

TencentDB for MariaDB

Operation Guide

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by the Tencent corporate group, including its parent, subsidiaries and affiliated companies, as the case may be. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Operation Guide

Precautions

Notes

Compatibility of TencentDB for MariaDB with MySQL 5.6

Instance Management

Renaming Instances

Assigning Instances to Projects

Changing Instance Specifications

Upgrading Database Engine Versions

Adjusting Deployed Node

Restarting an Instance

Isolating and Terminating Instances

Disaster Recovery Read-Only Instance

Account Management

Creating Account

Modifying Account Permissions

Configuring Read/Write Separation

Cloning Account

Resetting Account Password

Deleting Account

Read/Write Separation

Changing Networks

Backup and Rollback

Backup Mode

Downloading Backup Files

Decompressing Backups and Logs

Restoring Instances from Backup Files

Rolling back Databases

Modifying Data Replication Mode

Migrating Data

Importing Data to TencentDB for MariaDB Instances with DTS

Importing Data with mysqldump

Database Sync(Legacy)

Overview

Database Sync Tool IP Range

Connectivity Test

Security Management

Notes on Information Security

CAM

Overview

Policy Structure

Resource-level Permissions Supported

Console Examples

CAM-enabled Operations

Security Group Configuration

Transparent Data Encryption (TDE)

Monitoring and Alarms

Monitoring Feature

Alarming Feature

Killing Threads

Parameter Templates and Settings

Database Audit

Enabling Database Audit

Viewing Audit Logs

Modifying Log Retention Period

Operation Guide

Precautions

Notes

Last updated : 2024-01-11 15:28:37

Storage engine

TencentDB for MariaDB currently supports the following storage engines. We recommend that you use InnoDB as other ones may compromise the performance. You can run the `SHOW ENGINES` command to view the storage engines supported by the current database:

Storage Engine	Supported by TencentDB	Reason for Default Disablement
InnoDB	Yes (default)	-
MyISAM	Yes (you need to submit a ticket to enable it)	It may cause data inconsistency and affect sync efficiency
Memory	Yes	-
Merge	Yes	-
Archive	Yes (you need to submit a ticket to enable it)	It may cause data inconsistency and affect sync efficiency
Federated	No	It may cause security risks
CSV	Yes (you need to submit a ticket to enable it)	It may cause data inconsistency and affect sync efficiency
BLACKHOLE	Yes	-
MRG_MyISAM	Yes	-
PERFORMANCE_SCHEMA	Yes (you need to submit a ticket to enable it)	It may cause data inconsistency and affect sync efficiency

Precautions on upgrading a TencentDB instance

During the upgrade of a TencentDB instance, the instance will be disconnected for 1–30 seconds (for switch after upgrade); therefore, you need to get prepared and enable automatic reconnection between your application and database in advance to avoid service unavailability.

Definitions of primary and replica in the documentation

In TencentDB documentation, a replica server refers to the hot backup database server in the high-availability architecture. When the primary server fails, the replica server will be put into use in real time for continuous service.

The architecture of "1-primary-2-replicas" is recommended if strong sync is required

When you perform strong sync replication, the primary database will be hanged if it is disconnected from the replica database or if the replica database fails. If there is only one primary or replica database, the high-availability solution will be unavailable, because if only one single server provides the service, part of the data will be lost completely when a failure occurs.

Enabling public network access for a long time and using weak passwords may lead to security risks

A public network IP of the database is likely to be detected and scanned by malicious users if it is enabled for a long period. If you use a weak password (such as 12345678 or 1234abcd) in this case, your database may be at great security risk.

Notes on TencentDB for MariaDB rollback

TencentDB for MariaDB supports data rollback, but we recommend that you back up your key production data before performing rollback.

Data will be directly rolled back into a new pay-as-you-go instance.

Deleting the original instance will not affect the rollback instance.

Notes on TencentDB locking policy

TencentDB has a locking mechanism. If the storage capacity usage of your instance exceeds the threshold (which is usually 103–130%; the read-only threshold for TencentDB for MariaDB is 110% currently), the system will lock your instance which will become read-only. Therefore, we recommend that you periodically check the storage capacity usage and enable SMS reminder for usage in the TencentDB for MariaDB console. If you cannot expand the instance capacity timely due to financial reasons, you can submit a ticket to apply for temporary unlocking for 1–3 business days.

TencentDB for MariaDB failover

TencentDB for MariaDB adopts high-availability modes such as one primary and one replica or one primary and two replicas. When the primary database fails, TencentDB for MariaDB can switch to the replica database within 1 second (200 ms in average). However, during the switch, there may be a short period of up to 30 seconds when you cannot access the database (this period of time is used for failure detection and data sync). In this case, you need to set automatic reconnection between your application and the instance in advance to avoid service unavailability. The switch is imperceptible to the business (i.e., the IP and port remain unchanged and no human intervention is required). You only need to ensure that automatic reconnection is enabled for your business.

Things to do after TencentDB instance purchase:

After purchasing a TencentDB instance, you don't need to perform basic Ops tasks, such as high-availability maintenance, backup, and security patch, but you should keep in mind the following:

Check whether the CPU, IOPS, space, and number of connections of your TencentDB instance are sufficient. If they become insufficient, you need to optimize or upgrade the instance.

Check whether your TencentDB instance has performance issues, many slow or poor SQL statements, and excessive or missing indexes.

You cannot modify any data in the `mysql` , `information_schema` , `performance_schema` , and `sysdb` databases.

You cannot directly use SQL statements to configure accounts or grant permissions, which can only be done in the console. The following 19 common permissions are supported, while some rarely used permissions are not supported:

SELECT, INSERT, UPDATE, DELETE, CREATE, DROP, REFERENCES, INDEX, ALTER
CREATE TEMPORARY TABLES, LOCK TABLES, EXECUTE, CREATE VIEW, SHOW VIEW
CREATE ROUTINE, ALTER ROUTINE, EVENT, TRIGGER, SHOW DATABASES

TencentDB for MariaDB does not provide the root account

We recommend that you use a public network address only for routine maintenance, rather than for connecting business servers

Compatibility of TencentDB for MariaDB with MySQL 5.6

Last updated : 2024-01-11 15:28:38

Compatibility of TencentDB for MariaDB with Open-Source MariaDB

TencentDB for MariaDB is fully compatible with open-source MariaDB.

Compatibility of TencentDB for MariaDB with MySQL 5.6

As TencentDB for MariaDB is highly compatible with MySQL 5.6, you can directly use it with little or no changes to existing code, programs, drivers, or tools that are used with MySQL.

Data files and table definition files are binary compatible.

All client APIs and protocols are compatible.

All filenames, binary files, paths, and port numbers are the same.

All connectors (including those for PHP, Perl, Python, Java, .NET, Ruby, and MySQL) can be used in TencentDB for MariaDB with no modification required.

You can use a MySQL client to connect to TencentDB for MariaDB.

Incompatibility of TencentDB for MariaDB with MySQL 5.6

1. GTID incompatibility

`GTID` of TencentDB for MariaDB is incompatible with that of MySQL 5.6, i.e., MySQL cannot be used as a replica database of TencentDB for MariaDB.

2. Different default binlog configurations

Binlogs in TencentDB for MariaDB are in row format, while in native MySQL 5.6 and MariaDB earlier than 10.2.3, they are in statement format by default.

3. Row-based or command-based replication of the `CREATE TABLE ... SELECT` command

To ensure that the `CREATE TABLE ... SELECT` command can work properly in both row-based and command-based replication, this command in TencentDB for MariaDB will be converted to and executed as the `CREATE OR`

`REPLACE` command in a replica database. The advantage of this mechanism is that the replica database can run properly after recovery from downtime.

3.1. Default value deduction

When you create tables by using the `Create table ... Select from` statement, the differences between the default values of fields in `varchar(N)` type are as follows:

The fields do not have a default value in MariaDB 10.1.

The default value in MySQL 5.7 is `NULL`.

The default value in MySQL 5.5 or 5.6 is an empty string (`''`).

Default value of a decimal column: In MySQL 5.5 and 5.6, it is deduced to 0.00; in MariaDB 10.1, it is deduced to `NULL`.

Sample:

```
----- MySQL 5.5 -----
create table t1
select least(_latin1'a',_latin2'b',_latin5'c' collate latin5_turkish_ci) as f1;
show create table t1;
Table      Create Table
t1  CREATE TABLE `t1` (
  `f1` varchar(1) CHARACTER SET latin5 NOT NULL DEFAULT ''
) ENGINE=MyISAM DEFAULT CHARSET=latin1
----- MySQL 5.7 -----
create table t1
select least(_latin1'a',_latin2'b',_latin5'c' collate latin5_turkish_ci) as f1;
show create table t1;
Table      Create Table
t1  CREATE TABLE `t1` (
  `f1` varchar(1) CHARACTER SET latin5 DEFAULT NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
----- MariaDB 10.1* -----
create table t1
select least(_latin1'a',_latin2'b',_latin5'c' collate latin5_turkish_ci) as f1;
show create table t1;
Table      Create Table
t1  CREATE TABLE `t1` (
  `f1` varchar(1) CHARACTER SET latin5 NOT NULL
) ENGINE=MyISAM DEFAULT CHARSET=latin1
```

3.2. Differences in processing `select` statements in subqueries

In this statement: `SELECT a AS x, ROW(11, 12) = (SELECT MAX(x), 12), ROW(11, 12) IN (SELECT MAX(x), 12) FROM t1;`

In MySQL 5.5 and 5.6, the subquery `SELECT MAX(x), 12` is considered as `SELECT MAX(x), 12 from t1` if it is located after `in`; it is considered as `SELECT x, 12` if it is located after `=`, where "x" is the alias of `a` in the current row.

In MySQL 5.7 and MariaDB 10.1.*, the subquery `SELECT MAX(x), 12` always equals `SELECT x, 12`, where `x` is the alias of `a` in the current row.

Sample:

```
----- MySQL 5.5/5.6 -----
CREATE TABLE t1 (a INT);
INSERT INTO t1 VALUES (1), (2), (11);
SELECT a AS x, ROW(11, 12) = (SELECT MAX(x), 12), ROW(11, 12) IN (SELECT
MAX(x), 12) FROM t1;
x    ROW(11, 12) = (SELECT MAX(x), 12)    ROW(11, 12) IN (SELECT MAX(x), 12)
1    0    1
2    0    1
11   1    1

----- MariaDB 10.1.* or MySQL 5.7-----
-----
CREATE TABLE t1 (a INT);
INSERT INTO t1 VALUES (1), (2), (11);
SELECT a AS x, ROW(11, 12) = (SELECT MAX(x), 12), ROW(11, 12) IN (SELECT
MAX(x), 12) FROM t1;
x    ROW(11, 12) = (SELECT MAX(x), 12)    ROW(11, 12) IN (SELECT MAX(x), 12)
1    0    0
2    0    0
11   1    1
```

3.3. Processing of NULL for ALL and SOME

In MySQL 5.5, the `NULL` in `10 >= ALL (NULL, 1, 10)` or `1 <= ALL (NULL, 1, 10)` is skipped, i.e., it is considered as non-existent, because `NULL` is incomparable.

In MySQL 5.7 and TencentDB for MariaDB, `NULL` is an unknown value and the result in the comparisons above is unknown too; therefore, `NULL` will be returned.

3.4. alter table inplace operation

When `alter table` is used to change the sequence of columns only, the `inplace` algorithm can be used in TencentDB for MariaDB but not in MySQL.

When `inplace alter table` is executed in TencentDB for MariaDB, the result of running `show create table t1` will be the same as that of running `ALGORITHM=COPY` in MySQL.

4. Undefined behavior in MySQL and TencentDB for MariaDB

Undefined behavior is a feature of behavior that can be implemented through any method in MySQL or TencentDB for MariaDB, which may vary by version without the need to notify users or be specified. Implementation of undefined behaviors by MySQL and TencentDB for MariaDB may produce the same or different results.

For such same or different results in the current and future versions, TencentDB for MariaDB does not guarantee the results or ensure the same kernel optimization. For more information, see [MariaDB versus MySQL: Compatibility](#).

4.1. Case-insensitive sorting of character-type columns

Sorting (`order by` clause) of character-type columns is generally case-insensitive, which means that the order of fields with the same content but different letter cases will be undefined after sorting. You can use the `BINARY` keyword to force implement case-sensitive sorting, i.e., `ORDER BY BINARY column name`.

