT GENERATED ALWAYS AS (c2 + 2)STORED), ALGORITHM=COPY;
Adding a VIRTUAL derived column	ALTER TABLE t1 ADD COLUMN (c2 INT GENERATED ALWAYS AS (c1 + 1)VIRTUAL);
Dropping a table partition	ALTER TABLE {db}.tp DROP PARTITION p4;
Discarding a table partition	ALTER TABLE {db}.tp DISCARD PARTITION p2 TABLESPACE;
Importing a table partition	ALTER TABLE {db}.tp IMPORT PARTITION p2 TABLESPACE;

DDL Name	SQL Example
Truncating a table partition	ALTER TABLE {db}.tp TRUNCATE PARTITION p2;
Truncating a partitioned table	TRUNCATE {db}.tp;
Coalescing table partitions	ALTER TABLE {db}.tp_hash COALESCE PARTITION 2;
Reorganizing table partitions	ALTER TABLE {db}.tp REORGANIZE PARTITION p0,p1,p2,p3 INTO ( ... );
Exchanging partitions	ALTER TABLE {db}.tp EXCHANGE PARTITION p0 WITH TABLE {db}.tp2;
Removing a table partition	ALTER TABLE {db}.tp REMOVE PARTITIONING;
Using a REPLACE clause	CREATE OR REPLACE TABLE;

## 10.9.2 Synchronizing Data

### Scenarios

To analyze data in some databases of your GaussDB(for MySQL) instance, you can synchronize data from your GaussDB(for MySQL) instance to an HTAP instance for analysis.

### Prerequisites

- The parameters in your GaussDB(for MySQL) instance have been configured based on [Table 10-9](#).

**Table 10-9** Parameter description

Parameter	Value	How to Modify
default_authentication_plugin	mysql_native_password	<a href="#">Modifying a Parameter Template</a>
binlog_expire_logs_seconds	86400 <b>NOTE</b> It is recommended that the binlog retention period be greater than one day (= 86,400s = 60 (seconds) x 60 (minutes) x 24 (hours). This prevents incremental replication failures caused by a short binlog retention period.	<a href="#">Modifying a Parameter Template</a>


Parameter	Value	How to Modify
log_bin <b>NOTE</b> To use this parameter, ensure that the GaussDB(for MySQL) kernel version is earlier than 2.0.45.230900.	ON	<a href="#">How Do I Enable and View Binlog of My GaussDB(for MySQL) Instance?</a>
rds_global_sql_log_bin <b>NOTE</b> To use this parameter, ensure that the GaussDB(for MySQL) kernel version is 2.0.45.230900 or later.	ON	<a href="#">How Do I Enable and View Binlog of My GaussDB(for MySQL) Instance?</a>
binlog_format	ROW	Run the <b>SHOW VARIABLES;</b> command to check the parameter value. If you need to change the parameter value, contact the customer service personnel.
binlog_row_image	FULL	Run the <b>SHOW VARIABLES;</b> command to check the parameter value. If you need to change the parameter value, contact the customer service personnel.
log_bin_use_v1_row_events	OFF	Run the <b>SHOW VARIABLES;</b> command to check the parameter value. If you need to change the parameter value, contact the customer service personnel.

- Databases and tables have been created for your GaussDB(for MySQL) instance.

## Creating a Synchronization Task

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

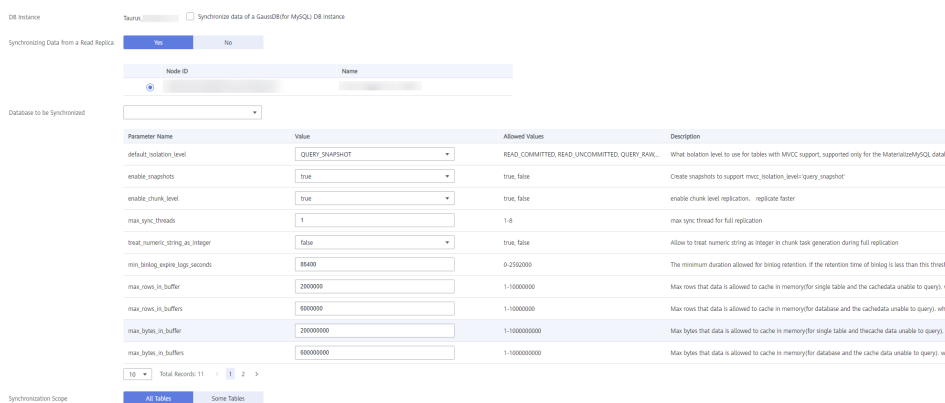
**Step 7** In the navigation pane on the left, choose **Data Synchronization**. On the displayed page, click **Create Synchronization Task**.

 **NOTE**

Before the synchronization, check the storage and connection status of the HTAP instance by referring to [Viewing Metrics](#).

**Step 8** Configure required parameters.

**Figure 10-11** Creating a synchronization task



- Currently, the databases whose name is Chinese cannot be synchronized.
- **Synchronizing Data from a Read Replica:** Select **Yes**: You need to select a read replica. Full data is synchronized from the selected read replica, preventing query pressure on the primary node during full synchronization. If there is only one read replica, this node is selected by default. During a full synchronization, ensure that the read replica is available, or you need to perform the synchronization again after the full synchronization fails.
- **Database to be Synchronized:** Select a database that the data will be synchronized to from the drop-down list. You can modify the database parameters as required.

**Table 10-10** Parameter description

Parameter	Default Value	Allowed Value	Description	Suggestion
-----------	---------------	---------------	-------------	------------

default_isolation_level	QUERY_SNAPSHOT	<ul style="list-style-type: none"> <li>- READ_COMMITTED</li> <li>- READ_UNCOMMITTED</li> <li>- QUERY_READ</li> <li>- QUERY_SNAPSHOT</li> </ul>	Controls the default isolation level. This parameter is suitable only for the MaterializeMySQL engine.	If the synchronized data will be modified later, set the parameter to <b>QUERY_SNAPSHOT</b> to accelerate queries, but generating snapshots will increase the data synchronization latency. If the synchronized data will not be modified later, set the parameter to <b>READ_COMMITTED</b> to reduce the data synchronization latency.
enable_snapshots	true	<ul style="list-style-type: none"> <li>- true</li> <li>- false</li> </ul>	Disables or enables snapshots. This parameter can be set to <b>true</b> only when <b>mvcc_isolation_level</b> is set to <b>query_snapshot</b> .	Set this parameter to <b>true</b> only when <b>default_isolation_level</b> is set to <b>QUERY_SNAPSHOT</b> .



enable_chunk_level	true	<ul style="list-style-type: none"> <li>- true</li> <li>- false</li> </ul>	Enables or disables chunk-level parallel replication.	If you want to quickly complete full synchronization, enable this function. However, during the synchronization, the read pressure of read replicas increases. If there is no time requirement on full synchronization, disable this function.
--------------------	------	---	---	--

max_sync_th reads	1	1 to vCPUs/2	Controls the maximum number of threads for full replication.	If you want to quickly complete full synchronization and HTAP instances deliver good performance, configure multiple threads. However, increasing this value will increase the read pressure of read replicas and affect the performance of HTAP instance. If there are no time requirements on the full synchronization, do not change this parameter value.
----------------------	---	--------------	--	---

max_rows_in_buffer	2000000	1 to 10000000	Controls the maximum number of rows of data that is allowed to cache in memory (for a single table and the cache data unable to query) When the number of rows exceeds the parameter value, the data will be materialized.	-
max_rows_in_buffers	6000000	1 to 10000000	Controls the maximum number of rows of data that is allowed to cache in memory (for databases and the cache data unable to query) When the number of rows exceeds the parameter value, the data will be materialized.	-

max_bytes_in_buffer	200000000	1~1000000000	Controls the maximum number of bytes of data that is allowed to cache in memory (for a single table and the cache data unable to query) When the number of bytes exceeds the parameter value, the data will be materialized.	-
max_bytes_in_buffers	600000000	1~1000000000	Controls the maximum number of bytes of data that is allowed to cache in memory (for databases and the cache data unable to query) When the number of bytes exceeds the parameter value, the data will be materialized.	-

max_flush_data_time	5000	1~10000	Controls how long that data is cached in the memory, in milliseconds.	If you want to shorten the data synchronization delay, set this parameter to a smaller value (at least 1000 ms). If the value is too small, the table may be fragmented, affecting the query performance. This parameter, <b>max_rows_in_buffer</b> , <b>max_rows_in_buffers</b> , <b>max_bytes_in_buffer</b> and <b>max_bytes_in_buffers</b> determine the time for flushing data to disks. If either of the conditions is met, the data is flushed to the disk.
treat_numeric_string_as_integer	false	- true - false	Whether numeric strings are treated as integers during block-level task generation during full synchronization.	-

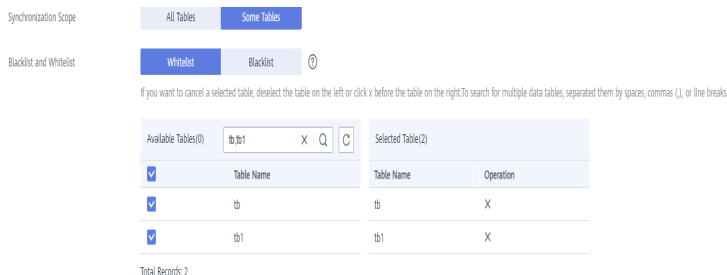
min_binlog_expire_logs_seconds	86400	0~2592000	Controls the minimum duration for storing binlogs. If the duration for storing binlogs is less than the value of this parameter, the synchronization will fail. <b>0</b> indicates that the duration is not limited.	-
--------------------------------	-------	-----------	--	---

- **Synchronization Scope:** Select **All Tables** or **Some Tables**.
- **Blacklist and Whitelist:** If **Synchronization Scope** is set to **Some Tables**, you need to configure tables for the blacklist or whitelist.

The tables in the whitelist or blacklist can be added or deleted. For details, see [Adding or Deleting Tables in the Blacklist or Whitelist](#).

**NOTE**


- The blacklist or whitelist cannot be empty.
- You need to configure table for either the whitelist or blacklist as required. If you select the whitelist, only the tables in the whitelist are synchronized. The tables in the blacklist cannot be synchronized.
- The tables to be synchronized must contain primary keys or a non-empty unique key. Otherwise, the tables cannot be synchronized to the HTAP instance.
- Extra disk space may be used during backend data combination and query. You are advised to reserve 50% free disk space for the system.
- To quickly find the desired tables to be added, you can enter multiple tables in the search box and click . Use commas (,), spaces (), and line breaks (\n) to separate multiple tables. The tables that you entered are selected by default in the **Available Tables** list and displayed in the **Selected Tables** list.



- **Table Synchronization:** Enable or disable table synchronization as required.
  - If you select **Enabled:**

- i. Select a table on the left and perform column operations, such as ORDER BY, PARTITION BY, SAMPLE BY, PRIMARY KEY, and TTL. For details about the syntax, see [Syntax](#).

**Figure 10-12** Table synchronization



Example:

- In ORDER BY(*COL1*, *COL2*), *COL1*, and *COL2* cannot be NULL.
- PARTITION BY toYYYYMM(*data\_time*);
- By default, PRIMARY KEY is the same as ORDER BY KEY and does not need to be specified.
- If SAMPLE BY is specified, ORDER BY must contain the content of SAMPLE BY, for example, **SAMPLE BY intHash32(UserID) ORDER BY (CounterID, EventDate, intHash32(UserID));**
- Do not set the TTL to a small value (at least one day), for example, **TTL time + INTERVAL 1 DAY**

**NOTE**

When configuring ORDER BY KEY, you need to ensure that the data is consistent after synchronization.


The value of PARTITION BY KEY cannot be too large, or the performance is affected.

- ii. After entering statements in **Column Operation** column, click **Verify** on the right of this area.
  - If you select **Disabled**, go to [Step 9](#).

**Step 9** After the settings are complete, click **Create Synchronization Task**.

**Step 10** On the details page, confirm the settings and click **Sync Now**. If you need to modify the information, click **Previous**.

**NOTE**

If you click **Previous** on the page or click  in the upper left corner of the page to return to the data synchronization page, a task to be synchronized will be generated. The status of the task is **Synchronization Stage: Waiting for synchronization**. To start the task, click **Synchronize** in the **Operation** column. To modify the task, see [Editing a Data Synchronization Task](#).

**Step 11** Click **Back to Synchronization List** to return to the data synchronization page. View details about the source database, destination database, status, and operations.

**Figure 10-13** Viewing task status




**NOTE**


If the status of a task is **Synchronization Stage: Incremental synchronization in progress**, the data synchronization is complete.

----End

## Viewing a Data Synchronization Task

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.


**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Data Synchronization**.

**Step 8** Locate a task and click **View** in the **Operation** column.

**Figure 10-14** Viewing a data synchronization task



**Step 9** View the detailed information about the task.


----End

## Editing a Data Synchronization Task

When the status of a task is **Synchronization Stage: Waiting for synchronization**, you can edit the task.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.




**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Data Synchronization**.

**Step 8** Locate a task and click **Edit** in the **Operation** column.

**Figure 10-15** Editing a data synchronization task



**Step 9** Configure required parameters.

**Step 10** After the settings are complete, click **Edit Synchronization Task**.

**Step 11** Confirm the settings and click **Sync Now**.

**Step 12** Click **Back to Synchronization List** to return to the data synchronization page. View details about the source database, destination database, status, and operations.

**Figure 10-16** Viewing task status




#### NOTE


If the status of a task is **Synchronization Stage: Incremental synchronization in progress**, the data synchronization is complete.

----End

## Deleting a Data Synchronization Task

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Data Synchronization**.

**Step 8** Locate a task and click **Delete** in the **Operation** column.

**Figure 10-17** Deleting a data synchronization task

**Step 9** In the displayed dialog box, click **Yes**.

----End

### 10.9.3 Adding or Deleting Tables in the Blacklist or Whitelist

You can add or delete tables in the blacklist or whitelist of an incremental data synchronization task.


#### Constraints

- Up to 50 tables can be selected in the whitelist or blacklist at a time. You are advised to select as small a number of tables as is appropriate based on workloads.
- The blacklist and whitelist cannot be switched.
- The blacklist or whitelist cannot be empty.
- You are advised to add or delete tables in the blacklist or whitelist during off-peak hours. If you perform this operation during peak hours, modifications cannot be updated. You can click **Stop** to stop the operation and reduce the number of tables to be added or deleted, or perform the operation during off-peak hours.

#### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.


**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.


**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Data Synchronization**, locate the synchronization task, and click **View** in the **Operation** column.

**Figure 10-18** Viewing a synchronization task

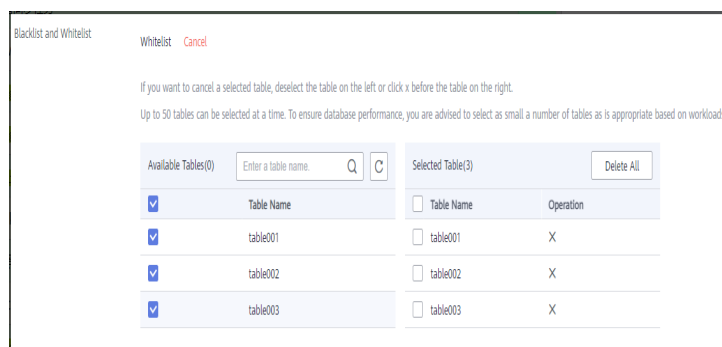
**Step 8** Click **Edit** next to **Blacklist and Whitelist**.

**Figure 10-19** Editing blacklist or whitelist



**Step 9** Add or delete tables in the whitelist or blacklist as required.

**Figure 10-20** Adding or deleting tables in the whitelist or blacklist




- To add a table to the whitelist or blacklist, select the table in the **Available Tables** list on the left.
- To delete a table from the whitelist or blacklist, locate the table in the **Selected Tables** list and click **X** in the **Operation** column. To delete multiple tables, select the desired tables and click **Delete All** in the upper right corner.
- To quickly find the desired tables to be added, enter multiple tables in the search box and click **Q**. Use commas (,), spaces (), and line breaks (\n) to separate multiple tables. The tables that you entered are selected by default in the **Available Tables** list and displayed in the **Selected Tables** list.

**Step 10** After the tables in blacklist and whitelist are modified, click **Modify Synch Task** at the lower right corner of the page.

**Step 11** Return to the data synchronization page and view the status of the task.

- When a table to be added contains a large amount of data, the message "Modification phase: Preparing data" is displayed in the **Status** column.
- You can click **View** to view the details about the modification. During the modification, you can click **Stop** to stop the modification.
- After the modification is complete, the status of the task becomes available. You can click **View** to view the modified results.

**Figure 10-21** Viewing the modified tables in the whitelist or blacklist

----End

## 10.10 Log Management

### 10.10.1 Viewing Slow Query Logs

Slow query logs record statements that exceed the value `slowlog_query_duration_ms` (1 second by default and its value cannot be modified). You can use these logs to identify and optimize the statements that are executing slowly.


HTAP instances support the following statement types:

- SELECT
- INSERT
- UPDATE
- DELETE
- CREATE
- ALTER
- DROP

#### Viewing Log Details

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.


**Step 6** Click the name of an HTAP instance to go to the **Basic Information** page.

**Step 7** In the navigation pane on the left, choose **Logs**.

**Step 8** On the **Slow Query Logs** page, view the slow query log details.

**Step 9** View slow query logs of different nodes in a given database and SQL statement types.

**Figure 10-22** Viewing slow query logs



- Enter a database name, click the drop-down list, and select your desired node.
- Click the drop-down list and select a SQL statement type (SELECT, INSERT, UPDATE, DELETE, CREATE, ALTER, or DROP).
- View logs generated in the last 15 minutes, last 30 minutes, last 1 hour, last 24 hours, last 7 days, or last 30 days. You can select **Custom time range** to set start time and end time to view the logs of the specified time range.

-----End

## Showing Original Log

### NOTE

After Show Original Log is enabled, original logs are displayed.

Original logs will be automatically deleted 30 days later. If the instance is deleted, its logs are also deleted.

**Step 1** On the **Instances** page, click the instance name to go to the **Basic Information** page.


**Step 2** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 3** Click the name of an HTAP instance to go to the **Basic Information** page.

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

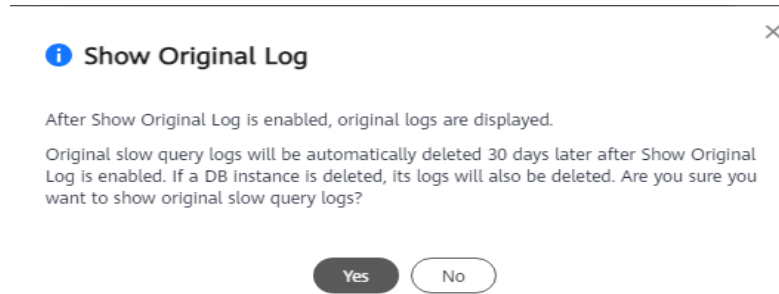
**Step 5** On the **Slow Query Logs** page, click next to **Show Original Log**.

**Figure 10-23** Enabling Show Original Log (1)




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

**Figure 10-24** Enabling Show Original Log (2)




----End

## Reporting Slow Log to LTS

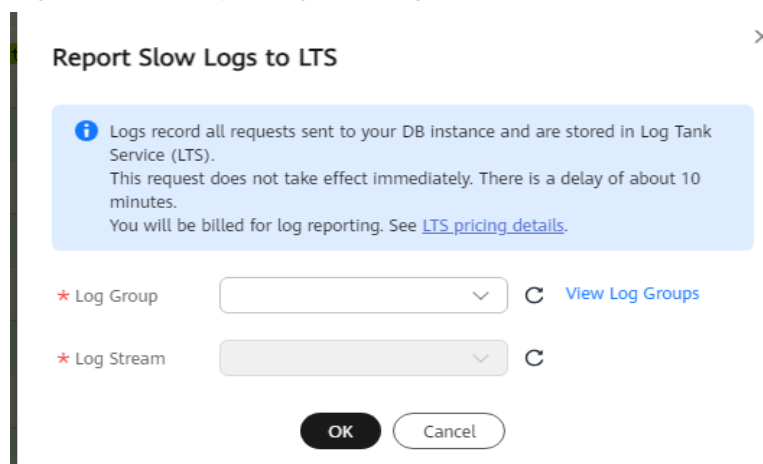
- Step 1** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 2** In the navigation pane on the left, choose **HTAP Analysis**.
- Step 3** Click the name of an HTAP instance to go to the **Basic Information** page.
- Step 4** In the navigation pane on the left, choose **Logs**.
- Step 5** On the **Slow Query Logs** page, click  next to **Report Slow Log to LTS**.

**Figure 10-25** Reporting slow logs to LTS (1)



- Step 6** Select an LTS log group and log stream and click **OK**.

**Figure 10-26** Reporting slow logs to LTS (2)



**NOTE**

- Audit logs record all requests sent to the HTAP instance and are stored in LTS.
- This function does not take effect immediately. There is a delay of about 10 minutes.
- You will be billed for this function. For pricing details, see [Billing Description](#).

----End

## 10.11 Connecting to an HTAP Instance

### 10.11.1 Connecting to an HTAP Instance Through DAS


Data Admin Service (DAS) enables you to manage DB instances on a web-based console, simplifying database management and improving efficiency and data security.

You can connect and manage HTAP instances through DAS. By default, you have remote login permissions. It is recommended that you use DAS to connect to the HTAP instances because this connection method is more secure and convenient than other methods.

#### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate a GaussDB(for MySQL) instance and click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **HTAP Analysis**.

**Step 6** In the instance list, locate an HTAP instance and click **Log In** in the **Operation** column.

**Figure 10-27** Logging in



Name/ID	DB Instance Type	Status	Billing Mode	Private IP Address	Database Port	Storage Type	Operation
htap-a151 29f1aa11db594be85...	Single	Available	Pay-per-use Created on Feb 04, 2024 16:46...		3306	Extreme SSD	<b>Log In</b>   New Metric   More >

**Step 7** On the displayed login page, enter the username and password and click **Log In**.

----End

## 10.11.2 Connecting to an HTAP Instance Through JDBC

### Prerequisites

- You are familiar with:
  - Computer basics
  - Java
  - JDBC knowledge
- You have downloaded official ClickHouse [JDBC driver](#)
- You can use the following command to connect to an HTAP instance through JDBC:

```
jdbc:clickhouse://<instance_ip>:<instance_port>/<database_name>
```

Parameter	Description
<i>&lt;instance_ip&gt;</i>	IP address of the instance. <b>NOTE</b> <ul style="list-style-type: none"><li>If you are accessing the instance through ECS, <i>instance_ip</i> is the private IP address of the instance. You can view the private IP address in the <b>Network Information</b> area on the <b>Basic Information</b>.</li><li>If you are accessing the instance through a public network, <i>instance_ip</i> indicates the EIP that has been bound to the instance. You can view the private IP address in the <b>Network Information</b> area on the <b>Basic Information</b>.</li></ul>
<i>&lt;instance_port&gt;</i>	Database port of the instance. The default port is <b>8123</b> . <b>NOTE</b> You can view the private IP address in the <b>Network Information</b> area on the <b>Basic Information</b> .
<i>&lt;database_name&gt;</i> >	Database name used for connecting to the instance.

### Sample Code

The following is an example of the Java code you could use for connecting to an HTAP instance. For details about other languages, see [ClickHouse official document](#).

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.Statement;
import java.sql.SQLException;
public class JDBCtest {
    static final String IP = "*.*.*.*"; //IP address of the instance
    //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 the code, set environment variables EXAMPLE_USERNAME_ENV and
    EXAMPLE_PASSWORD_ENV as needed.
```



```
static final String USER =System.getenv("EXAMPLE_USERNAME_ENV"); //username
static final String PASS = System.getenv("EXAMPLE_PASSWORD_ENV"); //password
public static void main(String[] args) {
    Connection conn = null;
    Statement stmt = null;
    String url = "jdbc:clickhouse://" + IP + ":8123";
    try {
        Class.forName("ru.yandex.clickhouse.ClickHouseDriver");
        conn = DriverManager.getConnection(url, USER, PASS);
        stmt = conn.createStatement();
        String sql = "show databases;";
        ResultSet rs = stmt.executeQuery(sql);
        int columns = rs.getMetaData().getColumnCount();
        for (int i = 1; i <= columns; i++) {
            System.out.print(rs.getMetaData().getColumnName(i));
            System.out.print("\t");
        }
        while (rs.next()) {
            System.out.println();
            for (int i = 1; i <= columns; i++) {
                System.out.print(rs.getObject(i));
                System.out.print("\t");
            }
        }
        rs.close();
        stmt.close();
        conn.close();
    } catch (SQLException se) {
        se.printStackTrace();
    } catch (Exception e) {
        e.printStackTrace();
    } finally {
        // release resource ....
    }
}
```

### 10.11.3 Connecting to an HTAP Instance Through HTTPS

#### Prerequisites

- You are familiar with:
  - Computer basics
  - Java
  - JDBC knowledge
- You have downloaded official ClickHouse [JDBC driver](#).
- SSL is enabled for an HTAP instance and SSL CA certificate has been downloaded. For details, see [Enabling SSL](#).

#### Precautions

Connecting to an HTAP instance through HTTPS increases the network connection response time and CPU usage. You are advised to use HTTPS only when the external network is used and encryption is required.

## Sample Code

The following is an example of the Java code you could use for connecting to an HTAP instance. For details about other languages, see [ClickHouse official document](#).

```
import java.sql.Connection;
import java.sql.DriverManager;
import java.sql.ResultSet;
import java.sql.Statement;
import java.sql.SQLException;
public void run() throws InterruptedException {
    final ClickHouseProperties clickHouseProperties = new ClickHouseProperties();
    clickHouseProperties.setSslRootCertificate("/user/CA.pem");
    clickHouseProperties.setSsl(true);
    clickHouseProperties.setSslMode("none");
    clickHouseProperties.setUser("test");
    clickHouseProperties.setPassword("123456Aa");
    clickHouseProperties.setSocketTimeout(2 * 3600 * 1000);
    final BalancedClickhouseDataSource dataSource = new BalancedClickhouseDataSource
("jdbc:<External network address, read/write internal network address>:<HTTPS port number>/
<Database name>?ssl=true", clickHouseProperties);
    try {
        final ClickHouseConnection conn = dataSource.getConnection();
        conn.createStatement().executeQuery("select now()");
    } catch (Throwable e) {
        e.printStackTrace();
    }
}
```

## 10.12 Developer Guide

### 10.12.1 Data Type Conversion

When data of GaussDB(for MySQL) instances is synchronized to HTAP instances, the data type will be converted. For details, see [Table 10-11](#).

**Table 10-11** Data type conversion

Data Type	GaussDB(for MySQL) Data Type	HTAP Data Type
NUMERTIC TYPES	TINYINT	Int8
	SMALLINT	Int16
	MEDIUMINT	Int32
	INTEGER/INT	Int32
	BIGINT	Int64
	DECIMAL/NUMERIC	Decimal
	FLOAT	Float32

	DOUBLE	Float64
	BIT	UInt64
DATE TIME TYPES	DATE	Date
	DATETIME	DateTime
	TIMESTAMP	DateTime
	TIME	String
	YEAR	UInt16
String Data Types	CHAR	String
	VARCHAR	String
	BINARY	FixedString
	VARBINARY	String
	BLOB	String
	TEXT	String
	ENUM	Enum
	SET	SET index
Spatial Data Types	GEOMETRY	String
	POINT	String
	LINestring	String
	POLYGON	String
	MULTIPOINT	String
	MULTILINestring	String
	MULTIPOLYGON	String
	GEOMETRYCOLLECTION	String
JSON Data Types	JSON	String

## 10.12.2 Syntax

HTAP supports the native ClickHouse syntax (case sensitive). For details, see [Official documents of ClickHouse](#).

Pay attention to the following:

When creating a data synchronization task, you can set **Table Synchronization to Enabled** and perform column operations on the management console. The

column operations for table synchronization are implemented through TABLE OVERRIDE. The database synchronization syntax is as follows:

```
CREATE DATABASE db_name ENGINE = MaterializedMySQL('host:port', 'db_name', 'user', 'password')  
[SETTINGS ] TABLE OVERRIDE t1 (ORDER BY col1), TABLE OVERRIDE t2 (ORDER BY col2)
```

TABLE OVERRIDE syntax:

```
TABLE OVERRIDE table_name ( function_name (col), another_function_name(col2))
```

Replace *function\_name* with ORDER BY, PARTITION BY, SAMPLE BY, PRIMARY KEY, or TTL.

For details about data synchronization, see [Data Synchronization](#).

### 10.12.3 Database Status

After databases are connected, run the following commands to check the database list and synchronization status:

- View the kernel version.

```
select version();
```

- Check database status.

```
show databases;
```

- Check synchronization status.

Querying binlog information will occupy the hourly query quota of GaussDB(for MySQL). You are advised to run this command at most once a minute.

```
show full slave status;
```

- Check snapshot status.

```
show snapshot status;
```

The **status** field is valid only when snapshots are enabled.

- **READY**: ready for query.
- **NOT\_READY**: not ready and cannot be used.
- **ERROR**: synchronization or update errors.
- **RECONNECTING**: Synchronization link disconnected.
- **PREPARING\_UPDATE**: ready for update.
- **UPDATING**: updating.

### 10.12.4 DB Engines

Databases synchronized from the GaussDB(for MySQL) instances use the MaterializedMySQL engine by default. Such database can only be queried but cannot be modified. If you want to use another DB engine, you can create a database and specify the required DB engine.

### 10.12.5 Table Engines

Tables to be synchronized use the ReplacingMultiVersionMergeTree engine by default.

# 11 Database Management

---

## 11.1 Creating a Database

### Scenarios

After your GaussDB(for MySQL) instance is created, you can create databases on it.


### Constraints

- This operation is not allowed when another operation is being performed on your DB instance.
- After a database is created, the database name cannot be changed.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Databases**. On the displayed page, click **Create Database**. In the displayed dialog box, enter a database name, select a character set, and authorize permissions for users. Then, click **OK**.

**Figure 11-1** Creating a database

**Create Database**

You can select up to 50 users at a time.

Database Name

Character Set  utf8mb4  utf8  latin1  gbk [Show](#)

User

User Not Authorized (0)		Authorized User (1)		
User	Host IP ...	Username	Permission	Operation
<input checked="" type="checkbox"/>	auto_user	%	<input checked="" type="radio"/> Read only <input type="radio"/> Read and write	x

Remarks

0/512

**OK** Cancel


- The database name can contain 1 to 64 characters. Only letters, digits, hyphens (-), and underscores (\_) are allowed. The total number of hyphens (-) cannot exceed 10.
- You can click **Show** to view all character sets.
- You can select one or more unauthorized users. If there are no unauthorized users, you can [create one](#).
- If you require fine-grained permissions control, log in to the database through the DAS console.


**Step 6** After the database is created, authorize or delete it on the **Databases** page. You can search for the desired database by character set and database name.

----End

## Creating a Database Through DAS

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance and click **Log In** in the **Operation** column.

**Step 5** On the displayed login page, enter the username and password and click **Log In**.

**Step 6** Choose **SQL Operations** > **SQL Query**, and enter the following command to create a database:

```
create database Database name;
```

```
----End
```

## APIs

- [Creating a Database](#)
- [Querying Databases](#)
- [Querying Available Database Character Sets](#)
- [Modifying Database Remarks](#)

# 11.2 Deleting a Database

## Scenarios

You can delete databases that you have created.

### NOTICE


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


## Constraints

This operation is not allowed when another operation is being performed on your DB instance.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Databases**. On the displayed page, locate the database that you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

```
----End
```

## APIs

- [Deleting Databases](#)
- [Creating a Database](#)



# 12 Account Management (Non-Administrator)

---

## 12.1 Creating a Database Account

### Scenarios

When you create a GaussDB(for MySQL) instance, account **root** is created at the same time by default. You can add other accounts as required.


### Constraints

This operation is not allowed when another operation is being performed on your DB instance.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Accounts**. On the displayed page, click **Create Account**. In the displayed dialog box, enter a username, authorize permissions for databases, enter a password, and confirm the password. Then, click **OK**.

- The username can consist of 1 to 32 characters. Only lowercase letters, digits, and underscores (\_) are allowed.
- You can select one or more unauthorized databases and authorize their permissions to the account. If there are no unauthorized databases, you can

**create ones.** You can also **modify the database permissions** after the account is created.

 **NOTE**

If you do not delete a database on the GaussDB(for MySQL) console but delete a database in other ways, permissions granted specifically for the database are not automatically deleted. They must be deleted manually. This is an open-source MySQL behavior. For details, see **DROP DATABASE Statement**.

- The password can consist of 8 to 32 characters, including at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%&^\*\_-=+?,()&). It cannot be the same as the username or the username spelled backwards.
- If you require fine-grained permissions control, log in to the database through the **DAS** console.


**Step 6** After the account is created, you can manage it on the **Accounts** page of the selected instance.

----End

## Creating a Database Account Through DAS

**Step 1** **Log in to the management console.**

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the target instance and click **Log In** in the **Operation** column.

**Step 5** On the displayed login page, enter the username and password and click **Log In**.

**Step 6** Choose **SQL Operations > SQL Query** and enter the following command:

```
create user Account name;
```

----End

## APIs

- [Creating a Database Account](#)
- [Querying Database Users](#)
- [Modifying Remarks of a Database User](#)

## 12.2 Resetting a Password for a Database Account

### Scenarios


You can reset passwords for the accounts you have created. To protect your instance against brute force cracking, change your password periodically, such as every three or six months.


## Constraints

This operation is not allowed when another operation is being performed on your DB instance.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Accounts**. On the displayed page, locate the target account and click **Reset Password** in the **Operation** column.

**Step 6** In the displayed dialog box, enter a new password, confirm the new password, and click **OK**.

- The password must consist of 8 to 32 characters and contain at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#\$%^\*\_-=+?,()&).
- The password you entered in the **Confirm Password** text box must be the same as that you entered in the **New Password** text box.
- It cannot be the username or the username spelled backwards.
- After the password is reset, the database will not be rebooted and permissions will not be changed.
- You can query password reset records on the CTS console. For details, see the [Cloud Trace Service User Guide](#).

**Step 7** If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

----End

## APIs

- [Changing Password of a Database User](#)
- [Querying Database Users](#)

## 12.3 Changing Permissions for Database Accounts

### Scenarios


You can authorize custom database users to specified databases and revoke permissions for authorized databases.


### Constraints

This operation is not allowed when another operation is being performed on your DB instance.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Accounts**. On the displayed page, locate the account and choose **More > Change Permission** in the **Operation** column.

**Step 6** In the displayed dialog box, select one or more unauthorized databases and authorize their permissions to the account. To delete a selected database, locate the database and click × in the **Operation** column.

----End

### APIs

- [Authorizing Permissions to a Database User](#)
- [Deleting Permissions of a Database User](#)

## 12.4 Deleting a Database Account

### Scenarios

You can delete database accounts you have created.

---

**NOTICE**

Deleted database accounts cannot be restored. Exercise caution when deleting an account.


---


## Constraints

This operation is not allowed when another operation is being performed on your DB instance.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Databases**. On the displayed page, locate the database that you want to delete and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

**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 obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

----End

## APIs

[Deleting a Database User](#)

# 13 Data Security

---

## 13.1 Resetting the Administrator Password

### Scenarios

If you forget the password of the administrator account, you can reset the password.

If an error occurs on the **root** account, for example, if your **root** account credentials are lost or deleted, you can restore the **root** account permissions through resetting the password.

You cannot reset the administrator password under the following circumstances:



- Your account is frozen.
- The database port is being changed.
- The instance status is **Creating**, **Restoring**, **Rebooting**, **Changing port**, **Changing instance specifications**, **Promoting to primary**, or **Abnormal**.

### Precautions

- If you have changed the administrator password of a DB instance, the passwords of the read replicas associated with the instance will also be changed accordingly.
- The time it takes for the new password to take effect depends on the amount of service data currently being processed by the primary node.
- To protect against brute force hacking and improve system security, change your password periodically, such as every three or six months.
- The instance may have been restored from a backup before you reset the administrator password.

### Method 1

- Step 1** [Log in to the management console.](#)

- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, locate the instance for which you want to change the password and choose **More** > **Reset Password** in the **Operation** column.
- 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 obtained verification code, and click **Verify** to close the page.
- For details about how to enable operation protection, see [Identity and Access Management User Guide](#).
- Step 6** Enter a new password and confirm the password.

---

**NOTICE**

Keep this password secure. If lost, the system cannot retrieve it.



---

The new password must consist of 8 to 32 characters and contain at least three of the following: uppercase letters, lowercase letters, digits, and special characters (~!@#%^\*\_-=+?,()&\$).

- To submit the new password, click **OK**.
- To cancel the reset, click **Cancel**.

----End

## Method 2

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the **DB Instance Information** area, click **Reset Password** in the **Administrator** field. In the displayed dialog box, enter and confirm the new password.

---

**NOTICE**

Keep this password secure. If lost, the system cannot retrieve it.

---

The new password must consist of 8 to 32 characters and contain at least three of the following: uppercase letters, lowercase letters, digits, and special characters (~!@#%^\*\_-=+?,()&\$).

----End

## 13.2 Changing a Security Group


### Scenarios

You can change the security group associated with your instance.


### Procedure



**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Network Information** area on the **Basic Information** page, click  in the **Security Group** field and select a new security group.

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

**Step 6** Click  in the upper right corner on the **Basic Information** page to view the result of the change. This process takes about 1 to 3 minutes.

----End

### APIs

[Modifying a Security Group](#)

## 13.3 Configuring Database Account Security

### Password Strength Requirements

For database password strength requirements on the management console, see the database configuration table in [Buying a DB Instance](#).

GaussDB(for MySQL) has a password security policy for newly created database users. Passwords must:

- Consist of at least eight characters.



- Contain at least three types of the following characters: uppercase letters, lowercase letters, digits, and special characters (~!@#%^\*\_-=+?,()&\$).

When you create instances, your password strength is checked. You can modify the password strength as user **root**. For security reasons, use a password that is at least as strong as the default password.

## Account Description

To provide O&M services, the system automatically creates system accounts when you create instances. These system accounts are unavailable to you.

---

### NOTICE

Deleting, renaming, and changing passwords or permissions for these accounts will cause the instance to run abnormally. Exercise caution when performing these operations.

---

- **rdsAdmin**: a management account with superuser permissions, which is used to query and modify instance information, rectify faults, migrate data, and restore data.
- **rdsRepl**: a replication account, which is used to synchronize data from the primary node to read replicas.
- **rdsBackup**: a backup account, which is used to back up data in the background.
- **rdsMetric**: a metric monitoring account, which is used by watchdog to collect database status data.
- **rdsProxy**: a database proxy account, which is used for authentication when the database is connected through the proxy address. This account is automatically created when you enable read/write splitting.

## 13.4 Configuring SSL

Secure Socket Layer (SSL) is an encryption-based Internet security protocol for establishing an encrypted link between a server and a client. It provides privacy, authentication, and integrity to Internet communications. SSL:

- Authenticates users and servers, ensuring that data is sent to the correct clients and servers.
- Encrypts data, preventing it from being intercepted during transmission.
- Ensures data integrity during transmission.

SSL is enabled by default. Enabling SSL increases the network connection response time and CPU usage, and you are advised to evaluate the impact on service performance before enabling SSL.

You can use a client to connect to your DB instance through a non-SSL or SSL connection.

- If SSL is enabled for your DB instance, you can connect to your DB instance using SSL, which is more secure.

- If SSL is disabled, you can only connect to your DB instance using a non-SSL connection.

---


#### NOTICE


Enabling or disabling SSL will cause the instance to be rebooted immediately and temporarily unavailable. You are advised to perform this operation during off-peak hours.

---

## Enabling SSL

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **DB Instance Information** area, click  in the **SSL** field.


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


**Step 7** On the **Basic Information** page, view the results.

----End

## Disabling SSL

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **DB Instance Information** area, click  in the **SSL** field.

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

**Step 7** On the **Basic Information** page, view the results.

----End

## APIs

### [Enabling or Disabling SSL](#)

## 13.5 Enabling TDE

Transparent Data Encryption (TDE) performs real-time I/O encryption and decryption on data files. Data is encrypted before being written to disks and is decrypted when being read from disks to memory. This effectively protects the security of databases and data files.

### Supported Regions

CN South-Guangzhou

### Constraints on Usage


- To enable TDE, submit a service ticket by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.
- To configure TDE, you must have the iam:agencies:createServiceLinkedAgencyV5 permission. If you do not have this permission, [create a custom policy](#).
- You need to enable Key Management Service (KMS) for your DB instance first. The data keys used for encryption are generated and managed by KMS. GaussDB(for MySQL) does not provide any keys or certificates required for encryption.
- To enable TDE, the kernel version of GaussDB(for MySQL) instances must be 2.0.47.231100 or later.
- Your DB instance must be billed at a pay-per-use or yearly/monthly basis.
- The instance type must be single-node or primary/standby deployment.
- TDE can be enabled only when a DB instance is created. After the instance is created, TDE cannot be enabled or disabled.
- TDE encrypts instance data, including full backups but excluding incremental backups.
- After TDE is enabled, the cryptographic algorithm cannot be changed later.
- Only instance-level encryption is supported.
- After TDE is enabled for a DB instance, you cannot:
  - Enable cross-region backup for the DB instance.
  - Restore the data of the DB instance to an existing DB instance.

### Procedure

**Step 1** Go to the [Buy DB Instance](#) page.

**Step 2** On the displayed page, set **TDE** to **Enabled** and select the corresponding cryptographic algorithm.

**Figure 13-1** Enabling TDE



**Step 3** After the DB instance is created, click the DB instance name to go to the **Basic Information** page and view the **TDE** field.

----End

# 14 Data Backups

---

## 14.1 Backup Principles


GaussDB(for MySQL) instances support automated and manual backups. You can periodically back up data in your DB instance, so if a DB instance becomes faulty or data is damaged, you can restore it using backups.

- **Automated backup**

You can click [Configure Some-Region Backup Policy](#) on the management console, and the system will automatically back up your instance data based on the time window and backup cycle you specify in the backup policy and will store the data for as long as you have configured the retention period for.

- Automated backups cannot be manually deleted. To delete them, you adjust the retention period specified in your [same-region backup policy](#). Retained backups (including full and incremental backups) will be automatically deleted at the end of the retention period.
- A full backup means that all data in your DB instance is backed up. In an incremental backup, only data that has changed within a certain period is backed up.

Incremental backups are created based on the most recent full backup, as shown in [Figure 14-1](#), so the most recent full backup that exceeds the retention period is still retained. For details, see the following example.

**Figure 14-1** Backup restoration

- **Manual backup**

Manual backups are user-initiated full backups of your DB instance. They are retained until you **delete them manually**.

Regularly backing up your DB instance is recommended, so if your DB instance becomes faulty or data is corrupted, you can restore it using backups.

## Backup Principles

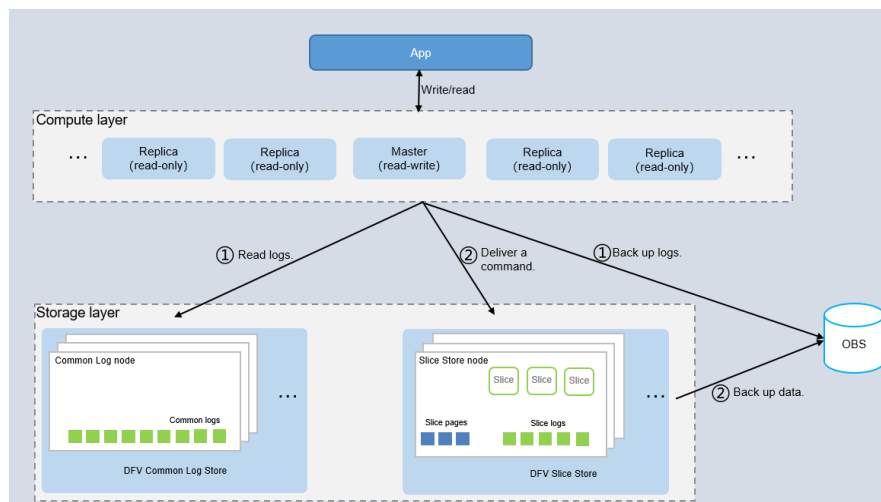
GaussDB(for MySQL) is based on Huawei's latest-generation DFV storage, which decouples storage from compute. The compute layer provides services for external systems and manages logs, and the storage layer stores the data. The storage layer consists of Common Log nodes and Slice Store nodes.

As shown in [Figure 14-2](#), the creation of backups involves in the compute layer and storage layer.

- The primary node at the compute layer reads the log content of the Common Log node at the storage layer and backs it up to OBS.
- The primary node at the compute layer sends a command for backing up data to the Slice Store node at the storage layer. The Slice Store node backs up data to OBS.

During the creation of a backup, the CPU usage and memory usage of the primary node of your instance increase slightly, but you will not notice anything at the storage layer. The final backup is stored in OBS as multiple data files and does not use up any of the disk space of the instance.

Figure 14-2 Backup principles



## 14.2 Backup Types

To ensure data reliability:

- GaussDB(for MySQL) supports multiple backup types. Automated backup is enabled by default. For details about automated backup, see [Configuring a Same-Region Backup Policy](#).
- GaussDB(for MySQL) periodically creates full backups and incremental backups for your instance. You can restore data to a specific point of time using a backup. For details, see [Restoring Instance Data to a Specific Point in Time](#).
- GaussDB(for MySQL) backs up data at the DB instance level.
- Automated and manual backup files cannot be downloaded.

Based on different dimensions, there are the following backup types:

- Full backup and incremental backup based on data volume
- Manual backup and automated backup based on backup methods
- Same-region backup and cross-region backup based on backup regions

**Table 14-1** Comparison of backup types

Backup Type	Enabled by Default	Retention Period	Description	How to View
Full backups	Yes	Full backups are retained till the retention period expires.	<ul style="list-style-type: none"><li>• A full backup is to back up all data of your DB instance in the current point of time.</li><li>• You can use a full backup to restore the complete data generated when the backup was created.</li><li>• Full backups include automated backups and manual backups.</li></ul>	Click the instance name. On the <b>Backups</b> page, click the <b>Full Backups</b> tab and view the backup size.
Incremental backups	Yes	Incremental backups are retained till the retention period expires.	<ul style="list-style-type: none"><li>• GaussDB(for MySQL) automatically backs up the updated data since the last automated or incremental backup was made every five minutes or after a certain amount of data is reached.</li><li>• Incremental backups are automated backups.</li></ul>	Click the instance name. On the <b>Backups</b> page, click the <b>Incremental Backups</b> tab and view the backup size.



Backup Type	Enabled by Default	Retention Period	Description	How to View
Same-region backups	Yes	1 to 732 days	Backup files are stored in the same region as the DB instance. Same-region backup (automated backup) is enabled by default and cannot be disabled.	<ul style="list-style-type: none"><li>• If cross-region backup is available: Click <b>Backups</b> in the left navigation pane. On the <b>Same-Region Backups</b> tab, view the backup size.</li><li>• If cross-region backup is not available: Click <b>Backups</b> in the left navigation pane and view the backup size.</li></ul>
Cross-region backups	No	1 to 1,825 days	GaussDB(for MySQL) can store backup files in a region different from the region where your instance is located. After you enable cross-region backup, the backup files are automatically stored in the region you specify.	Click <b>Backups</b> in the left navigation pane. On the <b>Cross-Region Backups</b> tab, view the backup size. <b>NOTE</b> To configure cross-region backup policies, contact customer service.
Automated backups	Yes	1 to 732 days	GaussDB(for MySQL) saves automated backups based on the retention period you specified.	-
Manual backups	Yes	Manual backups are always retained until you delete them manually.	Manual backups are user-initiated full backups of instances. They are retained until you delete them manually.	-

## 14.3 Backup Space and Billing

### Category

- Full backup: All data is backed up even if no data has changed since the last backup.
- Incremental backup: The system automatically backs up data that has changed since the last automated backup or incremental backup in binlogs every 5 minutes. The binlogs can be used to restore data to a specified point in time.
- Differential backup: The system backs up data that has changed since the most recent full backup or differential backup in to physical files. Physical files cannot be used for log replay.
- Billable space: Backup space that you are billed for.
- Logical space: Space occupied by full backups.
- Physical space: The amount of data that is backed up to OBS.

#### NOTE

After you purchase a DB instance, the logical space is the same as the physical space. When a backup starts in a backup chain, the physical space stores the data of the first full backup and subsequent differential backups.

### Backup Space Calculation Methods

There is a default backup chain (where there are seven backups). The first automated backup is a full backup, and subsequent automated backups are differential backups.

In a backup chain, the backup space is released only after all full backups and differential backups are deleted.

- Logical space: Total size of the logical space – Logical size of the expired backup file
- Physical space: Size of the first full backup file + total size of subsequent differential backup files
- Free space: There is free backup storage up to 100% of your provisioned database storage.

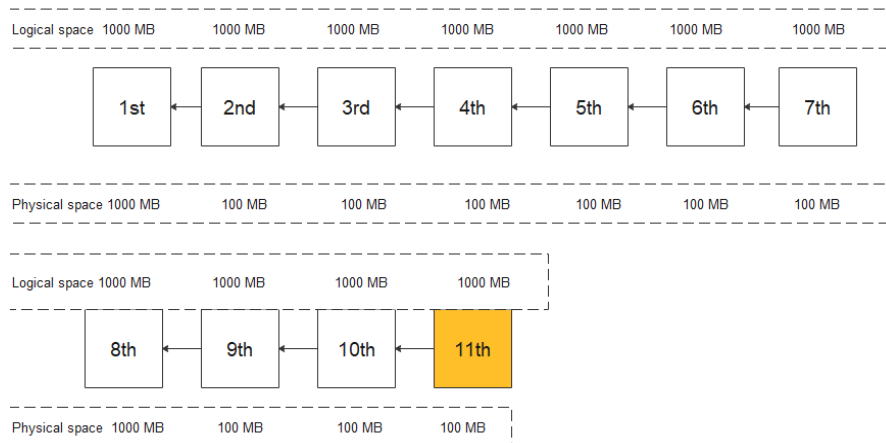
Billing is calculated as follows:

The system deducts the free space from either the logical space or the physical space, whichever is smaller.

### Example

A backup line contains seven backups by default. There are 11 backups shown in the following figure. The 1st backup to the 7th belong to one backup line and the 8th to the 11th belong to another.

**Figure 14-3 Backup example**



If there is 1,000 MB of backup space and the logical space is 1,000 MB each time, the physical space for the 1st backup is 1,000 MB. If the incremental data size is 100 MB each time, the physical space for the 2nd backup to 7th is 100 MB.

A backup line contains seven backups by default. The physical space for the 8th backup is 1,000 MB because it represents a new backup line.

Billed space includes the space of the two chains in the example.

Suppose that after the 11th backup was created, and the 1st, 2nd and 3rd backups expired and were automatically deleted. The size of each space is calculated as follows:

- Total logical space size = Total logical space size – Logical size of the expired backup file (1,000 MB x 11 - 3,000 MB = 8,000 MB in this example).
- Physical space: size of data backed up to OBS. In this example, physical space includes the sum of physical space on the two backup links: 1,000 MB + (100 MB x 6) + 1,000 MB + (100 MB x 3) = 2,900 MB
- Total billed space = Min (Total size of logical space, Total size of physical space) – free space, so the total billed space in this example = Min (8,000 MB, 2,900 MB) – 1,000 MB = 1,900 MB

## 14.4 Configuring a Same-Region Backup Policy

### Scenarios

When you create a GaussDB(for MySQL) instance, an automated backup policy is enabled by default and cannot be disabled. However, it can be modified after instance creation is complete. GaussDB(for MySQL) backs up data based on the automated backup policy you specify.

GaussDB(for MySQL) backs up data at the instance level. If a DB instance is faulty or data is damaged, you can still restore it using backups to ensure data reliability. 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.

After an automated backup policy is configured, full backups are created based on the time window and backup cycle specified in the policy. The time required for

creating a backup depends on how much data there is in the instance. Backups are stored for as long as you specified in the backup policy.

You do not need to configure incremental backup policies because the system automatically performs an incremental backup every 5 minutes. The generated incremental backups can be used to restore the database and table data to a specified point in time.


## Constraints

- Rebooting instances is not allowed during the creation of a full backup. Exercise caution when selecting a backup time window.
- When starting a full backup task, GaussDB(for MySQL) first tests connectivity to your instance. If the backup lock failed to be obtained from the DB instance, the verification fails and a retry is automatically performed. If the retry fails, the backup will fail.
- Performing a full backup may decrease instance throughput because it occupies node resources, especially disk bandwidth.

## Viewing or Modifying a Same-Region Backup Policy

**Step 1** [Log in to the management console.](#)

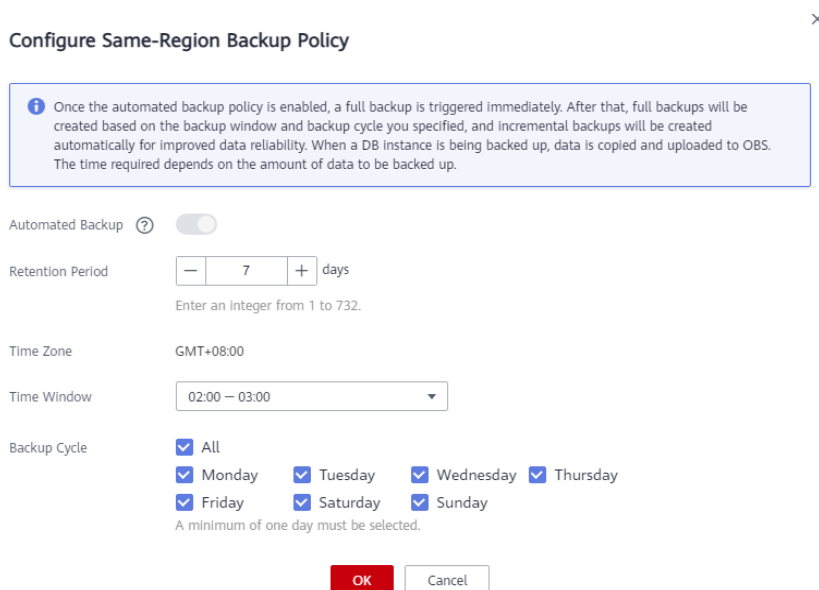
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** Choose **Backups** in the left navigation pane, click **Configure Same-Region Backup Policy**. You can view the configured backup policy. To modify the backup policy, adjust the parameter values as needed.

**Figure 14-4** Modifying backup policies



Configure Same-Region Backup Policy

**i** Once the automated backup policy is enabled, a full backup is triggered immediately. After that, full backups will be created based on the backup window and backup cycle you specified, and incremental backups will be created automatically for improved data reliability. When a DB instance is being backed up, data is copied and uploaded to OBS. The time required depends on the amount of data to be backed up.

Automated Backup

Retention Period  days  
Enter an integer from 1 to 732.

Time Zone GMT+08:00

Time Window

Backup Cycle  All  
 Monday  Tuesday  Wednesday  Thursday  
 Friday  Saturday  Sunday  
A minimum of one day must be selected.

**Table 14-2** Parameter description

Parameter	Description
Retention Period	<p>Number of days that your automated backups can be retained. The retention period is from 1 to 732 days and the default value is 7.</p> <ul style="list-style-type: none"><li>• Extending the retention period improves data reliability. You can configure the retention period if needed.</li><li>• If you shorten the retention period, the new backup policy takes effect for existing backups. Any backups (including full and automated backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.</li></ul> <p><b>NOTE</b> To extend the retention period to 3,660 days, contact customer service.</p>
Time Zone	The backup time is in UTC format. The backup time segment changes with the time zone during the switch between the DST and standard time.
Time Window	A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00 or 12:00-13:00.
Backup Cycle	By default, each day of the week is selected. You can change the backup cycle and must select at least one day of the week.

**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 **Retention Period** was set to 2:

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

**Step 6** Click **OK**.

----End

## APIs

- [Modifying an Automated Backup Policy](#)
- [Querying an Automated Backup Policy](#)

# 14.5 Configuring a Cross-Region Backup Policy

## Scenarios

GaussDB(for MySQL) can store backup files in the storage space that is in a different region from your DB instance for disaster recovery. If your DB instance in a region is faulty, you can use the backup files in another region to restore data to a new instance.

After you enable cross-region backup, the backup files are automatically stored in the region you specify. On the **Backups** page of the GaussDB(for MySQL) console, you can click **View Backup** in the **Operation** column and manage cross-region backup files.

Currently, cross-region backup can be enabled for up to 150 DB instances in a single region under a tenant. It is recommended that the data volume of a single DB instance be at most 2 TB. If the data volume is too large, the synchronization progress may be delayed.

To apply for the cross-region backup policy permission, submit a service ticket by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the management console.

## Supported Regions

To view regions where cross-region backup policy can be enabled, see [OBS Cross-Region Replication](#).

## Billing


Table 14-3 Billing

Specification Code	Pay-per-Use (USD/GB/Hour)
gaussdb.mysql.crossreg.backup.space	<ul style="list-style-type: none"><li>• Regions in China (excluding Hong Kong): \$0.000157 USD/GB/hour</li><li>• Hong Kong and regions outside China: \$0.00022 USD/GB/hour</li></ul>

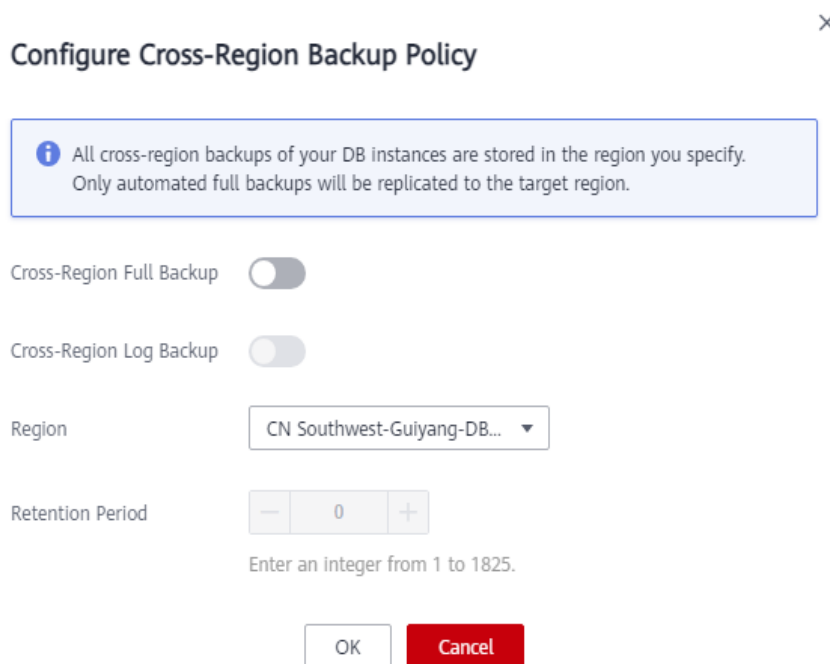
## Enabling or Modifying a Cross-Region Backup Policy

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Configure Cross-Region Backup Policy**. Configure parameters as shown in [Figure 14-5](#).

**Figure 14-5** Configuring a cross-region backup policy



- If you enable **Cross-Region Full Backup**, automated full backup files of the instance are stored in OBS in the region you specify.
- If you enable **Cross-Region Log Backup**, incremental backup files of the instance are stored in OBS in the region you specify. Enable cross-region full backup first.
- Cross-region backup files can be retained from 1 to 1,825 days.
- After cross-region log backup is enabled, you can restore an instance to a specified time point only after the next automated full backup replication is complete. The specified time point must be later than the time when the automated full backup is complete.

- Step 6** Click **OK**.
- Step 7** On the **Cross-Region Backups** tab of the **Backups** page, manage cross-region backup files.


**Figure 14-6** Cross-region backups

DB Instance Name/ID	DB Engine	Status	Source Backup Region	Target Backup Region	Retention Period	Operation
	GaussDB(for MySQL)	Available			1 day	Set Cross-Region Backup   More


- By default, all instances with cross-region backups are displayed.
  - To modify the cross-region backup policy, click **Set Cross-Region Backup** in the **Operation** column.
  - To view generated cross-region backup files, click **View Cross-Region Backup** in the **Operation** column. If an instance fails, you can use the cross-region backup files to restore data to a new instance.

Full or incremental backups can be resorted to a new DB instance. Select the backup you want to restore and click **Restore** in the **Operation** column.

**Figure 14-7** Restoring a full backup

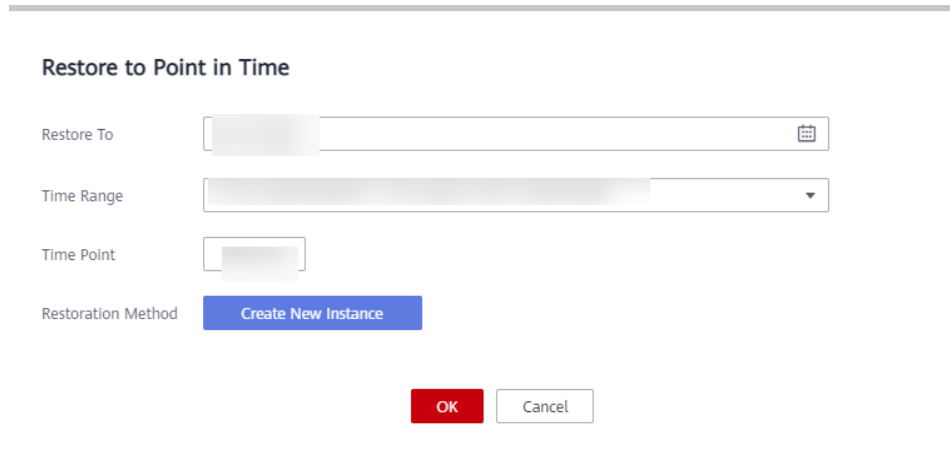


**Figure 14-8** Restoring an incremental backup



Incremental backups can be restored to a point of time. You need to select a time range, select or enter a time point within the acceptable range.

**Figure 14-9** Restoring an incremental backup to a point in time





- To view all cross-region backups, click **View All Backups**.  
To restore a backup, locate the backup and click **Restore** in the **Operation**.
- To return to the instance list, click **View All Instances**.

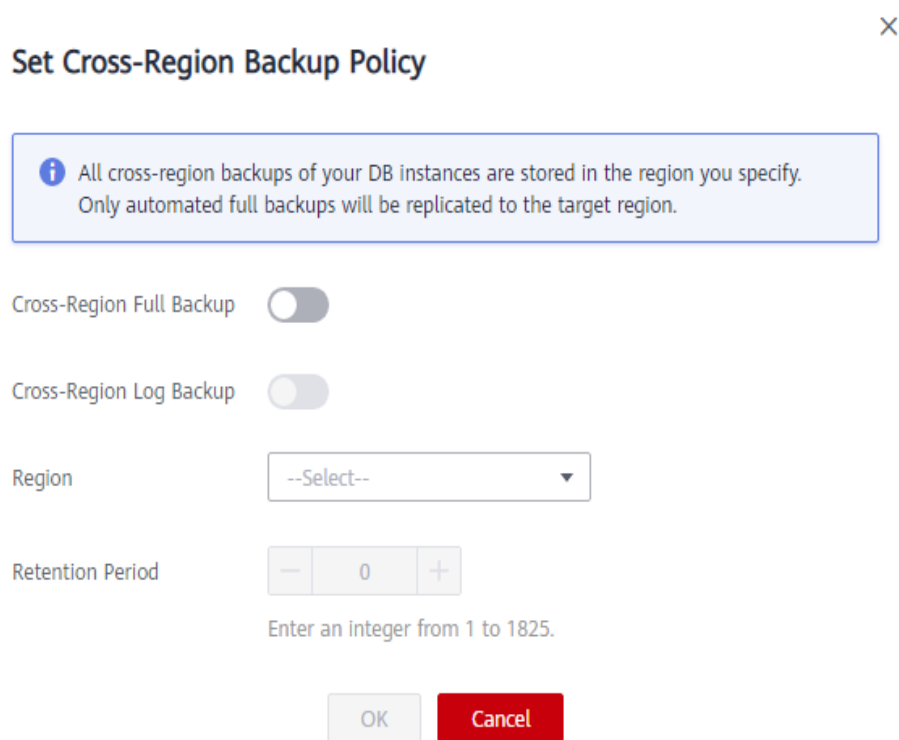
----End



## Disabling a Cross-Region Backup Policy

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Backups** page, click the **Cross-Region Backups** tab.
- Step 5** Locate the instance and click **Set Cross-Region Backup** in the **Operation** column. On the displayed page, disable the cross-region backup policy. Before disabling **Cross-Region Full Backup**, you need to disable **Cross-Region Log Backup**.

**Figure 14-10** Disabling a cross-region backup policy



- Step 6** Click **OK**.
- End

## 14.6 Enabling or Disabling Encrypted Backup

### Scenarios

GaussDB(for MySQL) can encrypt backups. After encrypted backup is enabled, a key is required, which is generated and managed by [Data Encryption Workshop \(DEW\)](#).


## Precautions

- To use encrypted backup, submit an application by choosing [Service Tickets > Create Service Ticket](#) in the upper right corner of the console.
- Only the backups generated after encrypted backup is enabled will be encrypted.
- After encrypted backup is disabled, new backup files will not be encrypted for storage. Backup files created before encrypted backup is disabled will not be decrypted.
- Currently, only the SM4 and AES\_256 key algorithms are supported. After encrypted backup is enabled, the key algorithm cannot be changed.
- The key cannot be disabled, deleted, or frozen while in use, or the encrypted backups cannot be used for restoration.
- Encrypted backups can be directly used to restore data on the management console. You do not need to manually decrypt backups.
- Once encrypted backup is enabled for your DB instance, data cannot be restored to an existing DB instance, even if encrypted backup is disabled later.
- Cross-region backup and encrypted backup cannot be both enabled.
- When encrypted backup is enabled for a DB instance, only the key of the corresponding enterprise project can be selected. To view keys in an enterprise project, see [Viewing a CMK](#).

## Enabling Encrypted Backup

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

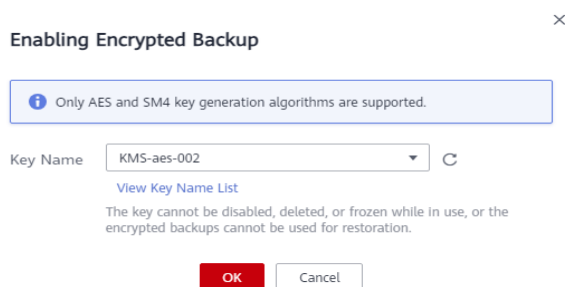
**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** On the **Backups** page, click  next to **Encrypted Backup**.

**Step 6** In the displayed dialog box, select a key name from the drop-down list and click **OK**.

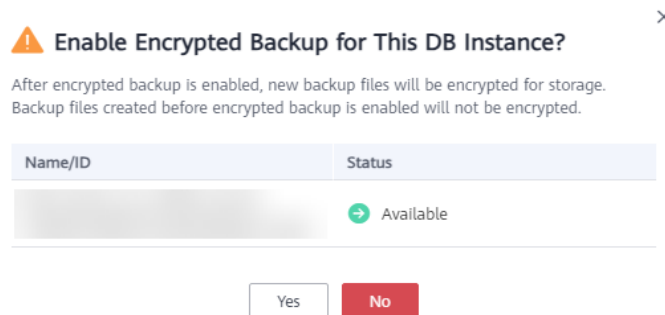
Only SM4 and AES\_256 key algorithms are supported.

**Figure 14-11** Selecting a key



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

**Figure 14-12** Enabling encrypted backup





**Step 8** Refresh the page and check whether encrypted backup is enabled.

----End

## Disabling Encrypted Backup

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.


**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** On the **Backups** page, click  next to **Encrypted Backup**.

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

**Figure 14-13** Disabling encrypted backup



----End

## 14.7 Creating a Manual Backup

### Scenarios

GaussDB(for MySQL) allows you to create manual backups for the available primary node of your instance. You can restore data from backups to ensure data reliability.

### Constraints

- You can create manual backups only when your account balance is no less than \$0 USD.
- The backup efficiency is in direct proportion to the instance data volume.
- The system verifies the connection to the DB instance when starting a full backup task. If the backup lock failed to be obtained from the DB instance, the verification fails and a retry is automatically performed. If the retry fails, the backup will fail.
- When an account is deleted, both automated and manual backups are deleted.


### Backup Clearing

When a DB instance is deleted, its automated backups are also deleted, but its manual backups are retained permanently until **you manually delete them**.

### Method 1

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

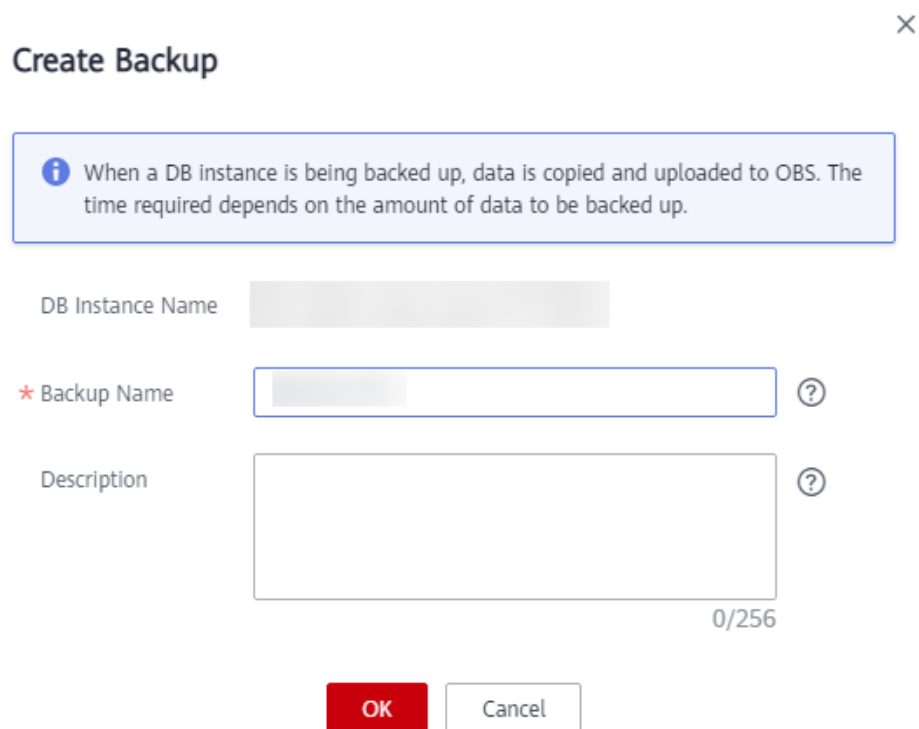
**Step 4** On the **Instances** page, locate the instance for which you want to create a manual backup and choose **More > Create Backup** in the **Operation** column.

**Step 5** In the displayed dialog box, enter a backup name and description. Then, click **OK**. If you want to cancel the backup creation task, click **Cancel**.

- The backup name must consist of 4 to 64 characters and start with a letter. It is case sensitive and can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (\_).
- The description can consist of up to 256 characters. It cannot contain carriage return characters and the following special characters:  
!<"='>&
- The time required for creating a manual backup depends on the data volume of the instance.
- When a DB instance is being backed up, data is copied and uploaded to OBS. The time required depends on the amount of data to be backed up.

Click  in the upper right corner of the page to check the instance status. If the instance status is **Available**, the backup creation is complete.

**Figure 14-14** Creating a manual backup



**Create Backup** ×

**i** When a DB instance is being backed up, data is copied and uploaded to OBS. The time required depends on the amount of data to be backed up.

DB Instance Name

\* Backup Name  ?

Description  ?

0/256


**OK** Cancel


**Step 6** View and manage the created backup on the **Backups** page.

----End

## Method 2

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

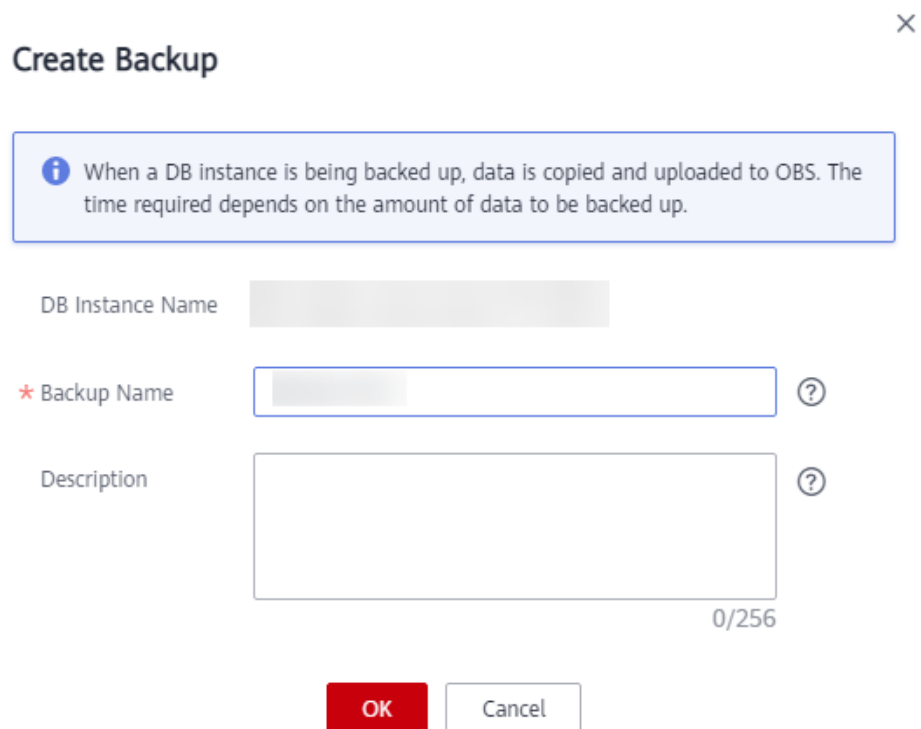
**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Create Backup**. In the displayed dialog box, enter a backup name and description, and click **OK**. If you want to cancel the backup creation task, click **Cancel**.

- The backup name must consist of 4 to 64 characters and start with a letter. It is case sensitive and can contain only uppercase letters, lowercase letters, digits, hyphens (-), and underscores (\_).

- The description can consist of up to 256 characters. It cannot contain carriage return characters and the special characters (!<"='>&).
- The time required for creating a manual backup depends on the data volume of the instance.
- When a DB instance is being backed up, data is copied and uploaded to OBS.

**Figure 14-15** Creating a manual backup



**Create Backup** ×

**i** When a DB instance is being backed up, data is copied and uploaded to OBS. The time required depends on the amount of data to be backed up.

DB Instance Name

\* Backup Name  ?

Description  ?

0/256

**Step 6** View and manage the created backup on the **Backups** page.

----End

## APIs

- [Creating a Manual Backup](#)
- [Querying Backups](#)

## 14.8 Exporting Backup Information

### Scenarios


You can export backup information of GaussDB(for MySQL) instances to an Excel file for further analysis. The exported information includes instance name, backup ID, backup name, backup type, backup time, status, size, and description.


### Constraints


Automated and manual backup files cannot be downloaded.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Backups**. On the displayed page, select the backups you want to export and click  to export the backup information.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Backups**. On the **Full Backups** page, select the backups you want to export and click **Export**.

- Currently, only the backup information displayed on the current page can be exported.
- The backup information is exported to an Excel file.

**Step 5** View the exported backup information.

----End

## 14.9 Deleting a Manual Backup

### Scenarios

You can delete manual backups to free up storage.

---

#### NOTICE


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

---

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Backups**. On the displayed page, locate the manual backup to be deleted and click **Delete** in the **Operation** column.

Alternatively, on the **Instances** page, click the instance name to go to the **Basic Information** page. On the **Backups** page, locate the backup you want to delete and click **Delete** in the **Operation** column.

The following backups cannot be deleted:

- Automated backups
- Backups that are being restored or created

**Step 5** Click **Yes**.

----**End**

## APIs

### [Deleting a Manual Backup](#)



# 15 Data Restorations

## 15.1 Restoring a DB Instance

If data is damaged or mistakenly deleted, you can restore it from backups.

Table 15-1 Restoring data

Scenario	Description
<a href="#">Restoring Instance Data to a Specific Point in Time</a>	You can restore instance data to a point in time. The data can be restored to a new DB instance, the original DB instances, and an existing DB instance.
<a href="#">Restoring Data to a DB Instance</a>	You can restore data to a new DB instance, the original DB instance, or an existing DB instance using automated or manual backups.
<a href="#">Restoring Table Data to a Specific Point in Time</a>	You can restore one or more tables in the database to a specified point of time.

## 15.2 Restoring Instance Data to a Specific Point in Time

### Scenarios

You can restore data of an instance to a specified point in time.

### Constraints


- Keep your account balance above zero so that backup data can be restored to a new DB instance. You will be billed for new instance specifications.

- Do not run the **reset master** command on instances within their lifecycle. Otherwise, an exception may occur during the point-in-time restore.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

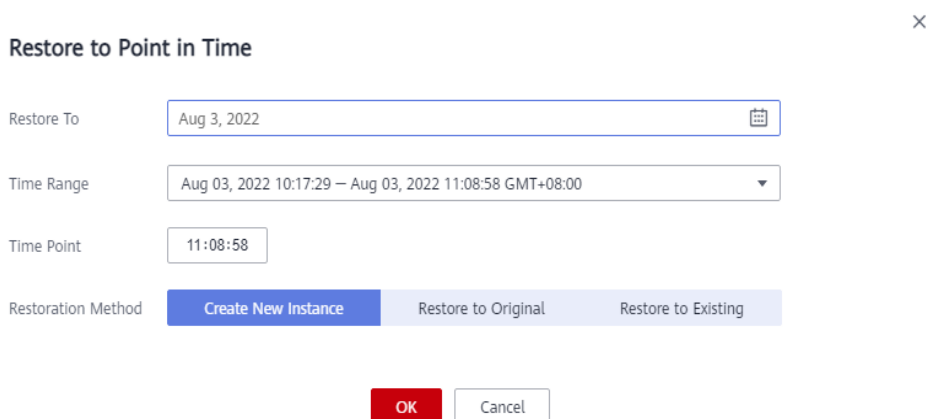
**Step 5** In the navigation pane on the left, choose **Backups**. On the displayed page, click **Restore to Point in Time**.

**Step 6** Select a time range, select or enter a time point within the acceptable range, and set **Restoration Method** to **Create New Instance**, **Restore to Original**, or **Restore to Existing**.

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

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

**Figure 15-1** Restoring to a point in time



Restore to Point in Time

Restore To: Aug 3, 2022

Time Range: Aug 03, 2022 10:17:29 – Aug 03, 2022 11:08:58 GMT+08:00

Time Point: 11:08:58

Restoration Method: **Create New Instance** | Restore to Original | Restore to Existing

OK Cancel

1. Select **Create New Instance** and click **OK**. On the **Create New Instance** page, configure parameters and click **Next**.
  - The region, DB engine and version of the new instance are the same as those of the original instance and cannot be changed.
  - The default database port is **3306**.
  - To synchronize database parameters of the original DB instance, select **Original DB instance parameter template**.

**NOTE**

- If the original DB instance is deleted, the database parameters of the original DB instance cannot be synchronized using backups.
  - When you synchronize the database parameters of the original DB instance, the following parameters cannot be synchronized and you need to **manually modify them** after the DB instance is restored.
    - innodb\_write\_io\_threads
    - innodb\_read\_io\_threads
    - max\_connections
    - innodb\_log\_buffer\_size
    - innodb\_parallel\_read\_threads
    - temptable\_max\_ram
    - threadpool\_size
    - innodb\_buffer\_pool\_size
    - innodb\_page\_cleaners
  - Other settings are the same as those of the original instance by default and can be modified. For details, see [Buying a DB Instance](#).
2. Select **Restore to Original** and click **Next**.
- Data on the original instance will be overwritten and the original DB instance will be unavailable during the restoration.

**Figure 15-2** Restoring to the original DB instance

**Restore to Point in Time** ×

Restore To: Aug 3, 2022

Time Range: Aug 03, 2022 10:17:29 – Aug 03, 2022 11:08:58 GMT+08:00

Time Point: 11:08:58

Restoration Method:  Create New Instance  Restore to Original  Restore to Existing

I acknowledge that after I select Restore to Original, data on the original databases will be overwritten and the original DB instance will be unavailable during the restoration.

**Next** Cancel

3. Select **Restore to Existing**, select the target DB instance, and click **Next**.
- Restoring data to an existing DB instance will overwrite data and password of user **root** on it and cause the existing DB instance to be unavailable during the restoration. DB instances will not be displayed unless they have the same DB engine type, version, and table name case sensitivity as the original DB instance.
  - The restored DB instance contains the data and account information of the original DB instance, but does not contain the parameter settings of the original DB instance.

- If the original password of the existing DB instance cannot be used to connect to the database after the restoration, you can reset the password.

**Figure 15-3** Restoring to an existing instance

**Restore to Point in Time** ×

Restore To: Aug 3, 2022

Time Range: Aug 03, 2022 10:17:29 – Aug 03, 2022 11:08:58 GMT+08:00

Time Point: 11:08:58

Restoration Method:  Create New Instance  Restore to Original  Restore to Existing

I acknowledge that restoring to an existing DB instance will overwrite data and root password on it and cause the existing DB instance to be unavailable during the restoration. DB instances will not be displayed unless they have same DB engine type, version, and table name case sensitivity as the original DB instance.

Enter a DB instance name or ID.

DB Instance Name/ID	Size	DB Engine
<input checked="" type="radio"/> [Redacted]	1.60 GB	GaussDB(for MySQL)
<input type="radio"/> [Redacted]	90.00 MB	GaussDB(for MySQL)

**Step 7** View the restoration results.

- **Create New Instance:** After the creation is complete, the instance status changes from **Creating** to **Available**. The new instance is independent from the original one and includes the data before the backup was created. If you want to offload read pressure from the primary node, create one or more read replicas for the new instance.

A full backup is triggered after the new instance is created.