Sample:

The sorting of the following samples in MySQL and TencentDB for MariaDB may be completely random:

```
mysql> SELECT email FROM t2 LEFT JOIN t1 ON kid = t2.id WHERE t1.id IS NULL
order by email;
```

```
+-----+
```

```
| email |
```

```
+-----+
```

```
| email |
```

```
| eMail |
```

```
| EMail |
```

```
+-----+
```

```
3 rows in set (0.00 sec)
```

4.2. Differences in processing `Auto_increment` field overflow

Undefined behavior specific to InnoDB:

In all auto-increment column lock modes (0, 1, 2), the auto-increment behavior will be considered undefined if you set the auto-increment column field to a negative value.

In all auto-increment column lock modes (0, 1, 2), the auto-increment behavior will be considered undefined if the value of the auto-increment column field is greater than the maximum integer that can be stored as integer type in the column.

Note:

Do not insert (incorrect) numbers into an auto-increment column.

4.3. Differences in statistics collection method for `FOUND_ROWS`

The returned value of `FOUND_ROWS()` will be accurate only when `UNION ALL` is used in the query statement.

If only `UNION` is used without `ALL`, TencentDB for MariaDB will remove duplicates in the statistics, while MySQL will retain duplicates. If the `UNION` query statement is used without the `LIMIT` clause, the

`SQL_CALC_FOUND_ROWS` keyword will be ignored, and the returned result of `FOUND_ROWS()` will be the number of rows in the temporary table created when `UNION` is executed.

4.4. Differences in locking sequence of the `LOCK TABLES` statement

The `LOCK TABLES` statement locks tables in the following method: first, all tables that need to be locked are sorted based on the internally-defined method; however, from user's perspective, the sorting order in MySQL and TencentDB for MariaDB is undefined. For example, if you write `LOCK TABLES t1, t2, t3`, TencentDB for MariaDB and MySQL will not lock the tables according to the sequence of `t1, t2, t3`.

This behavior is undefined in MySQL and TencentDB for MariaDB; therefore, they may use different methods to sort `t1, t2`, and `t3` and lock them based on the resulting sequence.

Therefore, you should not rely on locking sequence to ensure accuracy in your procedures or query code, as this may cause deadlock.

4.5. Timing for running the `RESET MASTER` statement

You cannot run `RESET MASTER` when any duplicate replica server is running; otherwise, the behaviors of primary and replica servers will be undefined (and not supported) in TencentDB for MariaDB and MySQL. Various predictable errors may or may not occur during execution of `RESET MASTER`. The official development teams of TencentDB for MariaDB and MySQL do not consider these errors as bugs and are not responsible for any errors that actually occur in this way.

4.6. Conversion of date and time types to year type

In MySQL 5.5, when variables in year and date types are compared, the date type will be converted to the year type. For example, "2011-01-01" will be converted to "2011".

In MySQL 5.7 and TencentDB for MariaDB, variables in date type will stay unchanged; therefore, comparisons with variables in year type are different.

Similarly, TencentDB for MariaDB cannot convert the time type to year type, while MySQL 5.6 uses the year part in the timestamp of the current session as the year value for every value in time type, which means that the year in the timestamp of the current session will be used every time a time-type value needs to be converted to year type.

4.7. Processing method of unknown characters

Character encoding conversion is implemented in different ways in different versions of MySQL and TencentDB for MariaDB. For example, if an encoding byte string is not recognized by `unhex`, an empty string ("") will be returned in MySQL 5.5/5.6/5.7, while a question mark character (?) will be returned in MariaDB 10.1.

The `UPDATE t1 SET a=unhex(code) ORDER BY code` statement assigns a value to field `a` in table `t1`; however, some of the assignment operations will fail as `unhex` can only recognize and convert byte strings within a certain range.

The default storage engine in MySQL 5.5 is MyISAM, which does not support transactions. Therefore, the statement will exit when it fails to assign a value to `a` in any row in `t1`; however, all assigned values will be still stored in `t1`.

The default storage engine in MySQL 5.7 is InnoDB; therefore, the transaction will be rolled back when the statement fails to assign a value to `a` in any row in `t1`, i.e., all assigned values will be rolled back as well.

The default storage engine in TencentDB for MariaDB is InnoDB. When `unhex` is unable to find a corresponding character for a byte string, a question mark (`?`, i.e., `0x3F`) will be returned; therefore, the operation will always succeed no matter whether the storage engine is InnoDB or MyISAM.

When inserting a hexadecimal string using the `insert into` statement, if the corresponding `utf8mb4` character is not found:

If the HEAP storage engine is used in MySQL 5.5/5.6, this unknown character will be ignored.

MariaDB 10.1 and MySQL 5.7 will use `0x3F` (i.e., question mark `"?"`) to replace this character.

For an invalidly encoded string field, MySQL with the InnoDB storage engine will directly return an error, while TencentDB for MariaDB will replace the field with `3F` and then insert it.

4.8. Precision of time type

```
SELECT CAST(CAST('10:11:12.098700' AS TIME) AS DECIMAL(20,6));
CAST(CAST('10:11:12.098700' AS TIME) AS DECIMAL(20,6))
...
```

When the above statements are used, the ways to process them are different in MySQL 5.5/5.6 and MySQL 5.7/MariaDB 10.1:

- In MySQL 5.5/5.6, `101112.098700` will be returned and the precision will still be retained.
- In MySQL 5.7 and MariaDB 10.1, `101112.000000` will be returned. This is because the statement does not specify a precision for ``TIME`` and the default precision of ``TIME`` is 0; therefore, the values after the decimal point of ``CAST('10:11:12.098700' AS TIME)`` will be lost.

You can use the following statement to keep the time precision.

```
SELECT CAST(CAST('10:11:12.098700' AS TIME(6)) AS DECIMAL(20,6));
```

```
+-----+
| CAST(CAST('10:11:12.098700' AS TIME(6)) AS DECIMAL(20,6)) |
+-----+
| 101112.098700 |
+-----+
```

>?The default precision of ``TIME`` is not consistent. If time precision is required, you should specify a precision of time for upgrade or migration.

```
CREATE TABLE t1(f1 TIME);
INSERT INTO t1 VALUES ('23:38:57');
SELECT TIMESTAMP(f1,'1') FROM t1;
```

In MySQL 5.5/5.6, `NULL` will be returned; in MariaDB 10.1 and MySQL 5.7, `2016-08-03 23:38:58` will be returned.

- If the first parameter of `TIMESTAMP()` is in time type, the returned value will be `NULL` as MySQL 5.5 cannot automatically convert it to timestamp type.
- In MySQL 5.7 and TencentDB for MariaDB, values in time type will be automatically converted to timestamp type, i.e., current date + entered time variable.

5. Appendix: TencentDB for MariaDB parameters and MySQL parameters

5.1. Different parameters with the same variable name

Parameters with the same variable name have the same main feature.

Parameter	MariaDB 10.1	MySQL 5.6
old_passwords	OFF	0
tmpdir	/tmp/5cXm2hHsWi/mysqld.1	/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-test/var/tmp/mysqld.1
version	10.1.9-MariaDB-log	5.6.31-log
slow_query_log_file	/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1/build_dongzhi/mysql-test/var/mysqld.1/mysqld-slow.log	/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-test/var/mysqld.1/mysqld-slow.log
table_definition_cache	400	1400
datadir	/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1/build_dongzhi/mysql-test/var/mysqld.1/data/	/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-test/var/mysqld.1/data/
pid_file		

```

<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/build_dongzhi/mysql-test/var/run/mysqld.1.pid</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-
test/var/run/mysqld.1.pid</td></tr>
<tr>
<td>max_seeks_for_key</td>
<td>4294967295</td>
<td>18446744073709500000</td></tr>
<tr>
<td>slave_load_tmpdir</td>
<td>/tmp/5cXm2hHsWi/mysqld.1</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-
test/var/tmp/mysqld.1</td></tr>
<tr>
<td>secure_file_priv</td>
<td>    /data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/build_dongzhi/mysql-test/var/</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-
test/var/</td></tr>
<tr>
<td>sql_mode</td>
<td>NO_AUTO_CREATE_USER,NO_ENGINE_SUBSTITUTION    </td>
<td>NO_ENGINE_SUBSTITUTION</td></tr>
<tr>
<td>ssl_cert</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1/mysql-
test/std_data/server-cert.pem</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/mysql-
test/std_data/server-cert.pem</td></tr>
<tr>
<td>ssl_ca</td>
<td>    /data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1/mysql-
test/std_data/cacert.pem</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/mysql-
test/std_data/cacert.pem</td></tr>
<tr>
<td>open_files_limit</td>
<td>1024</td>
<td>4161</td></tr>
<tr>
<td>binlog_checksum</td>
<td>NONE</td>
<td>CRC32</td></tr>
<tr>
<td>basedir</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6</td></tr>

```

```

<tr>
<td>query_alloc_block_size</td>
<td>16384</td>
<td>8192</td></tr>
<tr>
<td>innodb_max_dirty_pages_pct</td>
<td>75.000000</td>
<td>75</td></tr>
<tr>
<td>ssl_key</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-release1/mysql-
test/std_data/server-key.pem</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/mysql-
test/std_data/server-key.pem</td></tr>
<tr>
<td>myisam_sort_buffer_size</td>
<td>134216704</td>
<td>8388608</td></tr>
<tr>
<td>skip_name_resolve</td>
<td>ON</td>
<td>OFF</td></tr>
<tr>
<td>pseudo_thread_id</td>
<td>3</td>
<td>2</td></tr>
<tr>
<td>character_sets_dir</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/sql/share/charsets/</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/sql/share/charsets/</td>
</tr>
<tr>
<td>innodb_adaptive_flushing_lwm</td>
<td>10</td>
<td>10</td></tr>
<tr>
<td>myisam_recover_options</td>
<td>DEFAULT</td>
<td>OFF</td></tr>
<tr>
<td>performance_schema_max_statement_classes</td>
<td>179</td>
<td>168</td></tr>
<tr>
<td>innodb_version</td>
<td>5.6.26-74.0</td>

```



```

<td>5.6.31</td></tr>
<tr>
<td>max_write_lock_count</td>
<td>4294967295</td>
<td>18446744073709500000</td></tr>
<tr>
<td>thread_cache_size</td>
<td>0</td>
<td>9</td></tr>
<tr>
<td>innodb_checksum_algorithm</td>
<td>INNODB</td>
<td>innodb</td></tr>
<tr>
<td>optimizer_switch</td>
<td>
index_merge=on,<br>index_merge_union=on,<br>index_merge_sort_union=on,
<br>index_merge_intersection=on,<br>index_merge_sort_intersection=off,
<br>engine_condition_pushdown=off,<br>index_condition_pushdown=on,
<br>derived_merge=on,<br>derived_with_keys=on,<br>firstmatch=on,
<br>loosescan=on,<br>materialization=on,<br>in_to_exists=on,<br>semijoin=on,
<br>partial_match_rowid_merge=on,<br>partial_match_table_scan=on,
<br>subquery_cache=on,<br>mrr=off,<br>mrr_cost_based=off,<br>mrr_sort_keys=off,
<br>outer_join_with_cache=on,<br>semijoin_with_cache=on,
<br>join_cache_incremental=on,<br>join_cache_hashed=on,<br>join_cache_bka=on,
<br>optimize_join_buffer_size=off,<br>table_elimination=on,
<br>extended_keys=on,<br>exists_to_in=on</td>
<td>
index_merge=on,<br>index_merge_union=on,<br>index_merge_sort_union=on,
<br>index_merge_intersection=on,<br>engine_condition_pushdown=on,
<br>index_condition_pushdown=on,<br>mrr=on,<br>mrr_cost_based=on,
<br>block_nested_loop=on,<br>batched_key_access=off,<br>materialization=on,
<br>semijoin=on,<br>loosescan=on,<br>firstmatch=on,
<br>subquery_materialization_cost_based=on,<br>use_index_extensions=on
</td></tr>
<tr>
<td>timestamp</td>
<td>1471938276</td>
<td>1471937901</td></tr>
<tr>
<td>general_log_file</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/build_dongzhi/mysql-test/var/mysqld.1/mysqld.log</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-
test/var/mysqld.1/mysqld.log</td></tr>
<tr>
<td>myisam_stats_method</td>

```

```

<td>NULLS_UNEQUAL</td>
<td>>nulls_unequal</td></tr>
<tr>
<td>innodb_log_compressed_pages</td>
<td>OFF</td>
<td>ON</td></tr>
<tr>
<td>query_prealloc_size</td>
<td>24576</td>
<td>0</td></tr>
<tr>
<td>rand_seed2</td>
<td>297895171</td>
<td>0</td></tr>
<tr>
<td>rand_seed1</td>
<td>605568929</td>
<td>0</td></tr>
<tr>
<td>socket</td>
<td>/tmp/5cXm2hHsWi/mysql.1.sock</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/build_dongzhi/mysql-
test/var/tmp/mysql.1.sock</td></tr>
<tr>
<td>innodb_max_dirty_pages_pct_lwm</td>
<td>0.001</td>
<td>0</td></tr>
<tr>
<td>lc_messages_dir</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/build_dongzhi/sql/share/</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-
5.6/build_dongzhi/sql/share/</td></tr>
<tr>
<td>max_relay_log_size</td>
<td>1073741824</td>
<td>0</td></tr>
<tr>
<td>plugin_dir</td>
<td>/data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/lib/plugin/</td>
<td>/data/home/tdengine/dongzhi/src/mysql-server-5.6/lib/plugin/</td></tr>
<tr>
<td>thread_stack</td>
<td>294912</td>
<td>262144</td></tr>
</table>

```

5.2. Variables unique to TencentDB for MariaDB

```
- aria_block_size      8192
- aria_checkpoint_interval      30
- aria_checkpoint_log_activity  1048576
- aria_encrypt_tables      OFF
- aria_force_start_after_recovery_failures      0
- aria_group_commit      none
- aria_group_commit_interval      0
- aria_log_file_size      1073741824
- aria_log_purge_type      immediate
- aria_max_sort_file_size      9223372036853727232
- aria_page_checksum      ON
- aria_pagecache_age_threshold      300
- aria_pagecache_buffer_size      134217728
- aria_pagecache_division_limit      100
- aria_pagecache_file_hash_size      512
- aria_recover      NORMAL
- aria_repair_threads      1
- aria_sort_buffer_size      268434432
- aria_stats_method      nulls_unequal
- aria_sync_log_dir      NEWFILE
- aria_used_for_temp_tables      ON
- autoremoverelaylog      ON
- binlog_annotate_row_events      OFF
- binlog_commit_wait_count      0
- binlog_commit_wait_usec      100000
- binlog_optimize_thread_scheduling      ON
- deadlock_search_depth_long      15
- deadlock_search_depth_short      4
- deadlock_timeout_long      50000000
- deadlock_timeout_short      10000
- debug_no_thread_alarm      OFF
- default_master_connection
- default_regex_flags
- encrypt_binlog      OFF
- encrypt_tmp_disk_tables      OFF
- encrypt_tmp_files      OFF
- enforce_storage_engine
- expensive_subquery_limit      100
- extra_max_connections      20
- extra_port      0
- flush_relay_logs_for_strong_consistency      ON
- gtid_binlog_pos
- gtid_binlog_state
- gtid_current_pos
- gtid_domain_id      0
```

```
- gtid_ignore_duplicates      OFF
- gtid_seq_no                0
- gtid_slave_pos
- gtid_strict_mode           OFF
- histogram_size             0
- histogram_type             SINGLE_PREC_HB
- in_transaction             0
- innodb_adaptive_hash_index_partitions 1
- innodb_background_scrub_data_check_interval 3600
- innodb_background_scrub_data_compressed OFF
- innodb_background_scrub_data_interval 604800
- innodb_background_scrub_data_uncompressed OFF
- innodb_buf_dump_status_frequency 0
- innodb_buffer_pool_populate OFF
- innodb_cleaner_lsn_age_factor HIGH_CHECKPOINT
- innodb_compression_algorithm none
- innodb_corrupt_table_action assert
- innodb_default_encryption_key_id 1
- innodb_defragment          OFF
- innodb_defragment_fill_factor 0.900000
- innodb_defragment_fill_factor_n_recs 20
- innodb_defragment_frequency 40
- innodb_defragment_n_pages 7
- innodb_defragment_stats_accuracy 0
- innodb_disallow_writes      OFF
- innodb_empty_free_list_algorithm BACKOFF
- innodb_encrypt_log          OFF
- innodb_encrypt_tables       OFF
- innodb_encryption_rotate_key_age 1
- innodb_encryption_rotation_iops 100
- innodb_encryption_threads 0
- innodb_fake_changes         OFF
- innodb_fatal_semaphore_wait_threshold 600
- innodb_force_primary_key     OFF
- innodb_foreground_preflush   EXPONENTIAL_BACKOFF
- innodb_idle_flush_pct       100
- innodb_immediate_scrub_data_uncompressed OFF
- innodb_instrument_semaphores OFF
- innodb_kill_idle_transaction 0
- innodb_locking_fake_changes ON
- innodb_log_arch_dir         ./
- innodb_log_arch_expire_sec 0
- innodb_log_archive          OFF
- innodb_log_block_size       512
- innodb_log_checksum_algorithm INNODB
- innodb_max_bitmap_file_size 104857600
- innodb_max_changed_pages    1000000
```

```

- innodb_mtflush_threads      8
- innodb_prefix_index_cluster_optimization    OFF
- innodb_sched_priority_cleaner      19
- innodb_scrub_log      OFF
- innodb_scrub_log_speed      256
- innodb_show_locks_held      10
- innodb_show_verbose_locks      0
- innodb_simulate_comp_failures      0
- innodb_stats_modified_counter      0
- innodb_stats_traditional      ON
- innodb_track_changed_pages      OFF
- innodb_use_atomic_writes      OFF
- innodb_use_fallocate      OFF
- innodb_use_global_flush_log_at_trx_commit    ON
- innodb_use_mtflush      OFF
- innodb_use_stacktrace      OFF
- innodb_use_trim      OFF
- join_buffer_space_limit      2097152
- join_cache_level      2
- key_cache_file_hash_size      512
- key_cache_segments      0
- last_gtid
- log_slow_filter
admin,filesort,filesort_on_disk,full_join,full_scan,query_cache,query_cache_miss,tmp_table,tmp_table_on_disk
- log_slow_rate_limit      1
- log_slow_verbosity
- log_tc_size      24576
- loglevel      3
- max_long_data_size      4194304
- max_statement_time      0.000000
- mrr_buffer_size      262144
- myisam_block_size      1024
- mysql56_temporal_format      ON
- old_mode
- optimizer_selectivity_sampling_limit      100
- optimizer_use_condition_selectivity      1
- plugin_maturity      unknown
- progress_report_time      5
- query_cache_strip_comments      OFF
- relay_log_sync_threshold      134217728
- relay_log_sync_timeout      200
- relay_log_sync_txn_count      5
- replicate_annotate_row_events      OFF
- replicate_do_db
- replicate_do_table
- replicate_events_marked_for_skip      REPLICATE

```

```
- replicate_ignore_db
- replicate_ignore_table
- replicate_wild_do_table
- replicate_wild_ignore_table
- rowid_merge_buff_size      8388608
- rpl_semi_sync_master_enabled      OFF
- rpl_semi_sync_master_timeout      10000
- rpl_semi_sync_master_trace_level  32
- rpl_semi_sync_master_wait_no_slave      ON
- rpl_semi_sync_master_wait_point      AFTER_COMMIT
- rpl_semi_sync_slave_enabled      OFF
- rpl_semi_sync_slave_trace_level  32
- skip_parallel_replication      OFF
- skip_replication      OFF
- slave_current_parallel_transactions      0
- slave_ddl_exec_mode      IDEMPOTENT
- slave_domain_parallel_threads      0
- slave_max_parallel_transactions      0
- slave_parallel_max_queued      131072
- slave_parallel_mode      conservative
- slave_parallel_threads      0
- slave_run_triggers_for_rbr      NO
- sqlasyn      OFF
- sqlasyntimeout      10
- sqlasynwarntimeout      3
- strict_password_validation      ON
- thread_pool_high_prio_mode      transactions
- thread_pool_high_prio_tickets      4294967295
- thread_pool_idle_timeout      60
- thread_pool_max_threads      1000
- thread_pool_oversubscribe      3
- thread_pool_oversubscribe_parall      1
- thread_pool_size      8
- thread_pool_stall_limit      500
- use_stat_tables      NEVER
- userstat      OFF
- version_malloc_library      system
- version_ssl_library      OpenSSL 1.0.2d 9 Jul 2015
- wsrep_auto_increment_control      ON
- wsrep_causal_reads      OFF
- wsrep_certify_nonpk      ON
- wsrep_cluster_address
- wsrep_cluster_name      my_wsrep_cluster
- wsrep_convert_lock_to_trx      OFF
- wsrep_data_home_dir      /data/home/tdengine/dongzhi/src/tdsql-mariadb-10.1.9-
release1/build_dongzhi/mysql-test/var/mysqld.1/data/
- wsrep_debug_option
```

```

- wsrep_debug            OFF
- wsrep_desync           OFF
- wsrep_dirty_reads      OFF
- wsrep_drupal_282555_workaround  OFF
- wsrep_forced_binlog_format  NONE
- wsrep_gtid_domain_id   0
- wsrep_gtid_mode        OFF
- wsrep_load_data_splitting  ON
- wsrep_log_conflicts    OFF
- wsrep_max_ws_rows      131072
- wsrep_max_ws_size      1073741824
- wsrep_mysql_replication_bundle  0
- wsrep_node_address
- wsrep_node_incoming_address  AUTO
- wsrep_node_name
- wsrep_notify_cmd
- wsrep_on               OFF
- wsrep_osu_method        TOI
- wsrep_patch_version     wsrep_25.11
- wsrep_provider          none
- wsrep_provider_options
- wsrep_recover           OFF
- wsrep_replicate_myisam  OFF
- wsrep_restart_slave     OFF
- wsrep_retry_autocommit  1
- wsrep_slave_fk_checks   ON
- wsrep_slave_threads     1
- wsrep_slave_uk_checks   OFF
- wsrep_sst_auth
- wsrep_sst_donor
- wsrep_sst_donor_rejects_queries  OFF
- wsrep_sst_method        rsync
- wsrep_sst_receive_address  AUTO
- wsrep_start_position    00000000-0000-0000-0000-000000000000:-1
- wsrep_sync_wait         0

```

5.3. Variables unique to MySQL 5.6

```

- avoid_temporal_upgrade  OFF
- bind_address            *
- binlog_error_action      IGNORE_ERROR
- binlog_gtid_simple_recovery  OFF
- binlog_max_flush_queue_time  0
- binlog_order_commits    ON
- binlog_rows_query_log_events  OFF
- binlogging_impossible_mode  IGNORE_ERROR
- block_encryption_mode    aes-128-ecb

```

```

- core_file          ON
- disconnect_on_expired_password      ON
- end_markers_in_json      OFF
- enforce_gtid_consistency      OFF
- eq_range_index_dive_limit      1
- gtid_executed
- gtid_mode          OFF
- gtid_next          AUTOMATIC
- gtid_owned
- gtid_purged
- innodb_tmpdir
- log_bin_use_v1_row_events      OFF
- log_slow_admin_statements      OFF
- log_slow_slave_statements      OFF
- log_throttle_queries_not_using_indexes      0
- master_info_repository      FILE
- new          OFF
- optimizer_trace      enabled=off,one_line=off
- optimizer_trace_features
greedy_search=on,range_optimizer=on,dynamic_range=on,repeated_subselect=on
- optimizer_trace_limit      1
- optimizer_trace_max_mem_size      16384
- optimizer_trace_offset      -1
- relay_log_info_repository      FILE
- rpl_stop_slave_timeout      31536000
- server_id_bits      32
- server_uuid      9078a55d-6904-11e6-bfa9-ecf4bbcdc829
- sha256_password_private_key_path      private_key.pem
- sha256_password_public_key_path      public_key.pem
- show_old_temporals      OFF
- simplified_binlog_gtid_recovery      OFF
- slave_allow_batching      OFF
- slave_checkpoint_group      512
- slave_checkpoint_period      300
- slave_parallel_workers      0
- slave_pending_jobs_size_max      16777216
- slave_rows_search_algorithms      TABLE_SCAN, INDEX_SCAN
- table_open_cache_instances      1
- transaction_allow_batching      OFF

```


Instance Management

Renaming Instances

Last updated : 2024-01-11 15:28:37

Overview

This document describes how to rename a database instance in the TencentDB for MariaDB console.