- **Restore to Original and Restore to Existing:** When the instance status changes from **Restoring** to **Available**, the restoration is complete.

----End


## 15.3 Restoring Data to a DB Instance


### Scenarios

You can use an automated or manual backup to restore data to the point in time when the backup was created. The restoration is at the instance level.

## Procedure

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

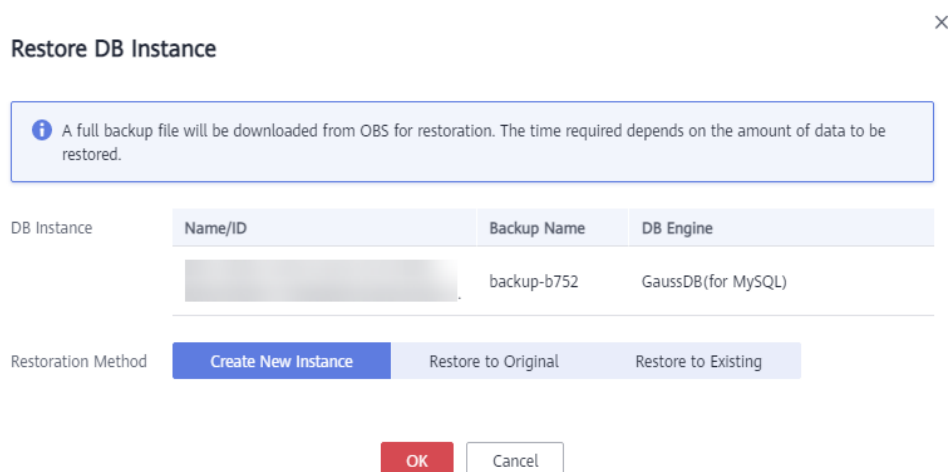
**Step 4** Select the backup to be restored using either of the ways:

In the navigation pane on the left, click **Backups**. On the **Backups** page, select the backup to be restored and click **Restore** in the **Operation** column.

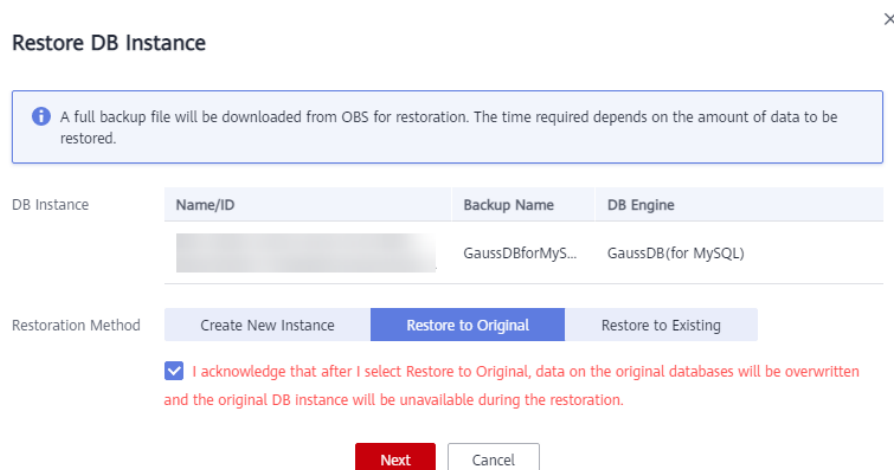
On the **Instances** page, click the instance name to go to the **Basic Information** page. On the **Backups** page, select the backup to be restored and click **Restore** in the **Operation** column.

**Step 5** Set **Restoration Method** to **Create New Instance**, **Restore to Original**, or **Restore to Existing** as needed.

**Figure 15-4** Restoring an instance from a backup



1. Select **Create New Instance** and click **OK**. On the **Create New Instance** page, configure parameters and click **Next**.
  - The region, DB engine and version of the new instance are the same as those of the original instance and cannot be changed.
  - The default database port is **3306**.
  - Other settings are the same as those of the original instance by default and can be modified. For details, see [Buying a DB Instance](#).
2. Select **Restore to Original** and click **Next**. In the displayed dialog box, click **OK**.
  - Data on the original instance will be overwritten and the original DB instance will be unavailable during the restoration.

**Figure 15-5** Restoring to the original DB instance

**Restore DB Instance** ×

**i** A full backup file will be downloaded from OBS for restoration. The time required depends on the amount of data to be restored.

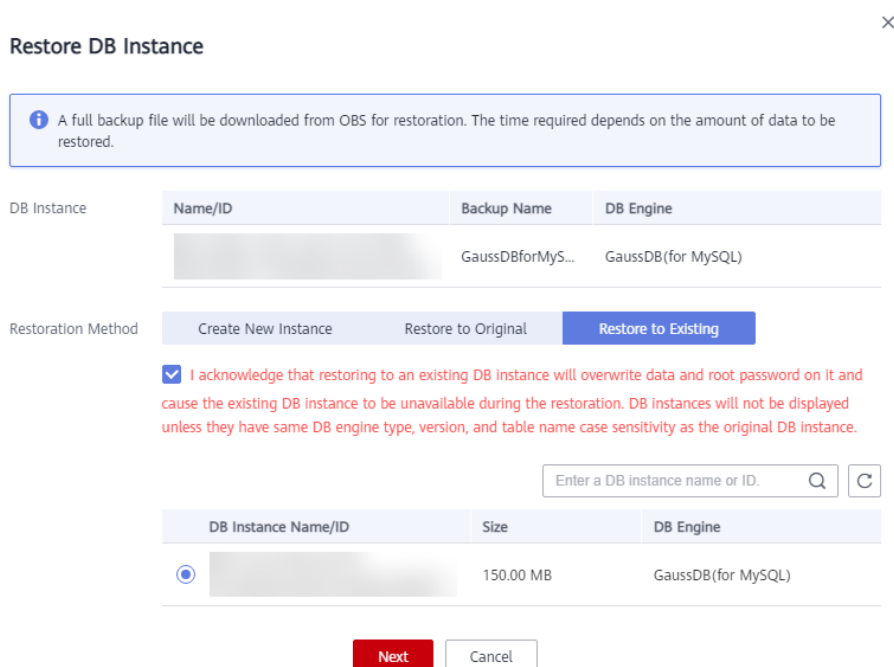
DB Instance	Name/ID	Backup Name	DB Engine
		GaussDBforMyS...	GaussDB(for MySQL)

Restoration Method: Create New Instance **Restore to Original** Restore to Existing

I acknowledge that after I select Restore to Original, data on the original databases will be overwritten and the original DB instance will be unavailable during the restoration.

Next Cancel

3. Select **Restore to Existing**, select the target DB instance, and click **Next**.
  - Restoring data to an existing DB instance will overwrite data and password of user **root** on it and cause the existing DB instance to be unavailable during the restoration. DB instances will not be displayed unless they have the same DB engine type, version, and table name case sensitivity as the original DB instance.
  - If the original password of the existing DB instance cannot be used to connect to the database, you can reset the password.

**Figure 15-6** Restoring to an existing instance

**Restore DB Instance** ×

**i** A full backup file will be downloaded from OBS for restoration. The time required depends on the amount of data to be restored.

DB Instance	Name/ID	Backup Name	DB Engine
		GaussDBforMyS...	GaussDB(for MySQL)

Restoration Method: Create New Instance Restore to Original **Restore to Existing**

I acknowledge that restoring to an existing DB instance will overwrite data and root password on it and cause the existing DB instance to be unavailable during the restoration. DB instances will not be displayed unless they have same DB engine type, version, and table name case sensitivity as the original DB instance.

Enter a DB instance name or ID.

DB Instance Name/ID	Size	DB Engine
<input checked="" type="radio"/>	150.00 MB	GaussDB(for MySQL)

Next Cancel

**Step 6** View the restoration results.

- **Create New Instance:** After the creation is complete, the instance status changes from **Creating** to **Available**. The new instance is independent from

the original one and includes the data before the backup was created. If you want to offload read pressure from the primary node, create one or more read replicas for the new instance.

A full backup is triggered after the new instance is created.

- **Restore to Original** and **Restore to Existing**: When the instance status changes from **Restoring** to **Available**, the restoration is complete.

----End

## APIs

- [Restoring Data to the Original Instance or an Existing Instance](#)
- [Querying the Restoration Time Range](#)

# 15.4 Restoring Table Data to a Specific Point in Time

## Scenarios

To ensure data integrity and reduce impact on the original instance performance, the system restores the full and incremental data at the selected time point to a temporary instance, exports the tables to be restored, and then restores the tables to the original instance. The time required depends on the amount of data to be backed up and restored on the instance. Restoring tables will not overwrite data in the instance.


## Constraints

- Tables that have foreign keys cannot be restored.
- Tables that have triggers cannot be restored.
- If the tables to be restored do not exist at the selected point in time, the restoration will fail.
- The DB instance cannot be rebooted or deleted, and the instance specifications cannot be modified.
- The number of tables to be restored must be no more than 20,000. If the number of tables to be restored exceeds 2,000, you can restore the instance to a point in time. For details, see [Restoring Instance Data to a Specific Point in Time](#).

## Procedure

**Step 1** [Log in to the management console](#).

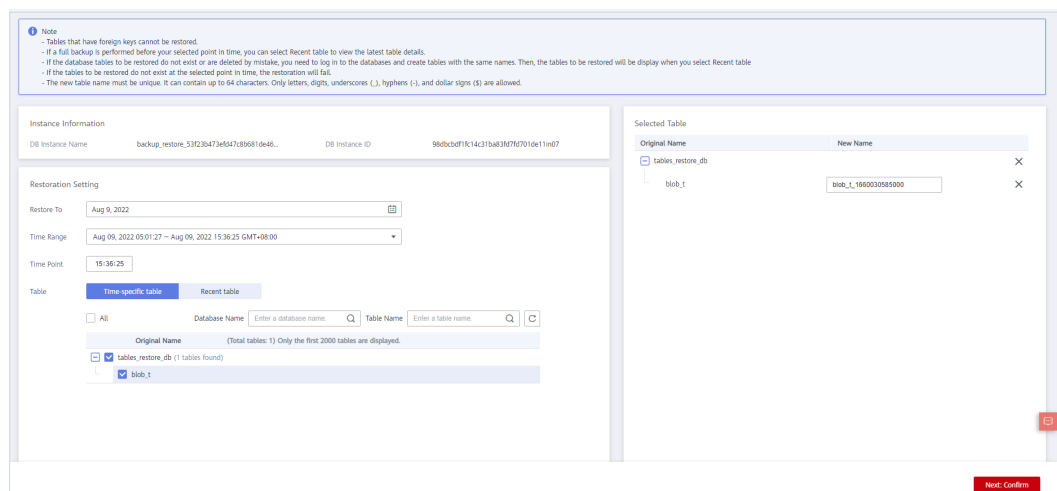
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

- Step 5** Choose **Backups** in the navigation pane on the left. On the **Full Backups** page, choose **More > Restore Table** above the backup list.
- Step 6** In the displayed dialog box, specify the restoration date, time range, time point, and tables to be restored, and click **Next: Confirm**.
- To facilitate your operations, you can search for the tables and databases to be restored.
  - After the restoration is complete, new tables with timestamps as suffixes are generated in the instance. You can rename the new tables. The new table name must be unique. It can contain up to 64 characters. Only letters, digits, underscores (\_), hyphens (-), and dollar signs (\$) are allowed.
  - **Time-specific table:** The tables to be restored are read from the latest full backup before the selected point in time. **Recent table:** The tables to be restored are read from the current point time.

**Figure 15-7** Restoring tables to a specified point in time



## NOTICE

- If a full backup is performed before your selected point in time, you can select **Recent table** to view the latest table details.
- If the database tables to be restored do not exist or are deleted by mistake, you need to log in to the databases and create tables with the same names. Then, the tables to be restored will be displayed when you select **Recent table**.
- Only specified tables are restored. Ensure that all tables to be restored are selected.

**Step 7** Confirm your settings and click **Restore Now**. If you need to modify your settings, click **Previous**.

**Step 8** On the **Instances** page, the instance status is **Restoring**. During the restoration process, services are not interrupted.

You can also view the progress and result of restoring tables to a specified point in time on the **Task Center** page.



After the restoration is successful, you can manage data in the tables as required.

----End

# 16 Connection Management

---

## 16.1 Binding and Unbinding an EIP

### Scenarios

By default, a GaussDB(for MySQL) instance is not publicly accessible (not bound with an EIP) after being created. You can bind an EIP to the instance for public accessibility and can unbind the EIP from the instance if needed.

---

#### NOTICE

- To ensure that the database can be accessed, the security group used by the database must allow access to the database port. For example, if the database port is **3306**, ensure that the security group allows access to the port **3306**.
  - You can purchase a new EIP if all available EIPs may have been bound to other applications.
- 

### Prerequisites


If an instance already has an EIP bound, that EIP needs to be unbound before a new one can be configured.


### Precautions

- You need to configure security group rules and enable specific IP addresses and ports to access your DB instance. Before accessing the DB instance, add an IP address or an IP address range that will access the DB instance to the inbound rule. For details, see [Configuring Security Group Rules](#).
- You can buy an EIP [on the console](#) and bind it to a DB instance. One EIP can be bound to only one DB instance. For pricing details, see [Price Calculator](#).

## Binding an EIP

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Network Information** area, click **Bind** in the **Public IP Address (EIP)** field.

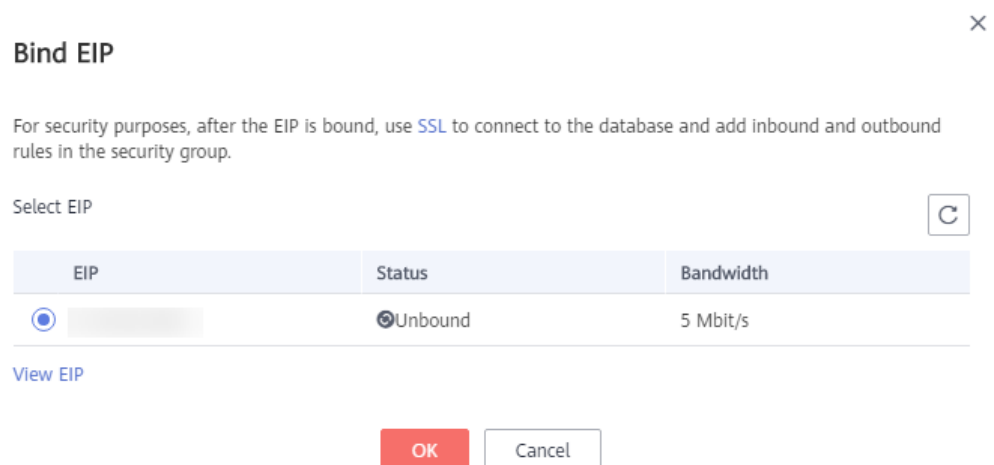
**Step 6** In the displayed dialog box, select an EIP and click **OK**.

If no EIPs are available, click **View EIP** to create an EIP on the network console. After the EIP is created, go back to the **Basic Information** page and bind the newly created EIP to the instance.

### NOTICE

You need to configure security group rules and enable specific IP addresses and ports to access the instance.

**Figure 16-1** Selecting an EIP



**Step 7** In the **Public IP Address (EIP)** field of the **Network Information** area, view the EIP that was bound.

To unbind the EIP from the instance, see [Unbinding an EIP](#).

-----End

## Unbinding an EIP

- Step 1** On the **Instances** page, click the instance that you want to unbind the EIP from. The **Basic Information** page is displayed.
- Step 2** In the **Network Information** area, click **Unbind** in the **Public IP Address (EIP)** field.
- Step 3** In the displayed dialog box, click **Yes**. To bind an EIP to the instance again, see [Binding an EIP](#).
- Step 4** If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

----End

## APIs

- [Binding an EIP](#)
- [Unbinding an EIP](#)

# 16.2 Changing a Database Port




## Scenarios

You can change the database port of a GaussDB(for MySQL) instance. The change will be applied to the ports of the primary node and read replicas.

## Constraints

The database port of a DB instance with read/write splitting enabled cannot be changed.



## Procedure

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the **Network Information** area, click  in the **Database Port** field.

**Figure 16-2** Changing a database port



The GaussDB(for MySQL) database port ranges from 1025 to 65534, excluding 5342, 5343, 5344, 5345, 12017, 20000, 20201, 20202, 33062, and 33071, which are reserved by the system.

- To submit the change, click .
  - In the dialog box, click **Yes**.
    - i. If you change the database port of an instance, the ports of the primary node and read replicas are changed accordingly and all of them are rebooted.
    - ii. This process takes about 1–5 minutes.
- To cancel the change, click .

**Step 6** View the results on the **Basic Information** page.

----End

## APIs

### Changing a Database Port

## 16.3 Applying for and Changing a Private Domain Name for a DB Instance

You can use a private network domain name to connect to a GaussDB(for MySQL) instance.

After your DB instance is created, you can change the private domain name as needed.


### Constraints

- To use a private domain name, submit an application by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the console.
- Domain Name Service (DNS) is deployed.
- Changing the private domain name will interrupt your database connection. To reconnect to the DB instance, change the connection address of your applications. The new private domain name is applied to the instance about 5 minutes after the change.

### Applying for a Private Domain Name

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.


**Step 5** In the **Network Information** area, click **Apply** in the **Private Domain Name** field.


**Step 6** In the **Private Domain Name** field, view the generated private domain name.

----End

## Changing a Private Domain Name

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Connection Information** area, click **Change** next to the **Private Domain Name** field.

**Step 6** In the displayed dialog box, enter a new domain name and click **OK**.

### NOTE

- Only the prefix of a private domain name can be modified.
- The prefix of a private domain name contains 8 to 63 characters, and can include only lowercase letters and digits.
- The new private domain name must be different from existing ones.

**Step 7** If you have enabled operation protection, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account and cloud product. For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

----End

## 16.4 Configuring and Changing a Private IP Address

### Scenarios

You can plan and change private IP addresses after migrating on-premises databases or other cloud databases to GaussDB(for MySQL).

### Constraints

After read/write splitting is enabled, the private IP address cannot be changed.

After a private IP address is changed, the domain name needs to be resolved again. This operation takes several minutes and may interrupt database connections. Therefore, you are advised to change a private IP address during off-peak hours.


## Configuring the Private IP Address of a DB Instance


When you buy an instance, select a VPC and subnet on the **Buy DB Instance** page. Then, a private IP address will be automatically assigned to your instance. You can also enter a private IP address.

### Procedure


You can change the private IP address of an existing instance.

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the **Network Information** area, click  next to the **Private IP Address** field.

**Step 6** In the displayed dialog box, enter a new private IP address and click **OK**.

An in-use IP address cannot be used as the new private IP address of the instance.

**Step 7** If you have enabled operation protection, click **Start Verification** in the displayed dialog box. On the displayed page, click **Send Code**, enter the obtained verification code, and click **Verify** to close the page.

For details about how to enable operation protection, see [Identity and Access Management User Guide](#).

----End

### APIs

#### [Changing a Private IP Address](#)

# 17 Parameter Template Management

---

## 17.1 Creating a Parameter Template

You can use database parameter templates to manage DB engine configurations. A database parameter template acts as a container for engine configuration values that can be applied to one or more instances.

---

**NOTICE**

Not all DB engine parameters can be changed in a custom parameter template.

---

If you want to use a custom parameter template, you just create a parameter template and select it when you create an instance or apply it to an existing instance. For details, see [Applying a Parameter Template](#).

If you already have a parameter template and want to include most of the custom parameters and values from that template in a new template, you can replicate that parameter template. For details, see [Replicating a Parameter Template](#).

The following are the key points you should know when using parameter templates:

- To change the parameters in a parameter template of the current instance, go to the **Parameters** page, change parameter values and save the changes. Dynamic parameter changes take effect immediately, but static parameter changes take effect only after you manually reboot the instance. This changes will apply only to the current instance. They will not affect other instances.
- To change the parameters in a parameter template, go to the **Parameter Templates** page and under **Custom Templates** tab, click the template name, change its parameter values and save the changes. Then, apply the changed parameter template to instances. Dynamic parameter changes take effect immediately, but static parameter changes take effect only after you manually reboot the instances.
- Inappropriate parameter settings may have unintended consequences, including degraded performance and system instability. Exercise caution when modifying database parameters and you need to back up data before




modifying parameters in a parameter template. Before applying parameter template changes to a production instance, you should try out these changes on a test instance.


**NOTE**

Each user can create up to 100 parameter templates.  
All GaussDB(for MySQL) engines share the parameter template quotas.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

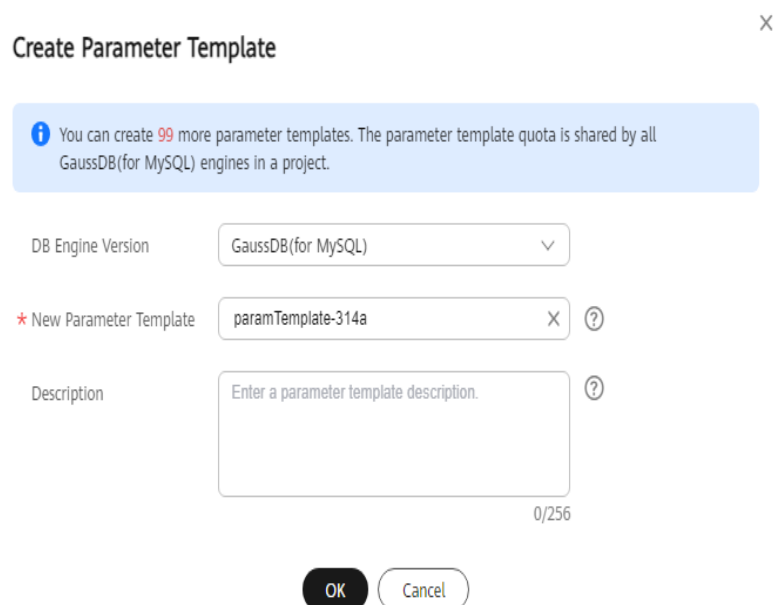
**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Parameter Templates**. On the **Parameter Templates** page, click **Create Parameter Template**.

**Step 5** In the displayed dialog box, configure required parameters and click **OK**.

- The DB engine is GaussDB(for MySQL).
- The template name must consist of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (\_), and periods (.).
- The description can consist of up to 256 characters. It cannot contain carriage return characters or special characters (>!<"&'=).

**Figure 17-1** Creating a parameter template



Create Parameter Template

**i** You can create 99 more parameter templates. The parameter template quota is shared by all GaussDB(for MySQL) engines in a project.

DB Engine Version GaussDB(for MySQL) ▾

\* New Parameter Template paramTemplate-314a X ?

Description Enter a parameter template description. ?

0/256

OK Cancel

----End

## APIs

- [Creating a Parameter Template](#)
- [Querying Parameter Templates](#)
- [Obtaining Details About a Parameter Template](#)

## 17.2 Modifying Parameters of a GaussDB(for MySQL) DB Instance

To optimize GaussDB(for MySQL) database performance, you can change parameters in a custom parameter template and then apply the updated template to your instances.

You can only change the values in custom parameter templates. You cannot change the values in default parameter templates.

The following are the key points you should know when using parameter templates:

- To change the parameters in a parameter template of the current instance, go to the **Parameters** page, change parameter values and save the changes. Dynamic parameter changes take effect immediately, but static parameter changes take effect only after you manually reboot the instance. This changes will apply only to the current instance. They will not affect other instances.
- To change the parameters in a parameter template, go to the **Parameter Templates** page and under **Custom Templates** tab, click the template name, change its parameter values and save the changes. Then, apply the changed parameter template to instances. Dynamic parameter changes take effect immediately, but static parameter changes take effect only after you manually reboot the instances.


### NOTE

GaussDB(for MySQL) has default parameter templates whose parameters cannot be changed. You can view these parameters by clicking the default parameter templates. If a custom parameter template is configured incorrectly and applied into instances, the instances may be unable to be rebooted. If this happens, you can refer to the settings used by a default parameter template.

## Modifying Parameter Template Parameters

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

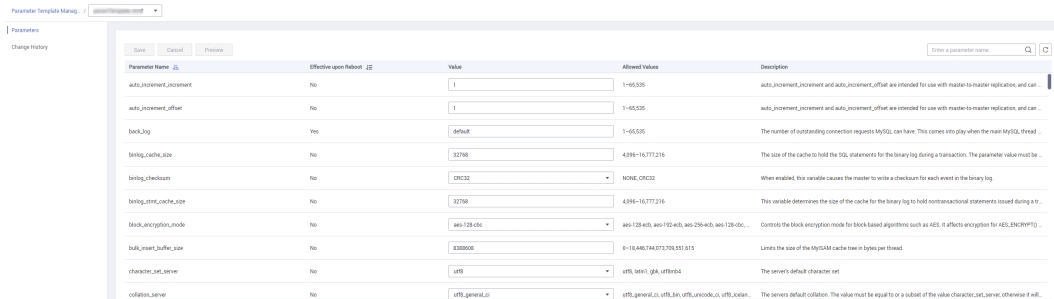
**Step 4** In the navigation pane on the left, choose **Parameter Templates**. On the **Custom Templates** page, click the parameter template you want to change.

**Step 5** Change parameters as needed.

You can save, cancel, or preview your changes.


**NOTICE**

**Figure 17-2** Changing parameters in a parameter template



- To save your changes, click **Save**.
- To cancel your changes, click **Cancel**.
- To preview your changes, click **Preview**.

**Figure 17-3** Previewing changes



**Step 6** After the parameters are changed, click **Change History** to view what changes have been made.

**NOTICE**


After you change parameters in a parameter template, some changes immediately take effect for the instance to which the parameter template applies, but other changes will take effect only after you apply the parameter templates to instances. For details about how to apply a parameter template, see [Applying a Parameter Template](#).

----End

## Modifying Instance Parameters

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, change parameters if needed.

You can save, cancel, or preview your changes.

---

#### NOTICE

After you change instance parameters, check the value in the **Effective upon Reboot** column.

- If **Yes** is displayed and the instance status on the **Instances** page is **Parameter change. Pending reboot**, you must reboot the instance for the changes to take effect.
- If **No** is displayed, the changes take effect immediately for the current instance.

- 
- To save your changes, click **Save**.
  - To cancel your changes, click **Cancel**.
  - To preview your changes, click **Preview**.

After parameters are changed, you can click **Change History** to view what changes have been made.

----End

## APIs


- [Modifying Parameters in a Parameter Template](#)
- [Querying Parameter Templates](#)
- [Obtaining Details About a Parameter Template](#)


## 17.3 Exporting Parameters

### Scenarios

You can export parameter template details (parameter names, values, and descriptions) of an instance to an EXCEL file for review and analysis.

### Procedure


- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.

- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Export** above the parameter list. In the displayed dialog box, enter the file name and click **OK**. You can export parameter template details (parameter names, values, and descriptions) of an instance to an EXCEL file for review and analysis.

**Figure 17-4** Exporting a parameter template

### Export Parameters

Export To

File Name  .CSV 

 **NOTE**

The file name can consist of 4 to 81 characters. It must start with a letter and contain only letters, digits, hyphens (-), and underscores (\_).



----End

## 17.4 Comparing Parameter Templates

### Scenarios

You can compare instance parameters with a parameter template to see the differences of parameter settings. You can also compare parameter templates to see the differences of parameter settings.

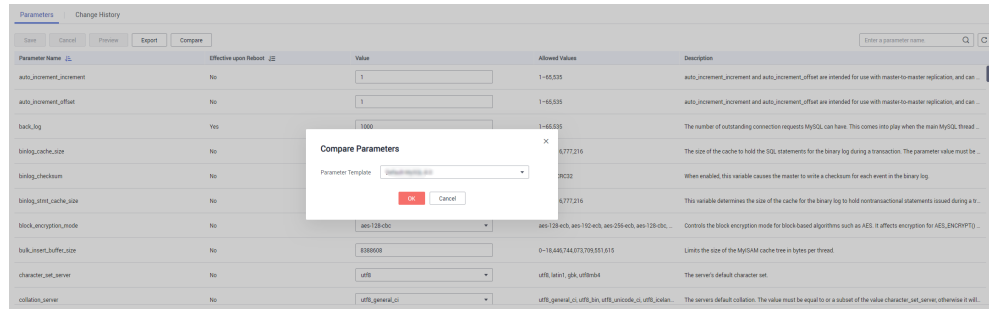
### Comparing Instance Parameters with a Parameter Template

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Compare** above the parameter list.

**Figure 17-5** Comparing instance parameters with those in a specified parameter template




**Step 6** In the displayed dialog box, select a parameter template that you want to compare with the current template and click **OK**.


- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

## Comparing Parameter Templates

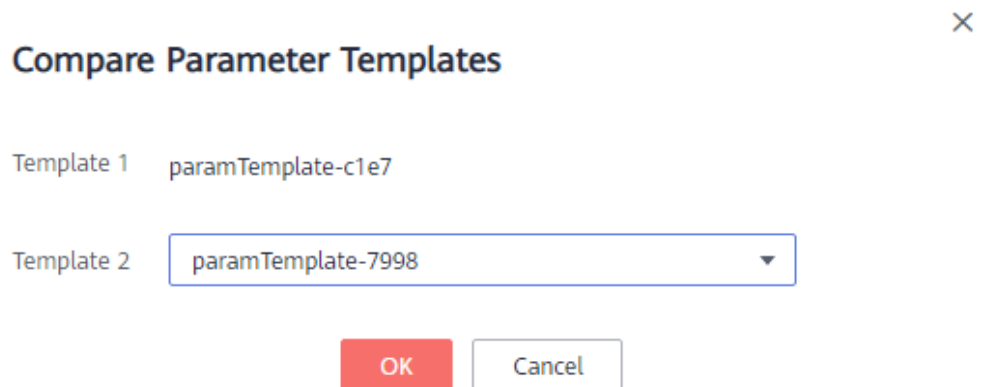
**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template to be compared and click **Compare** in the **Operation** column.

**Step 5** In the displayed dialog box, select a parameter template and click **OK**.

**Figure 17-6** Selecting a parameter template to be compared

- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

**Figure 17-7** Comparing parameter templates

Parameter Name	paramTemplate-b226	paramTemplate-5483-copy
auto_increment_offset	2	1

----End

## 17.5 Viewing Parameter Change History

### Scenarios


You can view the change history of instance parameters or custom parameter templates.


#### NOTE

If you did not make any change to a parameter template, the change history for the template is blank.

### Viewing Change History of Instance Parameters

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Parameters**. On the displayed page, click **Change History**.

**Figure 17-8** Viewing parameter change history



You can view the parameter names, original parameter values, new parameter values, modification statuses, modification time, application statuses, and application time.


You can apply the parameter template to instances if needed. For details, see [Applying a Parameter Template](#).

----End

## Viewing Change History of a Parameter Template

**Step 1** [Log in to the management console](#).


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Parameter Templates**. On the **Custom Templates** page, click the parameter template name.

**Step 5** On the displayed page, choose **Change History** in the navigation pane on the left.

**Figure 17-9** Viewing parameter change history



You can view the parameter names, original parameter values, new parameter values, modification statuses, and modification time.

----End

## 17.6 Replicating a Parameter Template

### Scenarios

You can replicate a parameter template you have created. If you already have 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 replicate the parameter template to generate a new parameter template for future use.


After the parameter template is replicated, the new template will be displayed about 5 minutes later.

Default parameter templates cannot be replicated, but you can create custom parameter templates based on those default templates.

## Replicating a Parameter Template of a DB Instance

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

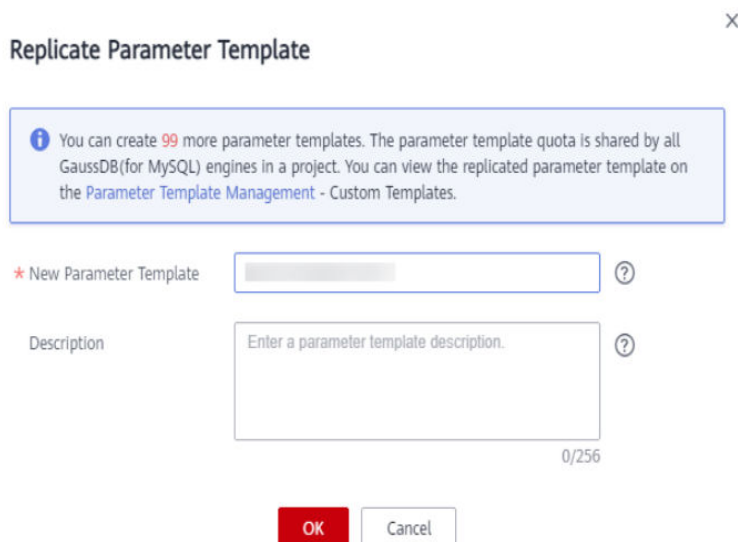
**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Parameters**. On the **Parameters** page, click **Replicate**.

**Step 6** In the displayed dialog box, configure required parameters and click **OK**.

**Figure 17-10** Replicating a Parameter Template of a DB Instance




- The template name consists of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (\_), and periods (.).
- The description consists of up to 256 characters. It cannot contain carriage returns or any of the following special characters:  
>!<"&'='


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

----End

## Replicating a Custom Parameter Template

**Step 1** Log in to the management console.

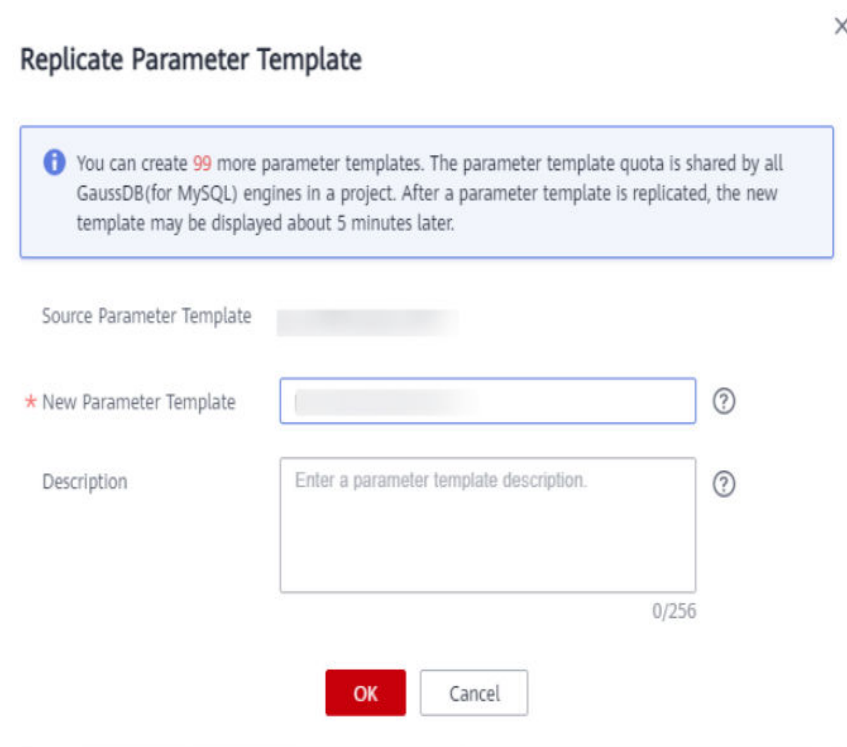
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template to be replicated and click **Replicate** in the **Operation** column.

**Step 5** In the displayed dialog box, configure required parameters and click **OK**.

**Figure 17-11** Replicating a Custom Parameter Template



**Replicate Parameter Template** ×

**i** You can create 99 more parameter templates. The parameter template quota is shared by all GaussDB(for MySQL) engines in a project. After a parameter template is replicated, the new template may be displayed about 5 minutes later.

Source Parameter Template

\* New Parameter Template  ?

Description  ?  
0/256

**OK**

- The template name consists of 1 to 64 characters. It can contain only uppercase letters, lowercase letters, digits, hyphens (-), underscores (\_), and periods (.).
- The description consists of up to 256 characters. It cannot contain carriage returns or any of the following special characters:  
>!<"&'='

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



----End

## 17.7 Resetting a Parameter Template

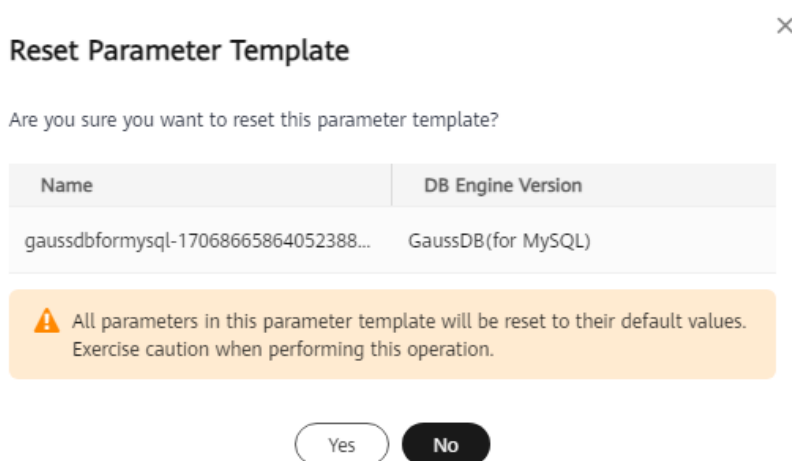
### Scenarios

You can reset all parameters in a custom parameter template to their default settings.

### Procedure

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template to be reset and choose **More > Reset** in the **Operation** column.
- Step 5** Click **Yes**.

**Figure 17-12** Confirming the reset



### NOTE

After you reset a parameter template, view the status of the instance to which the parameter template applies in the instance list. If the status is **Parameter change. Pending reboot**, you must reboot the instance.

----End

## 17.8 Applying a Parameter Template

### Scenarios


Changes to parameters in a custom parameter template do not take effect until the template is applied to instances.

- The parameter **innodb\_buffer\_pool\_size** is determined by the memory. Instances of different specifications have different value ranges. If this parameter value is out of range of the instance to which the parameter template is applied, the maximum value within the range is used.
- A parameter template can be applied only to instances of the same DB engine version.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Parameter Templates** page, apply a default template or a custom template to an instance:

- To apply a default template, click **Default Templates**, locate a parameter template and click **Apply** in the **Operation** column.
- To apply a custom template, click **Custom Templates**, locate a 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 to which the parameter template will be applied and click **OK**.

After the parameter template is applied, you can view its [application records](#).

----End

### APIs

[Applying a Parameter Template](#)


## 17.9 Viewing Application Records of a Parameter Template


### Scenarios

You can view the application records of a parameter template.

## Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

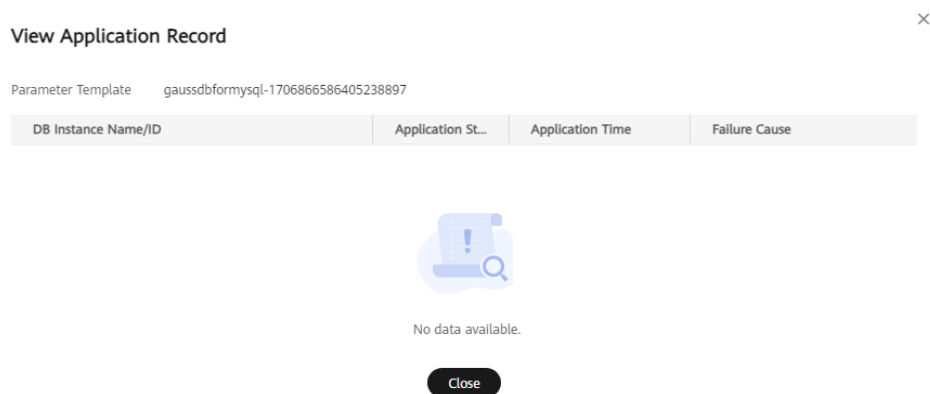
**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** Choose **Parameter Templates** in the navigation pane on the left:

- On the **Default Templates** page, locate a parameter template and click **View Application Record** in the **Operation** column.
- On the **Custom Templates** page, locate a parameter template and choose **More > View Application Record** in the **Operation** column.

You can view the name or ID of the instance the parameter template is applied to, as well as the application status, application time, and failure cause.

**Figure 17-13** Viewing application records



----End

## 17.10 Editing a Parameter Template Description

### Scenarios

You can edit the description of a parameter template you have created.





#### NOTE

You cannot edit the description of a default parameter template.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Parameter Templates** page, click **Custom Templates**. Locate the parameter template for which you want to edit the description and click  in the **Description** column.
- Step 5** Enter a new description and click  to submit or  to cancel the change.
- After the modification is successful, you can view the new description in the **Description** column.
  - The description contains up to 256 characters, and cannot contain carriage return characters and any of the following special characters:  
>!<"&'=  
  
----End

## APIs

- [Modifying Parameters in a Parameter Template](#)
- [Querying Parameter Templates](#)
- [Obtaining Details About a Parameter Template](#)

## 17.11 Deleting a Parameter Template

### Scenarios



You can delete a custom parameter template that is no longer needed.

---

#### NOTICE

- Deleted parameter templates cannot be recovered. Exercise caution when performing this operation.
  - Default parameter templates cannot be deleted.
- 

### Procedure

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- 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** In the displayed dialog box, click **Yes**.

----End

## APIs

### [Deleting a Parameter Template](#)

# 18 Metrics and Alarms

## 18.1 Introducing GaussDB(for MySQL) Metrics

### Function

You can monitor the status of your instances using Cloud Eye. The namespaces, descriptions, and dimensions of monitoring metrics of instances can be reported to Cloud Eye.

The monitoring interval can be 1 minute, 1 second, or 5 seconds. The default monitoring interval is 1 minute. To enable Monitoring by Seconds, contact customer service to apply for the required permissions.

### Namespace

SYS.GAUSSDB

### Monitoring Metrics Supported by Instances

**Table 18-1** Monitoring metrics supported by GaussDB(for MySQL) instances

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql001_cpu_util	CPU Usage	CPU usage of the monitored object	0–100%	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second



Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql002_mem_util	Memory Usage	Memory usage of the monitored object	0–100%	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql005_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql006_conn_count	Total Connections	Total number of connections that attempt to connect to the GaussDB(for MySQL) server	≥0 counts	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql007_conn_active_count	Current Active Connections	Number of active connections	≥0 counts	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql008_qps	QPS	Query times of SQL statements (including DDL, DML, SHOW, SET statements and storage procedures) per second	≥0 times/s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql009_tps	TPS	Execution times of submitted and rollback transactions per second	≥0 times/s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql010_innodb_buf_usage	Buffer Pool Usage	Ratio of used pages to total pages in the InnoDB buffer	0-1	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql011_innodb_buf_hit	Buffer Pool Hit Ratio	Ratio of read hits to read requests in the InnoDB buffer	0-1	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql012_innodb_buf_dirty	Buffer Pool Dirty Block Ratio	Ratio of dirty data to all data in the InnoDB buffer	0-100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql013_innodb_reads	InnoDB Read Throughput	Number of read bytes per second in the InnoDB buffer	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql014_innodb_writes	InnoDB Write Throughput	Bytes written to pages by InnoDB per second. GaussDB(for MySQL) writes data only to temporary tables	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql017_innodb_log_write_req_count	InnoDB Log Write Request Frequency	Number of InnoDB log write requests per second	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql019_innodb_log_writes	InnoDB Log Writes	Number of physical writes to the InnoDB redo log file	≥0 counts	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql020_temp_table_count	Temporary Tables	Number of temporary tables automatically created on disks when GaussDB(for MySQL) statements are executed	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql028_comdelete_count	DELETE Statements per Second	Number of DELETE statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql029_cominsert_count	INSERT Statements per Second	Number of INSERT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql030_cominsertselect_count	INSERT_SELECT Statements per Second	Number of INSERT_SELECT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql031_comreplace_count	REPLACE Statements per Second	Number of REPLACE statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql032_comreplaceselect_count	REPLACE_SELECT Statements per Second	Number of REPLACE_SELECT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql033_comdml_sel_count	SELECT Statements per Second	Number of SELECT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql034_comdml_upd_count	UPDATE Statements per Second	Number of UPDATE statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute 5 seconds 1 second
gaussdb_mysql035_innodb_del_row_count	Row Delete Frequency	Number of rows deleted from the InnoDB table per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql036_innodb_ins_row_count	Row Insert Frequency	Number of rows inserted into the InnoDB table per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql037_innodb_read_row_count	Row Read Frequency	Number of rows read from the InnoDB table per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql038_innodb_upd_row_count	Row Update Frequency	Number of rows updated into the InnoDB table per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql048_disk_used_size	Used Storage Space	Used storage space of the monitored object	0 GB-128 TB	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql060_rx_errors	Error Rate of Received Packets	Ratio of the number of error packets to the total number of received packets during the monitoring period	0–100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql061_rx_dropped	Loss Rate of Received Packets	Ratio of the number of lost packets to the total number of received packets during the monitoring period	0–100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql062_tx_errors	Error Rate of Sent Packets	Ratio of the number of error packets to the total number of sent packets during the monitoring period	0–100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql063_tx_dropped	Loss Rate of Sent Packets	Ratio of the number of lost packets to the total number of sent packets during the monitoring period	0–100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql072_conn_usage	Connection Usage	Percent of used GaussDB(for MySQL) connections to the total number of connections	0–100%	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql074_slow_queries	Slow Query Logs	Number of GaussDB(for MySQL) slow query logs generated per minute	≥0 counts/min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql077_replication_delay	Replication Delay	Delay between the primary node and read replicas <b>NOTE</b> This metric is used only for read replicas.	≥ 0s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql104_dfv_write_delay	Storage Write Delay	Average delay of writing data to the storage layer in a specified period	≥0 ms	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql105_dfv_read_delay	Storage Read Delay	Average delay of reading data from the storage layer in a specified period	≥0 ms	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql106_innodb_row_lock_current_waits	InnoDB Row Locks	Number of row locks being waited by operations on the InnoDB table <b>NOTE</b> If there are DDL statements, long transactions, or slow SQL statements, the number of row locks being waited may increase.	≥0 locks/s	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql107_comdm_l_ins_and_ins_sel_count	INSERT and INSERT_SELECT Statements per Second	Number of INSERT and INSERT_SELECT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql108_com_commit_count	COMMIT Statements per Second	Number of COMMIT statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql109_com_rollback_count	ROLLBACK Statements per Second	Number of ROLLBACK statements executed per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql110_innodb_bufpool_reads	InnoDB Storage Layer Read Requests per Second	Number of times that InnoDB reads data from the storage layer per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql111_innodb_bufpool_read_requests	InnoDB Read Requests per Second	Number of InnoDB read requests per second	≥0 counts /s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql114_innodb_bufpool_read_ah_ead	InnoDB Bufpool Read Ahead	Number of pages read into the InnoDB buffer pool by the read-ahead background thread	≥0 counts	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql115_innodb_bufpool_read_ah_ead_evicted	InnoDB Bufpool Read Ahead Evicted	Number of pages read into the InnoDB buffer pool by the read-ahead background thread that were subsequently evicted without having been accessed by queries	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql116_innodb_bufpool_read_ah_ead_rnd	InnoDB Bufpool Read Ahead Rnd	Number of random read-aheads initiated by InnoDB	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql117_innodb_pages_read	InnoDB Pages Read	Number of pages read from the InnoDB buffer pool by operations on InnoDB tables	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql118_innodb_pages_written	InnoDB Pages Written	Number of pages written by operations on InnoDB tables	≥0 counts	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql119_disk_usage_ratio	Disk Usage	Disk usage of the monitored object	0–100%	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql120_innodb_buffer_pool_bytes_data	Total Bytes of Buffer Pool	Total number of bytes in the InnoDB buffer pool containing data	≥0 bytes	GaussDB(for MySQL) instances	1 minute



Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql121_innodb_row_lock_time	Row Lock Time	Total time spent in acquiring row locks for InnoDB tables	≥0 ms	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql122_innodb_row_lock_waits	Row Lock Waits	Number of times operations on InnoDB tables had to wait for a row lock	≥0 counts /min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql123_sort_range	Sorts Using Ranges	Number of sorts that were done using ranges	≥0 counts /min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql124_sort_rows	Sorted Rows	Number of sorted rows	≥0 counts /min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql125_sort_scan	Sorts by Scanning Tables	Number of sorts that were done by scanning tables.	≥0 counts /min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql126_table_open_cache_hits	Hits for Open Tables Cache Lookups	Number of hits for open tables cache lookups	≥0 counts /min	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql127_table_open_cache_misses	Misses for Open Tables Cache Lookups	Number of misses for open tables cache lookups	≥0 counts /min	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb_mysql128_long_transactions_count	Long-Running Transactions	Number of long transactions that are not closed	≥0 counts	GaussDB(for MySQL) instances	150s
gaussdb_mysql342_iostat_iops_write	I/O Write IOPS	I/O write IOPS	≥0 counts/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql344_iostat_iops_read	I/O Read IOPS	I/O read IOPS	≥0 counts/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql346_iostat_throughput_write	I/O Write Bandwidth	Disk write bandwidth per second	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql348_iostat_throughput_read	I/O Read Bandwidth	Disk read bandwidth per second	≥0 bytes/s	GaussDB(for MySQL) instances	1 minute
gaussdb_mysql371_taurus_binlog_total_file_counts	Binlog Files	Number of GaussDB(for MySQL) binlog files	≥0	GaussDB(for MySQL) instances	5 minutes
gaussdb_mysql378_create_temp_tbl_per_min	Temporary Tables Created per Minute	Number of temporary tables automatically created on disks per minute when GaussDB(for MySQL) statements are executed	≥0 counts/min	GaussDB(for MySQL) instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
gaussdb386_undo_spaces_tx_count	Existing Transactions in Undo Space	Number of transactions that are not cleared in the undo space.	$\geq 0$	GaussDB(for MySQL) instances	30s

## Monitoring Metrics Supported By Proxy

**Table 18-2** GaussDB(for MySQL) Proxy metrics

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