Note:

Renaming an instance does not change the private IP of the database or affect database connections.

After the instance is renamed, its project and network remain unchanged.

If an instance is in another task flow (such as upgrade or initialization), it cannot be renamed.

Directions

1. Log in to the [TencentDB for MariaDB console](#), locate an instance in the instance list, and click the



icon next to its name. You can also click an instance name/ID in the instance list to access the instance details page, and click the



icon next to the instance name in the **Basic Info** section.

2. In the pop-up window, modify the instance name and click **OK**.

Note:

An instance cannot be renamed to an existing database instance name.

Modify instance name×

Instance Name *

✓

Up to 60 chars, including Chinese characters, letters, digits, dashes, underscores, and periods

Confirm

Cancel

Assigning Instances to Projects

Last updated : 2024-01-11 15:28:37

TencentDB for MariaDB supports assigning instances to different projects for management.

Project is a resource assignment method defined by Tencent Cloud for teams. You can assign different resources to different teams based on your organizational structure, and this assignment method is called project in Tencent Cloud. Read-only replicas and disaster recovery instances are the associated instances of the primary instance and should be in the same project as the primary instance.

Assigning and moving database instances across projects will not affect the services provided by the instances.

Users need to specify projects to which the instances belong when purchasing them. The default project is "Default Project".

Assigned instances can be reassigned to other projects through the **Switch to another project** feature in the [console](#).

The screenshot displays the Tencent Cloud console interface for a TencentDB for MariaDB instance. The top navigation bar includes a 'Switch to another project' button, which is highlighted with a red box. Below the navigation bar, the 'Instance Details' page is shown, featuring a sidebar with 'Basic Info' and a main area with an 'Instance Architecture Diagram'.

Basic Info

- Instance Name: tdsq-kbsrcfaj
- Instance ID: tdsq-kbsrcfaj
- Running Status: Running
- Instance Type: Primary Instance
- Instance Version: Standard Edition (1 primary-1 replica)
- Region: South China (Guangzhou)
- Private IPv4 Address: [Redacted] Security Group
- Private Port: 3306
- Public IPv4 Address: Enable
- Private IPv6 Address: --
- Public IPv6 Address: This instance
- Network: Classic Net
- Project: DEFAULT PROJECT (with a red box around the project name and a 'Modify Project' button next to it)
- Character Set: UTF8
- Tag: --

Instance Architecture Diagram

The diagram shows the instance architecture for South China (Guangzhou). It includes a primary instance (tdsq-kbsrcfaj) and a replica server (tdsq-dtjh2urr). The primary instance is connected to a proxy, which is connected to the replica server. The replica server is connected to a Cold Standby Center. The diagram also shows a disaster recovery instance (tdsq-dtjh2urr) connected to a Cold Standby Center. The diagram is labeled 'South China (Guangzhou)' and 'Guangzhou Zone 3'.

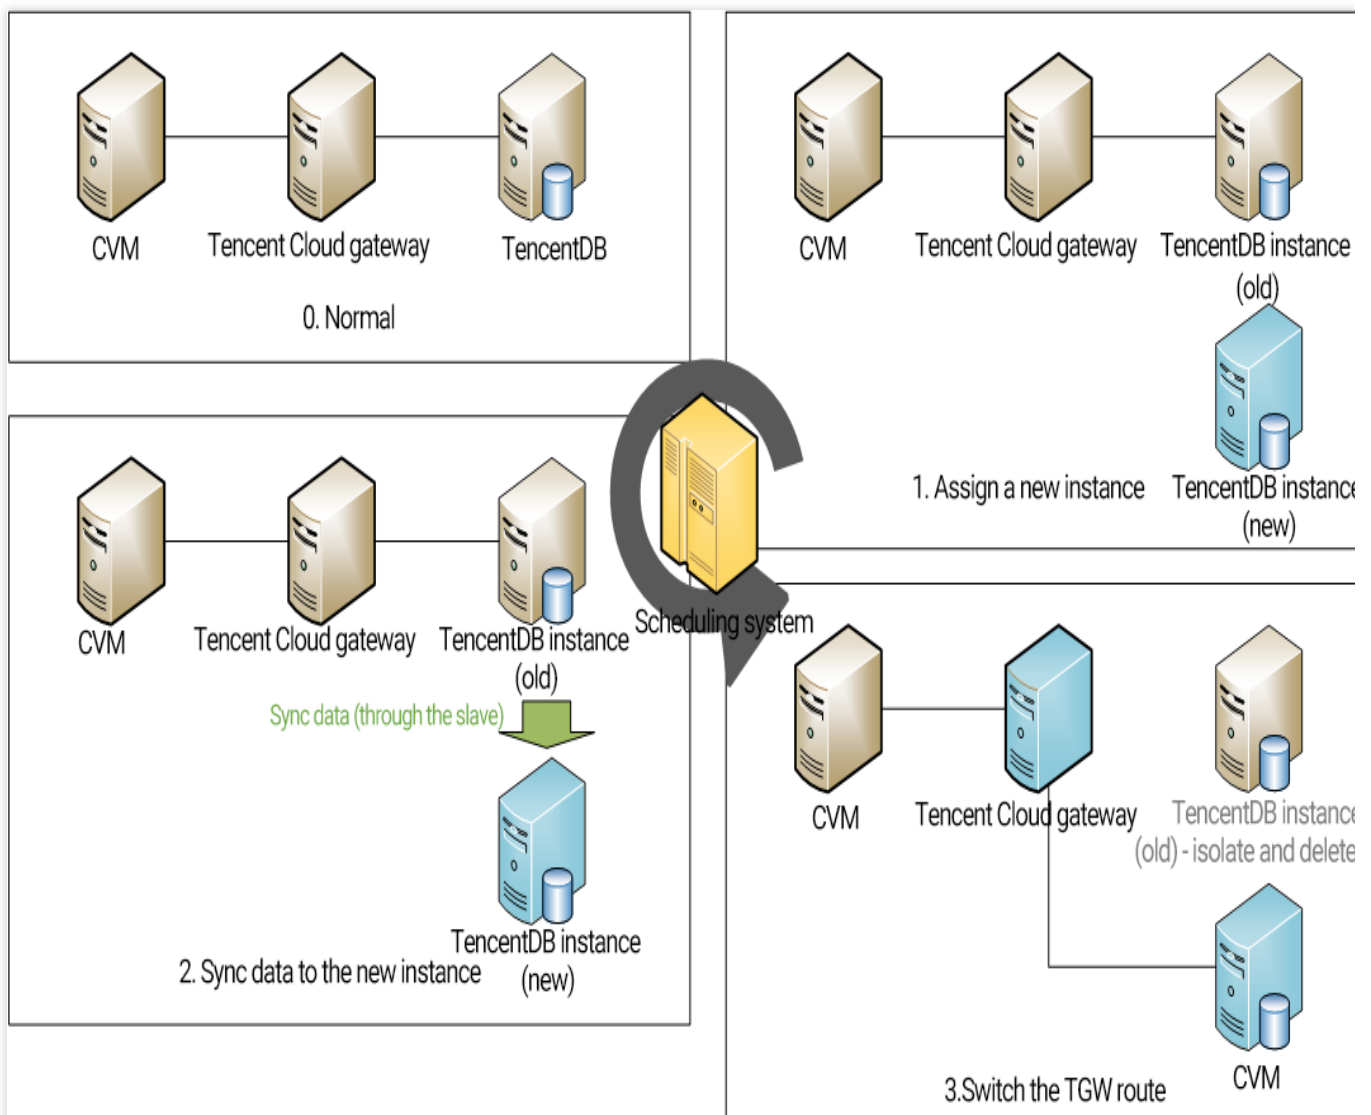
Changing Instance Specifications

Last updated : 2024-01-11 15:28:38

How the configuration adjustment works

After you click **Adjusted Configuration** in the console, the Tencent Cloud Ops management system will generally perform the following steps:

1. Assign a new instance (the "new instance") based on the required configuration.
2. Sync the data and configuration of the instance to be adjusted (the "old instance") to the new instance.
3. After the sync is completed, switch the route in the Tencent Cloud gateway to the new instance for continued use.



Changing instance specifications

Note:

You can still use the old instance as usual during the adjustment, such as importing or exporting data.

The name, access IP, and access port of the instance remain unchanged after the adjustment.

When the adjustment is completed, the database will be disconnected for several seconds. We recommend that you implement an automatic reconnection feature in your program.

During the adjustment, try avoiding operations such as modifying global parameters, instance name, or user password of the database.

1. Log in to the [TencentDB for MariaDB console](#), find the target instance in the instance list, and select **More > Adjust Configurations** in the **Operation** column.
2. On the **Adjust Configurations** page, select the target specification, disk capacity, and switch time and click **Adjust Configurations**.

TencentDB for MariaDB supports scheduled switch for configuration adjustment, which allows you to switch the database to its new specification at a specified time (usually during off-peak hours).

Note:

Scheduled switch

: You can choose to switch the database to its new configuration at a specified time, which is usually during off-peak hours and must be within 72 hours.

Generally, the switch time has a deviation of about 15 minutes, as there may be high amounts of write requests to large transactions, which will affect the data sync progress. In this case, the system will first guarantee sync between the new and old instances instead of performing the scheduled switch.

To ensure a successful switch, you can select the option for retry upon failure, and the system will try switching again two hours after a switch failure.

Configuration adjustment fees

If you upgrade a database instance, the price difference between original and upgraded specifications is deducted from your account. If the account balance is insufficient, you need to top it up. The upgraded instance will be billed by the new specification.

Upgrading Database Engine Versions

Last updated : 2024-01-11 15:28:37

Upgrading the Major Version of the Database Engine

Upgrade of the major version of the database engine refers to such upgrade as upgrading MariaDB 10.0 to 10.1. Currently, major version upgrade is not supported. If needed, you are recommended to purchase an instance on the new version. After testing is completed, migrate data from the old instance to the new one and switch the business system.

Upgrading the Version of the Database Cluster Module

Upgrading feature modules such as Agent and Proxy in the cluster to their latest version does not involve business system compatibility issues. If needed, please [submit a ticket](#) for application and specify "instance region, instance ID, upgrade requirement, and scheduled switch time". The upgrade in this case works in the same way as upgrading instance specification. For more information, please see [How Upgrade Works](#).

Adjusting Deployed Node

Last updated : 2024-01-11 15:28:37

This document describes how to adjust deployed nodes in the TencentDB for MariaDB console. You can add replica nodes to enjoy cross-region replica support, reduce the execution pressure, and increase the read speed. You can also remove unnecessary replica nodes to reduce the redundant performance costs during idle hours.

Note:

You can still use the old instance as usual during the adjustment.

The name, access IP, and access port of an instance will remain the same after the adjustment; however, the SQL passthrough ID (Setid) will change.

When the adjustment is completed, the database will be disconnected for several seconds. We recommend that you implement an automatic reconnection feature in your program.

During the adjustment, try avoiding operations such as modifying global parameters, instance name, or user password of the database.

Adjusting the node deployment region

1. Log in to the [TencentDB for MariaDB console](#) and click the target instance ID in the instance list to enter the instance details page.
2. In **Availability Info > Deployment Mode** on the **Instance Details** page, click **Change Deployment Mode**.
3. On the **Change Deployment Mode** page, select the target deployment mode, and select the regions of the primary and replica nodes in the drop-down lists.

Note:

Target Deployment Mode: You can select **Single-AZ** or **Multi-AZ**. In single-AZ mode, the region of replica nodes must be the same as that of the primary node. In multi-AZ mode, replica nodes can be in any regions.

Adding/Removing replica nodes

1. Log in to the [TencentDB for MariaDB console](#) and click the target instance ID in the instance list to enter the instance details page.
2. In **Availability Info > Deployment Mode** on the **Instance Details** page, click **Change Deployment Mode**.
3. On the **Change Deployment Mode** page, click **Add Replica Node** to add up to five replica nodes.

Note:

Delete: Click it to remove existing replica nodes. If there is only one replica node, it cannot be removed.

Scheduled switch: You can choose to switch the database to its new configuration at a specified time, which is usually during off-peak hours and must be within 72 hours.

Generally, the switch time has a deviation of about 15 minutes, as there may be high amounts of write requests to large transactions, which will affect the data sync progress. In this case, the system will first guarantee sync between the new and old instances instead of performing the scheduled switch.

To ensure a successful switch, you can select the option for retry upon failure, and the system will try switching again two hours after a switch failure.

Restarting an Instance

Last updated : 2024-01-11 15:28:38

This document describes how to restart an instance in the console.

Overview

Instance restart is a common maintenance method for TencentDB for MariaDB. It is similar to restarting a local database.

Notes

Preparation for restart: during the restart, the instance cannot provide services. Therefore, before the restart, please ensure that TencentDB for MariaDB has stopped accepting business requests. During the restart, dirty pages will be generated if the business write volume is high. In this case, the restart may fail in order to shorten the business interruption.

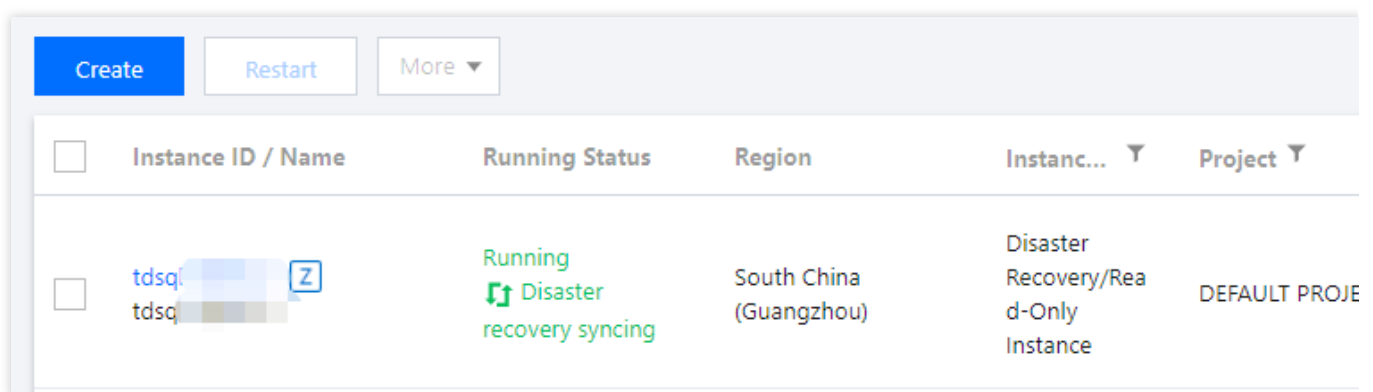
Restart method: you are recommended to restart an instance by following the steps provided by Tencent Cloud instead of running the restart command on the instance.

Restart time: generally, it takes only a few minutes to restart an instance.

Physical instance features: restarting an instance does not change its physical features or private IP.

Directions

1. Log in to the [MariaDB Console](#), select one or more instances from the instance list and click **Restart** at the top.



2. In the pop-up dialog box, check that all information is correct, and click **Confirm** to restart a single instance or multiple instances in batches.

Isolating and Terminating Instances

Last updated : 2024-01-11 15:28:37

Isolating Instances

An instance can be isolated when you no longer use it. Once isolated, the instance can neither be used nor accessed (but is not eliminated yet), and will be moved to the recycle bin, where you can restore or terminate it or it will be automatically terminated upon expiration. Even though the instance is isolated, the space occupied by its resources is not freed up, and it still has the minimal data replicas.

You can log in to the [console](#), select the pay-as-you-go instance from the instance list, click **Terminate/Return** to return it manually. After the instance is returned, it is in the **Isolated* status and will be retained for 3 days, during which it cannot be accessed but can be restored in the recycle bin.

You can log in to the [console](#), select the dedicated cluster instance in the instance list, click **Terminate/Return** to return it manually. After the instance is returned, it is in the **Isolated* status and will be retained for 3 days, during which it cannot be accessed. To restore it, you can do so in the recycle bin list.

After an instance is returned, once its status has changed to **Isolated**, it will no longer generate fees.

Note:

After an instance is isolated, its IP will be released, and you may not get back the original IP after the instance is restored.

After an instance is isolated, you cannot upgrade it, modify its parameters, create or modify an account for it, roll it back, or rename it.

Directions

1. Log in to the [MariaDB console](#). In the instance list, select an instance, and click **More > Terminate/Return** at the top.
2. In the pop-up dialog box, indicate your consent and click **OK**.

Terminate Instance

×

You've selected **1 instance in total**. [View Details](#) ▼


After the instance is completely terminated, **the data will not be recovered**. Please back up the instance data in advance.

After the instance is completely terminated, the IP resources are released at the same time. If this instance has associated with one or more disaster recovery/read-only instances:

- The disaster recovery/read-only instance will stop the sync connection and automatically promote to primary instance

Refund after the instance is completely terminated:

- The amount refunded without any reason will be refunded to the original payment account in 5 days
- The normal self-refund amount will be returned to your Tencent Cloud account by the proportion of the cash and voucher amount paid for the purchase.
- For orders from promotional reward channel, the refund will be charged 25% of their actual cash payment amount.
- These types of orders do not support self-service refunds, please submit a ticket to request a refund.

☒ I have read and agreed to [Termination Rules](#) 

ConfirmCancel

Then the instance becomes "Isolated" and is moved to the recycle bin.

Restoring Instances

An isolated instance can be restored to its normal running status, which may take several minutes. The restored instance may have a new IP rather than the original IP before isolation.

Directions

1. Log in to the [TencentDB for MariaDB console](#), locate the instance in the recycle bin list, and click **Restore/Start up**.
2. In the pop-up window, click **OK**.

Terminating Instances

If you don't need an instance anymore, you can return it. Once returned, it is in the "Isolated" status and moved to the recycle bin, where it will be automatically terminated upon expiration, or you can click **Eliminate Now** to completely terminate it.

Notes

After an instance is terminated, its data will not be recoverable. You need to back up the data in advance.

After an instance is terminated, its IP resources will be released simultaneously, and its disaster recovery instance will stop the sync connection and will be automatically promoted to primary instance.

After an instance is terminated completely, the refund procedures are as detailed below:

For instances that meet the 5-day no-questions-asked refund policy, the payment will be returned to your Tencent Cloud account.

For normal instances, the payment will be returned to your Tencent Cloud account by the proportion of the cash and gift cards paid for the purchase.

For an order placed from promotion rewarding channels, 25% of the actual cash payment amount will be deducted from the refund amount as service charges. Currently, self-service refund is unavailable for such kind of orders, you can [submit a ticket](#) to apply for the refund.

Disaster Recovery Read-Only Instance

Last updated : 2024-01-11 15:28:38

This document describes how to create and manage disaster recovery read-only instances in the console.

Overview

TencentDB for MariaDB provides cross-AZ/region disaster recovery read-only instances to enhance your capacity to deliver continuous services at low costs while improving data reliability for applications with greater service continuity, data reliability, and compliance requirements.

Note:

Disaster recovery read-only instance costs the same as the primary instance. For detailed pricing, see [Pricing](#)

Use Cases

Remote disaster recovery: To ensure data security, you can use disaster recovery instances to back up your business and data in multiple regions. In the event that an instance becomes unavailable due to an AZ/region failure, you can quickly switch to a cross-AZ/region disaster recovery instance to minimize the impact on your business.

Nearby access: You can use an instance in a specific AZ as the primary instance and those in other AZs/regions as read-only instances, which provides users with nearby access, remote read capabilities, and improved access speed.

Multi-region deployment: MariaDB instance can be deployed across multiple regions. When an MariaDB instance experiences network fluctuations or unavailability in an AZ/region, it can be switched to another AZ/region based on business needs.

Features

Disaster recovery read-only instances provide separate database connection addresses for read-only access. They can be used for nearby access and data analysis at a lower cost of device redundancy.

A primary instance can create multiple disaster recovery read-only instances that can be deployed in different regions and AZs.

Disaster recovery read-only instances support high-availability(1 primary-1 replica and 1 primary-2 replica) architecture, which helps avoid single point of failure for databases.

If the primary instance fails, the disaster recovery read-only instance can be activated in seconds to provide full read/write capability.

Data in a disaster recovery read-only instance is synced over a private network, which has lower latency and greater stability than a public network.

The traffic of data sync over a private network is currently free of charge during the promotion period. If fees will be charged for it, we will inform you in advance.

Feature Limits

Disaster recovery read-only instance do not support parameter setting and account management features.

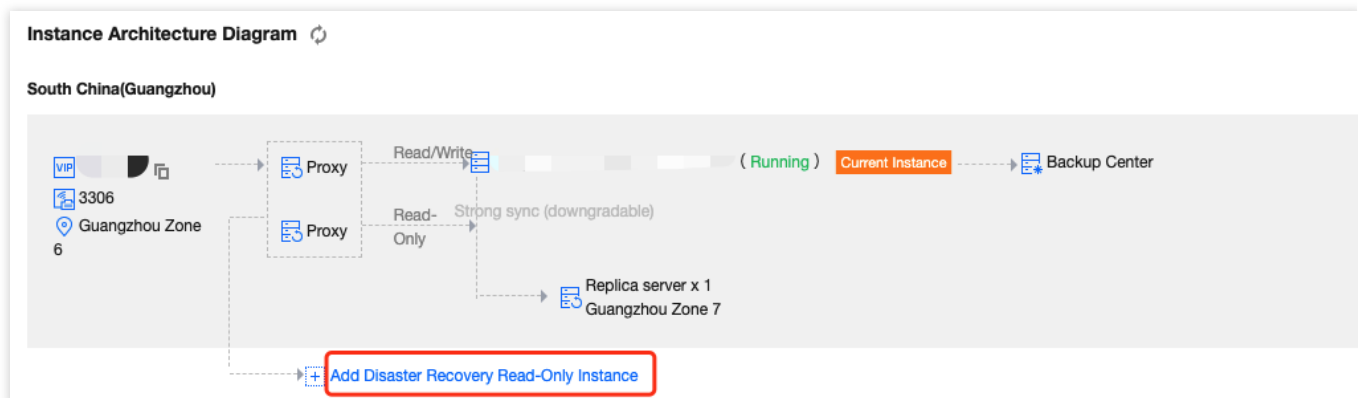
Database version of disaster recovery read-only instance is the same as that of the primary instance by default.

Instance specification and disk size should be greater than or equal to that of the primary instance.

Directions

Create a disaster recovery read-only instance

1. Log in to the [MariaDB console](#), click an instance ID in the instance list, and enter the instance management page.
2. In instance architecture diagram on the instance details page, click **Add Disaster Recovery Read-Only Instance**, and enter instance purchase page.



3. On the purchase page, set billing mode, region, and other basic info of disaster recovery read-only instance, and click **Buy Now**.

Note:

The time required to complete the creation depends on the amount of data, and no operations can be performed on the primary instance in the console during the creation. It's recommended to do so at an appropriate time.

Only the entire instance data can be synced. Make sure that the disk space is sufficient.









Make sure that the primary instance is in the running status and no tasks are executing, otherwise the sync task may fail.