rds_proxy_frontend_connections	Frontend Connections	Number of connections between applications and the proxy	$\geq 0$ counts	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_backend_connections	Backend Connections	Number of connections between the proxy and GaussDB(for MySQL) databases	$\geq 0$ counts	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_average_response_time	Average Response Time	Average response time	$\geq 0$ ms	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_query_per_seconds	QPS	Query times of SQL statements	$\geq 0$ counts	GaussDB(for MySQL) proxy instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
rds_proxy_read_query_proportions	Read Proportion	Proportion of read requests to total requests	0–100%	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_write_query_proportions	Write Proportion	Proportion of write requests to total requests	0–100%	GaussDB(for MySQL) proxy instances	1 minute
rds001_cpu_util	CPU Usage	CPU usage of the monitored object	0–100%	GaussDB(for MySQL) proxy instances	1 minute
rds002_mem_util	Memory Usage	Memory usage of the monitored object	0–100%	GaussDB(for MySQL) proxy instances	1 minute
rds004_bytes_in	Network Input Throughput	Incoming traffic in bytes per second	≥0 bytes/s	GaussDB(for MySQL) proxy instances	1 minute
rds005_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second	≥0 bytes/s	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_frontend_connection_creation	Front-End Connections Created per Second	Number of connections created per second between the database proxy and applications	≥ 0 counts	GaussDB(for MySQL) proxy instances	1 minute
rds_proxy_multi_statement_query	Multi-Statement Queries per Second	Number of multi-statements executed in transactions per second	≥ 0 counts	GaussDB(for MySQL) proxy instances	1 minute

Metric ID	Metric	Metric Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
rds_proxy_transaction_query	Transaction Queries per Second	Number of SELECT statements executed in transactions per second	≥ 0 counts	GaussDB(for MySQL) proxy instances	1 minute

## Dimension

**Table 18-3** Monitoring metric dimension

Key	Value
gaussdb_mysql_instance_id	GaussDB(for MySQL) instance ID.
gaussdb_mysql_node_id	GaussDB(for MySQL) node ID.
dbproxy_instance_id	GaussDB(for MySQL) proxy instance ID
dbproxy_node_id	GaussDB(for MySQL) proxy node ID

## 18.2 Viewing Monitoring Metrics

### 18.2.1 Viewing Instance Monitoring Metrics

#### Scenarios

Cloud Eye monitors status of your instances. You can view the monitoring metrics of instances on the management console. With these metrics, you can identify periods of high resource usage. You can also check error logs or slow query logs to optimize database performance.

#### Prerequisites

- Instances are running properly.  
Monitoring metrics of the instances that are faulty or have been deleted cannot be displayed on the Cloud Eye console, but you can view them after the instances are rebooted or restored to be available.


 NOTE

If an instance has been faulty for 24 hours, Cloud Eye considers that it does not exist and deletes it from the monitoring object list. You need to manually clear the alarm rules created for the instance.


- Instances have kept running properly for about 10 minutes.  
For a newly created instance, you need to wait for a while before viewing the monitoring metrics.

## Viewing Monitoring Metrics of Nodes

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.

**Step 3** In the navigation pane on the left, choose **Cloud Service Monitoring > GaussDB(for MySQL)**.

**Step 4** Click  in the front of the instance. Locate a node and click **View Metric** in the **Operation** column.

You can also perform the following operations to switch to the Cloud Eye console:

- On the **Instances** page, click the instance name to go to the **Basic Information** page. In the upper right corner of the page, click **View Metric** to go to the Cloud Eye console and view monitoring metrics.
- In the **Node List** area of the **Basic Information** page, locate the primary node or a read replica and click **View Metric** in the **Operation** column to go to the Cloud Eye console and view the monitoring metrics.

**Step 5** View monitoring metrics of the node. Click **Select Metric** in the upper right corner. In the displayed dialog box, select the metrics to be displayed and sort them by dragging them at desired locations.


Cloud Eye can monitor performance metrics from the last 1 hour, last 3 hours, last 12 hours, last 24 hours or last 7 days.

----End

## Viewing Real-Time Instance Monitoring Metrics

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Advanced O&M**.

**Step 6** Under **Real-Time Monitoring**, view real-time monitoring data such as CPU usage, memory usage, SELECT statements per second, DELETE statements per second, and INSERT statements per second.

You can also click **View details** to view more metrics on the Cloud Eye console.

----End

## 18.2.2 Viewing Proxy Monitoring Metrics


### Prerequisites

The instance is running properly.


Read/write splitting has been enabled for the instance. For details, see [Create a Proxy Instance](#).

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.

**Step 3** In the navigation pane on the left, choose **Cloud Service Monitoring > Database Proxy Service**.

**Step 4** In the instance list, click  in the front of the target instance. Locate a node and click **View Metric** in the **Operation** column.

**Step 5** View monitoring metrics of the node. Click **Select Metric** in the upper right corner. In the displayed dialog box, you can select the metrics to be displayed and sort them by dragging them at desired locations.

Cloud Eye can monitor performance metrics from the last 1 hour, last 3 hours, last 12 hours, last 24 hours or last 7 days.

----End

## 18.3 Configuring Alarm Rules

### 18.3.1 Creating Alarm Rules for a DB Instance


#### Scenarios

You can create alarm rules for GaussDB(for MySQL) to customize the monitored objects and notification policies and stay aware of the GaussDB(for MySQL) instance statuses.


The GaussDB(for MySQL) alarm rules include alarm rule names, services, dimensions, monitored objects, metrics, alarm thresholds, monitoring period, and whether to send notifications.

## Creating Alarm Rules for Instances


**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.

**Step 3** In the navigation pane on the left, choose **Cloud Service Monitoring > GaussDB(for MySQL)**.

**Step 4** In the instance list, click  in the front of the instance. Locate a node and click **Create Alarm Rule** in the **Operation** column.

**Figure 18-1** Creating an alarm rule for an instance



**Step 5** On the **Create Alarm Rule** page, configure parameters as needed. For details, see [Creating an Alarm Rule](#).

1. Configure the alarm rule name and description.

**Table 18-4** Parameter description

Parameter	Parameter description
Name	Specifies the name of the alarm rule. The system generates a random name, but you can change it if needed. Example value: <b>alarm-b6al</b>
Description	(Optional) Provides supplementary information about the alarm rule.

2. Configure alarm content parameters.

**Table 18-5** Parameter description

Parameter	Parameter description
Method	Select an associated template, use an existing template or create a custom template as required. <ul style="list-style-type: none"><li>– Modifying the template will also modify its associated alarm rules.</li><li>– If you select <b>Configure manually</b>, you can configure <b>Alarm Policy</b> and <b>Alarm Severity</b> as required.</li></ul>



Parameter	Parameter description
Template	Select the template to be used. You can select a default alarm template or create a custom template. For details about how to create a custom template, see <a href="#">Creating a Custom Template</a>
Alarm Policy	Specifies the policy for triggering an alarm. 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.

3. Configure alarm notification parameters.

**Figure 18-2** Configuring an alarm notification



**Table 18-6** Parameter description

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email or text message, or through HTTP/HTTPS request to servers.
Notification Recipient	You can select a notification group or topic subscription as required.
Notification Group	Specifies the notification group that needs to send alarm notifications.

Parameter	Description
Notification Object	Specifies the object an alarm notification is to be sent to. You can select the account contact or a topic. <ul style="list-style-type: none"> <li>Account contact is the mobile phone number and email address provided for registration.</li> <li>Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see <a href="#">Creating a Topic</a> and <a href="#">Adding Subscriptions</a>.</li> </ul>
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If you set <b>Validity Period</b> to <b>08:00-20:00</b> , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification. You can select <b>Generated alarm</b> (when an alarm is generated), <b>Cleared alarm</b> (when an alarm is cleared), or both.

4. Configure the enterprise project and tag.

Figure 18-3 Configuring advanced settings




Table 18-7 Parameter description


Parameter	Description
Enterprise Project	Specifies the enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule. For details about how to create an enterprise project, see <a href="#">Creating an Enterprise Project</a> .
Tag	Adding tags helps you better identify and manage your DB instances.


**Step 6** Click **Create**. The alarm rule is created.

----End

## Creating Alarm Rules for Metrics

**Step 1** [Log in to the management console.](#)


**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

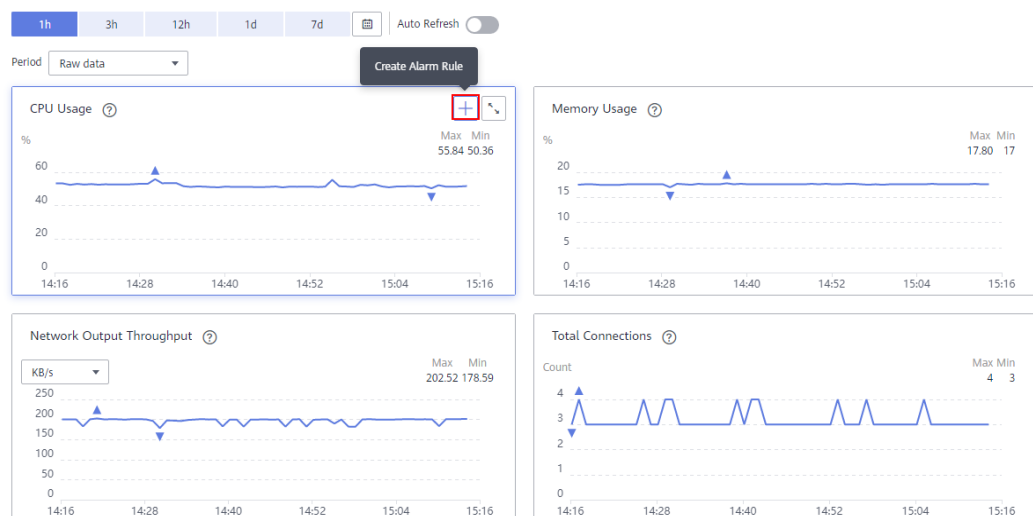
**Step 4** On the **Instances** page, locate the instance and click **View Metric** in the **Operation** column to go to the Cloud Eye console.

Alternatively, go to the Cloud Eye console using either of the following methods:

- On the displayed **Basic Information** page, click **View Metric** in the upper right corner.
- In the **Node List** area of the **Basic Information** page, locate a node and click **View Metric** in the **Operation** column.

**Step 5** Locate the monitoring metric that you want to create an alarm for and click  in the upper right corner of the metric.

**Figure 18-4** Creating an alarm rule for a metric



**Step 6** On the **Create Alarm Rule** page, configure parameters as needed. For details, see [Creating an Alarm Rule](#).

1. Configure the alarm rule name and description.

**Table 18-8** Parameter description

Parameter	Parameter description
Name	Specifies the name of the alarm rule. The system generates a random name, but you can change it if needed.  Example value: <b>alarm-b6al</b>

Parameter	Parameter description
Description	(Optional) Provides supplementary information about the alarm rule.

2. Configure alarm content parameters.

**Table 18-9** Parameter description

Parameter	Description
Method	Specifies the method for triggering an alarm. If you select <b>Configure manually</b> , you can configure <b>Alarm Policy</b> and <b>Alarm Severity</b> as required.
Alarm Policy	Specifies the policy for triggering an alarm.

3. Configure alarm notification parameters.

**Figure 18-5** Configuring an alarm notification



**Table 18-10** Parameter description

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email or text message, or through HTTP/HTTPS request to servers.
Notification Recipient	You can select a notification group or topic subscription as required.
Notification Group	Specifies the notification group that needs to send alarm notifications.

Parameter	Description
Notification Object	Specifies the object an alarm notification is to be sent to. You can select the account contact or a topic. <ul style="list-style-type: none"> <li>Account contact is the mobile phone number and email address provided for registration.</li> <li>Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see <a href="#">Creating a Topic</a> and <a href="#">Adding Subscriptions</a>.</li> </ul>
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If you set <b>Validity Period</b> to <b>08:00-20:00</b> , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification. You can select <b>Generated alarm</b> (when an alarm is generated), <b>Cleared alarm</b> (when an alarm is cleared), or both.

4. Configure the enterprise project and tag.

Figure 18-6 Configuring advanced settings




Table 18-11 Parameter description

Parameter	Description
Enterprise Project	Specifies the enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule. For details about how to create an enterprise project, see <a href="#">Creating an Enterprise Project</a> .
Tag	Adding tags helps you better identify and manage your DB instances.

**Step 7** Click **Create**.

----End



## 18.3.2 Creating Alarm Rules for a Proxy Instance

### Scenarios

You can configure alarm rules for a database proxy instance to customize the monitored objects and notification policies and stay aware of the proxy instance statuses.

The alarm rules include alarm rule names, alarm thresholds, monitoring period, and whether to send notifications.

### Procedure

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.
- Step 3** In the navigation pane on the left, choose **Cloud Service Monitoring > Database Proxy Service**.
- Step 4** In the instance list, click  in the front of the instance. Locate a node and click **Create Alarm Rule** in the **Operation** column.
- Step 5** On the **Create Alarm Rule** page, configure parameters as needed. For details, see [Creating an Alarm Rule](#).
  1. Configure the alarm rule name and description.

**Table 18-12** Parameter description

Parameter	Parameter description
Name	Specifies the name of the alarm rule. The system generates a random name, but you can change it if needed. Example value: <b>alarm-b6al</b>
Description	(Optional) Provides supplementary information about the alarm rule.

2. Configure alarm content parameters.

**Table 18-13** Parameter description

Parameter	Parameter description
Method	Select an associated template, use an existing template or create a custom template as required. <ul style="list-style-type: none"> <li>– Modifying the template will also modify its associated alarm rules.</li> <li>– If you select <b>Configure manually</b>, you can configure <b>Alarm Policy</b> and <b>Alarm Severity</b> as required.</li> </ul>
Template	Select the template to be used. You can select a default alarm template or create a custom template. For details about how to create a custom template, see <a href="#">Creating a Custom Template</a>
Alarm Policy	Specifies the policy for triggering an alarm. 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.

## 3. Configure alarm notification parameters.

**Figure 18-7** Configuring an alarm notification

Alarm Notification 

\* Notification Recipient Notification group Topic subscription

\* Notification Group  

If you create notification group, you must click refresh to make it available for selection. After you create the notification group, click Add Notification Object in the Operation column of the notification group list to add notification objects.

\* Notification Window Daily  -  GMT+08:00 

\* Trigger Condition  Generated alarm  Cleared alarm


**Table 18-14** Parameter description

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email or text message, or through HTTP/HTTPS request to servers.
Notification Recipient	You can select a notification group or topic subscription as required.

Parameter	Description
Notification Group	Specifies the notification group that needs to send alarm notifications.
Notification Object	Specifies the object an alarm notification is to be sent to. You can select the account contact or a topic. <ul style="list-style-type: none"> <li>Account contact is the mobile phone number and email address provided for registration.</li> <li>Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see <a href="#">Creating a Topic</a> and <a href="#">Adding Subscriptions</a>.</li> </ul>
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If you set <b>Validity Period</b> to <b>08:00-20:00</b> , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification. You can select <b>Generated alarm</b> (when an alarm is generated), <b>Cleared alarm</b> (when an alarm is cleared), or both.

4. Configure the enterprise project and tag.

**Figure 18-8** Configuring advanced settings



**Table 18-15** Parameter description

Parameter	Description
Enterprise Project	Specifies the enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule. For details about how to create an enterprise project, see <a href="#">Creating an Enterprise Project</a> .
Tag	Adding tags helps you better identify and manage your DB instances.



**Step 6** Click **Create**.

----End

## 18.4 Event Monitoring

### 18.4.1 Introducing Event Monitoring

Event monitoring provides reporting, query, and alarm functions for event data. You can create alarm rules for both system events and custom events. When specific events occur, Cloud Eye generates alarms for you.

Events are key operations on GaussDB(for MySQL) that are stored and monitored by Cloud Eye. You can view events to see operations performed by specific users on specific resources, such as deleting a read replica or changing instance specifications.

Event monitoring provides an API for reporting custom events (abnormal events or important change events) generated by services to Cloud Eye.

Event monitoring is enabled by default. You can view monitoring details about system events and custom events. For details about system events, see [Events Supported by Event Monitoring](#).

### 18.4.2 Viewing Event Monitoring Data

#### Scenarios


In event monitoring, you can query system events that are automatically reported to Cloud Eye and custom events reported to Cloud Eye through the API. You can create alarm rules for both system events and custom events. When specific events occur, Cloud Eye generates alarms for you.


Event monitoring is enabled by default.

You can view the event monitoring data.

#### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.


**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, locate the instance and click **View Metric** in the **Operation** column to go to the Cloud Eye console.

Alternatively, go to the Cloud Eye console using either of the following methods:

- On the displayed **Basic Information** page, click **View Metric** in the upper right corner.

- In the **Node List** area of the **Basic Information** page, locate a node and click **View Metric** in the **Operation** column.

**Step 5** Click  to return to the Cloud Eye console.

**Step 6** In the navigation pane on the left, choose **Event Monitoring**.

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

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

**Step 7** Locate an event and click **View Event** in the **Operation** column to view details about a specific event.

----End


## 18.4.3 Creating Alarm Rules for Event Monitoring

### Scenarios

You can create alarm rules for event monitoring.

### Procedure

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.

**Step 3** In the navigation pane on the left, choose **Event Monitoring**.


**Step 4** On the event list page, click **Create Alarm Rule** in the upper right corner.

**Step 5** On the displayed page, configure parameters if needed.

**Table 18-16** Parameter description

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

Parameter	Description
Event Source	Specifies the service the event is generated for. Example value: <b>GaussDB(for MySQL)</b> .
Monitoring Scope	Specifies the monitoring scope for event monitoring.
Method	Specifies the event creation method.
Alarm Policy	Events indicate the instantaneous operations users performed on system resources, such as login and logout. For details about events supported by Event Monitoring, see <a href="#">Events Supported by Event Monitoring</a> . You can select a trigger mode and alarm severity as needed.

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

**Table 18-17** Alarm notification parameters

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email or text message, or through HTTP/HTTPS request to servers.
Notification Object	Specifies the object an alarm notification is to be sent to. You can select the account contact or a topic. <ul style="list-style-type: none"><li>Account contact is the mobile phone number and email address provided for registration.</li><li>Topic is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see <a href="#">Creating a Topic</a> and <a href="#">Adding Subscriptions</a>.</li></ul>
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If you set <b>Validity Period</b> to <b>08:00-20:00</b> , Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering the alarm notification.

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

----End

## 18.4.4 Events Supported by Event Monitoring

**Table 18-18** GaussDB(for MySQL)

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
GaussDB(for MySQL)	Incremental backup failure	TaurusIncrementalBackupInstanceFailed	Major	The network between the instance and the management plane or the OBS is disconnected, or the backup environment created for the instance is abnormal.	Submit a service ticket.	Backup jobs fail.
	Read replica creation failure	addReadOnlyNodesFailed	Major	The quota is insufficient or underlying resources are exhausted.	Check the read replica quota. Release resources and create read replicas again.	Read replicas fail to be created.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	DB instance creation failure	createInstanceFailed	Major	The quota is insufficient or underlying resources are exhausted.	Check the instance quota. Release resources and create instances again.	Instances fail to be created.
	Read replica promotion failure	activeStandBySwitchFailed	Major	The read replica fails to be promoted to the primary node due to network or server failures. The original primary node takes over services quickly.	Submit a service ticket.	The read replica fails to be promoted to the primary node.
	Instance specifications change failure	flavorAlterationFailed	Major	The quota is insufficient or underlying resources are exhausted.	Submit a service ticket.	Instance specifications fail to be changed.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Faulty DB instance	TaurusInstanceRunningStatusAbnormal	Major	The instance process is faulty or the network between the instance and the DFV storage is disconnected.	Submit a service ticket.	Services may be affected.
	DB instance recovered	TaurusInstanceRunningStatusRecovered	Major	The instance is recovered.	Observe the service running status.	None.
	Faulty node	TaurusNodeRunningStatusAbnormal	Major	The node process is faulty or the network between the node and the DFV storage is disconnected.	Observe the instance and service running status es.	A read replica may be promoted to the primary node.
	Node recovered	TaurusNodeRunningStatusRecovered	Major	The node is recovered.	View the node running status.	None.
	Read replica deletion failure	TaurusDeleteReadOnlyNodeFailed	Major	The network between the management plane and the read replica is disconnected or the VM fails to be deleted from IaaS.	Submit a service ticket.	Read replica fails to be deleted.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Password reset failure	TaurusResetInstancePasswordFailed	Major	The network between the management plane and the instance is disconnected or the instance is abnormal.	Check the instance status and try again. If the fault persists, submit a service ticket.	Passwords fail to be reset for instances.
	DB instance reboot failure	TaurusRestartInstanceFailed	Major	The network between the management plane and the instance is disconnected or the instance is abnormal.	Check the instance status and try again. If the fault persists, submit a service ticket.	Instances fail to be rebooted.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Restoration to new DB instance failure	TaurusRestoreToNewInstanceFailed	Major	The instance quota is insufficient, underlying resources are exhausted, or the data restoration logic is incorrect.	If the new instance fails to be created, check the instance quota, release resources, and try to restore to a new instance again. In other cases, submit a service ticket.	Backup data fails to be restored to new instances.
	EIP binding failure	TaurusBindEIPToInstanceFailed	Major	The binding task fails.	Submit a service ticket.	EIPs fail to be bound to instances.



Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	EIP unbinding failure	TaurusUnbindEIPFromInstanceFailed	Major	The unbinding task fails.	Submit a service ticket.	EIPs fail to be unbound from instances.
	Parameter modification failure	TaurusUpdateInstanceParameterFailed	Major	The network between the management plane and the instance is disconnected or the instance is abnormal.	Check the instance status and try again. If the fault persists, submit a service ticket.	Instance parameters fail to be modified.
	Parameter template application failure	TaurusApplyParameterGroupToInstanceFailed	Major	The network between the management plane and instances is disconnected or the instances are abnormal.	Check the instance status and try again. If the fault persists, submit a service ticket.	Parameter templates fail to be applied to instances.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Full backup failure	TaurusBackupInstanceFailed	Major	The network between the instance and the management plane or the OBS is disconnected, or the backup environment created for the instance is abnormal.	Submit a service ticket.	Backup jobs fail.
	Read replica promotion	TaurusActiveStandbySwitched	Major	When the primary node is faulty, a read replica is promoted to the primary node.	Check the instance status. If the fault persists, submit a service ticket.	Services are intermittently interrupted.
	Instance read-only	NodeReadOnlyMode	Major	The instance supports only query operations.	Submit a service ticket.	After the instance becomes read-only, write requests cannot be processed.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Instance read/write	NodeReadWriteMode	Major	The instance can process both write and read requests.	Submit a service ticket.	None.
	Instance DR switchover	DisasterSwitchOver	Major	If an instance is faulty and unavailable, a switchover is performed to ensure that the instance continues to provide services.	Contact technical support.	The database connection is intermittently interrupted. The DR instance is promoted to primary to provide services.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Database process restarted	TaurusDatabaseProcessRestarted	Major	The database process is stopped due to insufficient memory or high load.	Log in to the Cloud Eye console. Check whether the memory usage increases sharply or the CPU usage is too high for a long time. You can increase the specifications or optimize the service logic.	When the database process is suspended, workloads on the node are interrupted. In this case, the HA service automatically restarts the database process and attempts to recover the workloads.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Connection failure between proxy instance and DB instance	proxy_connection_failure_to_db	Major	The database proxy failed to establish a new connection with the primary node of a DB instance, and it may fail to establish a new connection with a read replica. The DB instance or proxy instance is overloaded, or the network between them is abnormal.	Change values of related parameters based on metrics (Connections, Active Connections, and CPU Usage) of the DB instance and proxy instance. If the metrics are normal, submit a service ticket.	Service requests accessed through the proxy instance are interrupted.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Connection failure between database proxy and read replica	proxy_connection_failure_to_replica	General	The proxy instance failed to establish a new connection with a read replica. The read replica is overloaded, or the network between the proxy instance and read replica is abnormal.	Change values of related parameters based on metrics (Connections, Active Connections, and CPU Usage) of the read replica. If the metrics are normal, submit a service ticket.	Read requests accessed through the proxy instance are interrupted.

Event Source	Event Name	Event ID	Alarm Severity	Description	Handling Suggestion	Impact
	Proxy instance access to DB instance failure	proxy_connection_failure_cause_security_group	Major	No rules in the security group allow the proxy instance to access the DB instance.	Add the proxy instance address to the rules of the security group.	Service requests accessed through the proxy instance are interrupted.

## 18.5 Configuring Monitoring by Seconds

GaussDB(for MySQL) supports Monitoring by Seconds. You can set the monitoring interval to 1 second or 5 seconds to view the metric values.

### Billing

Pricing is listed on a per-hour basis, but bills are calculated based on the actual usage duration. GaussDB(for MySQL) provides monitoring every 60 seconds for free.

**Table 18-19** Price details


Region	Monitoring Interval	Pay-per-Use (USD/Hour)
CN East-Shanghai1, CN North-Beijing4, CN South-Guangzhou, CN Southwest-Guiyang1, CN North-Ulanqab1, and CN South-Guangzhou-InvitationOnly	1s	0.024
	5s	0.012

Region	Monitoring Interval	Pay-per-Use (USD/ Hour)
AP-Singapore, AP-Jakarta, RU-Moscow2, CN-Hong Kong, AP-Bangkok, and TR-Istanbul	1s	0.032
	5s	0.016
LA-Sao Paulo1	1s	0.054
	5s	0.027

## Enabling Monitoring by Seconds

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **DBA Assistant** > **Real-Time Diagnosis**. The **Dashboard** page is displayed.

**Step 6** Click **Performance**.

**Step 7** In the upper part of the page, click **Enable Monitoring by Seconds**.

**Step 8** In the displayed dialog box, click  next to **Monitoring by Seconds**, select a collection interval, and click **OK**.

After you enable this function, monitoring data will be reported and displayed by the second after about five minutes.


**Step 9** In the navigation pane on the left, click **Advanced O&M** > **Real-Time Monitoring** to view the monitoring metric data.

- View the current data collection period in the upper part of the page.
- Monitoring by Seconds supports the following metrics: CPU usage, memory usage, SELECT statements per second, DELETE statements per second, and INSERT statements per second. You can click **View details** to view more metrics. For details about the metrics, see [Introducing GaussDB\(for MySQL\) Metrics](#).
- If you need to change the collection period, see [Modifying Collection Interval](#).

----End



## Disabling Monitoring by Seconds

- Step 1** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 2** In the navigation pane on the left, choose **DBA Assistant > Real-Time Diagnosis**. The **Dashboard** page is displayed.
- Step 3** Click **Performance**.
- Step 4** In the upper part of the page, click **Enable Monitoring by Seconds**.
- Step 5** In the displayed dialog box, click  next to **Monitoring by Seconds** and click **OK**.

After you disable this function, monitoring data will be reported and displayed by the minute after about five minutes.

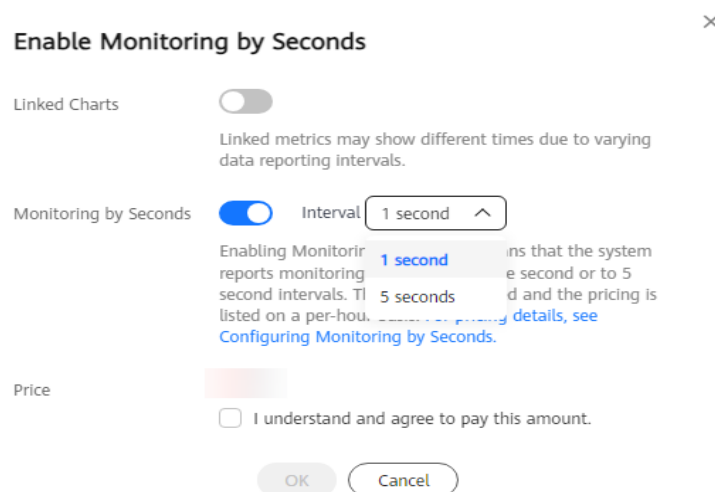
----End

## Modifying Collection Interval

- Step 1** On the **Instances** page, click the instance name to go to the **Basic Information** page.
- Step 2** In the navigation pane on the left, choose **DBA Assistant > Real-Time Diagnosis**. The **Dashboard** page is displayed.
- Step 3** Click **Performance**.
- Step 4** In the upper part of the page, click **Enable Monitoring by Seconds**.
- Step 5** Select the monitoring interval and click **OK**.

Monitoring data will be reported based on the new collection interval about 5 minutes later.

**Figure 18-9** Modifying the collection interval



----End

## APIs

- [Configuring the Monitoring By Seconds Function](#)
- [Querying the Configuration of Monitoring by Seconds](#)

# 19 Interconnection with CTS

## 19.1 Key Operations Supported by CTS

Cloud Trace Service (CTS) records operations related to GaussDB(for MySQL) for further querying, auditing, and backtracking. [Table 19-1](#) lists the supported operations.

**Table 19-1** GaussDB(for MySQL) operations recorded by CTS

Operation	Resource Type	Trace Name
Creating a DB instance	instance	createInstance
Creating a read replica	instance	addNodes
Deleting a read replica	instance	deleteNode
Rebooting a DB instance	instance	restartInstance
Changing a database port	instance	changeInstancePort
Changing a security group	instance	modifySecurityGroup
Promoting a read replica to the primary node	instance	instanceSwitchOver
Binding or unbinding an EIP	instance	setOrResetPublicIP
Deleting a DB instance	instance	deleteInstance
Renaming a DB instance	instance	renameInstance
Changing a failover priority	instance	modifyPriority
Creating a database	instance	createDatabase

Operation	Resource Type	Trace Name
Creating a database account	instance	createDatabaseUser
Resetting a password	instance	resetPassword
Deleting a database	instance	dropDatabase
Deleting a database account	instance	dropDatabaseUser
Changing the password of a database user	instance	modifyDatabaseUserPwd
Restoring data to a new DB instance	instance	restoreInstance
Enabling read/write splitting	instance	openProxy
Disabling read/write splitting	instance	closeProxy
Assigning read weights	instance	setProxyWeight
Changing the CPU and memory specifications of an instance	instance	resizeFlavorOrVolume
Configuring monitoring by seconds	instance	openSecondExtend
Upgrading a minor version	instance	upgradeVersion
Adding a tag	instance	addInstanceTags
Authorizing database user permissions	instance	grantDatabaseUser
Revoking database user permissions	instance	revokeDatabaseUser
Applying for a private domain name	instance	createDnsName
Modifying a private domain name	instance	modifyDnsName
Changing the routing policy of a proxy instance	instance	modifyProxyRouteMode
Changing the port of a proxy instance	instance	modifyProxyPort

Operation	Resource Type	Trace Name
Applying for a private domain name for a database proxy instance	instance	proxyCreateDns
Changing a private domain name for a database proxy instance	instance	modifyProxyDnsName
Deleting a private domain name for a database proxy instance	instance	deleteProxyDnsName
Deleting database proxy nodes	instance	reduceProxy
Creating a backup	backup	createManualSnapshot
Configuring an automated backup policy	backup	setBackupPolicy
Deleting a backup	backup	deleteManualSnapshot
Creating a parameter template	parameterGroup	createParameterGroup
Modifying parameters in a parameter template	parameterGroup	updateParameterGroup
Deleting a parameter template	parameterGroup	deleteParameterGroup
Replicating a parameter template	parameterGroup	copyParameterGroup
Resetting a parameter template	parameterGroup	resetParameterGroup
Comparing parameter templates	parameterGroup	compareParameterGroup
Applying a parameter template	parameterGroup	applyParameterGroup

## 19.2 Viewing Tracing Events

### Scenarios



After CTS is enabled, operations on cloud resources are recorded. You can view the operation records of the last 7 days on the CTS console.

This section describes how to query the operation records of last 7 days on the CTS console.

 NOTE

Before using CTS, you need to enable it. For details, see [Enabling CTS](#).

## Procedure

- Step 1** [Log in to the management console](#).
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** In the upper left corner of the page, click  and choose **Management & Governance > Cloud Trace Service**.
- Step 4** In the navigation pane on the left, choose **Trace List**.
- Step 5** Filter conditions to query traces.

**Table 19-2** Filtering criteria

Filtering Criteria	Description
Time Range	In the upper right corner, choose <b>Last 1 hour</b> , <b>Last 1 day</b> , or <b>Last 1 week</b> , or specify a custom time range.
Trace Type	Select <b>Management</b> or <b>Data</b> <ul style="list-style-type: none"><li>• Management traces record details about creating, configuring, and deleting cloud service resources in your tenant account.</li><li>• Data traces record operations on data, such as data upload and download.</li></ul> <b>NOTE</b> <ul style="list-style-type: none"><li>- If you select <b>Data</b> for <b>Trace Type</b>, you can only filter traces by tracker.</li><li>- The trace list does not record queries.</li></ul>
Trace Source	Select a trace source as needed.
Resource Type	Select a resource type as needed.
Search By	If you select <b>Resource ID</b> for <b>Search By</b> , you need to enter a resource ID.
Operator	Select a specific operator from the drop-down list.
Trace Status	Select <b>All trace statuses</b> , <b>Normal</b> , <b>Warning</b> , or <b>Incident</b> .

- Step 6** View the events that meet the search criteria.
- Step 7** Click an event name. Details about the event are displayed in the dialog box on the right.
- Step 8** Click **Export** in the upper left corner of the list. CTS exports traces collected in the past seven days to a CSV file. The CSV file contains all information related to the traces.

For details about key fields in the CTS trace structure, see sections "[Trace Structure](#)" and "[Example Traces](#)" in the *Cloud Trace Service User Guide*.

----End

# 20 Log Management

---

## 20.1 Enabling or Disabling Log Reporting

GaussDB(for MySQL) log management allows you to view database-level logs, including error logs and slow SQL query logs.

### Scenarios

If you enable log reporting for your DB instance, new logs generated for the instance will be uploaded to [Log Tank Service \(LTS\)](#) for management.

### Precautions

- You will be billed for this function.
- Ensure that there are available LTS log groups and log streams in the same region as your instance.


### Constraints

- Error logs and slow query logs cannot share the same log stream.
- You can bind a new structuring template to an error log stream or slow log query stream, but once selected, the log stream type cannot be changed.
- If a structuring template has been bound to a log stream, ensure that the template type is the same as the log type when you select the log stream. For example, if an error log template has been bound to a log stream, the log stream cannot be used for slow query logs.

### Enabling Log Reporting in Batches

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.



**Step 4** In the navigation pane, choose **Log Reporting**.

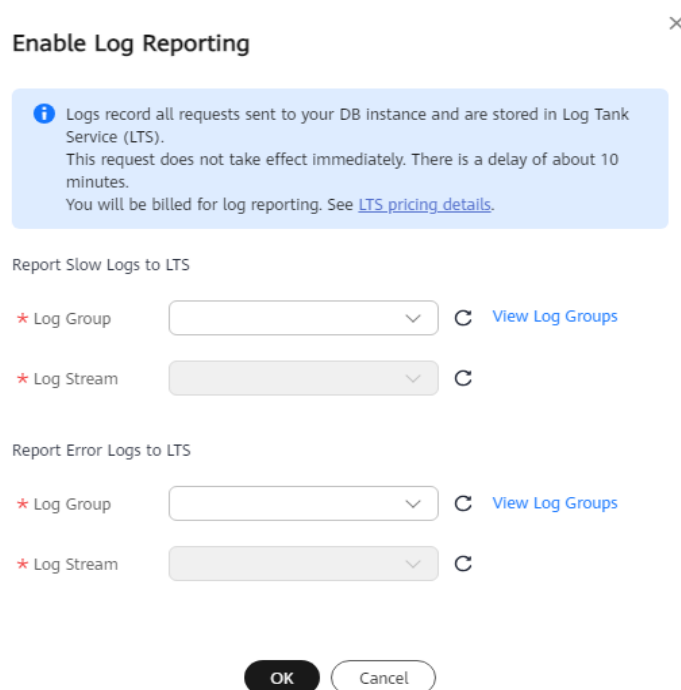
**Step 5** Select one or more instances and click **Enable Log Reporting**.

**Step 6** In the displayed dialog box, select a log group and log stream, and click **OK**.

 **NOTE**

- Error logs and slow query logs cannot share the same log stream.
- This request does not take effect immediately. There is a delay of about 10 minutes.
- You can only enable either error log reporting to LTS or slow log reporting to LTS.
- Audit logs record all requests sent to your DB instance and are stored in LTS.

**Figure 20-1** Enabling log reporting



**Enable Log Reporting** ×

**i** Logs record all requests sent to your DB instance and are stored in Log Tank Service (LTS). This request does not take effect immediately. There is a delay of about 10 minutes. You will be billed for log reporting. See [LTS pricing details](#).

Report Slow Logs to LTS

\* Log Group  ⌵ ⌵ [View Log Groups](#)

\* Log Stream  ⌵ ⌵

Report Error Logs to LTS

\* Log Group  ⌵ ⌵ [View Log Groups](#)


\* Log Stream  ⌵ ⌵


**OK** **Cancel**

----End

## Disabling Log Reporting

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

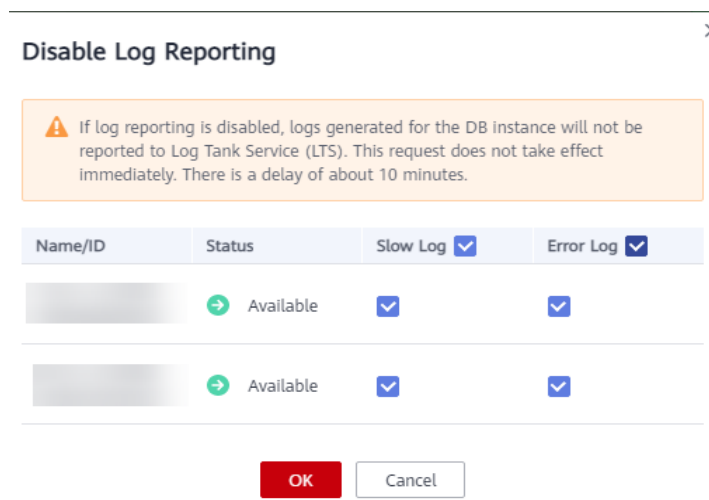
**Step 4** In the navigation pane, choose **Log Reporting**.


**Step 5** Disable log reporting in either of the following ways:

 NOTE

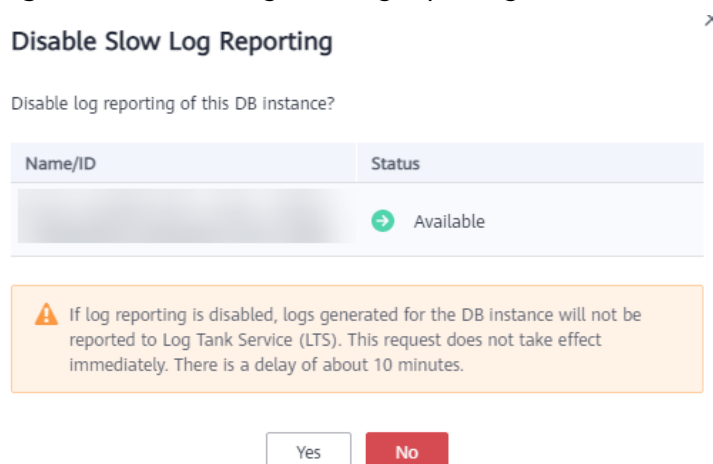
- If log reporting is disabled, logs generated for the DB instance will not be reported to LTS.
- This request is not applied immediately. There is a delay of about 10 minutes.
- Disabling log reporting for multiple instances in batches
  - a. Select one or more instances and click **Disable Log Reporting**.
  - b. In the displayed dialog box, click **OK**.

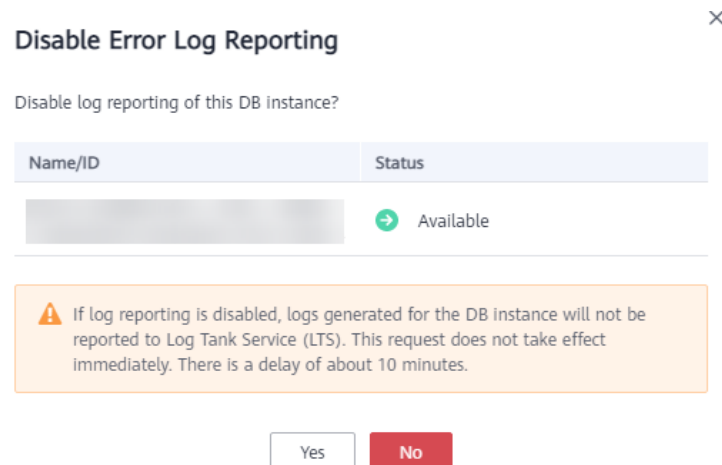
**Figure 20-2** Disabling log reporting



- Disabling log reporting for a DB instance
  - a. Locate a DB instance and click  in the **Report Error Log to LTS** column or **Report Slow Log to LTS** column.
  - b. In the displayed dialog box, click **Yes**.

**Figure 20-3** Disabling slow log reporting



**Figure 20-4** Disabling error log reporting


----End


## 20.2 Viewing Error Logs

Error logs contain logs generated during the database running. These can help you analyze problems with the database.

### Viewing Log Details

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Logs**.

**Step 6** On the **Error Logs** page, view error logs of different nodes, at different log levels, and within a specified time range.

Click the drop-down list in the upper right corner, and select a node name and a log level as needed.

The levels of error logs include ALL, INFO, WARNING, ERROR, FATAL and NOTE.

Click  and specify a time period.


Figure 20-5 Viewing error logs


Time	Log Level	Description
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-013003] [InnoDB] Information of replica 1: write bin 26497056, recycle bin 26497056, low_limit_no 24836
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-018231] [InnoDB] Delete one record in innodb_membership table, server_id: 1
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-013682] [InnoDB] Replica 1 expired.
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-013582] [InnoDB] Information of replica 2: write bin 26497056, recycle bin 26497056, low_limit_no 24836
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-018231] [InnoDB] Delete one record in innodb_membership table, server_id: 2
Jan 23, 2024 10:14:05 GMT+08:00	WARNING	[MY-013682] [InnoDB] Replica 2 expired.
Jan 23, 2024 10:12:06 GMT+08:00	WARNING	[MY-010008] [Server] CA certificate cA/csr.pem is self signed.
Jan 23, 2024 10:12:06 GMT+08:00	WARNING	[MY-011018] [Server] Skip updating information_schema metadata in InnoDB read-only mode.
Jan 23, 2024 10:12:06 GMT+08:00	WARNING	[MY-010009] [Server] Skip re-populating collations and character sets tables in read-only mode.
Jan 23, 2024 10:12:06 GMT+08:00	WARNING	[MY-018235] [InnoDB] Failed to set priority of log reader thread to -20

----End