4. Return to instance list after payment, initialize the instance, and you can proceed to the subsequent operations.

Manage disaster recovery read-only instances

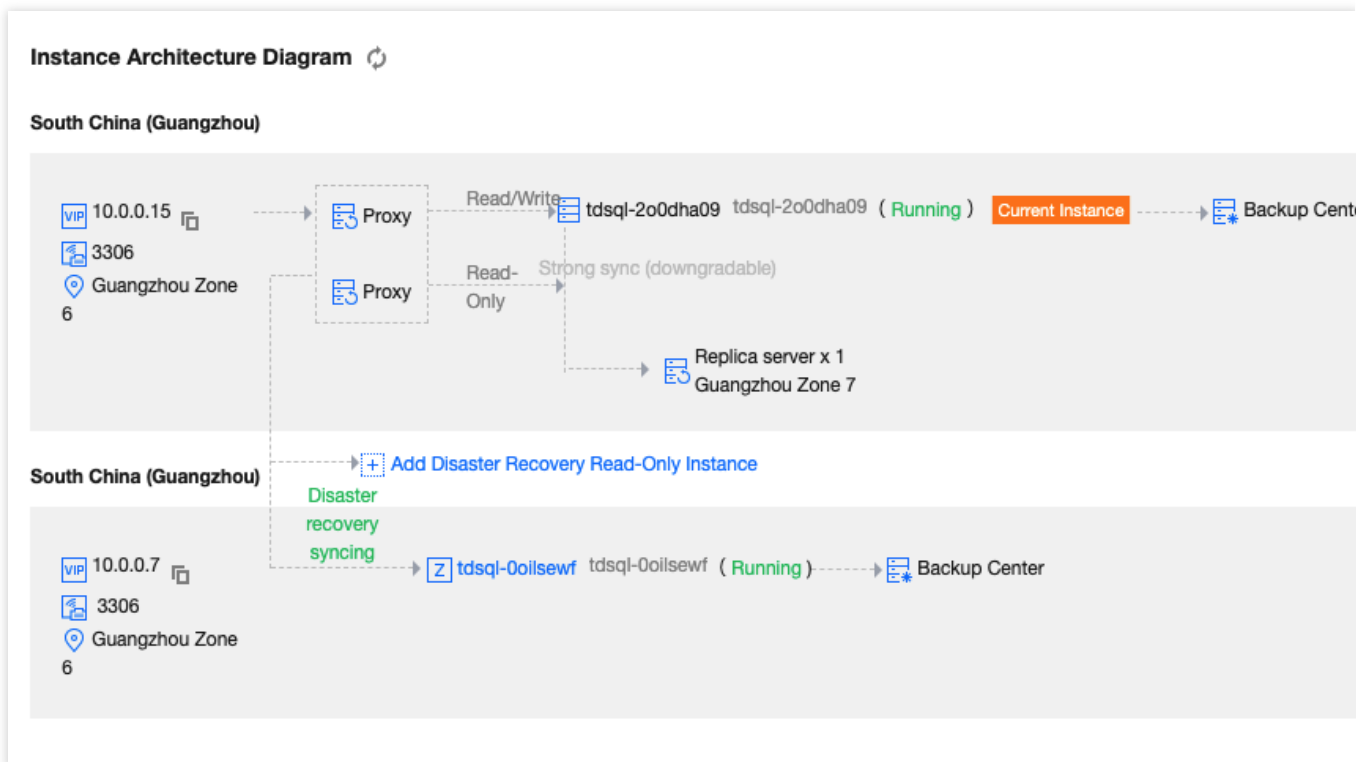
View disaster recovery read-only instances

You can view disaster recovery read-only instances from the region where they reside, and filter them out in the instance list.

<input type="checkbox"/>	Instance ID/Name	Monitoring/Status	AZ	Instance Type ▾	Project ▾
<input type="checkbox"/>	 	 Creating  Creating disaster recovery sync	Guangzhou Zone 6	<input type="checkbox"/> All <input type="checkbox"/> Primary (Dedicated) <input type="checkbox"/> Primary Instance <input type="checkbox"/> Disaster Recovery Read-Only Instance <input type="checkbox"/> Disaster Recovery Read-Only Instance (Dedicated)	
<input type="checkbox"/>	 	 Running  Creating disaster recovery sync	Guangzhou Zone 6	<input type="button" value="OK"/> <input type="button" value="Reset"/>	

View the relationship between the primary instance and the disaster recover read-only instance

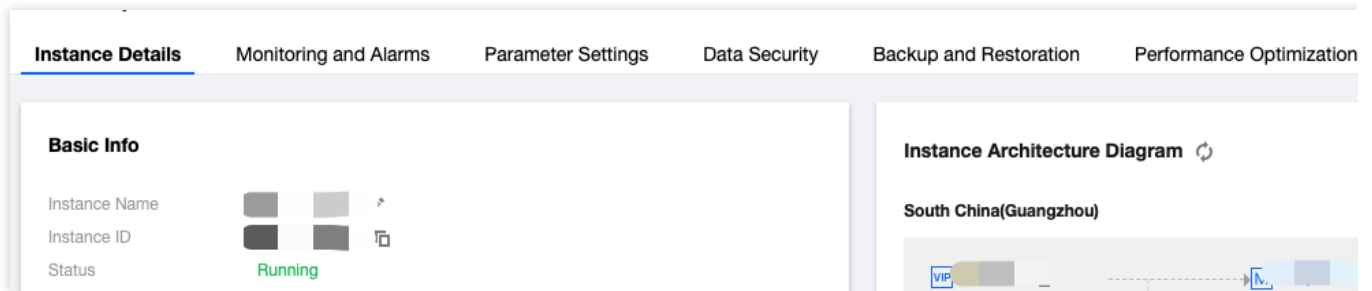
In instance architecture diagram on the instance details page, you can view the relationship between the primary instance and the disaster recover read-only instance.



Disaster recovery read-only instance feature

Disaster recovery read-only instance provides instance details, system monitoring, data security, backup and

recovery, and performance optimization features.



Promote a disaster recovery read-only instance to the primary instance

You can promote a disaster recovery read-only instance to primary instance in the console as needed.

1. Log in to the [TencentDB for MariaDB console](#), select the target disaster recovery read-only instance, click an instance name in the instance list to enter the instance management page.
2. Click **Promote to Primary Instance** in the top-right corner to promote the disaster recovery read-only instance to primary instance. After the promotion, the sync link with the primary instance will be disconnected, so that the promoted instance can get data write capability and full MariaDB functionality.

Note:

The disconnected sync link cannot be reconnected. You must exercise caution with this operation.

Account Management

Creating Account

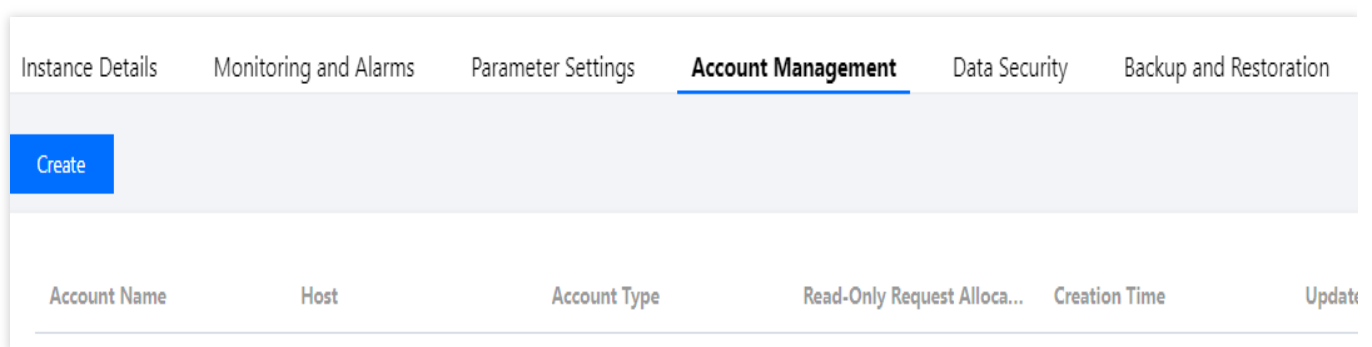
Last updated : 2024-01-11 15:28:38

Overview

This document describes how to create a TencentDB for MariaDB account in the console to manage and connect to the database instance.

Directions

1. Log in to the [TencentDB for MariaDB console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select **Account Management** and click **Create Account**.



3. In the pop-up dialog box, enter the account name, host, and password. After confirming that everything is correct, click **Next**.

Account ID: It must contain 1-32 letters, digits, or symbols, and start with a letter.

Host: It can be an IP and contain `%`.

Password: It must contain 8-32 lowercase letters, uppercase letters, digits, and symbols (`()~!@#$%^&*~+=_ | { }` `[] : < > , . ? /`), and cannot start with a slash (/).

Maximum connections: If left empty or `0` is passed in, the "max_connections" parameter will take effect.

Create



Account Name *

The account name contain 1-32 letters, digits, or symbols (_), and must start with a letter.

Create as Read-Only Account

☐ Yes ☒ No

If yes, you can set the parameters of the read-only account after clicking OK.

Host *

It is in the format of an IP ending with %. You can also enter % or 127.0.0.1.

Set Password *



The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols (0~!@#\$%^&*+=_[]:;<>.,?/). It cannot start with a slash (/).

Confirm Password *



The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols (0~!@#\$%^&*+=_[]:;<>.,?/). It cannot start with a slash (/).

Maximum Connections

If left empty or `0` is passed in, the "max_connections" parameter will take effect.

Remarks

Up to 256 characters for remarks

[Confirm and Go Next](#)[Cancel](#)

4. To create a read-only account, you need to [configure read/write separation](#) for it. Confirm the information you enter, click **OK**.

-If **Primary Server** is selected, read from the primary server when the delay of replica server times out.

If **Report Errors** is selected, report an error when all replica servers are delayed.

If **Read Only from Replica Server** is selected, ignore the replica delay and always read from replica server (generally used to fetch binlog for sync).

If your instance architecture is 1-primary-1-replica, select **Read Only from Replica Server** carefully to prevent high-load tasks like large transactions from affecting backup tasks and replica server availability.

Read-Only Account Settings

Account Namecycloneli

Host%

Read-Only Request Allocation Policy *

☒ Primary Server

☐ Report Errors

☐ Read Only from Replica Server

If "Primary Server" is selected, read from the primary server when the delay of replica server times out.

If "Report Errors" is selected, report errors for the replica delay.

If "Read Only from Replica Server" is selected, ignore the delay parameter and always read from replica server (generally used to fetch binlog for sync).

If your instance architecture is 1-source-1-replica, please select "Read Only from Replica Server" carefully to avoid the impact on the backup tasks and availability of the replica server by high-load tasks such as large transactions.

Specified Read-Only Replica Server *

☒

When the specified read-only replica server is enabled, the read-only requests will be automatically disconnected and will not be switched to another replica server if the source-replica delay exceeds the delay parameter.

When the specified read-only replica server is disabled, another available replica server will be automatically selected if the source-replica delay exceeds the delay parameter.

Read-Only Replica Server Delay Parameter *

–

10

+

sec

If the replica server delay exceeds this parameter value, the replica server is considered faulty. We recommend that you set this parameter to a value larger than 10.

OK

Cancel

5. In the **Modify Permissions** pop-up window, grant permissions as needed and click **Modify**. To discard the changes, click **Cancel Modification**.

Account name: cycloneli@%

Modify the database permissions for one or more objects

Global Privileges

☒ ALTER

☐ ALTER ROUTINE

☐ CREATE

☐ CREATE ROUTINE

☐ CREATE TEMPORARY TABLES

☐ CREATE VIEW

☐ DELETE

☐ DROP

☐ EVENT

☐ EXECUTE

☐ INDEX

☐ Select All

Objects selected and permissions modified

Global Privileges

☒ ALTER

☐ ALTER ROUTINE

☐ CREATE

☐ CREATE ROUTINE

☐ CREATE TEMPORARY TABLES

☐ CREATE VIEW

☐ DELETE

☐ DROP

☐ EVENT

☐ EXECUTE

☐ INDEX

Refresh

Reset

Modify

Cancel Modification

Related APIs

API Name	Description
CreateAccount	Creates an account

Modifying Account Permissions

Last updated : 2024-01-11 15:28:38

Overview

You can grant global/object-level privileges for TencentDB for MariaDB accounts in the console.

Directions

1. Log in to the [TencentDB for MariaDB console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select the **Account Management** tab, find the account for which to modify the permission, and click **Modify Permissions**.

Account Name	Host	Account Type	Read-Only Request Alloca...	Creation Time	Update Time	Remarks	Operation
cycloneli	%	General Account	None	2022-12-19 16:40:14	2022-12-19 16:40:14		Modify Permissions Clone Account More

3. In the pop-up dialog box, select or deselect permissions and click **OK** to complete the modification.

Global Privileges: Grant permissions to all databases in the instance.

Object-Level Privileges: Grant permissions to certain databases in the instance.