## Reporting Error Logs to LTS

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Logs**.

**Step 6** On the **Error Log** page, click  next to **Report Error Log to LTS**.

**Step 7** Select an LTS log group and log stream and click **OK**.

Figure 20-6 Reporting error logs to LTS

✕

**Enable Error Log Reporting to LTS**

**i** Logs record all requests sent to your DB instance and are stored in Log Tank Service (LTS). This request does not take effect immediately. There is a delay of about 10 minutes. You will be billed for log reporting. See [LTS pricing details](#).

\* Log Group  View Log Groups

\* Log Stream

----End

## APIs

### Querying Error Logs

## 20.3 Viewing Slow Query Logs

### Scenarios

Slow query logs record statements that exceed **long\_query\_time** (10 seconds by default). You can view log details and statistics to identify statements that are executing slowly and optimize the statements.

GaussDB(for MySQL) supports the following statement types:

- SELECT
- INSERT
- UPDATE
- DELETE
- CREATE
- ALTER
- DROP

### Parameter Description


**Table 20-1** Parameters related to slow queries

Parameter	Description
long_query_time	Specifies how many seconds a SQL query has to take to be recorded in slow query logs. The default value is 10s. You are advised to set this parameter to 1s. The lock wait time is not calculated into the query time.
log_queries_not_using_indexes	Specifies whether to record the slow query that without indexes. The default value is <b>OFF</b> .
log_throttle_queries_not_using_indexes	Specifies the SQL statement that can be written to the slow query log every minute. The default value is <b>0</b> .

### Viewing Log Details

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Logs**.

**Step 6** On the **Slow Query Logs** page, view the slow query log details.

You can [download slow query logs](#).


**Step 7** View slow query logs of different nodes in a given database and SQL statement types. In the upper right corner of the page:

Enter a database name, click the drop-down list, and select your desired node.

Click the drop-down list and select a SQL statement type (SELECT, INSERT, UPDATE, DELETE, CREATE, ALTER, or DROP).

Click  and specify a time period.

**Figure 20-7** Viewing slow query logs



----End

## Enabling Show Original Log


### NOTE

By default, SQL statements are displayed anonymously. If **Show Original Log** is enabled, SQL statements in the logs will be displayed in plaintext.


Logs displayed in plaintext will be automatically deleted 30 days later. If a DB instance is deleted, its related logs will also be deleted.

**Step 1** [Log in to the management console](#).

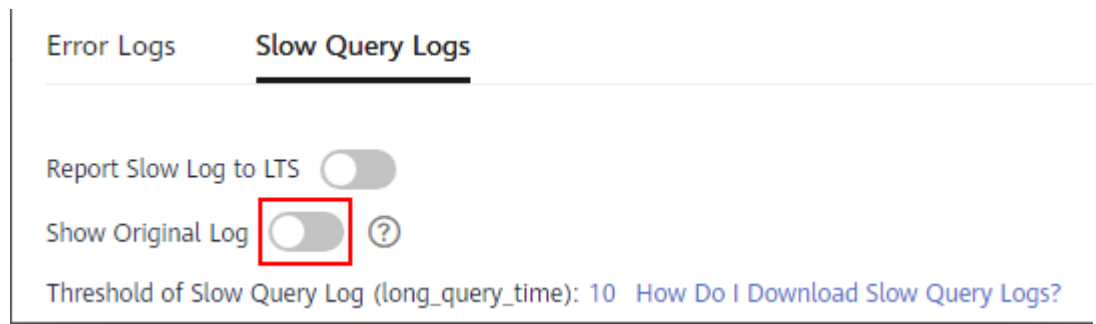
**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Logs**. On the **Slow Query Logs** tab, click  next to **Show Original Log**.


**Figure 20-8** Enabling Show Original Log




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

## Reporting Slow Logs to LTS

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

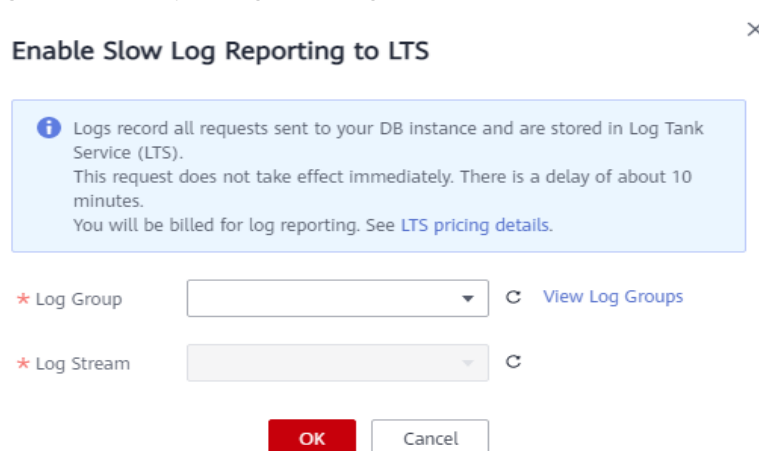
**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane on the left, choose **Logs**.

**Step 6** On the **Slow Query Logs** page, click  next to **Report Slow Log to LTS**.

**Step 7** Select an LTS log group and log stream and click **OK**.

**Figure 20-9** Reporting slow logs to LTS



----End

## APIs

### Querying Slow Query Logs

## 20.4 Enabling or Disabling SQL Explorer

Enabling SQL Explorer will allow GaussDB(for MySQL) to store all SQL statement logs for analysis.

You can enable SQL Explorer on the [DAS](#) console.

### Constraints

SQL Explorer cannot record all data. It has the following constraints:

- Some data cannot be recorded if a buffer overrun occurs.
- If the size of a SQL statement exceeds the value of **rds\_sql\_tracer\_max\_record\_size**, the statement is not recorded by default.

#### NOTE

**rds\_sql\_tracer\_max\_record\_size** controls the maximum size of a SQL statement. To change its value, see [Modifying Parameters of a GaussDB\(for MySQL\) DB Instance](#).



# 21 DBA Assistant

---

## 21.1 Function Overview

DBA Assistant provides you with a range of database O&M functions, making it easy to diagnose database problems, locate faults, analyze and optimize database performance.

DBA Assistant consists of the following modules:

### Dashboard

**Dashboard** shows the status of your instance, including alarms, resource usages, and key performance metrics. DBA Assistant diagnoses instance health using operational data analytics and intelligent algorithms, and provides you with solutions and suggestions for handling detected exceptions. For details, see [Dashboard](#).

### Sessions

On the **Sessions** page, you can view current session statistics of your instance, identify abnormal sessions, and kill the sessions. For details, see [Sessions](#).

### Performance

The **Performance** page displays key metrics of your instance and provides metric comparison between different days. You can keep track of metric changes and detect exceptions in a timely manner. Monitoring by Seconds helps accurately locate faults. For details, see [Performance](#).

### Slow Query Log



The **Slow Query Log** page displays slow queries within a specified time period. Slow query logs are collected by user, IP address, SQL template, and other keywords, sort statistics, and identify sources of slow SQL statements. For details, see [Slow SQL Analysis](#).

## SQL Explorer

After **Collect All SQL Statements** is enabled, you can gain a comprehensive insight into SQL statements on the **SQL Explorer** page. For details, see [SQL Insights](#) and [Concurrency Control](#).

## 21.2 Dashboard

### Alarms


- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the DB instance name.
- Step 5** In the navigation pane, choose **DBA Assistant > Real-Time Diagnosis**.
- Step 6** On the **Dashboard** page, view instance alarms provided by Cloud Eye.

You can customize alarm rules by adjusting alarm policies and severities for key metrics, such as CPU usage and disk usage. To view alarm details, click the number next to an alarm severity.

**Figure 21-1** Alarm statistics



**Figure 21-2** Alarm list




----End

## Health

In the **Health** area, you can view real-time health diagnosis results of your instance. By default, the diagnosis results of high vCPU utilization, memory bottleneck, high-frequency slow SQL, and lock wait are displayed.

For abnormal metrics, click **Diagnose** to view diagnosis details and suggestions. For details, see [Table 21-1](#).

**Figure 21-3** Health diagnosis



**Table 21-1** Health diagnosis and suggestions

Health Item	Exception Trigger Condition
High vCPU utilization	<p>Either of the following conditions is met:</p> <ul style="list-style-type: none"> <li>After you configure alarm rules on Cloud Eye, an alarm is reported, indicating the CPU usage is high.</li> <li>The CPU usage exceeds 95% for more than 2.5 minutes within 5 minutes.</li> </ul>
Memory bottleneck	<p>Either of the following conditions is met:</p> <ul style="list-style-type: none"> <li>After you configure alarm rules on Cloud Eye, an alarm is reported, indicating the memory usage is high.</li> <li>The memory usage exceeds 95% within 5 minutes.</li> </ul>
High-frequency slow SQL	<p>Either of the following conditions is met:</p> <ul style="list-style-type: none"> <li>After you configure alarm rules on Cloud Eye, an alarm is reported, indicating there are too many slow logs.</li> <li>There are more than 100 slow logs for 5 consecutive minutes.</li> </ul>
Lock wait	<p>After you configure alarm rules on Cloud Eye, any of the following alarms is reported:</p> <ul style="list-style-type: none"> <li>Too long row lock time</li> <li>Too many InnoDB row locks</li> <li>Too many row lock waits</li> </ul>

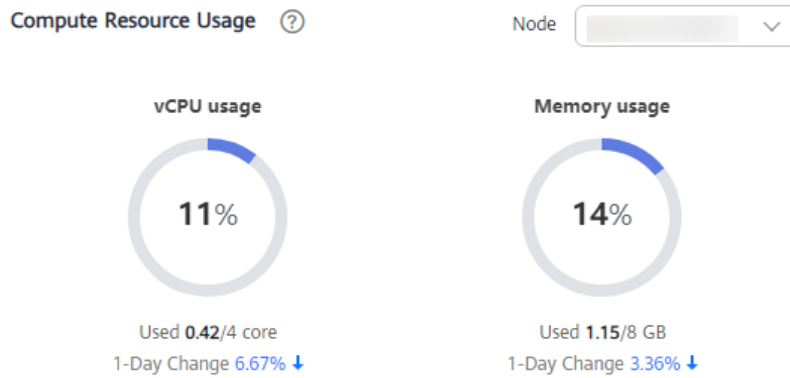
**NOTE**

- For details about how to configure alarm rules, see [Creating Alarm Rules for a DB Instance](#).
- For details about monitoring metrics, see [Introducing GaussDB\(for MySQL\) Metrics](#).

## Compute Resource Usage

In the **Compute Resource Usage** area, the vCPU usage and memory usage are displayed by default. The displayed values are the average values within 5 minutes.


Figure 21-4 Compute resource usage



## Storage Resource Usage

In the **Storage Resource Usage** area, the storage usage, disk read IOPS, and disk write IOPS are displayed by default. The displayed values are the average values within 5 minutes.


Figure 21-5 Storage resource usage



## Key Performance Metrics

In the **Key Performance Metrics** area, the CPU usage & slow query logs, connections, memory utilization, and disk reads/writes in the last hour are displayed by default. The displayed values are real-time values.

Figure 21-6 Key performance metrics



## 21.3 Sessions

On the **Sessions** page, you can view current session statistics of your instance, identify abnormal sessions, and kill the sessions. For details, see [Sessions](#).

## 21.4 Performance

### Creating Alarm Rules

Click **Create Alarm Rule**: You can create alarm rules for GaussDB(for MySQL) to customize the monitored objects and notification policies and stay aware of the GaussDB(for MySQL) instance statuses.


### Enabling Monitoring by Seconds


Click **Enable Monitoring by Seconds**: You can configure linked charts and enable Monitoring by Seconds.

- **Linked Charts**: Enabling it means that you can view all metrics at the same time.
- **Monitoring by Seconds**: Enabling it means that the system reports monitoring data precise to the second or to 5 second intervals. This function is billed and the pricing is listed on a per-hour basis. For details, see [Configuring Monitoring by Seconds](#).

### Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

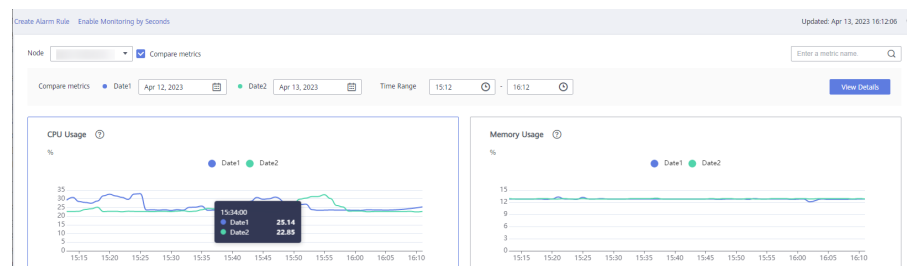
**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the DB instance name.

**Step 5** In the navigation pane, choose **DBA Assistant > Real-Time Diagnosis**.

**Step 6** Click the **Performance** tab to view your instance metric trends within the same time range on different days.

**Figure 21-7** Viewing performance metrics



----End


## 21.5 Slow SQL Analysis


### Scenarios

Slow SQL Analysis displays a chart of SQL statements that are taking too long to execute and allows you to sort slow SQL statements by multiple dimensions, such as by user, host, or SQL template. It helps you quickly identify bottlenecks and improve instance performance.

### Viewing Slow Query Logs

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the DB instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **DBA Assistant > Historical Diagnosis**.

**Step 6** Click the **Slow Query Log** tab.

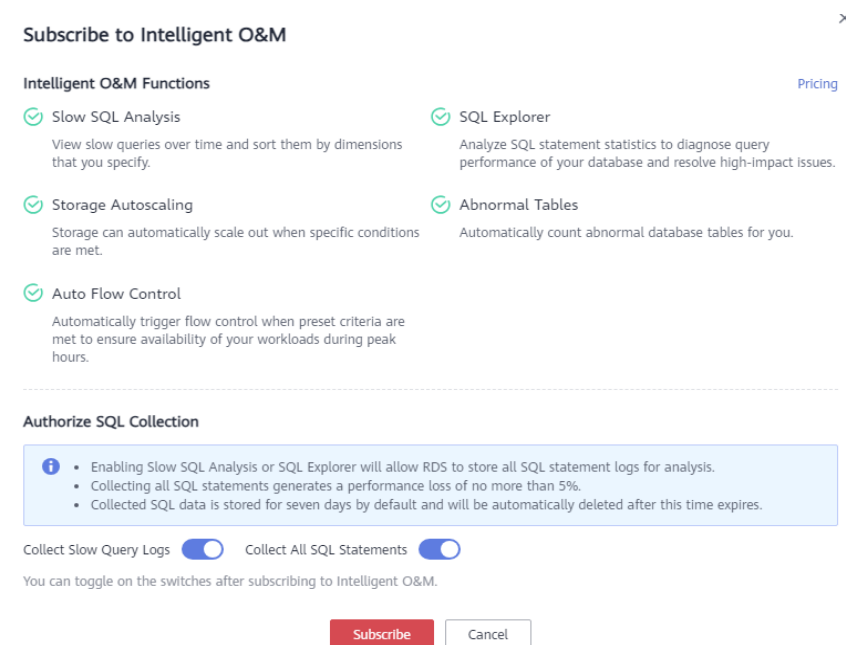
#### NOTE

Slow SQL Analysis needs to be purchased separately. To use this function, subscribe to Intelligent O&M first.

Only the data of the last hour is displayed if Intelligent O&M is not subscribed. The data will be automatically deleted after one hour. After Intelligent O&M is subscribed, data can be stored for up to 30 days. For details, see [Slow Query Log Storage](#).

**Step 7** Click **Subscribe**. In the displayed dialog box, you can learn about Intelligent O&M functions and pricing.


**Figure 21-8** Subscribe to Intelligent O&M




**Step 8** After subscribing to Intelligent O&M, view slow queries over time of your instance.

**Step 9** View slow queries and slow log archive history for the last 1 hour, last 3 hours, last 12 hours, or a custom time period (spanning no more than one day).


In the **Top 5 Slow Query Logs** area, logs can be displayed by the usernames and client IP addresses.




**Step 10** In the **Details** area, click  next to **Collect Slow Query Logs** to view slow query log details and template statistics. If you want to disable it, click **Manage Log** in the upper right corner and toggle off the switch.

**Figure 21-9** Viewing slow query log details (**Collect Slow Query Logs** disabled)



**Figure 21-10** Viewing slow query log details (**Collect Slow Query Logs** enabled)


- In the **Details** area:
  - Filter slow query logs by database, client IP address, or user.
  - Locate a SQL statement and click **Concurrency Control** in the **Operation** column. For details, see [Concurrency Control](#).
  - Click . In the displayed dialog box, select an OBS bucket and click **OK** to export slow query logs to the OBS bucket. Up to 100,000 logs can be exported.

If no OBS bucket is available, click **Create**. In the displayed dialog box, enter an OBS bucket name, and click **OK**.

Creating an OBS bucket is free, but you will be billed for storing data in the bucket. For pricing details, see [OBS Pricing Details](#).

A bucket name:

- Cannot be the same as that of any existing bucket.
  - Can contain 3 to 63 characters. Only lowercase letters, numbers, hyphens (-), and periods (.) are allowed.
  - Cannot start or end with a period (.) or hyphen (-), and cannot contain two consecutive periods (.) or contain a period (.) and a hyphen (-) adjacent to each other.
  - Cannot be an IP address.
  - If the bucket name contains a period (.), certificate-based verification is required when you use the name to access an OBS bucket or object.
- After the log details are exported, you can click **View Export List** to view export records. You can also download the details to your local PC for analysis.

- In the **Statistics** area:
  - Click **View Sample** in the **Operation** to view the sample of the SQL template.
  - Click . In the displayed dialog box, select an OBS bucket and click **OK** to export slow query logs to the OBS bucket. Up to 100,000 logs can be exported.

If no OBS bucket is available, click **Create**. In the displayed dialog box, enter an OBS bucket name, and click **OK**.

Creating an OBS bucket is free, but you will be billed for storing data in the bucket. For pricing details, see [OBS Pricing Details](#).

A bucket name:



- Cannot be the same as that of any existing bucket.
  - Can contain 3 to 63 characters. Only lowercase letters, numbers, hyphens (-), and periods (.) are allowed.
  - Cannot start or end with a period (.) or hyphen (-), and cannot contain two consecutive periods (.) or contain a period (.) and a hyphen (-) adjacent to each other.
  - Cannot be an IP address.
  - If the bucket name contains a period (.), certificate-based verification is required when you use the name to access an OBS bucket or object.
- After the templates are exported, you can click **Export Slow Query Logs** to view export records. You can also download the details to your local PC for analysis.

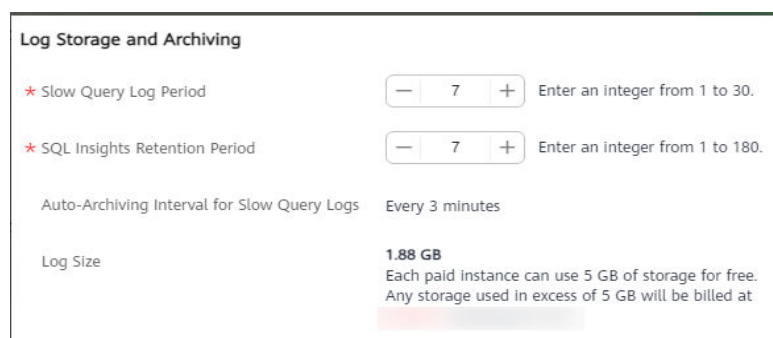
----End

## Slow Query Log Storage

After **Collect Slow Query Logs** is enabled, SQL text content will be stored in OBS for analysis.

- If you have subscribe intelligent O&M, click **Log Settings** in the upper right corner.
  - **Slow Query Log Period:** The default value is 7. The value ranges from 1 to 30. After the period expires, the logs are automatically deleted.
  - **SQL Insights Retention Period:** The default value is 7. The value ranges from 1 to 180.
  - **Log Size:** Each paid instance can use 5 GB of storage for slow query logs for free. Any storage used in excess of 5 GB will be billed on a pay-per-use basis.

**Figure 21-11** Log storage and archiving (Intelligent O&M subscribed)



- If you do not subscribe intelligent O&M, click **Log Settings** in the upper right corner.
  - **Slow Query Log Period:** The default value is 1 hour and cannot be changed. After the period expires, the logs are automatically deleted.

- **SQL Insights Retention Period:** 1 hour

**Figure 21-12** Log storage and archiving (Intelligent O&M not subscribed)



Log Storage and Archiving	
Slow Query Log Period	1 hour
SQL Insights Retention Period	1 hour
Auto-Archiving Interval for Slow Query Logs	Every 3 minutes

## 21.6 Top SQL

### Scenarios

After **Collect All SQL Statements** is enabled, you can gain a comprehensive insight into SQL statements on the **SQL Explorer** page. Top SQL helps you locate exceptions.

### Procedure

- Step 1** [Log in to the management console.](#)
- Step 2** Click  in the upper left corner and select a region and a project.
- Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.
- Step 4** On the **Instances** page, click the DB instance name to go to the **Basic Information** page.
- Step 5** In the navigation pane, choose **DBA Assistant > Historical Diagnosis**.
- Step 6** Choose **SQL Explorer > Top SQL**.

#### NOTE

Top SQL needs to be purchased separately. To use this function, subscribe to Intelligent O&M first.

Only the data of the last hour is displayed if Intelligent O&M is not subscribed. The data will be automatically deleted after one hour. If Intelligent O&M is subscribed, you can configure how long that top SQL statements are stored for (spanning no more than one day).

- Step 7** Click **Subscribe**. In the displayed dialog box, you can learn about Intelligent O&M functions and pricing.

×

### Subscribe to Intelligent O&M

**Intelligent O&M Functions** [Pricing](#)

<p><input checked="" type="checkbox"/> <b>Slow SQL Analysis</b> View slow queries over time and sort them by dimensions that you specify.</p> <p><input checked="" type="checkbox"/> <b>Storage Autoscaling</b> Storage can automatically scale out when specific conditions are met.</p> <p><input checked="" type="checkbox"/> <b>Auto Flow Control</b> Automatically trigger flow control when preset criteria are met to ensure availability of your workloads during peak hours.</p>	<p><input checked="" type="checkbox"/> <b>SQL Explorer</b> Analyze SQL statement statistics to diagnose query performance of your database and resolve high-impact issues.</p> <p><input checked="" type="checkbox"/> <b>Abnormal Tables</b> Automatically count abnormal database tables for you.</p>
---	--

---

**Authorize SQL Collection**

**i**

- Enabling Slow SQL Analysis or SQL Explorer will allow RDS to store all SQL statement logs for analysis.
- Collecting all SQL statements generates a performance loss of no more than 5%.
- Collected SQL data is stored for seven days by default and will be automatically deleted after this time expires.


Collect Slow Query Logs     Collect All SQL Statements

You can toggle on the switches after subscribing to Intelligent O&M.

SubscribeCancel

**Step 8** After subscribing to Intelligent O&M, view top SQL statements of your instance.

**Step 9** View execution durations of the top SQL statements in the last 1 hour, last 3 hours, last 6 hours, or a custom time period (spanning no more than one day).

- Click a point in time or drag to select a time period to view the SQL statistics of a SQL template.
- Click  to export information about all top SQL templates in the list. To use this export function, subscribe to Intelligent O&M.
- Locate a SQL template and click **Details** to view the total execution times, average rows scanned, average execution duration, and the like.
- Locate a SQL template and click **Concurrency Control** in the **Operation** column. For details, see [Concurrency Control](#).

----End

## 21.7 SQL Insights

### Scenarios

The SQL Insights function allows you to not only query all executed SQL statements, but also analyze and search for the tables that are accessed and updated most frequently, and the SQL statements that have the longest lock wait, helping you quickly identify exceptions.

### Constraints


- You need to enable **Collect All SQL Statements** before using SQL Insights.
- After **Collect All SQL Statements** is disabled, new SQL statements will not be collected anymore and the collected SQL data will be deleted.

- Some data cannot be recorded if a buffer overrun occurs.
- If the length of a SQL statement exceeds the value of **rds\_sql\_tracer\_max\_record\_size**, the statement is not recorded by default. To configure the parameter value, see [Modifying Parameters of a GaussDB\(for MySQL\) DB Instance](#).

## Procedure

**Step 1** [Log in to the management console](#).

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the DB instance name to go to the **Basic Information** page.

**Step 5** In the navigation pane, choose **DBA Assistant > Historical Diagnosis**.

**Step 6** Click **SQL Explorer** and then **SQL Insights**.

**Step 7** Click  next to **Collect All SQL Statements**.

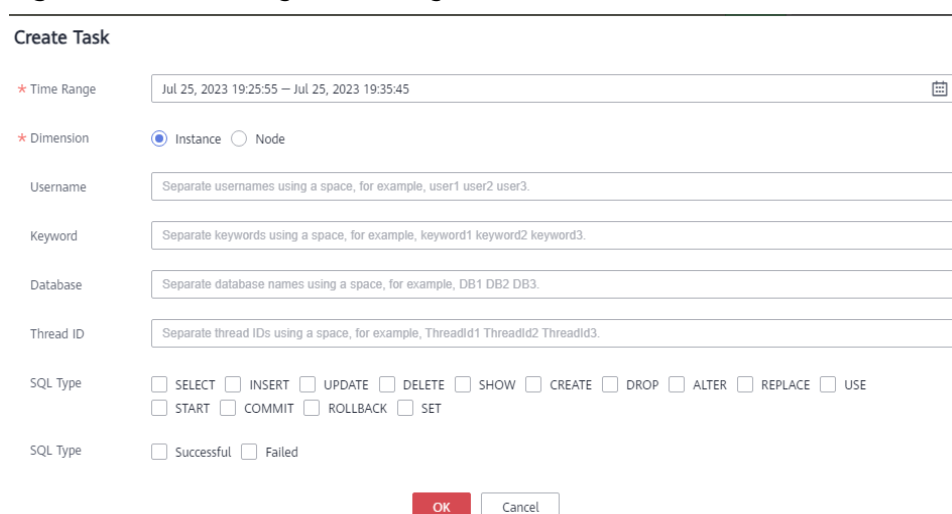
### NOTE

- Collecting all SQL statements generates a performance loss of no more than 5%.
- To disable this function, click **Log Settings** in the upper right corner, toggle off the **Collect All SQL Statements** switch, and click **OK**.

**Step 8** Click **Create Task** and specify **Time Range**, **Dimension**, **Username**, **Keyword**, **Database**, **Thread ID**, **SQL Type**, and **Execution Status**.

You can set **Dimension** to **Instance** or **Node**. When **Node** is selected, you can view the SQL logs of deleted nodes.

**Figure 21-13** Creating a SQL insights task



**Create Task**

\* Time Range: Jul 25, 2023 19:25:55 - Jul 25, 2023 19:35:45

\* Dimension:  Instance  Node

Username: Separate usernames using a space, for example, user1 user2 user3.

Keyword: Separate keywords using a space, for example, keyword1 keyword2 keyword3.

Database: Separate database names using a space, for example, DB1 DB2 DB3.

Thread ID: Separate thread IDs using a space, for example, ThreadId1 ThreadId2 ThreadId3.

SQL Type:  SELECT  INSERT  UPDATE  DELETE  SHOW  CREATE  DROP  ALTER  REPLACE  USE  
 START  COMMIT  ROLLBACK  SET

SQL Type:  Successful  Failed

**OK** Cancel

**Step 9** Click **OK**.

**Step 10** In the task list, click **Details** in the **Operation** column to view task details.

**Step 11** Specify conditions such as **Time Range**, **User**, **Keyword**, **Database** and click **Query**. The selected time range must be after the time when the new task is added.

----End

## 21.8 Concurrency Control

### Scenarios

SQL statement concurrency control aims to keep GaussDB(for MySQL) instances running stably even if there is a sudden increase in concurrent SQL statements.

### Constraints


- This function is available only to the GaussDB(for MySQL) instances that meet the following requirements:
  - 2.0.28.40 > kernel version ≥ 2.0.28.15
  - Kernel version ≥ 2.0.29.1
- Each SQL concurrency control rule can contain up to 128 keywords.
- The keywords in a rule cannot contain \t, \r, and \n, and cannot be a backslash (\) or a single null character ("").
- Spaces at the start, end of or in the middle of a keyword are ignored.
- The SQL concurrency control rule cannot end with a tilde (~).
- Keywords in a concurrency control rule are sorted in a specific order, and the system will match them from first to last. For example, if one rule contains the keyword **a~and~b**, the system only matches **xxx a>1 and b>2**.
- Each SQL concurrency control rule applies to only the SQL statements that your database received after the rule is created.
- If different rules are created for the primary node and read replicas of a DB instance, the rules still apply to the primary node and read replicas after their roles are switched over.
- If a SQL statement matches multiple concurrency control rules, only the most recently created rule is applied.
- SQL statements that have been executed before a concurrency control rule is added are not counted.
- The total length of all rules for SELECT, UPDATE, or DELETE statements and the **Concurrency** value in each rule cannot exceed 1024 bytes.
- If you add too many SQL concurrency control rules for your instance, the execution of SELECT, UPDATE, or DELETE statements will slow down.
- SQL concurrency control rules are applied based on prefix match. For example, if the concurrency control rule is **SELECT~COUNT~t1**, SQL statements **SELECT COUNT(\*) FROM t1** and **SELECT COUNT(\*) FROM t1 LIMIT 1** will both be intercepted.

- After concurrency control is triggered, an execution error is reported on the service side, indicating that query execution was interrupted. The error code is ERROR 1317 (70100).
- This function controls how many statements can run at the same time. However, it does not limit concurrency for:
  - system catalog
  - Queries where no database data is involved, such as **select sleep(xxx)**
  - Account **root**
  - SQL statements in stored procedures, triggers, and functions

## Procedure

**Step 1** Log in to the management console.

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the DB instance name to go to the **Basic Information** page.

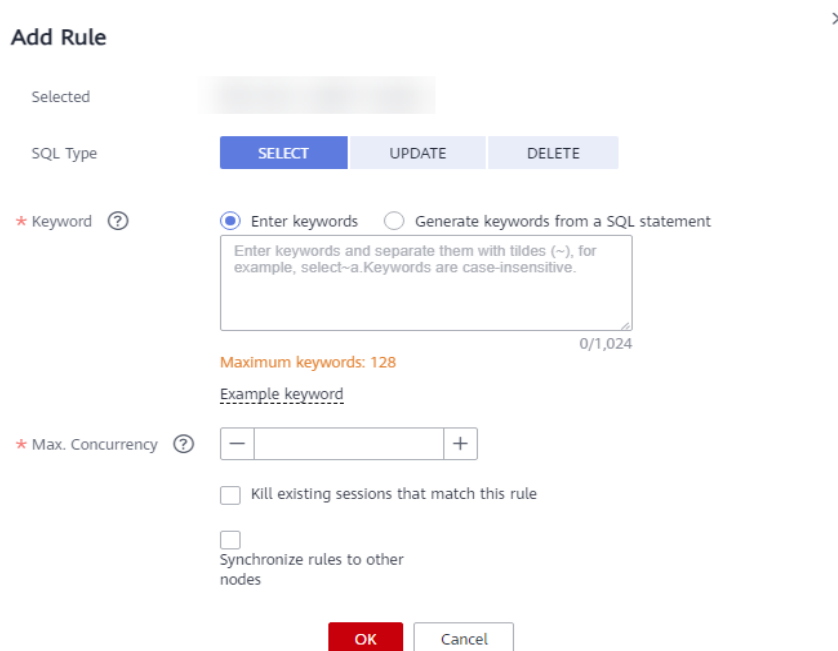
**Step 5** In the navigation pane, choose **DBA Assistant > Historical Diagnosis**.

**Step 6** Choose **SQL Explorer > Concurrency Control**.

**Step 7** On the displayed page, enable **Concurrency Control**.

**Step 8** Click **Add Rule**. In the displayed dialog box, specify **SQL Type**, **Keyword**, and **Max. Concurrency**.

**Figure 21-14** Adding a rule



**Add Rule** ×

Selected

SQL Type:  SELECT  UPDATE  DELETE

\* Keyword ?  Enter keywords  Generate keywords from a SQL statement

Enter keywords and separate them with tildes (~), for example, select~a. Keywords are case-insensitive.

0/1,024

Maximum keywords: 128

Example keyword

\* Max. Concurrency ?  -  +

Kill existing sessions that match this rule

Synchronize rules to other nodes

- **Keyword:** You can enter keywords or copy an existing SQL statement to the text box and click **Generate Keyword**.  
**Keyword:** Take **select~a** as an example. **select** and **a** are two keywords contained in a concurrency control rule. The keywords are separated by a tilde (~). In this example, the rule restricts the execution of only the SQL statements containing keywords **select** and **a**.
- **Max. Concurrency:** SQL statements that meet the specified SQL type and keyword and exceed the value of **Max. Concurrency** will not be executed.
- If you select **Kill existing sessions that meet this rule**, the sessions that meet the rule will be killed.
- If you select **Synchronize rules to other nodes**, the new rules can be synchronized to other nodes in the same instance.

**Step 9** Confirm the settings and click **OK**.

**Step 10** If a concurrency control rule is not required, select the rule and click **Delete** above the rule list. In the displayed dialog box, click **OK**.

----End


# 22 Task Center


## 22.1 Viewing a Task

You can view the progresses and results of instant and scheduled tasks on the **Task Center** page.

### Viewing an Instant Task

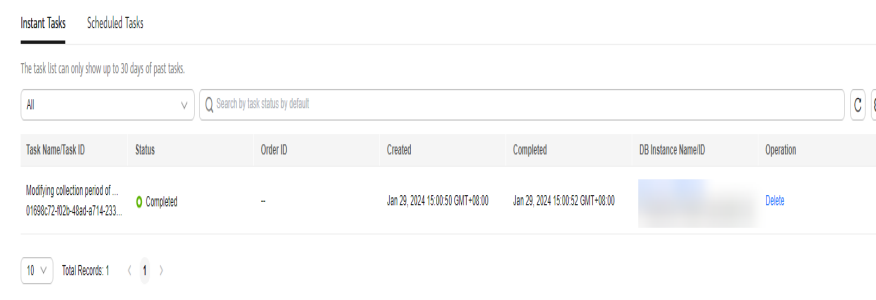
**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Task Center**. Locate the task and view its details on the displayed **Instant Tasks** page.

**Figure 22-1** Viewing an instant task



The screenshot shows the 'Instant Tasks' page. At the top, there are tabs for 'Instant Tasks' and 'Scheduled Tasks'. Below the tabs, there is a message: 'The task list can only show up to 30 days of past tasks.' There is a search bar with a dropdown menu set to 'All' and a search icon. Below the search bar is a table with the following columns: Task Name/Task ID, Status, Order ID, Created, Completed, DB Instance Name/ID, and Operation. The table contains one row with the following data: Task Name/Task ID: 'Modifying collection period of... 01689c72-902e-48ea-a714-233...', Status: 'Completed', Order ID: '-', Created: 'Jan 29, 2024 15:00:50 GMT+08:00', Completed: 'Jan 29, 2024 15:00:52 GMT+08:00', DB Instance Name/ID: (blurred), and Operation: 'Delete'. At the bottom of the table, there is a pagination bar showing '10' records per page and 'Total Records: 1'.

Task Name/Task ID	Status	Order ID	Created	Completed	DB Instance Name/ID	Operation
Modifying collection period of... 01689c72-902e-48ea-a714-233...	Completed	-	Jan 29, 2024 15:00:50 GMT+08:00	Jan 29, 2024 15:00:52 GMT+08:00		Delete

- Click the **All** drop-down list box in the upper part to view the task execution progress and status in a specified period. The default period is all time. The task list shows tasks that have been executed in the past 30 days.
- Click the filter box in the upper part to query the desired instant tasks by task name and task status.



- Task status: **Runing, Completed, and Failed**
- Task name:
  - Creating a GaussDB(for MySQL) instance
  - Creating a GaussDB(for MySQL) read replica
  - Rebooting a GaussDB(for MySQL) instance
  - Changing a GaussDB(for MySQL) instance port
  - Promoting a GaussDB(for MySQL) read replica to the primary node
  - Binding an EIP to a GaussDB(for MySQL) instance
  - Unbinding an EIP from a GaussDB(for MySQL) instance
  - Changing the instance name for a GaussDB(for MySQL) instance
  - Changing a security group for a GaussDB(for MySQL) instance
  - Deleting a GaussDB(for MySQL) instance
  - Deleting a GaussDB(for MySQL) read replica
  - Changing the specifications of a GaussDB(for MySQL) instance
  - Restoring to a new GaussDB(for MySQL) instance
  - Changing private IP address
  - Modifying collection period of Monitoring by Seconds
  - Adding database proxy nodes
  - Deleting database proxy nodes
  - Enabling database proxy
  - Disabling database proxy
  - Changing IP address of a proxy instance
  - Changing proxy instance specifications
  - Enabling or disabling SSL
  - Changing consistency level of a proxy instance
  - Changing read weights of nodes
  - Restoring to an existing DB instance
  - Restoring tables to a point in time
  - Creating a database


- Deleting a database
- Creating a database account
- Deleting a database account
- Changing the password of a database user
- Changing the host IP address of a database user
- Authorizing database user permissions
- Deleting database user permissions
- Rebooting a node
- Changing read/write splitting address
- Changing a node name
- Increasing specifications of a serverless instance
- Decreasing specifications of a serverless instance
- Changing the port of a proxy instance
- Applying for a private domain name for a proxy instance
- Changing the private domain name of a proxy instance
- Deleting the private domain name of a proxy instance
- Changing the routing policy of a proxy instance
- Enabling or disabling SSL for a proxy instance
- Applying for a private domain name for the DB instance
- Changing the private domain name of the DB instance

----End


## Viewing a Scheduled Task

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Task Center**. On the **Scheduled Tasks** page, view the task progress and results.

**Figure 22-2** Viewing a scheduled task

- You can enter the instance ID or task status in the search box to determine the desired task and view the task creation time and execution time.  
Task status: **Runing, Completed, Failed, Canceled, To be executed, and To be authorized.**
- Click the **All** drop-down list box in the upper part to view the task execution progress and status in a specified period. The default period is all time.

----End

## APIs

- [Obtaining Information About a Task with a Specified ID](#)
- [Obtaining Instant Tasks](#)
- [Obtaining Scheduled Tasks](#)

## 22.2 Deleting a Task Record

You can delete the task records no longer need to be displayed. The deletion only deletes the task records. It does not delete the instances or terminate tasks in progress.


### NOTICE

Deleted task records cannot be recovered. Exercise caution when performing this operation.

### Deleting an Instant Task Record

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases > GaussDB(for MySQL)**.

**Step 4** In the navigation pane on the left, choose **Task Center**. On the displayed **Instant Tasks** page, locate the task record to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.


You can delete instant task records with the following statuses:


- Completed
- Failed

----End

## Deleting a Scheduled Task Record

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** Choose **Task Center** in the navigation pane on the left. On the **Scheduled Tasks** page, locate the task record to be deleted and check whether the task record status is **To be executed** or **To be authorized**.

- If yes, go to [Step 5](#).
- If no, go to [Step 6](#).

**Step 5** Click **Cancel** in the **Operation** column. In the displayed dialog box, click **Yes** to cancel the task record. Then, click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes** to delete the task record.

**Step 6** Click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes** to delete the task record.

You can delete scheduled task records with the following statuses:

- Completed
- Failed
- Canceled
- To be authorized

----End

## APIs

- [Canceling a Scheduled Task](#)
- [Deleting a Task Record](#)

# 23 Managing Tags

---


## Scenarios


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

- You are advised to configure predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key.
- Each instance can have up to 20 tags.

## Adding a Tag

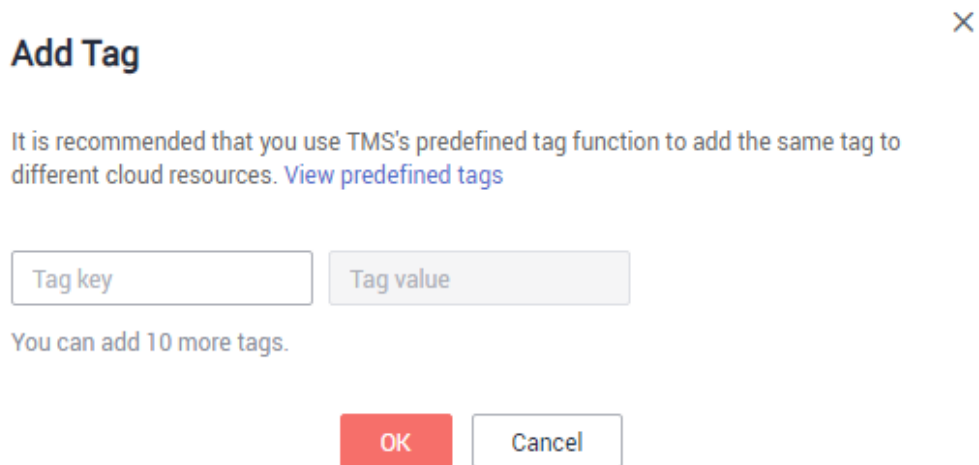
**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name to go to the **Basic Information** page.

**Step 5** On the **Tags** page, click **Add Tag**. In the displayed dialog box, enter a tag key and a tag value, and click **OK**.

**Figure 23-1** Adding a tag

**Add Tag** ×

It is recommended that you use TMS's predefined tag function to add the same tag to different cloud resources. [View predefined tags](#)

Tag key  Tag value

You can add 10 more tags.

- When you enter a tag key and value, the system automatically displays all tags (including predefined tags and resource tags) associated with all instances except the current one.
- The tag key must be unique and must consist of 1 to 36 characters. Only letters, digits, hyphens (-), and underscores (\_) are allowed.
- The tag value can be empty or consist of 1 to 43 characters. Only letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed.


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

----End

## Editing a Tag

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name.

**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.
- The tag value can be empty or consist of 1 to 43 characters. Only letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed.


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

----End

## Deleting a Tag

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Click  in the upper left corner of the page, choose **Databases** > **GaussDB(for MySQL)**.

**Step 4** On the **Instances** page, click the instance name.

**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** View that the tag is no longer displayed on the **Tags** page.

----End

## APIs

- [Querying Resource Tags](#)
- [Querying Project Tags](#)
- [Adding or Deleting Tags in Batches](#)

# 24 Managing Quotas


## Scenarios

A quota is a limit on the quantity or capacity of a certain type of service resources available to you. Examples of GaussDB(for MySQL) quotas include the maximum number of DB instances that you can create. Quotas are put in place to prevent excessive resource usage.

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

## Viewing Quotas


**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.

**Step 3** Choose **Resources > My Quotas** in the upper right corner of the page.

The **Quota** page is displayed.

**Figure 24-1** Viewing quotas




**Step 4** View the used and total quotas of each type of resources.

----End



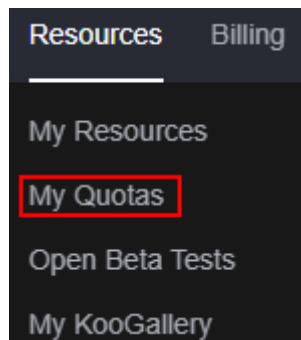
## Increasing Quotas

**Step 1** [Log in to the management console.](#)

**Step 2** Click  in the upper left corner and select a region and a project.


**Step 3** In the upper right corner of the console page, choose **Resources > My Quotas**.

**Figure 24-2** Viewing quotas



**Step 4** In the upper right corner of the page, click **Increase Quota**.

**Figure 24-3** Increasing quotas



Service	Resource Type	Used Quota	Total Quota
Elastic Cloud Server	ECS	1	20
	VCPU	1	88
	Memory (MB)	1,024	1,024.0

**Step 5** On the **Create Service Ticket** page, configure parameters as required.

In the **Problem Description** area, fill in the content and reason for quota adjustment.

**Step 6** After all necessary parameters are configured, select the agreement and click **Submit**.

----End

## APIs

- [Querying the Instance Quotas of a Tenant](#)
- [Querying the Resource Quotas of a Specified Enterprise Project](#)
- [Configuring Resource Quotas for a Specified Enterprise Project](#)
- [Modifying the Resource Quotas of a Specified Enterprise Project](#)

# 25 Kernel Functions


## 25.1 Parallel Query

### 25.1.1 Overview

Parallel query (PQ) reduces the processing time of analytical queries to satisfy the low latency requirements of enterprise-class applications. It distributes a query task to multiple CPU cores for computation to shorten the query time. Theoretically, the performance improvement of parallel query is positively correlated with the number of CPU cores. The more CPU cores are used, the higher the performance improvement is.

The following figure shows the count(\*) process for a table based on parallel query. Table data is divided into blocks and distributed to multiple cores for parallel computing. Each core processes some data to obtain an intermediate count(\*) result, and all the intermediate results are aggregated to obtain the final result.

Figure 25-1 Parallel query



## Prerequisites

The engine version of GaussDB(for MySQL) is MySQL 8.0.22 or later.

## Scenarios

Parallel query is mainly suitable for SELECT statements to query large tables, multiple tables, and a large amount of data. This feature does not benefit extremely short queries.

- Lightweight analysis  
The SQL statements for report queries are complex and time-consuming. Parallel query can improve the efficiency of a single query.
- More available system resources  
Parallel query requires more system resources. You can enable parallel query to improve resource utilization and query efficiency only when the system has a large number of CPUs, low I/O loads, and sufficient memory resources.
- Frequent data queries  
For data-intensive queries, you can use parallel query to improve query processing efficiency, ease network traffic, and reduce pressure on compute nodes.

### NOTE

Both read replicas and primary nodes support parallel query. Parallel query consumes a lot of computing resources (such as CPU and memory). To ensure instance stability, parallel query does not take effect on primary nodes of instance 2.0.42.230600 and later versions by default. To use parallel query, contact customer service.

## 25.1.2 Usage

### 25.1.2.1 Application Scenarios

Parallel query is suitable for the following scenarios:

- Full table scans, index scans, index range scans, index reverse scans, index point queries, and index pushdown.
- Single-table queries, multi-table joins, views, subqueries, and partial CTE queries.
- Multiple JOIN algorithms, including BNL Join, BKA Join, HASH Join, Nested Loop Join, Semi Join, Anti Join, and Outer join.
- Multiple subqueries, including conditional subqueries, scalar subqueries, some correlated subqueries, non-correlated subqueries, and derived tables.
- Multiple data types, including Integer, Character, Floating Point, and Time.
- Arithmetic expressions (+, -, \*, %, /, |, and &), conditional expressions (<, <=, >, >=, <>, Between/And, In), logical operations (Or, And, Not), and common functions (Character, Integer, and Time), and aggregation functions (COUNT/SUM/AVG/MIN/MAX)

 **NOTE**

The count aggregate function can be executed concurrently only when **innodb\_parallel\_select\_count** is disabled.

- Non-partitioned table queries, and queries for a single partition of partitioned tables
- Order By, Group By/Distinct, Limit/Offset, Where/Having, and Column Projection
- UNION/UNION ALL queries
- EXPLAIN statements to view execution plans. Traditional Explain statements, Explain format=tree, and Explain format=json are included.

### 25.1.2.2 Constraints

Currently, there are some restrictions on parallel queries, but they will be gradually resolved in the future.

### Unsupported Statements

Parallel queries are not suitable for:

- Non-SELECT statements
- Window functions
- Triggers
- Prepared Statements
- Spatial indexes
- System tables, temporary tables, and non-InnoDB tables
- Full-text indexes
- Stored procedures
- Subqueries that cannot be converted to semi-joins
- SQL statements that do not meet the **only\_full\_group\_by** rules
- Index Merge statements
- HASH JOIN operations, during which data overflows to disks
- Lock queries, such as SERIALIZABLE isolation level, FOR UPDATE, or SHARE LOCK
- Recursive queries
- With rollup
- Statements with keyword HIGH\_PRIORITY
- No line of data returned in the execution result. (The execution plan shows: Zero limit, Impossible WHERE, Impossible HAVING, No matching min/max row, Select tables optimized away, Impossible HAVING noticed after reading const tables, or no matching row in const table)
- Columns with type ZEROFILL. Its column values can be optimized to constants.
- Generated columns, BLOB, TEXT, JSON, and GEOMETRY
- Spatial functions (such as SP\_WITHIN\_FUNC)

- DISTINCT clauses in Aggregate functions, such as SUM(DISTINCT), AVG(DISTINCT), and COUNT(DISTINCT)
- GROUP\_CONCAT
- JSON\_ARRAYAGG/JSON\_OBJECTAGG
- User-defined functions
- STD/STDDEV/STDDEV\_POP
- VARIANCE/VAR\_POP/VAR\_SAMP
- BIT\_AND, BIT\_OR and BIT\_XOR
- set\_user\_var
- RAND functions with parameters
- json\_\* (such as json\_length and json\_type)
- st\_distance
- get\_lock
- is\_free\_lock, is\_used\_lock, release\_lock, and release\_all\_locks
- sleep
- xml\_str
- weight\_string
- REF functions (VIEW\_REF, OUTER\_REF, AGGREGATE\_REF)
- SHA, SHA1, SHA2, and MD5
- row\_count
- User-related functions (such as user, current\_user, session\_user, system\_user)
- extractvalue
- GeomCollection, GeometryCollection, LineString, MultiLineString, MultiPoint, MultiPolygon, Polygon
- MASTER\_POS\_WAIT
- Spatial relationship functions, such as MBRContains, MBRCoveredBy, MBR Covers, MBRDisjoint, MBREquals, MBRIntersects, MBROverlaps, MBRTouches, MBRWithin
- Point
- PS\_CURRENT\_THREAD\_ID()
- PS\_THREAD\_ID(CONNECTION\_ID())
- WAIT\_FOR\_EXECUTED\_GTID\_SET
- WAIT\_UNTIL\_SQL\_THREAD\_AFTER\_GTIDS
- UNCOMPRESS (COMPRESS ( ))
- STATEMENT\_DIGEST\_TEXT
- Functions BINARY and CONVERT
- Functions starting with ST\_

## Incompatible with the Serial Execution Results

The execution results of parallel queries may be incompatible with that of serial queries.

- Number of errors or alarms

If an error or alarm message is displayed during serial queries, the error or alarm message will be displayed in each worker thread during the parallel queries. As a result, the total number of error or alarm messages increases.

```
mysql> SELECT dt1 = 99991231235959.999999 AS a, dt2 = 99991231235959.999999 AS b FROM t7;
+-----+
| a | b |
+-----+
| 0 | 0 |
| 0 | 0 |
| 0 | 1 |
+-----+
3 rows in set, 2 warnings (0.00 sec)

mysql> show warnings;
+-----+
| Level | Code | Message
+-----+
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
+-----+
2 rows in set (0.00 sec)

mysql> set force_parallel_execute=ON;
Query OK, 0 rows affected (0.00 sec)

mysql> SELECT dt1 = 99991231235959.999999 AS a, dt2 = 99991231235959.999999 AS b FROM t7;
+-----+
| a | b |
+-----+
| 0 | 0 |
| 0 | 0 |
| 0 | 1 |
+-----+
3 rows in set, 12 warnings (0.01 sec)

mysql> show warnings;
+-----+
| Level | Code | Message
+-----+
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
| Warning | 1441 | Datetime function: datetime field overflow
| Warning | 1292 | Incorrect datetime value: '99991231235959.999999' for column 'dt1' at row 1
+-----+
12 rows in set (0.00 sec)
```

- Precision

During the parallel queries, if there is a function type in a SELECT statement, additional stored procedures will be generated in the intermediate results. As a result, compared with serial queries, the precision of the floating point part in parallel queries may be different, and the final result may be slightly different.

```
mysql> create table tb(double_col double);
Query OK, 0 rows affected (0.08 sec)

mysql> insert into tb values (-1.7976931348623157e308),(-1.7976931348623157e308);
Query OK, 2 rows affected (0.02 sec)
Records: 2 Duplicates: 0 Warnings: 0

mysql> select sum(double_col) from tb;
+-----+
| sum(double_col) |
+-----+
| 0 |
+-----+
1 row in set (0.00 sec)

mysql> set force_parallel_execute=ON;
Query OK, 0 rows affected (0.00 sec)

mysql> select sum(double_col) from tb;
+-----+
| sum(double_col) |
+-----+
| -1.7976931348623157e308 |
+-----+
1 row in set (0.02 sec)
```

- Truncation

During the parallel queries, if there is a function type in a SELECT statement, additional stored procedures will be generated in the intermediate results. In this process, the calculation result of the function needs to be cached, and data truncation may occur (generally due to data type conversion, for example, covering a floating-point value to a character string). As a result, the final result is different from the serial queries.

- Sequence of result sets

Because tasks are executed by multiple worker threads during parallel queries, the sequence of the returned result set may not be consistent with that of serial queries. In the case of a query with LIMIT, this problem is more likely to occur. If fields of GROUP BY are invisible characters, the sequence of the returned result set is also different.