Account name: cycloneli@%

Modify the database permissions for one or more objects

Global Privileges

Object Level Privilege

☒ ALTER

☐ ALTER ROUTINE

☐ CREATE

☐ CREATE ROUTINE

☐ CREATE TEMPORARY TABLES

☐ CREATE VIEW

☐ DELETE

☐ DROP

☐ EVENT

☐ EXECUTE

☐ INDEX

☐ Select All

Objects selected and permissions modified

Global Privileges

☒ ALTER

☐ ALTER ROUTINE

☐ CREATE

☐ CREATE ROUTINE

☐ CREATE TEMPORARY TABLES

☐ CREATE VIEW

☐ DELETE

☐ DROP

☐ EVENT

☐ EXECUTE

☐ INDEX

Refresh

Reset

Modify

Cancel Modification

Related APIs

API Name	Description
DescribeAccountPrivileges	Queries account permission
GrantAccountPrivileges	Sets account permission

Configuring Read/Write Separation

Last updated : 2024-01-11 15:28:38

Read/Write Separation Overview

TencentDB for MariaDB supports read/write separation by default. Each replica in the primary/replica architecture can be read-only. If multiple replicas are configured, read requests will be automatically assigned to low-load replicas by the gateway cluster (TProxy).

Read/Write Separation Based on Read-Only Accounts

A read-only account has only the read permission to read data from the replica server (or read-only instances) in a database cluster by default. In the [TencentDB for MariaDB console](#), you can set a read-only account and a read policy on the **Account Management** tab of the instance management page.

Create✕

Account Name *

Account Name

The account name contain 1-32 letters, digits, or symbols (_-), and must start with a letter.

Create as Read-Only Account

☐ Yes ☒ No


If yes, you can set the parameters of the read-only account after clicking OK.

Host *

Host


It is in the format of an IP ending with %. You can also enter % or 127.0.0.1.

Set Password *

Enter the password 

The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols ([~!@#\$%^&*+=_[]:;<>.,/?/). It cannot start with a slash (/).

Confirm Password *

Please enter confirm pas: 

The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols ([~!@#\$%^&*+=_[]:;<>.,/?/). It cannot start with a slash (/).

Maximum Connections

—

0

+

If left empty or `0` is passed in, the "max_connections" parameter will take effect.

Remarks

Enter remarks

Up to 256 characters for remarks

Confirm and Go Next

Cancel

In read-only account settings, you can set **Read-Only Request Allocation Policy** to define the read policy when a

replica failure (or long delay) occurs. The **Read-Only Replica Server Delay Parameter** defines the data sync delay time and is used together with **Read-Only Request Allocation Policy**.

-If **Primary Server** is selected, read from the primary server when the delay of replica server times out.

If **Report Errors** is selected, report an error when all replica servers are delayed.

If **Read Only from Replica Server** is selected, ignore the replica delay and always read from replica server (generally used to fetch binlog for sync).

The **Read-Only Replica Server Delay Parameter** defines the data sync delay time.

Allocation Policy	Read-Only Replica Server Is Specified	No Read-Only Replica Server Is Specified
Primary Server	Read from the primary server when the actual delay exceeds the delay parameter.	Read from the other replica servers when the actual delay exceeds the delay parameter, and switch to the primary server when the delay of the replica server occurs.
Report Errors	An error will be reported when the actual delay exceeds the delay parameter.	Read from the other replica servers first when the actual delay exceeds the delay parameter, and an error will be reported when the delay of the replica server occurs.
Read Only from Replica Server	Read from the specified read-only replica	-

>?>- To modify the settings of read-only account, select **More** > **Modify Read-Only Request Allocation Policy** in the **Operation** column.>- If your instance architecture is 1-primary-1-replica, the read-only separation feature can only be used for low-load read-only tasks. Avoid high-load tasks such as large transactions, as they affect the backup tasks and availability of the replica server.

For example, if you design a transaction system, the following configuration items are recommended:

Core transaction module: set a regular account which has read/write permission.

Balance query module: Set a read-only account that reads from the replica by default, configure the request allocation policy so that the primary will be read from upon replica failure, and set the delay parameter to below 10 seconds in order to ensure primary/replica performance and data consistency of user inquiries.

Batch query module: Set a read-only account that reads from the replica by default, configure the request allocation policy so that an error will be reported upon replica failure, and set the delay parameter to above 30 seconds in order to avoid affecting the primary performance.

In addition, the Multi-thread Asynchronous Replication (MAR) mechanism is to return a response immediately after data is written to a replica transaction log. In this case, replica table data may not be updated; therefore, delay will occur.

Read/Write Separation Based on Comments

Add `/*slave*/` field before each SQL to be executed by replica, and add `-c` parameter after `mysql` to resolve the annotation `mysql -c -e "/*slave*/sql"`, to automatically distribute the read request to replica. Examples are shown below:

```
//Read data from primary//
select * from emp order by sal, deptno desc;
//Read data from replica//
/*slave*/ select * from emp order by sal, deptno desc;
```

Note:

Only "read data from replica" (SELECT) is supported rather than other operations. Non-SELECT statements will fail.

`-c` parameter needs to be added after `mysql` to resolve the annotation.

`/*slave*/` must be lowercase, and no spaces are needed before and after the statement.

If the MAR mechanism is affected due to replica exception, the read operation on replica is automatically switched to that on primary.

Read-only Instance (Remote Read-only Instance)

If the above read/write separation schemes do not meet your needs, TencentDB for MariaDB provides [read-only instance](#) for you. A read-only instance is an independent database instance that does not participate in high-availability switch of the original primary instance and is used only for improving read performance.

Cloning Account

Last updated : 2024-01-11 15:28:38

Overview

You can clone a database account in the TencentDB for MariaDB console, and retain its original account password to provide different permissions.

Directions

1. Log in to the [TencentDB for MariaDB console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the database management page, select the **Account Management** tab, find the account for which to reset the password, and click **Clone Account**.

Account Name	Host	Account Type	Read-Only Request Alloca...	Creation Time	Update Time	Remarks	Operation
cycloneli	%	General Account	None	2022-12-19 16:40:14	2022-12-19 16:40:14		Modify Permissions Clone Account More

3. In the pop-up window, enter the primary server IP , account name, and password (the name and password can be the same as that of the original account), then click **Confirm and Go Next**.

Clone Account ✕

Account Name *

cycloneli

The account name contain 1-32 letters, digits, or symbols (_), and must start with a letter.

☒ Clone Original Account Password

Create as Read-Only Account

☐ Yes ☒ No

If yes, you can set the parameters of the read-only account after clicking OK.

Host *

%

It is in the format of an IP ending with %. You can also enter % or 127.0.0.1.

Set Password *

Enter the password

The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols (0~!@#\$%^&*~+=_[]:;<>.,?/). It cannot start with a slash (/).

Confirm Password *

Please enter confirm pas:

The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols (0~!@#\$%^&*~+=_[]:;<>.,?/). It cannot start with a slash (/).

Maximum Connections

-

0

+

If left empty or `0` is passed in, the "max_connections" parameter will take effect.

Remarks

Enter remarks

Up to 256 characters for remarks

Confirm and Go Next

Cancel

4. Return to the account management page to view the cloned account.

Related APIs

API Name	Description
CloneAccount	Clones an account

Resetting Account Password

Last updated : 2024-01-11 15:28:38

Overview

If you forgot your database account password or need to modify it while using TencentDB for MariaDB, you can reset it in the console.

Note:

We recommended that you regularly reset the password at least once every three months for the sake of data security.

Directions

1. Log in to the [TencentDB for MariaDB console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select **Account Management** tab, find the account for which to reset the password, and select **More > Reset Password**.

Account Name	Host	Account Type	Read-Only Request Alloca...	Creation Time	Update Time	Remarks	Operation
cycloneli	%	General Account	None	2022-12-19 16:40:14	2022-12-19 16:40:14		Modify Permissions Clone Account More

[Reset Passv](#)
[Delete Acco](#)


3. In the pop-up window, enter the **New Password** and **Confirm Password** and click **OK**.

Note:

To avoid the risks caused by creating, modifying, and deleting account information, we recommend that you configure [access management](#), and reset the password with caution.

Reset Password



 To avoid your business risks caused by creating, modifying, or deleting account info, it is recommended to carefully control the function menu permissions (API keyword: *Account*) when you configure access management.

Instance Name

Account Name

Host

%

Set Password *

Enter the password



The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols ({}~!@#\$%^&*~+=_[]: <>.,?/). It cannot start with a slash (/).

Confirm Password *

Please enter confirm pas



The password must contain 8-32 characters in all of the following four types: lowercase letters, uppercase letters, digits, and symbols ({}~!@#\$%^&*~+=_[]: <>.,?/). It cannot start with a slash (/).

OK

Cancel

Related APIs

API Name	Description
ResetAccountPassword	Resets account password

Deleting Account

Last updated : 2024-01-11 15:28:38

Overview

This document describes how to delete a TencentDB for MariaDB instance in the console.

Note:

A database account cannot be recovered once deleted. Ensure that the account is no longer in use and proceed with caution.

Directions

1. Log in to the [TencentDB for MariaDB console](#). In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select **Account Management** tab, find the account for which to reset the password, and select **More > Delete Account**.

Account Name	Host	Account Type	Read-Only Request Alloca...	Creation Time	Update Time	Remarks	Operation
cycloneli	%	General Account	None	2022-12-19 16:40:14	2022-12-19 16:40:14		Modify Permissions Clone Account More
							Reset Passv Delete Acco

3. In the pop-up dialog box, confirm that everything is correct and click **OK**.

Related APIs

API Name	Description
DeleteAccount	Deletes an account

Read/Write Separation

Last updated : 2024-01-11 15:28:38

Read/Write Separation Overview

TencentDB for MariaDB supports read/write separation by default. Each replica in the primary/replica architecture can be read-only. If multiple replicas are configured, read requests will be automatically assigned to low-load replicas by the gateway cluster (TProxy).

Read/Write Separation Based on Read-only Accounts

A read-only account has only the read permission and reads data from the replica server (or read-only replicas) in a database cluster by default. In the [TencentDB for MariaDB console](#), you can set a read-only account and a read policy on the **Account Management** tab of the instance management page.

Create



Account Name *

Account name:

Account ID must be a combination of 1-32 chars comprised of digits, letters and special chars. It should start with a letter. The special chars are underscores and hyphens.

Create as read-only account

☒ Yes ☐ No

If yes, you can set the parameters of the read-only account after clicking OK.

Host *

Host:

It is in the format of an IP ending with %. You can also enter % or 127.0.0.1.

Set Password *

Enter the password



The password contains 8-32 characters, which must be comprised of lowercase letters, uppercase letters, digits, and symbols (()~!@#\$\$%^&*+ =_{}[]:;<>.,?/), but cannot start with a slash (/).

Confirm Password *

Please enter confirm password



The password contains 8-32 characters, which must be comprised of lowercase letters, uppercase letters, digits, and symbols (()~!@#\$\$%^&*+ =_{}[]:;<>.,?/), but cannot start with a slash (/).

Remarks

Enter remarks

Up to 256 chars for remarks

Confirm and Go Next

Cancel

In read-only account settings, you can set **Read-only Request Allocation Policy** to define the read policy when a

replica failure (or long delay) occurs. The **Read-only Replica Delay Parameter** defines the data sync delay time and is used together with **Read-only Request Allocation Policy**.

Read-only Account Settings

×

Account Name

test

Host

%

Read-Only Request Allocation Policy *

☒ Primary Server
 ☐ Report Errors
 ☐ Read Only from Secondary Server

If "Primary Server" is selected, read from the primary server when the delay of secondary server times out.

If "Report Errors" is selected, report errors for the secondary delay.

If "Read Only from Secondary Server" is selected, ignore the delay parameter and always read from secondary server (generally used to fetch binlog for sync).

Read-Only Secondary Server Delay Parameter *

−

10

+

sec

If the secondary server delay exceeds this parameter value, the secondary server is considered faulty. It is recommended to set this parameter to a value larger than 10.

OK

Cancel

For example, if you design a transaction system, the following configuration items are recommended:

Core transaction module: set a regular account which has read/write permission.

Balance inquiry module: set a read-only account that reads from the replica by default, configure the request allocation policy so that the primary will be read from upon replica failure, and set the delay parameter to below 10 seconds in order to ensure primary/replica performance and data consistency of user inquiries.