```
mysql> select a,count(*) from t group by a;
+-----+-----+
| a      | count(*) |
+-----+-----+
| 0      | 32768    |
| 1      | 32768    |
| 2      | 32768    |
| 3      | 32768    |
| 4      | 32768    |
| 5      | 32768    |
| 6      | 32768    |
| 7      | 32768    |
| 8      | 32768    |
| 9      | 32768    |
+-----+-----+
10 rows in set (6.37 sec)

mysql> set force_parallel_execute=ON;
Query OK, 0 rows affected (0.00 sec)

mysql> select a,count(*) from t group by a;
+-----+-----+
| a      | count(*) |
+-----+-----+
| 4      | 32768    |
| 5      | 32768    |
| 6      | 32768    |
| 7      | 32768    |
| 8      | 32768    |
| 9      | 32768    |
| 0      | 32768    |
| 1      | 32768    |
| 2      | 32768    |
| 3      | 32768    |
+-----+-----+
10 rows in set (4.35 sec)
```

- UNION ALL result sets

UNION ALL ignores sort operators. The sequence of the returned result set in parallel execution may be different from that in non-parallel execution. In the case of a query with LIMIT, the result sets are different.

### 25.1.2.3 Parameters and Variables

#### NOTE

Parallel query is in the open beta test (OBT) phase. You are advised to use it in the test environment.

### Supported Parameters and Variables

**Table 25-1** System Parameters

Parameter	Level	Description
force_parallel_execute	Global, Session	Whether to enable parallel query. If this parameter is set to <b>ON</b> , parallel query is enabled. Value: <b>ON</b> or <b>OFF</b> Default value: <b>OFF</b>
parallel_max_threads	Global	Maximum number of active threads allowed for parallel execution. If the number of active threads in the current system exceeds the value of this parameter, parallel execution cannot be enabled for new queries. Value range: <b>0-4294967295</b> Default value: <b>64</b>
parallel_default_degree	Global, Session	Default parallelism degree for parallel execution. If the parallelism degree is not specified in query statements, this parameter value is used. Value range: <b>0-1024</b> Default value: <b>4</b> .
parallel_cost_threshold	Global, Session	The cost threshold for enabling parallel execution. If the parallel execution cost of query statements exceeds the value of this parameter, the optimizer enables parallel execution. Value range: <b>0-4294967295</b> Default value: <b>1000</b>
parallel_queue_timeout	Global, Session	Waiting time of the parallel execution. If the waiting time exceeds the value of this parameter, new queries will be executed in single-thread mode. Value range: <b>0-4294967295</b> Default value: <b>0</b>



Parameter	Level	Description
parallel_memory_limit	Global	Maximum available memory for parallel execution. When the memory used for parallel execution exceeds the value of this parameter, new SQL queries will not be executed in parallel. Value range: <b>0-4294967295</b> Default value: <b>104857600</b>

**Table 25-2** Status variables

Variable	Level	Description
PQ_threads_running	Global	Total number of concurrent threads that are running.
PQ_memory_used	Global	Total memory used for parallel execution.
PQ_threads_refused	Global	Total number of queries that fail to be executed in parallel due to the limit on the total number of threads.
PQ_memory_refused	Global	Total number of queries that fail to be executed in parallel due to the limit on the total memory.

### 25.1.2.4 Enabling or Disabling Parallel Query

You can enable or disable parallel query by configuring system parameters in the console or using hints in SQL statements.

#### Method 1: Configuring system parameters in the console

Log in to the console and go to the **Parameters** page to change the following parameters:

**force\_parallel\_execute**: determines whether to forcibly enable parallel execution.

**parallel\_default\_dop**: indicates the parallelism degree for parallel execution. It controls the number of concurrent threads.

**parallel\_cost\_threshold**: indicates the cost threshold for enabling parallel execution.

These parameters can be modified at any time. The modifications will take effect immediately and you do not need to reboot the instance.

For example, if you want to forcibly enable parallel execution, set the parallelism degree to **4**, and set the minimum execution cost to **0**, configure the parameters as follows:

```
SET force_parallel_execute=ON
```

```
SET parallel_default_dop=4
```

```
SET parallel_cost_threshold=0
```

## Method 2: Using hints in SQL statements

Hints can be used to control whether a single statement is executed in parallel. Parallel execution is disabled by default, but you can use hints to enable parallel execution for specific SQL statements. You can also use hints to disable parallel execution of specified SQL statements.

### Enabling parallel execution:

Enabling parallel execution: **SELECT /\*+ PQ() \*/... FROM...**

Enabling parallel execution and setting the parallelism degree to **8**: **SELECT /\*+ PQ(8) \*/... FROM...**

Enabling parallel execution and set the parallel-executed table to **t1**: **SELECT /\*+ PQ(t1) \*/... FROM...**

Enabling parallel execution, set the parallel-executed table to **t1**, and set the parallelism degree to **8**: **SELECT /\*+ PQ(t1 8) \*/... FROM...**

#### NOTE

SELECT is followed by PQ (*Hints*). Otherwise, the hints do not take effect. **dop** indicates the parallelism degree of a parallel query and its value ranges from **1** to **min(parallel\_max\_threads, 1024)**.

When the **dop** value exceeds the normal range, parallel query does not take effect.

**Disabling parallel execution:** When parallel query is enabled, use the **NO\_PQ** to disable parallel execution of a single SQL statement.

```
SELECT /*+ NO_PQ */ ... FROM ...
```

#### NOTE

**NO\_PQ** (*Hints*) takes precedence over **PQ** (*Hints*). If a SQL statement contains **NO\_PQ** (*Hints*), the SQL statement will not be executed concurrently even if **PQ** (*Hints*) is configured.

## 25.1.2.5 Viewing the Statuses of Parallel Query

### Query Statement Statuses

Run **show status like "%PQ %"** to display the current statuses of query statements executed in parallel, as shown in [Figure 25-2](#).

Figure 25-2 Status

```
mysql> show status like "%PQ%";
+-----+-----+
| Variable_name | Value |
+-----+-----+
| PQ_memory_refused | 0 |
| PQ_memory_used | 0 |
| PQ_threads_refused | 0 |
| PQ_threads_running | 0 |
+-----+-----+
```

## Parallel Execution Plans

Use EXPLAIN to display the parallel execution plans of the query statements, as shown in [Figure 25-3](#).

Figure 25-3 Parallel execution plan

```
mysql> explain select /*+ PQ(4) */ n_name, sum(l_extendedprice * (1 - l_discount)) as revenue from customer, orders, lineitem, supplier, nation, region where c_custkey = o_custkey and l_orderkey = o_orderkey and l_suppkey = s_suppkey and c_nationkey = s_nationkey and s_nationkey = n_nationkey and n_regionkey = r_regionkey and r_name = 'MIDDLE EAST' and o_orderdate >= date '1994-01-01' and o_orderdate < date '1994-01-01' + interval '1' year group by n_name order by revenue desc;
```

id	select_type	table	partitions	type	possible_keys	key	key_len	ref	rows	filtered	Extra
1	SIMPLE	<gather>	NONE	ALL	NONE	NONE	NONE	NONE	10	100.00	Parallel execute (4 workers), tpch.supplier
2	SIMPLE	region	NONE	ALL	PRIMARY	NONE	NONE	NONE	5	20.00	Using where; Using temporary; Using filesort
2	SIMPLE	supplier	NONE	index	PRIMARY,SUPPLIER_FK1	SUPPLIER_FK1	4	NONE	10	100.00	Using index; Using join buffer (Block Nested Loop)
2	SIMPLE	nation	NONE	eq_ref	PRIMARY,NATION_FK1	PRIMARY	4	tpch.supplier_S_NATIONKEY	1	20.00	Using where
2	SIMPLE	customer	NONE	ref	PRIMARY,CUSTOMER_FK1	CUSTOMER_FK1	4	tpch.supplier_S_NATIONKEY	6	100.00	Using index
2	SIMPLE	orders	NONE	ref	PRIMARY,ORDERS_FK1	ORDERS_FK1	4	tpch.customer_C_CUSTKEY	15	11.11	Using where
2	SIMPLE	lineitem	NONE	ref	PRIMARY	PRIMARY	4	tpch.orders_O_ORDERKEY	4	10.00	Using where

rows in set, 1 warning (0.02 sec)

### NOTE

Compared with a traditional execution plan, a parallel execution plan has one more row of records. In the first row of the query result, the parallel-executed tables and parallelism degree are displayed.

## 25.1.3 Performance Tests

### 25.1.3.1 Test Method

Use the TPC-H test tool to test the performance of 22 parallel queries.

Test tool: TPC-H

Specifications: 32 vCPUs and 256 GB of memory

Kernel version: 2.0.26.1

Concurrent threads: 16

Data volume: 100 GB

### Procedure

#### Step 1 Generate test data.

1. Download the shared source code in the TPC-H test from <https://github.com/electrum/tpch-dbggen>.
2. Find the **makefile.suite** file, modify its contents as follows, and save the modifications:

```
CC = gcc
# Current values for DATABASE are: INFORMIX, DB2, TDAT (Teradata)
# SQLSERVER, SYBASE, ORACLE
# Current values for MACHINE are: ATT, DOS, HP, IBM, ICL, MVS,
# SGI, SUN, U2200, VMS, LINUX, WIN32
# Current values for WORKLOAD are: TPCH
DATABASE= SQLSERVER
MACHINE = LINUX
WORKLOAD = TPCH
```

3. In the root directory of the source code, run the following command to compile and generate the data tool dbgen:

```
make -f makefile.suite
```

4. Run the following command to generate 100 GB data.

```
./dbgen -s 100
```

- Step 2** Log in to the target GaussDB(for MySQL) instance, create a database, and run the following command to create a table:

```
CREATE TABLE nation ( N_NATIONKEY INTEGER NOT NULL,
                       N_NAME CHAR(25) NOT NULL,
                       N_REGIONKEY INTEGER NOT NULL,
                       N_COMMENT VARCHAR(152));
CREATE TABLE region ( R_REGIONKEY INTEGER NOT NULL,
                       R_NAME CHAR(25) NOT NULL,
                       R_COMMENT VARCHAR(152));
CREATE TABLE part ( P_PARTKEY INTEGER NOT NULL,
                    P_NAME VARCHAR(55) NOT NULL,
                    P_MFGR CHAR(25) NOT NULL,
                    P_BRAND CHAR(10) NOT NULL,
                    P_TYPE VARCHAR(25) NOT NULL,
                    P_SIZE INTEGER NOT NULL,
                    P_CONTAINER CHAR(10) NOT NULL,
                    P_RETAILPRICE DECIMAL(15,2) NOT NULL,
                    P_COMMENT VARCHAR(23) NOT NULL );
CREATE TABLE supplier ( S_SUPPKEY INTEGER NOT NULL,
                         S_NAME CHAR(25) NOT NULL,
                         S_ADDRESS VARCHAR(40) NOT NULL,
                         S_NATIONKEY INTEGER NOT NULL,
                         S_PHONE CHAR(15) NOT NULL,
                         S_ACCTBAL DECIMAL(15,2) NOT NULL,
                         S_COMMENT VARCHAR(101) NOT NULL);
CREATE TABLE partsupp ( PS_PARTKEY INTEGER NOT NULL,
                         PS_SUPPKEY INTEGER NOT NULL,
                         PS_AVAILQTY INTEGER NOT NULL,
                         PS_SUPPLYCOST DECIMAL(15,2) NOT NULL,
                         PS_COMMENT VARCHAR(199) NOT NULL );
CREATE TABLE customer ( C_CUSTKEY INTEGER NOT NULL,
                          C_NAME VARCHAR(25) NOT NULL,
                          C_ADDRESS VARCHAR(40) NOT NULL,
                          C_NATIONKEY INTEGER NOT NULL,
                          C_PHONE CHAR(15) NOT NULL,
                          C_ACCTBAL DECIMAL(15,2) NOT NULL,
                          C_MKTSEGMENT CHAR(10) NOT NULL,
                          C_COMMENT VARCHAR(117) NOT NULL);
CREATE TABLE orders ( O_ORDERKEY INTEGER NOT NULL,
                       O_CUSTKEY INTEGER NOT NULL,
                       O_ORDERSTATUS CHAR(1) NOT NULL,
                       O_TOTALPRICE DECIMAL(15,2) NOT NULL,
                       O_ORDERDATE DATE NOT NULL,
                       O_ORDERPRIORITY CHAR(15) NOT NULL,
                       O_CLERK CHAR(15) NOT NULL,
                       O_SHIPPRIORITY INTEGER NOT NULL,
                       O_COMMENT VARCHAR(79) NOT NULL);
CREATE TABLE lineitem ( L_ORDERKEY INTEGER NOT NULL,
                        L_PARTKEY INTEGER NOT NULL,
                        L_SUPPKEY INTEGER NOT NULL,
                        L_LINENUMBER INTEGER NOT NULL,
                        L_QUANTITY DECIMAL(15,2) NOT NULL,
                        L_EXTENDEDPRICE DECIMAL(15,2) NOT NULL,
                        L_DISCOUNT DECIMAL(15,2) NOT NULL,
                        L_TAX DECIMAL(15,2) NOT NULL,
                        L_RETURNFLAG CHAR(1) NOT NULL,
                        L_LINestatus CHAR(1) NOT NULL,
                        L_SHIPDATE DATE NOT NULL,
                        L_COMMITDATE DATE NOT NULL,
                        L_RECEIPTDATE DATE NOT NULL,
                        L_SHIPINSTRUCT CHAR(25) NOT NULL,
                        L_SHIPMODE CHAR(10) NOT NULL,
                        L_COMMENT VARCHAR(44) NOT NULL);
```

- Step 3** Run the following command to import the generated data to the table:

```
load data INFILE '/path/customer.tbl' INTO TABLE customer FIELDS TERMINATED BY '|';
load data INFILE '/path/region.tbl' INTO TABLE region FIELDS TERMINATED BY '|';
load data INFILE '/path/nation.tbl' INTO TABLE nation FIELDS TERMINATED BY '|';
load data INFILE '/path/supplier.tbl' INTO TABLE supplier FIELDS TERMINATED BY '|';
load data INFILE '/path/part.tbl' INTO TABLE part FIELDS TERMINATED BY '|';
load data INFILE '/path/partsupp.tbl' INTO TABLE partsupp FIELDS TERMINATED BY '|';
load data INFILE '/path/orders.tbl' INTO TABLE orders FIELDS TERMINATED BY '|';
load data INFILE '/path/lineitem.tbl' INTO TABLE lineitem FIELDS TERMINATED BY '|';
```

**Step 4** Create an index for the table.

```
alter table region add primary key (r_regionkey);
alter table nation add primary key (n_nationkey);
alter table part add primary key (p_partkey);
alter table supplier add primary key (s_suppkey);
alter table partsupp add primary key (ps_partkey,ps_suppkey);
alter table customer add primary key (c_custkey);
alter table lineitem add primary key (l_orderkey,l_linenumber);
alter table orders add primary key (o_orderkey);
```


**Step 5** Obtain 22 query statements from <https://github.com/dragansah/tpch-dbgen/tree/master/tpch-queries> and perform corresponding operations.

----End

### 25.1.3.2 Test Results

Based on 16-thread parallel execution, the performance of 17 query statements is greatly improved. The query speed of all statements is improved by more than 10 times on average. The following figure shows the TPC-H performance test results.

**Figure 25-4** Test results



## 25.2 Near Data Processing

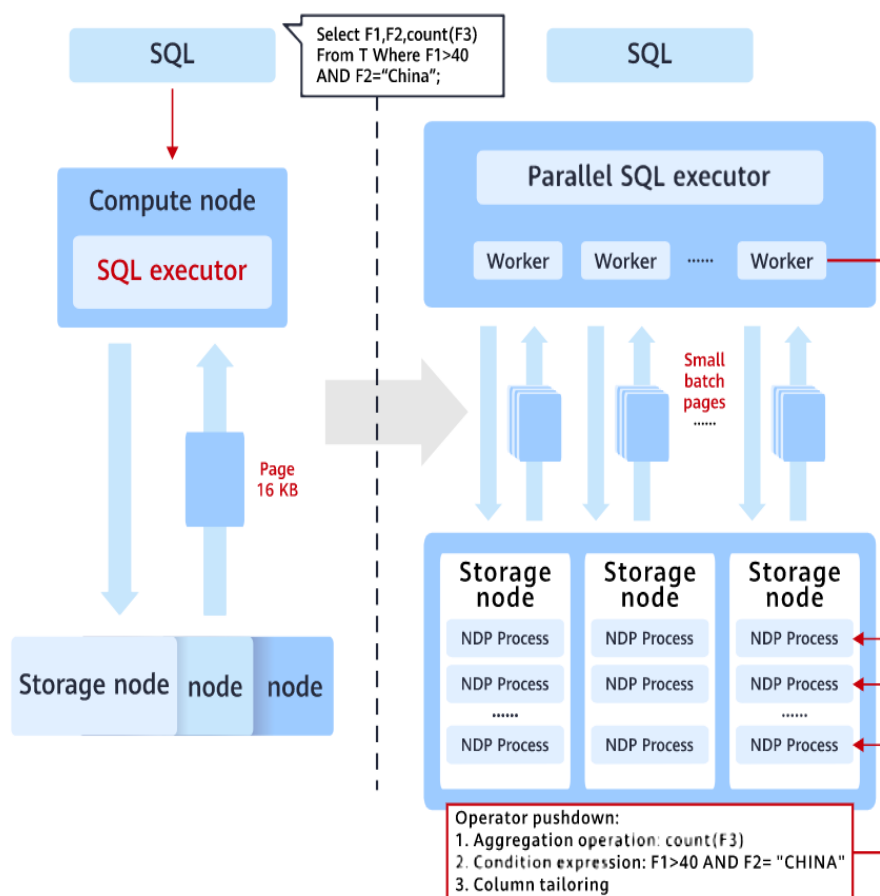
### 25.2.1 NDP Overview

Near Data Processing (NDP) is a computing pushdown solution to improve data query efficiency. For data-intensive queries, operations such as column extraction, aggregation calculation, and condition filtering are pushed down to multiple nodes on a distributed storage layer for parallel execution. This reduces query processing pressure on compute nodes, improves parallel processing capabilities, and saves network traffic.

## Architecture

GaussDB(for MySQL) uses an architecture with decoupled storage and compute to reduce network traffic. Based on this architecture, NDP is used to accelerate data queries. Without NDP, all raw data needs to be transmitted from storage nodes to compute nodes for query processing. NDP pushed the most I/O-intensive and CPU-intensive query tasks down to storage nodes. Only the required columns and filtered rows or aggregated results are sent back to compute nodes, greatly reducing network traffic. Additionally, parallel processing across storage nodes reduces the CPU usage of compute nodes and improves the query efficiency.

NDP is integrated with parallel query. Batch pages are prefetched to realize the entire process in parallel. The query execution efficiency is greatly improved.



## Application Scenarios

In terms of functions, NDP is suitable for the following scenarios:

### 1. Projection

Column pruning: Only the fields required by a query statement are sent to the compute node.

### 2. Aggregate

Typical aggregation operations include COUNT, SUM, AVG, MAX, MIN, and GROUP BY. Only the aggregated results (not all tuples) are sent to the query engine. COUNT (\*) is the most common.

### 3. SELECT - WHERE clause for filtering

Common condition expressions are COMPARE(>=,<=,<,>==), BETWEEN, IN, AND/OR, and LIKE.

A filter expression is executed on the storage nodes. Only the rows that meet the conditions are sent to the compute node.

## Application Constraints

1. InnoDB tables.
2. Tables with rows in the COMPACT or DYNAMIC format.
3. Primary keys or B-tree indexes. Hash and full-text indexes are not supported.
4. SELECT statements among the DML statements. INSERT INTO SELECT statements and SELECT statements that will lock rows (such as SELECT FOR SHARE/UPDATE) are not supported.
5. Expressions with numeric, log, time, or partial string types (CHAR and VARCHAR). The utf8mb4 and utf8 character sets are supported.
6. Expression predicates with comparison operators (<,>=,<=,>=,!=), IN, NOT IN, LIKE, NOT LIKE, BETWEEN AND, and AND/OR.

## Parameter Description

Table 25-3 Parameter description

Parameter	Level	Description
ndp_mode	Global <b>NOTE</b> <ul style="list-style-type: none"><li>• To enable NDP at the global level, contact technical support for assistance.</li><li>• NDP is in the test phase. There are 10 test users in total.</li></ul>	NDP switch. Value: <b>off</b> or <b>on</b> Default value: <b>off</b>

## 25.3 DDL Optimization

### 25.3.1 Parallel DDL

Traditional DDL is designed based on a single core and traditional disks. It takes a long time to perform DDL operations on large tables and the latency is too high. For example, when creating secondary indexes, DDL operations with high latency block subsequent DML queries that depend on new indexes.

GaussDB(for MySQL) supports parallel DDL. When database hardware resources are idle, you can use parallel DDL to accelerate DDL execution, preventing subsequent DML operations from being blocked and shortening the DDL operation window.

## Constraints

- This function is supported when the kernel version is 2.0.45.230900 or later.
- This function is only suitable for BTREE secondary indexes.
- This function is not suitable for primary key indexes, spatial indexes, and fulltext indexes. If an SQL statement for concurrently creating indexes contains a primary key index, spatial index, or fulltext index, the client will receive an alarm indicating that the operation does not support concurrent index creation. The statement is executed in single-thread index creation mode. Assume that multiple threads are specified when a primary key index is modified. An alarm will also be reported and the index is created through a single thread.

## Enabling Parallel DDL

Table 25-4 Parameter description

Parameter	Level	Description
innodb_rds_parallel_index_creation_threads	Global, Session	<ul style="list-style-type: none"> <li>• Number of threads for concurrently creating indexes.</li> <li>• If the value is greater than 1, concurrent creation is performed. Otherwise, single-thread creation is performed.</li> <li>• Default value: <b>8</b>. You are advised to set the value to be half of the number of CPU cores and be at most the value of <b>innodb_rds_parallel_index_creation_threads_max</b>.</li> </ul>

## Example

1. Prepare a sysbench table with 100 million data records.

Figure 25-5 Viewing table information

```
mysql> show create table sbtest1;
+-----+-----+
| Table | Create Table |
+-----+-----+
| sbtest1 | CREATE TABLE `sbtest1` (
  `id` int NOT NULL AUTO_INCREMENT,
  `k` int NOT NULL DEFAULT '0',
  `c` char(120) NOT NULL DEFAULT '',
  `pad` char(60) NOT NULL DEFAULT '',
  PRIMARY KEY (`id`),
  KEY `k_1` (`k`)
) ENGINE=InnoDB AUTO_INCREMENT=10000001 DEFAULT CHARSET=utf8mb4 COLLATE=utf8mb4_0900_ai_ci
```

2. Create an index in the **k** field of the table.  
Create an index for the **k** field in the table. If a single thread is used to create the index by default, it should take 146.82 seconds.



**Figure 25-6** Creating an index using a single thread

```
mysql> alter table sbtest1 add index idx_s(k);  
Query OK, 0 rows affected (2 min 26.82 sec)  
Records: 0 Duplicates: 0 Warnings: 0
```

3. Set `innodb_rds_parallel_index_creation_threads = 4` to use four threads to create the index.

It should take 38.72 seconds to create the index, 3.79 times faster than with a single thread.

**Figure 25-7** Creating an index using multiple threads

```
mysql> set innodb_rds_parallel_index_creation_threads = 4;  
Query OK, 0 rows affected (0.00 sec)  
  
mysql> alter table sbtest1 add index idx_p4(k);  
Query OK, 0 rows affected (38.72 sec)  
Records: 0 Duplicates: 0 Warnings: 0
```

4. Assume that a primary key index needs to be modified. Even if multiple threads are specified, a warning will be received and the index is created using just a single thread.

**Figure 25-8** Modifying a primary key index

```
mysql> set innodb_rds_parallel_index_creation_threads = 4;  
Query OK, 0 rows affected (0.00 sec)  
  
mysql> alter table sbtest1 add primary key(id,k);  
Query OK, 0 rows affected, 1 warning (10 min 11.36 sec)  
Records: 0 Duplicates: 0 Warnings: 1  
  
mysql> show warnings;  
+-----+-----+-----+  
| Level | Code | Message  
+-----+-----+-----+  
| Warning | 7519 | InnoDB: Creating or rebuilding PK in parallel is disallowed. |  
+-----+-----+-----+  
1 row in set (0.00 sec)
```

## 25.3.2 DDL Fast Timeout

For some specific DDL operations, you can configure their metadata lock (MDL) waiting time, preventing subsequent DML operations from being blocked.

### Constraints

- The kernel version is 2.0.45.230900 or later.
- Currently, the following DDL operations are supported: ALTER TABLE, CREATE INDEX, and DROP INDEX.

### Enabling DDL Fast Timeout

**Table 25-5** Parameter description

Parameter	Level	Description
-----------	-------	-------------

rds_ddl_lock_wait_timeout	Global, Session	<p>Defines how long that a DDL operation waits for a lock in the current session or global sessions.</p> <ul style="list-style-type: none"><li>• Value range: <b>1</b> to <b>31536000</b> (s). Default value: <b>31536000</b>, indicating that the function is disabled.</li><li>• The actual lock wait timeout for DDL operations is the smaller value between <b>lock_wait_timeout</b> and this parameter value.</li><li>• The actual table lock timeout during DDL execution at the InnoDB layer is the minimum value of <b>innodb_lock_wait_timeout</b> and this parameter value. Row locks are not considered.</li></ul>
---------------------------	-----------------	---

## Example

1. Start a client and add a lock for tables.

Figure 25-9 Adding a lock

```
mysql>
mysql>
mysql> flush tables with read lock;
Query OK, 0 rows affected (0.01 sec)

mysql>
```

2. Run the following command to check the status of the DDL fast timeout function

```
show variables like "%rds_ddl_lock_wait_timeout%";
```

Figure 25-10 Querying the status of the DDL fast timeout function

```
mysql> show variables like "rds_ddl_lock_wait_timeout";
+-----+-----+
| Variable_name | Value |
+-----+-----+
| rds_ddl_lock_wait_timeout | 31536000 |
+-----+-----+
1 row in set (0.02 sec)
```

As shown in the preceding figure, the value of **rds\_ddl\_lock\_wait\_timeout** is **31536000** (default value). The function is disabled. The subsequent operations will wait for a long time.

```
mysql> set rds_ddl_lock_wait_timeout=31536000;
Query OK, 0 rows affected (0.00 sec)

mysql>
mysql>
mysql> alter table lzk.t_lzk drop index indexa;
```

To enable function, referring to 3.

3. Run the following command to set `rds_ddl_lock_wait_timeout`.  
**set rds\_ddl\_lock\_wait\_timeout=1;**

Figure 25-11 Configuring parameters

```
mysql> show variables like "%rds_ddl_lock_wait_timeout%";
+-----+
| Variable_name | Value |
+-----+
| rds_ddl_lock_wait_timeout | 31536000 |
+-----+
1 row in set (0.01 sec)

mysql> set rds_ddl_lock_wait_timeout=1;
Query OK, 0 rows affected (0.00 sec)

mysql> show variables like "%rds_ddl_lock_wait_timeout%";
+-----+
| Variable_name | Value |
+-----+
| rds_ddl_lock_wait_timeout | 1 |
+-----+
1 row in set (0.01 sec)
```

4. Run the following command to create an index. It is found that the DDL operation times out quickly.  
**alter table lzk.t\_lzk drop index indexa;**

Figure 25-12 Creating an index

```
mysql> alter table lzk.t_lzk drop index indexa;
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting transaction
mysql> █
```

## 25.4 Statement Outline

During the runtime of a MySQL DB instance, the execution plan of a SQL statement often changes, causing database instability. To resolve the issue, GaussDB(for MySQL) provides the Statement Outline function, which uses MySQL optimizer and index hints to stabilize plan execution. GaussDB(for MySQL) also provides a group of management interfaces (**dbms\_outln package**) for easy use.

### Prerequisites

The kernel version of GaussDB(for MySQL) is 2.0.42.230600 or later.

## Precautions

1. Statement Outline is disabled by default. To enable it, contact customer service.
2. If Statement Outline is disabled, the performance is not affected. If there are a large number of rules after Statement Outline is enabled, the performance deteriorates.

## Description

Statement Outline supports the optimizer hints and index hints of MySQL 8.0.

- **Optimizer hints**  
Optimizer hints are classified into Global-Level Hint, Table-Level Hint, Index-Level Hint and Join-Order Hints based on the scope (query blocks) and Hint objects. For details, see [Optimizer Hints](#).
- **Index hints**  
Index hints provide the optimizer with information about how to select indexes during query processing without changing the optimizer's policy. There are three common index hints: USE INDEX hint, IGNORE INDEX hint, and FORCE INDEX hint. For details, see [Index Hints](#).

## outline Table

GaussDB(for MySQL) has a built-in system table (**outline**) to store hints. This table is automatically created when the system is started. The SQL statements for creating the table are as follows.

```
CREATE TABLE `mysql`.`outline` (  
  `id` bigint(20) NOT NULL AUTO_INCREMENT,  
  `Schema_name` varchar(64) COLLATE utf8_bin DEFAULT NULL,  
  `Digest` varchar(64) COLLATE utf8_bin NOT NULL,  
  `Digest_text` longtext COLLATE utf8_bin,  
  `Type` enum('IGNORE INDEX','USE INDEX','FORCE INDEX','OPTIMIZER') CHARACTER SET utf8 COLLATE  
utf8_general_ci NOT NULL,  
  `Scope` enum('','FOR JOIN','FOR ORDER BY','FOR GROUP BY') CHARACTER SET utf8 COLLATE  
utf8_general_ci DEFAULT '',  
  `State` enum('N','Y') CHARACTER SET utf8 COLLATE utf8_general_ci NOT NULL DEFAULT 'Y',  
  `Position` bigint(20) NOT NULL,  
  `Hint` text COLLATE utf8_bin NOT NULL,  
  PRIMARY KEY (`id`)  
) ENGINE=InnoDB  
DEFAULT CHARSET=utf8 COLLATE=utf8_bin STATS_PERSISTENT=0 COMMENT='Statement outline'
```

For details about the parameter description, see the following table.

**Table 25-6** Parameter description

Parameter	Description
Id	ID of the outline table.
Schema_name	Database name.
Digest	64-byte hash string calculated from <b>Digest_text</b> during the hash calculation.

Digest_text	Digest of the SQL statement.
Type	In optimizer hints, the value is <b>OPTIMIZER</b> . In index hints, the value can be <b>USE INDEX</b> , <b>FORCE INDEX</b> , or <b>IGNORE INDEX</b> .
Scope	This field is required only for index hints. Its value can be: <ul style="list-style-type: none"> <li>• <b>FOR GROUP BY</b></li> <li>• <b>FOR ORDER BY</b></li> <li>• <b>FOR JOIN</b></li> <li>• An empty string</li> </ul> <b>NOTE</b> If this field is set to an empty string, it indicates all types of index hints.
State	Whether Statement Outline is enabled. Its value can be: <ul style="list-style-type: none"> <li>• <b>N</b></li> <li>• <b>Y</b> (default value)</li> </ul>
Position	<ul style="list-style-type: none"> <li>• <b>Optimizer hints</b> Sequence number of the keyword in query blocks on which the hint is applied. Its value starts from <b>1</b>. All optimizer hints must be applied to the query block.</li> <li>• <b>Index hints</b> Sequence number of the table on which the hint is applied. Its value starts from <b>1</b>.</li> </ul>
Hint	<ul style="list-style-type: none"> <li>• <b>Optimizer hints</b> A complete hint string, for example, <b>/*+ MAX_EXECUTION_TIME(1000) */</b></li> <li>• <b>Index hints</b> A list of index names, for example, <b>ind_1,ind_2</b></li> </ul>

## Statement Outline Management

There are six local storage rules to manage Statement Outline.

- **add\_optimizer\_outline**

Adding optimizer hints

- **Syntax**

```
dbms_outln.add_optimizer_outline(<Schema_name>,<Digest>,<Query_block>,<Hint>,<Query>);
```

- **NOTE**

You can set either *Digest* or *Query* (original SQL statement). If you set *Query*, DBMS\_OUTLN calculates **Digest** and **Digest\_text**. You are advised to set *Query* directly.

- **Description**

Parameter	Mandatory	Type	Description
<i>Schema_name</i>	Yes	VARCHAR	Name of the database to which the statement belongs.  This parameter can be set to <b>NULL</b> or left blank, the statement cannot be matched.
<i>Digest</i>	No	VARCHAR	Hash value of the statement.  You can set this parameter or <b>Query</b> . If you do not want to set it to a specific value, set it to an empty string.
<i>Query_block</i>	Yes	INT	Position of the object to which the hint applies.  Value range: Greater than or equal to 1
<i>Hint</i>	Yes	VARCHAR	Hint name.
<i>Query</i>	No	VARCHAR	SQL statement.  <ul style="list-style-type: none"> <li>You can set either this parameter or <b>Digest</b>. If you do not want to set it to a specific value, set it to an empty string.</li> <li>If both of them are set, check whether <b>Digest</b> and <b>Query</b> match. If they do not match, the parameter verification fails and the execution fails.</li> </ul>

– **Example**

```

call dbms_outln.add_optimizer_outline('outline_db', '', 1, /*+ MAX_EXECUTION_TIME(1000) */, 'select * from t1 where id = 1');
call dbms_outln.add_optimizer_outline('outline_db', '', 1, /*+ SET_VAR(foreign_key_checks=OFF) */, 'select * from t1 where id = 1');

mysql> explain select * from t1 where id = 1;
+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered | Extra |
+-----+
| 1 | SIMPLE | NULL | NULL | NULL | NULL | NULL | NULL | NULL | NULL | no matching row in const table |
+-----+
1 row in set, 1 warning (0.01 sec)

mysql> show warnings;
+-----+
| Level | Code | Message |
+-----+
| Note | 1003 | /*+ select#1 */ select /*+ MAX_EXECUTION_TIME(1000) SET_VAR(foreign_key_checks=OFF) */ NULL AS 'id',NULL AS 'col1',NULL AS 'col2' from 'outline_db`.`t1' where multiple equal(1, NULL) |
+-----+
1 row in set (0.00 sec)

```

- **add\_index\_outline**

## Adding index hints

- **Syntax**

```
dbms_outln.add_index_outline(<Schema_name>,<Digest>,<Position>,<T
ype>,<Hint>,<Scope>,<Query>);
```

 **NOTE**

You can set either *Digest* or *Query* (original SQL statement). If you set *Query*, DBMS\_OUTLN calculates *Digest* and *Digest\_text*. You are advised to set *Query* directly.

- **Description**

Parameter	Mandatory	Type	Description
<i>Schema_name</i>	Yes	VARCHAR	Name of the database to which the statement belongs.  This parameter can be set to <b>NULL</b> or left blank, the statement cannot be matched.
<i>Digest</i>	No	VARCHAR	Hash value of the statement.  Set either this parameter or <i>Query</i> . If you do not want to set it to a specific value, set it to an empty string.
<i>Position</i>	Yes	INT	Position of the table to which the index hint applies in the statement.  The value must be greater than or equal to <b>1</b> .
<i>Type</i>	Yes	ENUM	Hint type. Its value can be: <ul style="list-style-type: none"> <li>• <b>OPTIMIZER</b></li> <li>• <b>USE INDEX</b></li> <li>• <b>FORCE INDEX</b></li> <li>• <b>IGNORE INDEX</b></li> </ul>
<i>Hint</i>	Yes	VARCHAR	Hint name or index name set. Use commas (,) to separate multiple index names.

<i>Scope</i>	Yes	ENUM	Hint scope. Its value can be: <ul style="list-style-type: none"> <li>• <b>FOR GROUP BY</b></li> <li>• <b>FOR ORDER BY</b></li> <li>• <b>FOR JOIN</b></li> <li>• An empty string</li> </ul>
<i>Query</i>	No	VARCHAR	SQL statement. <ul style="list-style-type: none"> <li>• You can select either or <i>Digest</i>. If you do not want to set it to a specific value, set it to an empty string.</li> <li>• If both of them are set, check whether <b>Digest</b> and <b>Query</b> match. If they do not match, the parameter verification fails and the execution fails.</li> </ul>

– **Example**

```
call dbms_outln.add_index_outline('outline_db', '', 1, 'USE INDEX', 'ind_1', '', 'select * from t1 where t1.col1 =1 and t1.col2 ='xpchild'');
```

• **preview\_outline**

Querying the status of the SQL statement matching the statement outline, which can be used for manual verification.

– **Syntax**

```
dbms_outln.preview_outline(<Schema_name>,<Query>);
```

– **Description**

Parameter	Mandatory	Data Type	Description
<i>Schema_name</i>	Yes	VARCHAR	Database name.
<i>Query</i>	Yes	VARCHAR	SQL statement.

– **Example**

```
mysql> call dbms_outln.preview_outline('outline_db', 'select * from t1 where t1.col1 =1 and t1.col2 ='xpchild');
+-----+-----+-----+-----+-----+-----+-----+-----+
| SCHEMA | DIGEST | BLOCK_TYPE | BLOCK_NAME | BLOCK | HINT |
+-----+-----+-----+-----+-----+-----+-----+-----+
| outline_db | b4369611be7ab2d27c85897632576a04bc08f50b928a1d735b62d0a140628c4c | TABLE | t1 | 1 | USE INDEX ('ind_1') |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

• **show\_outline**

Displaying the in-memory hit rate of statement outline

– **Syntax**

```
dbms_outln.show_outline();
```



- **Example**

```
mysql> call dbms_outln.view_outline();
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID | OWNER | OBJECT | TYPE | SCOPE | POS | HINT | HIT | OVERFLOW | OBJECT_TEXT |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 11 | OUTLINE_01 | 3586cc31fa7e320302a6c3f6d780a45891874c8d186861c087462d46 | OPTIMIZE | 1 | /* SET_VAR(foreign_key_checks=OFF) */ | 1 | 0 | SELECT * FROM 't1' WHERE 't1' = ? |
| 12 | OUTLINE_01 | 3586cc31fa7e320302a6c3f6d780a45891874c8d186861c087462d46 | OPTIMIZE | 1 | /* MAX_EXECUTION_TIME(1000) */ | 2 | 0 | SELECT * FROM 't1' WHERE 't1' = ? |
| 14 | OUTLINE_01 | 69e4ef36e4e0403048c0f782a24e113a208079e172a1320f232af72a6f | OPTIMIZE | 1 | /* BIG_TABLE */ | 1 | 0 | SELECT 't1' - 't2' - 't3' FROM 't1', 't2' |
| 15 | OUTLINE_01 | 5a726a8080f8f7808f9f2a24e113a208079e172a1320f232af72a6f | OPTIMIZE | 2 | /* QB_NAME(sales) */ | 2 | 0 | SELECT * FROM 't1' WHERE 't1' - 't2' IN ( SELECT 't1' FROM 't2' ) |
| 16 | OUTLINE_01 | 5a726a8080f8f7808f9f2a24e113a208079e172a1320f232af72a6f | OPTIMIZE | 1 | /* USE_INDEX(sales MATERIALIZATION, DUPS(OUT)) */ | 2 | 0 | SELECT * FROM 't1' WHERE 't1' - 't2' IN ( SELECT 't1' FROM 't2' ) |
| 18 | OUTLINE_01 | 84180513a7a03d7c85f932278a80c8f80828a1d7f35d328a4802d46 | USE INDEX | 1 | End_1 | 3 | 0 | SELECT * FROM 't1' WHERE 't1' - 't2' = ? AND 't1' - 't2' = ? |
| 19 | OUTLINE_01 | 281734170805f78a1320f232af72a6f9f2a24e113a208079e172a1320f232af72a6f | USE INDEX | 2 | End_2 | 2 | 0 | SELECT * FROM 't1', 't2' WHERE 't1' - 't2' = 't2' - 't2' AND 't1' - 't2' = ? |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
7 rows in set (0.00 sec)
```

HIT and OVERFLOW description

- i. **HIT** indicates the number of times that the statement outline finds the destination query block or table.
- ii. **OVERFLOW** indicates the number of times that statement outline does not find the destination query block or table.

- **del\_outline**

Deleting a statement outline from the memory and table.

- **Syntax**

```
dbms_outln.del_outline(<id>);
```

- **Description**

Parameter	Mandatory	Type	Description
<i>id</i>	Yes	INT	Statement outline ID, which is the value in the <b>id</b> column in the <b>mysql.outline</b> table. The value cannot be left blank.

- **Example**

```
mysql> call dbms_outln.del_outline(1000);
Query OK, 0 rows affected, 2 warnings (0.00 sec)

mysql> show warnings;
+-----+-----+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+-----+-----+
| Warning | 7521 | Statement outline 1000 is not found in table |
| Warning | 7521 | Statement outline 1000 is not found in cache |
+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

Note: If the statement outline to be deleted does not exist, the system reports a warning. You can run the **show warnings;** command to view the warning content.

- **flush\_outline**

If you modify the statement outline in the **outline** table, you need to make the statement outline take effect again.

- **Syntax**

```
dbms_outln.flush_outline();
```

- **Example**

```
update mysql.outline set Position = 1 where Id = 18;
call dbms_outln.flush_outline();
```

## Function Verification

To check whether the statement outline takes effect, perform the following steps:

- Use the `preview_outline` interface.

```
mysql> call dbms_outln.preview_outline('outline_db', 'select * from t1 where t1.col1 =1 and t1.col2 = "xpcchild"');
+-----+-----+-----+-----+-----+-----+
| SCHEMA | DIGEST | BLOCK_TYPE | BLOCK_NAME | BLOCK | HINT |
+-----+-----+-----+-----+-----+-----+
| outline_db | b4369611be7ab2d27c85897632576a04bc08f50b928a1d735b62d0a140628c4c | TABLE | t1 | 1 | USE INDEX (`ind_1`) |
+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

- Run the **EXPLAIN** command.

```
mysql> explain select * from t1 where t1.col1 =1 and t1.col2 = "xpcchild";
+-----+-----+-----+-----+-----+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered | extra |
+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | t1 | | index | ref | ind_1 | 19 | | 1 | 100.00 | Using where |
+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

mysql> show warnings;
+-----+-----+-----+-----+-----+-----+
| Level | Code | Message |
+-----+-----+-----+-----+-----+-----+
| Note | 1000 | 77 selects '1' selects '1' select 'outline_db.' '1' '1' AS 'id', 'outline_db.' '1' '1' AS 'col1', 'outline_db.' '1' '1' AS 'col2' from 'outline_db.' '1' USE INDEX (`ind_1`) USE INDEX (`ind_1`) USE INDEX (`ind_1`) where ((('outline_db.' '1' '1' AS 'id') = 1) and (('outline_db.' '1' '1' AS 'col2') = 'xpcchild')) |
+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

## 25.5 Idle Transaction Disconnection

### 25.5.1 Function

If an idle transaction is not committed for a long time, its rollback will consume database resources and performance. If a large number of idle transactions are not committed and not rolled back for a long time, the performance loss to a database is severe especially during peak hours. GaussDB(for MySQL) can proactively terminate idle transactions. Different parameters are used to control different types of transactions. When idle transactions timed out, they are automatically rolled back and disconnected.

#### NOTE

This function is supported when the kernel version is 2.0.39.230300 or later.

### 25.5.2 Parameter Description

```
mysql> show variables like '%idle%';
+-----+-----+-----+-----+
| Variable_name | Value |
+-----+-----+-----+-----+
| idle_readonly_transaction_timeout | 0 |
| idle_transaction_timeout | 0 |
| idle_write_transaction_timeout | 0 |
+-----+-----+-----+-----+
```

**Table 25-7** Parameter description

Parameter	Level	Description
idle_readonly_transaction_timeout	global, session	Time in seconds that the server waits for idle read-only transactions before killing the connection. If this parameter is set to <b>0</b> , there is not timeout threshold for idle read-only transactions.
idle_transaction_timeout	global, session	Time in seconds that the server waits for common idle transactions before killing the connection. If this parameter is set to <b>0</b> , there is not timeout threshold for common idle transactions.
idle_write_transaction_timeout	global, session	Time in seconds that the server waits for idle read/write transactions before killing the connection. If this parameter is set to <b>0</b> , there is not timeout threshold for idle read/write transactions.

The parameters **idle\_readonly\_transaction\_timeout** and **idle\_write\_transaction\_timeout** have higher priorities than the parameter **idle\_transaction\_timeout**.

**Figure 25-13** Read-only transactions

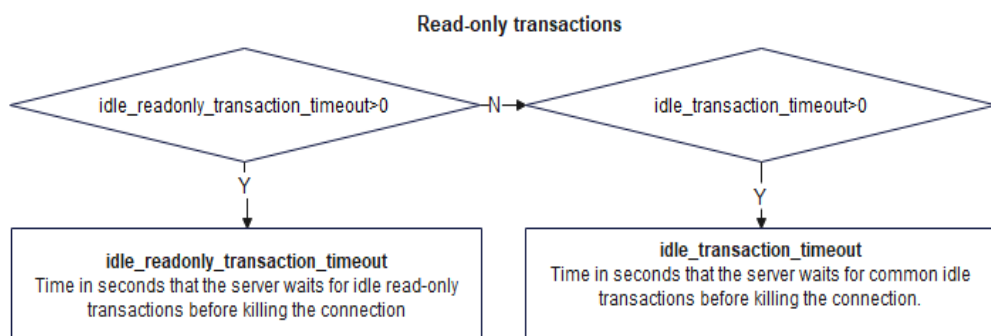
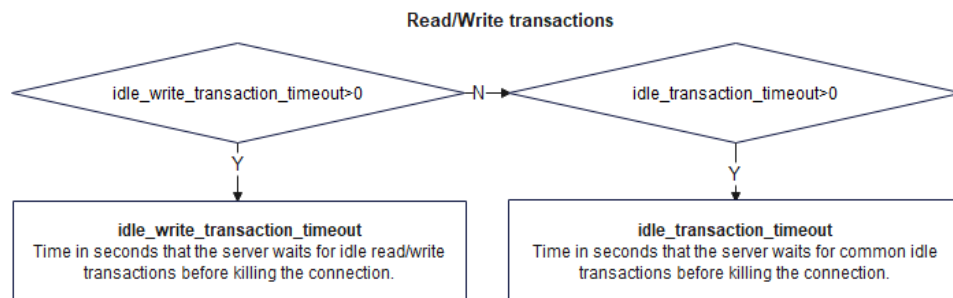


Figure 25-14 Read/Write transactions



### 25.5.3 Example

1. Set **idle\_transaction\_timeout** to **10**, **idle\_readonly\_transaction\_timeout** to **0**, and **idle\_write\_transaction\_timeout** to **0**.

- Read-only transactions

```
mysql> begin;
Query OK, 0 rows affected (0.00 sec)
```

Wait for 10 seconds and run a query statement again. The following information is displayed.

```
mysql> select * from t1;
ERROR 2013 (HY000): Lost connection to MySQL server during query
```

- Read/Write Transaction

Run the **begin** statement to start a transaction and run a query statement. The following information is displayed.

```
mysql> select * from t1;
+-----+
| col_int |
+-----+
|      1 |
+-----+
1 row in set (0.00 sec)
mysql> begin;
Query OK, 0 rows affected (0.00 sec)
```

```
mysql> insert into t1 values(2);
Query OK, 1 row affected (0.00 sec)
```

Wait for 10 seconds and run a query statement again. The following information is displayed.

```
mysql> select * from t1;
ERROR 2013 (HY000): Lost connection to MySQL server during query
```

Reconnect the transaction to the database and run a query statement. If the following information is displayed, the transaction has been rolled back.

```
mysql> select * from t1;
+-----+
| col_int |
+-----+
|      1 |
+-----+
1 row in set (0.00 sec)
```

2. Set **idle\_write\_transaction\_timeout** to **15**.

- Read/Write transactions

Run the **begin** statement to start a transaction and run a query statement. The following information is displayed.

```
mysql> select * from t1;
+-----+
| col_int |
+-----+
|      2 |
+-----+
1 row in set (0.00 sec)

mysql> begin;
Query OK, 0 rows affected (0.00 sec)

mysql> insert into t1 values(3);
Query OK, 1 row affected (0.00 sec)
```

Wait for 15 seconds and run a query statement again. The following information is displayed.

```
mysql> select * from t1;
ERROR 2013 (HY000): Lost connection to MySQL server during query
```

Reconnect the transaction to the database and run a query statement. If the following information is displayed, the transaction has been rolled back.

```
mysql> select * from t1;
+-----+
| col_int |
+-----+
|      2 |
+-----+
1 row in set (0.01 sec)
```

### 3. Set **idle\_readonly\_transaction\_timeout** to 15.

- Read-only transactions

```
mysql> begin;
Query OK, 0 rows affected (0.00 sec)
```

Wait for 15 seconds and run a query statement again. The following information is displayed.

```
mysql> select * from t1;
ERROR 2013 (HY000): Lost connection to MySQL server during query
```

## 25.6 LIMIT...OFFSET Pushdown

### 25.6.1 Function

In MySQL Community Edition, If you use  $LIMIT(M)$  and  $OFFSET(P)$  in a SELECT statement, the engine layer returns all rows that meet the WHERE condition to the SQL layer for processing. The SQL layer skip  $P$  rows of data and returns  $N$  rows of data. When a secondary index needs to access the columns in the primary table, the engine layer returns the table to obtain all required column information. If OFFSET value ( $P$ ) is much greater than the LIMIT value ( $M$ ), the engine layer sends a large amount of data to the SQL layer for processing.

In GaussDB(for MySQL), If you use  $LIMIT(M)$  and  $OFFSET(P)$  in a SELECT statement, data is pushed down to the engine layer for processing, speeding up queries.

## 25.6.2 Usage

**Table 25-8** Parameter description

Parameter	Level	Description
optimizer_switch	Global, Session	<p>Enables or disables query optimization. After this parameter is enabled, you can configure <b>offset_pushdown</b> to enable or disable LIMIT OFFSET pushdown.</p> <ul style="list-style-type: none"> <li>● <b>ON</b>: enabled</li> <li>● <b>OFF</b>: (default value): disabled</li> </ul>

You can also add a HINT clause to enable or disable LIMIT OFFSET pushdown.

- **OFFSET\_PUSHDOWN**(*table\_name*): enabled
- **NO\_OFFSET\_PUSHDOWN**(*table\_name*): disabled

Example:

Take a schema as an example in a TPC-H test. After LIMIT OFFSET pushdown is enabled using the parameter or the HINT clause, **Using limit-offset pushdown** is displayed in the **Extra** column when you run EXPLAIN SQL to view an execution plan.

- Enabling LIMIT OFFSET pushdown by configuring **offset\_pushdown**

```
mysql> EXPLAIN SELECT * FROM lineitem LIMIT 1000000,10;
+-----+
+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered |
Extra |
+-----+
+-----+
| 1 | SIMPLE | lineitem | NULL | ALL | NULL | NULL | NULL | NULL | 59281262 | 100.00 |
Using offset pushdown |
+-----+
+-----+
1 row in set, 1 warning (0.00 sec)
```

- Enabling LIMIT OFFSET pushdown by adding a HINT clause

```
mysql> EXPLAIN SELECT /*+ OFFSET_PUSHDOWN() */ * FROM lineitem LIMIT 1000000,10;
+-----+
+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered |
Extra |
+-----+
+-----+
| 1 | SIMPLE | lineitem | NULL | ALL | NULL | NULL | NULL | NULL | 59281262 | 100.00 |
Using offset pushdown |
+-----+
+-----+
1 row in set, 1 warning (0.00 sec)

mysql> EXPLAIN SELECT /*+ NO_OFFSET_PUSHDOWN() */ * FROM lineitem LIMIT 1000000,10;
+-----+
+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered | Extra |
+-----+
+-----+
```

```

+-----+
| 1 | SIMPLE | lineitem | NULL | ALL | NULL | NULL | NULL | NULL | 59281262 | 100.00 |
NULL |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
1 row in set, 1 warning (0.00 sec)

```

### 25.6.3 Performance Tests

- Run following SQL statement (Q1) with no predicate conditions to access the primary table.

```

mysql> EXPLAIN SELECT * FROM lineitem LIMIT 1000000,10;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered |
Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | lineitem | NULL | ALL | NULL | NULL | NULL | NULL | 59281262 | 100.00 |
Using offset pushdown |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

```

- Run following SQL statement (Q2) with predicate conditions to access the secondary index (including the index range conditions). Information about other columns needs to be obtained from the table.

```

mysql> EXPLAIN SELECT * FROM lineitem WHERE L_partkey > 10 AND L_partkey < 200000 LIMIT
5000000, 10;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref |
rows | filtered | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | lineitem | NULL | range | i_L_partkey_suppkey,i_L_partkey | i_L_partkey | 4 |
NULL | 10949662 | 100.00 | Using offset pushdown; Using index condition |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

```


- Run following SQL statement (Q3) with predicate conditions and ORDER BY to sort data by index.

```

mysql> EXPLAIN SELECT * FROM lineitem WHERE L_partkey > 10 AND L_partkey < 200000 ORDER BY
L_partkey LIMIT 5000000, 10;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref |
rows | filtered | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | lineitem | NULL | range | i_L_partkey_suppkey,i_L_partkey | i_L_partkey | 4 |
NULL | 10949662 | 100.00 | Using offset pushdown; Using index condition |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

```

The following figure describes the performance of Q1, Q2, and Q3 when LIMIT OFFSET pushdown is enabled and disabled in the TPC-H benchmark (10 scale).

**Figure 25-15** Performance comparison

## 25.7 Conversion of IN Predicates Into Subqueries

### 25.7.1 Function

To execute complex queries efficiently, the GaussDB(for MySQL) optimizer can convert some big IN predicates into IN subqueries. The conversion happens if the following conditions are met:

- The GaussDB(for MySQL) kernel version is 2.0.42.230600 or later.
- The number of elements in the **IN** list exceeds the value of **rds\_in\_predicate\_conversion\_threshold**.

#### Overview

In MySQL Community Edition, if column IN (const1, const2, ..... ) is executed and there is an index on the column, the optimizer usually performs a range scan. The parameter **range\_optimizer\_max\_mem\_size** controls the memory available to the range optimizer. If there are many elements in the IN list and the used memory exceeds the parameter value, the range scan will fail and the query performance deteriorates. To solve this problem, you can increase the parameter value to expand the memory that can be used. However, the memory is at the session level. It means that each session occupies the same memory, so the instance may be out of memory. Even if the range optimizer can be used, if the number of elements in the IN list exceeds the **eq\_range\_index\_dive\_limit** value, index statistics, instead of index dive is used. This may cause inaccurate estimation and performance rollback. After IN predicates into subqueries, the optimizer will continue to consider whether to convert the IN clause into a semijoin to improve performance. A specific conversion process is as follows.

```
select ... from lineitem where L_partkey in (...)
```

====>

```
select ... from lineitem where L_partkey in  
(select tb_col_1 from (values (9628136),(19958441),...) tb)
```



## 25.7.2 Scenarios

### Supported Query Statements

- SELECT
- INSERT ... SELECT
- REPLACE ... SELECT
- PREPARED STMT and views

### Constraints

- Only the constant IN LIST (including statements that do not involve table queries, such as NOW() and ?) is supported.
- Stored procedures, functions, and triggers are not supported.
- NOT IN is not supported. Statements where indexes cannot be used are not supported.

## 25.7.3 Usage

You can use the **rds\_in\_predicate\_conversion\_threshold** parameter to convert IN predicates into subqueries.

#### NOTE

The default value is **0**, indicating the conversion is disabled. To configure this parameter, contact customer service.

**Table 25-9** Parameter description

Parameter	Level	Description
rds_in_predicate_conversion_threshold	Global	Controls the minimum number of elements in the value list of an IN predicate that triggers its conversion to an IN subquery.

### Example:

- Query before conversion:

```
mysql> explain select * from t where a in (1,2,3,4,5);
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | partitions | type | possible_keys | key | key_len | ref | rows | filtered |
Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | t | NULL | ALL | idx1 | NULL | NULL | NULL | 5 | 100.00 | Using where |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set, 1 warning (0.00 sec)

mysql> explain format=tree select * from t where a in (1,2,3,4,5);
+-----+
| EXPLAIN |
```

```

+-----+
| -> Filter: (t.a in (1,2,3,4,5)) (cost=0.75 rows=5)
|   -> Table scan on t (cost=0.75 rows=5)
|
+-----+
1 row in set (0.01 sec)

```

- Query after conversion:

```

mysql> set rds_in_predicate_conversion_threshold=3;
Query OK, 0 rows affected (0.00 sec)

```

```

mysql> explain select * from t where a in (1,2,3,4,5);

```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table      | partitions | type | possible_keys | key      | key_len | ref |
| rows | filtered | Extra      |            |      |               |         |         |    |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE     | t          | NULL       | ALL | idx1          | NULL    | NULL   | NULL |
| 5 | 100.00    | Using where |            |      |               |         |         |     |
| 1 | SIMPLE     | <in_predicate_2> | NULL      | eq_ref | <auto_distinct_key> | <auto_distinct_key> | |
| 5 | test.t.a   | 1 | 100.00 | IN-list converted |
+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set, 1 warning (0.00 sec)

```

```

mysql> explain format=tree select * from t where a in (1,2,3,4,5);

```

```

+-----+
|
| EXPLAIN
|
+-----+
|
+-----+
| -> Nested loop inner join (cost=2.50 rows=5)
|   -> Filter: (t.a is not null) (cost=0.75 rows=5)
|     -> Table scan on t (cost=0.75 rows=5)
|       -> Single-row index lookup on <in_predicate_2> using <auto_distinct_key> (a=t.a) (cost=0.27
|         rows=1)
|
+-----+

```

EXPLAIN returns the execution plan. There is **<in\_predicate\_\*** (\* indicates a number) in the table column. It means that the table is a temporary table that stores all data in the IN query.

You can also view **in\_to\_subquery\_conversion** information in the optimize trace.

```

| explain format=tree select * from t where a in (1,2,3,4,5) | {
| "steps": [
|   {
|     "join_preparation": {
|       "select#": 1,
|       "steps": [
|         {
|           "IN_uses_bisection": true
|         },
|         {
|           "in_to_subquery_conversion": {
|             "item": "(`t`.`a` in (1,2,3,4,5))",
|             "steps": [
|               {
|                 "creating_tmp_table": {
|                   "tmp_table_info": {

```

```
"table": "intermediate_tmp_table",
"columns": 1,
"row_length": 5,
"key_length": 5,
"unique_constraint": false,
"makes_grouped_rows": false,
"cannot_insert_duplicates": true,
"location": "TempTable"
}
},
```

## 25.7.4 Performance Tests

sysbench is used to perform a benchmark test.

1. Prepare 10 million data records.

```
sysbench /usr/share/sysbench/oltp_read_only.lua --tables=1 --report-interval=10 --table-size=10000000 --mysql-user=root --mysql-password=123456 --mysql-host=127.0.0.1 --mysql-port=3306 --mysql-db=sbtest --time=300 --max-requests=0 --threads=200 prepare
```

2. Run a statement where there are 10,000 elements in IN list.

```
select count(*) from sbtest1 where id/k in (... ...);
```

The following table lists the performance comparison.

**Table 25-10** Performance data

Method	Function Enabled	Function Disabled (Not Suitable for range_opt)	Performance Comparison
Statements using indexes	0.09	2.48	Improved by 26.5 times

## 25.8 Diagnosis on Large Transactions

Large transactions affect the health and stability of DB instances. In typical scenarios, long rollbacks of large transactions prolong the upgrade and specification change time. GaussDB(for MySQL) provides diagnosis on large transactions. When there is a large transaction, an alarm is generated to notify you to submit the transaction in a timely manner.

### Prerequisites

- The kernel version is 2.0.39.230300 or later.
- The related parameter is configured based on the following conditions:
  - If the kernel version is earlier than 2.0.45.230900, set the value of **log\_bin** is **ON**. To view and modify the parameter value, see [Modifying Parameters of a GaussDB\(for MySQL\) DB Instance](#).
  - If the kernel version is 2.0.45.230900 or later, set the value of **rds\_global\_sql\_log\_bin** to **ON**.

## Usage

1. Configure the parameter **rds\_warn\_max\_binlog\_cache\_size** as required.

**Table 25-11** Parameter description

Parameter	Level	Description
rds_warn_max_binlog_cache_size	global	Controls the maximum binlog cache size for a transaction. If the size in a transaction exceeds the parameter value, a WARNING message is reported.  Default value: <b>18446744073709547520</b>  Value range: <b>4096</b> to <b>18446744073709547520</b>

To prevent multiple WARNING messages from being sent to the client, a WARNING message can be sent to the client once for each statement in a transaction.

In this example, **rds\_warn\_max\_binlog\_cache\_size** is set to **40960** (40 KB).

```
mysql> CREATE TABLE t1 (
-> a longtext
-> ) DEFAULT CHARSET=latin1;
Query OK, 0 rows affected (0.12 sec)

mysql> show variables like 'rds_warn_max_binlog_cache_size';
+-----+-----+
| Variable_name          | Value |
+-----+-----+
| rds_warn_max_binlog_cache_size | 40960 |
+-----+-----+
1 row in set (0.01 sec)

mysql> begin;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t1 VALUES (REPEAT('a',20000));
Query OK, 1 row affected (0.01 sec)

mysql> INSERT INTO t1 VALUES (REPEAT('a',20000));
Query OK, 1 row affected (0.00 sec)

mysql> INSERT INTO t1 VALUES (REPEAT('a',20000));
Query OK, 1 row affected, 1 warning (0.00 sec)

mysql> show warnings;
+-----+-----+
| Level | Code | Message |
+-----+-----+
+-----+-----+
| Warning | 4008 | Recommend you to INSERT/UPDATE/DELETE rows in batches by multiple transactions. The current transaction required more than 'rds_warn_max_binlog_cache_size' (40960) bytes of storage. Which shall cause replication latency. Please commit it. |
+-----+-----+
```

```

+-----+-----
+-----+-----
+-----+-----
-----+
1 row in set (0.00 sec)

mysql> select count(*) from t1;
+-----+
| count(*) |
+-----+
|      3 |
+-----+
1 row in set (0.01 sec)

mysql> commit;
Query OK, 0 rows affected (0.01 sec)

mysql> select count(*) from t1;
+-----+
| count(*) |
+-----+
|      3 |
+-----+
1 row in set (0.01 sec)

mysql> INSERT INTO t1 VALUES (REPEAT('a',50000));
Query OK, 1 row affected, 1 warning (0.01 sec)

mysql> show warnings;
+-----+-----
+-----+-----
+-----+-----
| Level | Code | Message          |
+-----+-----
+-----+-----
+-----+-----
| Warning | 4008 | Recommend you to INSERT/UPDATE/DELETE rows in batches by multiple
transactions. The current transaction required more than 'rds_warn_max_binlog_cache_size' (40960)
bytes of storage. Which shall cause replication latency. |
+-----+-----
+-----+-----
1 row in set (0.00 sec)

```

2. Check the binlog cache size of the transactions in the current connection.

```

mysql> CREATE TABLE t1 (
->  a longtext
-> ) ENGINE=InnoDB DEFAULT CHARSET=latin1;
Query OK, 0 rows affected (0.10 sec)

mysql> SHOW STATUS LIKE 'Rds_binlog_trx_cache_size';
+-----+-----+
| Variable_name          | Value |
+-----+-----+
| Rds_binlog_trx_cache_size | 0     |
+-----+-----+
1 row in set (0.04 sec)

mysql> begin;
Query OK, 0 rows affected (0.00 sec)

mysql> INSERT INTO t1 VALUES (REPEAT('a',20000));
Query OK, 1 row affected (0.01 sec)

mysql> SHOW STATUS LIKE 'Rds_binlog_trx_cache_size';
+-----+-----+
| Variable_name          | Value |
+-----+-----+
| Rds_binlog_trx_cache_size | 20150 |
+-----+-----+

```

```
1 row in set (0.05 sec)
mysql> commit;
Query OK, 0 rows affected (0.00 sec)

mysql> SHOW STATUS LIKE 'Rds_binlog_trx_cache_size';
+-----+-----+
| Variable_name      | Value |
+-----+-----+
| Rds_binlog_trx_cache_size | 0     |
+-----+-----+
1 row in set (0.09 sec)
```

3. Check the binlog cache size of transactions in all connections.

```
mysql> SHOW GLOBAL STATUS LIKE 'rds_binlog_trx_cache_size';
+-----+-----+
| Variable_name      | Value |
+-----+-----+
| Rds_binlog_trx_cache_size | 40300 |
+-----+-----+
1 row in set (0.05 sec)
```

# A Change History

Released On	Description
2024-02-23	<p>This issue is the forty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added the parameter <b>min_binlog_expire_logs_seconds</b> in <a href="#">Synchronizing Data</a>.</li><li>• Added <a href="#">Viewing Slow Query Logs</a>. HTAP instances support slow query logs.</li><li>• Added <a href="#">Adding or Deleting Tables in the Blacklist or Whitelist</a>.</li><li>• Allowed you to enter tables that you want to add in the whitelist or blacklist in <a href="#">Synchronizing Data</a>.</li><li>• Added <a href="#">Configuring SSL</a> and <a href="#">Connecting to an HTAP Instance Through HTTPS</a>. HTAP instances support SSL data encryption.</li></ul>
2024-01-23	<p>This issue is the forty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Changed read/write splitting address to proxy address.</li><li>• Changed the description for <b>innodb_parallel_select_count</b> in <a href="#">Changing vCPUs and Memory of a DB Instance</a>.</li><li>• Allowed you to set <b>Scaling Type</b> to <b>Number of read replicas</b> for yearly/monthly instances. For details, see <a href="#">Configuring Auto Scaling Policies</a>.</li><li>• Changed HTAP Read Replicas to HTAP Analysis.</li></ul>
2024-01-15	<p>This issue is the forty-third official release, which incorporates the following change:</p> <p>Modified the billing description in <a href="#">Billing</a>.</p>
2023-12-29	<p>This issue is the forty-second official release, which incorporates the following changes:</p> <p>Added <a href="#">Enabling TDE</a>.</p>

Released On	Description
2023-12-05	<p>This issue is the forty-first official release, which incorporates the following changes:</p> <p>Added the following kernel functions:</p> <ul style="list-style-type: none"><li>● <a href="#">Parallel DDL</a></li><li>● <a href="#">DDL Fast Timeout</a></li><li>● <a href="#">Idle Transaction Disconnection</a></li><li>● <a href="#">LIMIT...OFFSET Pushdown</a></li><li>● <a href="#">Conversion of IN Predicates Into Subqueries</a></li><li>● <a href="#">Diagnosis on Large Transactions</a></li></ul>
2023-11-03	<p>This issue is the fortieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>● Added the metric "Existing Transactions in Undo Space".</li><li>● Optimized <a href="#">Billing</a>.</li><li>● Allowed you to delete database proxy nodes.</li><li>● Added the parameter <b>Subnet</b> in <a href="#">Create a Proxy Instance</a>.</li><li>● Added <a href="#">HTAP Analysis</a>.</li></ul>
2023-10-12	<p>This issue is the thirty-ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>● Added <a href="#">Auto-Renewing a DB Instance</a>.</li><li>● Added <a href="#">Migration Solution Overview</a>.</li></ul>
2023-08-31	<p>This issue is the thirty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>● Added <a href="#">Slow SQL Analysis</a>.</li><li>● Added <a href="#">Top SQL</a>.</li><li>● Added <a href="#">SQL Insights</a>.</li><li>● Added <a href="#">Concurrency Control</a>.</li><li>● Optimized the content in <a href="#">Changing a Security Group</a>. After read/write splitting is enabled, you can change the security group of the DB instance.</li></ul>



Released On	Description
2023-07-24	<p>This issue is the thirty-seventh official release, which incorporates the following change:</p> <ul style="list-style-type: none"><li>• Added <b>Billing</b> for auto scaling policies.</li><li>• Added <b>feature comparison</b> among different backup types.</li><li>• Optimized <b>Backup Principles</b>.</li><li>• Added the restriction on the number of tables to be restored in <b>Restoring Table Data to a Specific Point in Time</b>.</li><li>• Added <b>Constraints</b> and <b>Backup Clearing</b> in <b>Creating a Manual Backup</b>.</li><li>• Added <b>Constraints</b> in <b>Configuring a Same-Region Backup Policy</b>.</li></ul>
2023-07-18	<p>This issue is the thirty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <b>Applying for and Changing a Private Domain Name for a DB Instance</b>.</li><li>• Added <b>Changing the Compute Range of a Serverless Instance</b>.</li><li>• Added <b>Enabling or Disabling Encrypted Backup</b>.</li><li>• Added <b>Configuring a Routing Policy</b>.</li><li>• Added <b>Changing the Port of a Proxy Instance</b>.</li><li>• Allowed you to delete and change the private domain name for a proxy instance in <b>Using a Private Domain Name for a Proxy Instance</b>.</li></ul>
2023-06-20	<p>This issue is the thirty-fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Allowed you to configure auto scaling policies for yearly/monthly DB instances in <b>Configuring Auto Scaling Policies</b>.</li><li>• Added two events <b>Instance DR switchover</b> and <b>Database process restarted</b> in <b>Events Supported by Event Monitoring</b>.</li></ul>

Released On	Description
2023-05-22	<p>This issue is the thirty-fourth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Allowed you to reduce the storage for a DB instance to 10 GB in <a href="#">Changing Storage of a Yearly/Monthly DB Instance</a>.</li><li>• Extended the backup retention period to 3,660 days in <a href="#">Configuring a Same-Region Backup Policy</a>.</li><li>• Supported changing the specifications of yearly/monthly instances during the maintenance period in <a href="#">Changing vCPUs and Memory of a DB Instance</a>.</li><li>• Added a function: When a DB instance is restored to a specified point in time, the original DB instance parameters can be automatically synchronized in <a href="#">Restoring Instance Data to a Specific Point in Time</a>.</li></ul>
2023-04-20	<p>This issue is the thirty-third official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Configuring Connection Pools</a>.</li><li>• Added <a href="#">DBA Assistant</a>.</li><li>• Added remarks during database creation. For details, see <a href="#">Creating a Database</a>.</li><li>• Allowed you to add remarks during database account creation in <a href="#">Creating a Database Account</a>.</li></ul>
2023-03-31	<p>This issue is the thirty-second official release, which incorporates the following change:</p> <p>Added <a href="#">Upgrading the OS of a DB Instance</a>.</p>
2023-03-20	<p>This issue is the thirty-first official release, which incorporates the following change:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Changing a Node Name</a>.</li><li>• Optimized <a href="#">Rebuilding a Deleted Instance from Recycle Bin</a>. You can configure recycle bin policies.</li></ul>
2023-02-23	<p>This issue is the thirtieth official release, which incorporates the following change:</p> <p>Modified <a href="#">Creating a Database</a>. The database name can contain hyphens (-).</p>
2023-01-30	<p>This issue is the twenty-ninth official release, which incorporates the following change:</p> <p>Added <a href="#">Rebooting a Proxy Instance</a>.</p>

Released On	Description
2022-12-30	<p>This issue is the twenty-eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <b>Modifying Proxy Instance Parameters</b>.</li><li>• Added <b>Configuring Multi-Statement Processing Mode</b>.</li><li>• Added <b>Upgrading the Kernel Version of a Proxy Instance</b>.</li><li>• Added <b>Enabling or Disabling Log Reporting</b>. New logs generated by instances are uploaded to LTS.</li><li>• Optimized <b>Viewing Error Logs</b>. Error logs can be reported to LTS.</li><li>• Optimized <b>Viewing Slow Query Logs</b>. Slow logs can be reported to LTS.</li><li>• Optimized <b>Rebooting a Node</b>.</li><li>• Optimized <b>Creating a Read Replica</b>. The description of the failover priority is added.</li></ul>
2022-12-20	<p>This issue is the twenty-seventh official release, which incorporates the following change:</p> <p>Optimized <b>Introducing GaussDB(for MySQL) Metrics</b>.</p>
2022-12-13	<p>This issue is the twenty-sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <b>Backup Principles</b>.</li><li>• Added <b>Backup Space and Billing</b>.</li><li>• Optimized <b>Create a Proxy Instance</b>. Constraints are added.</li><li>• Optimized <b>Introducing Read/Write Splitting</b>. Constraints are added.</li></ul>
2022-11-21	<p>This issue is the twenty-fifth official release, which incorporates the following change:</p> <p>Optimized <b>Introducing Read Replicas</b>. Constraints are added.</p>
2022-11-14	<p>This issue is the twenty-fourth official release, which incorporates the following change:</p> <p>Optimized <b>Enabling or Disabling SQL Explorer</b>. Constraints are added.</p>
2022-10-31	<p>This issue is the twenty-third official release, which incorporates the following change:</p> <p>Added <b>Configuring Transaction Splitting</b>.</p>

Released On	Description
2022-10-26	<p>This issue is the twenty-second official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Rebooting a Node</a>.</li><li>• Optimized <a href="#">Configuring a Same-Region Backup Policy</a>. The policy of automatically deleting full backup files is added.</li><li>• Optimized <a href="#">Upgrading a Minor Version</a>. Precautions are added.</li><li>• Optimized <a href="#">Changing vCPUs and Memory of a DB Instance</a>. Specifications can be reduced and expanded.</li></ul>
2022-07-30	<p>This issue is the twenty-first official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Restoring Table Data to a Specific Point in Time</a>.</li><li>• Added <a href="#">Database Management</a>.</li><li>• Added <a href="#">Account Management (Non-Administrator)</a>.</li><li>• Updated <a href="#">Changing Storage of a Yearly/Monthly DB Instance</a>. The storage of yearly/monthly instances can be reduced.</li><li>• Optimized the description in <a href="#">Changing Consistency Level</a>.</li><li>• Updated <a href="#">Viewing Slow Query Logs</a>. Slow logs can be reported to LTS.</li></ul>
2022-06-30	<p>This issue is the twentieth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Optimized <a href="#">Rebooting a DB Instance</a>. You can reboot your instance during the maintenance window.</li><li>• Optimized <a href="#">Restoring Instance Data to a Specific Point in Time</a>. Data can be restored to the original DB instance and an existing DB instance.</li><li>• Optimized <a href="#">Restoring Data to a DB Instance</a>. Data can be restored to the original DB instance and an existing DB instance.</li><li>• Billed for <a href="#">Monitoring by Seconds</a>.</li><li>• Optimized <a href="#">Viewing Slow Query Logs</a>. You can enable <a href="#">Show Original Log</a>.</li></ul>
2022-05-19	<p>This issue is the nineteenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added notice in <a href="#">Promoting a Read Replica to the Primary Node</a>.</li><li>• Added <a href="#">Introducing Consistency Levels</a>.</li><li>• Added <a href="#">Changing Consistency Level</a>.</li></ul>

Released On	Description
2022-04-21	This issue is the eighteenth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Added <a href="#">Changing the Specifications of a Proxy Instance</a>.</li><li>• Added <a href="#">Changing the IP Address of a Proxy Instance</a>.</li><li>• Updated <a href="#">Introducing GaussDB(for MySQL) Metrics</a>.</li><li>• Added <a href="#">Performance Tests</a>.</li></ul>
2022-03-04	This issue is the sixteenth official release, which incorporates the following change: Added <a href="#">Configuring and Changing a Private IP Address</a> .
2021-11-24	This issue is the fifteenth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Allowed you to <a href="#">enable multiple proxy instances</a>.</li><li>• Added <a href="#">Usage Rules</a>.</li></ul>
2021-10-31	This issue is the fourteenth official release, which incorporates the following changes: Added <a href="#">Unsubscribing a Yearly/Monthly Read Replica</a> .
2021-09-31	This issue is the thirteenth official release, which incorporates the following change: Added the metric Disk Usage.
2021-08-31	This issue is the twelfth official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Updated the units of supported monitoring metrics.</li><li>• Optimized the content for changing a security group.</li><li>• Added <a href="#">Configuring SSL</a>.</li><li>• Added <a href="#">Changing a DB Instance Description</a>.</li><li>• Added read replica promotion.</li><li>• Optimized the content of the section "Creating a GaussDB(for MySQL) Custom Policy".</li><li>• Added <a href="#">Enabling or Disabling SQL Explorer</a>.</li></ul>
2021-06-10	This issue is the eleventh official release, which incorporates the following change: Updated events supported by event monitoring.

Released On	Description
2021-05-31	<p>This issue is the tenth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Supported scaling down the vCPUs and memory of an instance.</li><li>• Added <a href="#">Changing the Billing Mode from Yearly/Monthly to Pay-per-Use</a>.</li><li>• Added <a href="#">Upgrading a Minor Version</a>.</li><li>• Added monitoring metrics in <a href="#">Introducing GaussDB(for MySQL) Metrics</a>.</li></ul>
2021-03-30	<p>This issue is the ninth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Optimized the overall structure of <i>Working with GaussDB(for MySQL)</i>.</li><li>• Optimized the content for viewing monitoring metrics of DB instances.</li><li>• Optimized constraints in <a href="#">Create a Proxy Instance</a>.</li><li>• Optimized constraints in <a href="#">Parameter Template Management</a>.</li></ul>
2020-12-30	<p>This issue is the eighth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Viewing Proxy Monitoring Metrics</a>.</li><li>• Added <a href="#">Creating Alarm Rules for a Proxy Instance</a>.</li><li>• Added <a href="#">Event Monitoring</a>.</li></ul>
2020-10-30	<p>This issue is the seventh official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Rebuilding a Deleted Instance from Recycle Bin</a>.</li><li>• Added <a href="#">Managing Tags</a>.</li></ul>
2020-08-30	<p>This issue is the sixth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Added <a href="#">Database Proxy (Read/Write Splitting)</a>.</li><li>• Added <a href="#">Exporting Backup Information</a>.</li><li>• Added <a href="#">Parameter Template Management</a>.</li></ul>
2020-07-30	<p>This issue is the fifth official release, which incorporates the following changes:</p> <ul style="list-style-type: none"><li>• Supported the setting of the maintenance window for GaussDB(for MySQL) instances.</li><li>• Added tags for GaussDB(for MySQL) metrics.</li></ul>

Released On	Description
2020-04-24	This issue is the fourth official release, which incorporates the following change: Updated the content related to billing management, database connection, database instance management, read replicas, database security, backup and restoration, and logs.
2020-03-13	This issue is the third official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Supported restoration to a specific point in time for GaussDB(for MySQL).</li><li>• Added three monitoring metrics for GaussDB(for MySQL): Replication Delay, Storage Write Delay, and Storage Read Delay.</li></ul>
2020-02-24	This issue is the second official release, which incorporates the following changes: <ul style="list-style-type: none"><li>• Supported renewal of yearly/monthly GaussDB(for MySQL) instances.</li><li>• Supported billing mode changes from pay-per-use to yearly/monthly for GaussDB(for MySQL) instances.</li><li>• Supported unsubscription of yearly/monthly GaussDB(for MySQL) instances.</li><li>• Supported storage scaling up for GaussDB(for MySQL) instances billed at a yearly/monthly basis.</li><li>• Supported the following operations on GaussDB(for MySQL) read replicas: creating and managing read replicas, promoting a read replica to the primary node, and delete pay-per-use read replicas.</li></ul>
2019-09-03	This issue is the first official release.