Batch inquiry module: set a read-only account that reads from the replica by default, configure the request allocation policy so that an error will be reported upon replica failure, and set the delay parameter to above 30 seconds in order to avoid affecting the primary performance.

In addition, the Multi-thread Asynchronous Replication (MAR) mechanism is to return a response immediately after

data is written to a replica transaction log. In this case, replica table data may not be updated; therefore, delay will occur.

Read/Write Separation Based on Comments

Add `/slave/` field before each SQL to be executed by replica, and add `-c` parameter after `mysql` to resolve the annotation `mysql -c -e "/*slave*/sql"`, to automatically distribute the read request to replica. Examples are shown below:

```
//Read data from primary//
select * from emp order by sal, deptno desc;
//Read data from replica//
/*slave*/ select * from emp order by sal, deptno desc;
```

Note:

Only "read data from replica" (SELECT) is supported rather than other operations. Non-SELECT statements will fail.

`-c` parameter needs to be added after `mysql` to resolve the annotation.

`/*slave*/` must be lowercase, and no spaces are needed before and after the statement.

If the MAR mechanism is affected due to replica exception, the read operation on replica is automatically switched to that on primary.

Read-only Replicas (Remote Read-only Replicas)

If the above read/write separation schemes do not meet your needs, TencentDB for MariaDB provides read-only replicas for you. A read-only replica is an independent database instance that does not participate in high-availability switch of the original primary instance and is used only for improving read performance.

Changing Networks

Last updated : 2024-01-11 15:28:38

This document describes how to change the instance network type and modify the instance access address.

Note:

Modifying the network configurations of an instance is highly risky. Please do so only during off-peak hours. After modification, unless assigned to another service, the original IP will remain valid for another 24 hours. We recommend modifying your business configuration accordingly as soon as possible.

Modifying the Private IP

You can modify the private IP of a TencentDB instance in VPC.

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance details page.
2. In the **Basic Info** section, click



next to **Private IP** to modify it. You can do so only when the current subnet has available IPs.

Basic Info

Instance Name	<div></div>
Instance ID	<div></div>
Running Status	Running
Instance Type	Disaster Recovery/Read-Only Instance
Instance Version	Standard Edition (1 primary-1 replica)
Region	South China (Guangzhou)
Private IPv4 Address	<div></div> Security Group
Private Port	3306 <div></div>
Public IPv4 Address	Enable
Private IPv6 Address	--
Public IPv6 Address	This instance is not supported currently
Network	Classic Network Switch to VPC
Project	DEFAULT PROJECT <div></div>
Character Set	UTF8 <div></div>
Tag	-- <div></div>

3. In the pop-up dialog box, modify the private IP and click **Confirm**.

Switching between VPC Subnets

You can switch an instance between VPC subnets.

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance details page.
2. In the **Basic Info** section, click **Change Subnet** next to **Network**.
3. In the pop-up dialog box, select a subnet, select **Auto-assign IP** or **Specify IP**, and click **Confirm**.

Note:

The original VIP address remains valid for another 24 hours after the VPC is modified. Please modify your business IP address accordingly within 24 hours.

Because the product supports an intra-city active-active architecture, you are recommended to choose a VPC subnet in the same region as your business server or primary node.

Change Subnet

Changing subnet may cause the change of the instance IP. The original IP will become invalid after 24 hours. Modify the instance IP on the client in time.

Select a subnet

VPC1 test0723 | Guangzhou Zone 3  CIDR 192.168.252.0/24
253 subnet IPs in total, with 242 available

To change the network, please go to the console to [Create VPC](#) or [Create Subnet](#)

In the current network environment, only devices in the VPC1 VPC can access this database instance.

It is recommended to select the same subnet as that of the business server or database instance primary node

- ☒ Auto-assign IP
☐ Specify IP

Confirm

Cancel

Switching between VPCs

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance details page.
2. In the **Basic Info** section, click **Change Subnet** next to **Network**.
3. In the pop-up dialog box, select a VPC, select **Auto-assign IP** or **Specify IP**, and click **Confirm**.

Note:

The original VIP address remains valid for another 24 hours after the VPC is modified. Please modify your business IP address accordingly within 24 hours.

Because the product supports an intra-city active-active architecture, you are recommended to choose a VPC subnet in the same region as your business server or primary node.

Switching from a Classic Network to a VPC

You can switch an instance from classic network to VPC.

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance details page.
2. In the **Basic Info** section, click **Switch to VPC** next to **Network**.

Basic Info

Instance Name	<div></div>
Instance ID	<div></div>
Running Status	Running
Instance Type	Disaster Recovery/Read-Only Instance
Instance Version	Standard Edition (1 primary-1 replica)
Region	South China (Guangzhou)
Private IPv4 Address	<div></div> Security Group
Private Port	<div></div>
Public IPv4 Address	Enable
Private IPv6 Address	--
Public IPv6 Address	This instance is not supported currently
Network	Classic Network Switch to VPC
Project	DEFAULT PROJECT <div></div>
Character Set	UTF8 <div></div>
Tag	-- <div></div>

3. In the pop-up dialog box, select a VPC, select **Auto-assign IP** or **Specify IP**, and click **Confirm**.

Note:

The switch from classic network to VPC is irreversible.

After the switch, VPC access will take effect immediately. The original classic network access will be retained for 24 hours; therefore, other instances associated to the TencentDB for MySQL instance should be migrated to VPC within 24 hours so as to guarantee uninterrupted access.

Because the product supports an intra-city active-active architecture, you are recommended to choose a VPC subnet in the same region as your business server or primary node.

Backup and Rollback

Backup Mode

Last updated : 2024-01-11 15:28:38

TencentDB for MariaDB supports full backup and incremental backup, and backups are compressed with LZ4. For more information on how to use LZ4, see [Decompressing Backups and Logs](#).

Backup Type

Full backup

You can set the backup retention period and backup execution time for full backups, which is set to seven days by default (30 days for finance cloud).

Incremental backup

Incremental backup is implemented based on binlogs, which are generated in real time. The binlogs occupy some disk space and are periodically uploaded to the TencentDB backup system.

Custom Backup Time



1. Log in to the [TencentDB for MariaDB console](#). Click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. Click **Backup and Restoration** in the instance management page
3. On **Backup and Restoration > Backup and Log Settings** page, you can set the storage period and backup execution time.

Storage time: Data and log backups can be retained for 1 to 365 days. Default value: 7 days.

Backup execution time: It can be set to any time period in hours.

Note:

Log backup is enabled by default and cannot be disabled. Logs include error logs, slow logs, and transaction logs (binlogs).

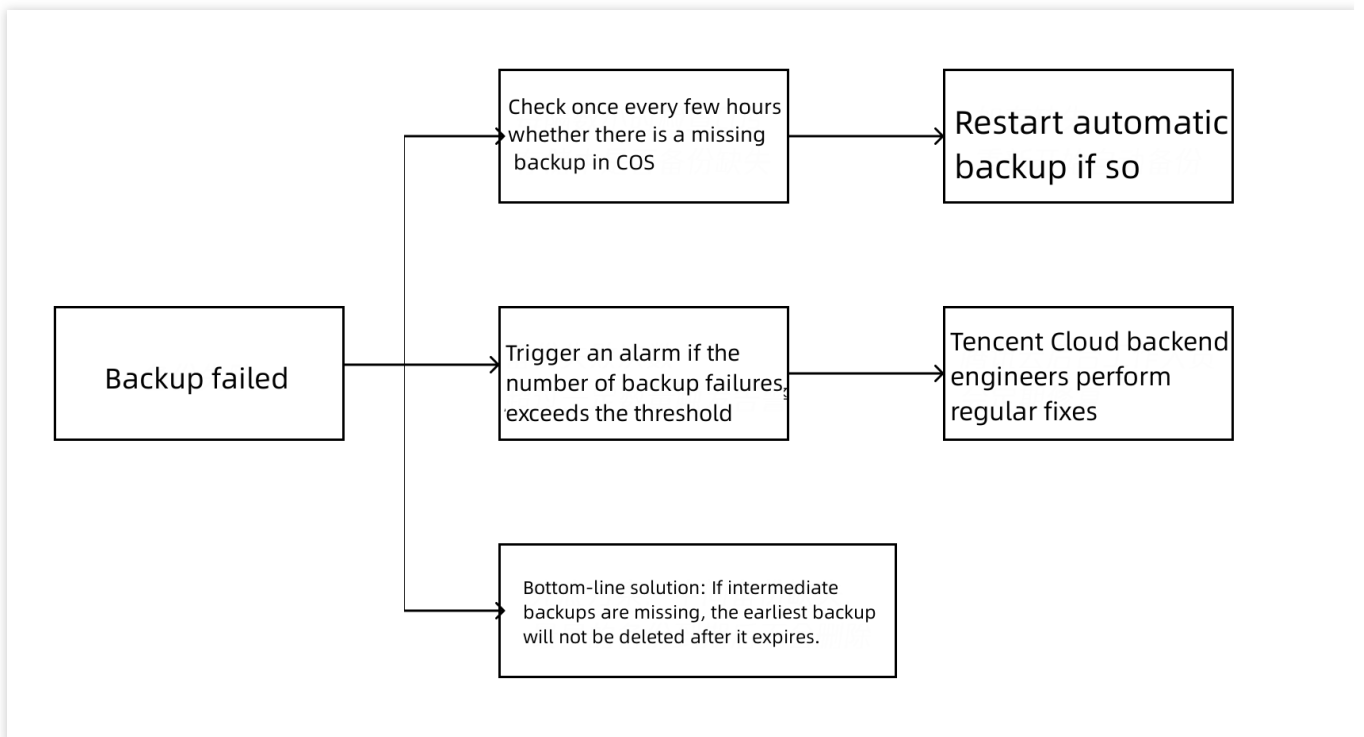
Instance Details	Monitoring and Alarms	Parameter Configuration	Account Management	Data Security	Backup and Restore	Performance Optimization
Cold Backup List	Temp Instance	Binlog List	Backup and Log Settings			
Item	Storage Time		Backup Execution Time		Notes	
Backup and Log Storage Days	30 days 		00:00~23:59 		50% of the instance capacity is provided as the storage space for free	

?There are three handling methods for backup failures:

Automatic fix: The database system checks once every few hours whether there is a missing backup in COS and starts another automatic backup if so.

Alarm and check: After the number of backup failures exceeds the threshold, an alarm will be triggered on the Tencent Cloud backend, and backend engineers will troubleshoot the issue.

Bottom-line solution: If there are so many backup failures that no backups are available, the database system will ensure that at least one backup is retained; that is, if intermediate backups are missing, the earliest backup will not be deleted after it expires.



Downloading Backup Files

Last updated : 2024-01-11 15:28:38

You can download the cold backup data and binlogs in the TencentDB for MariaDB console.

Downloading a cold backup or binlog

1. Log in to the [TencentDB for MariaDB console](#). Click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. Select **Backup and Restoration** > **Cold Backup List** or **Binlog List**
3. Select the file to be downloaded and click **Download** in the **Operation** column.
4. In the pop-up dialog box, click **Get Download Address** to get the download address in a VPC.
5. Log in to CVM (Linux system) under the VPC where the database resides as instructed in [Customizing Linux CVM Configurations](#) and run the `wget` command to download the file.

Note:


Download from Public Network: Enable this option in **Download Settings** on the **Database Backup** page. Then, you can directly copy the download link to a browser for download.

Download from Private Network: Access the instance in the VPC and use the `wget` command for download: `wget -O <custom name.log> '<backup file download address>'`.

The address is valid for 15 minutes. Refresh the page to get a new one after expiration.

Download

Notes

1. To ensure data security, **Each address is valid for 15 minutes**. Upon expiration, click "Get Download Address" again to get a new one.
2. Download over public network: Enable it on "Database Backup" > "Download Settings", and copy the download address to the browser for download.
3. Download over private network: Run the `wget` command (format: `wget -O <custom filename.log> '<backup download address>'`) for download in a VPC.
4. View and download [Help Documentation](#) 

Get Download Address

Close

Downloading a slow log

1. Log in to the [TencentDB for MariaDB console](#). Click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. Select **Performance Optimization > Slow Log** tab
3. Select the file to be downloaded and click **Download** in the **Operation** column.

Note:

If the file size is 0 KB, no slow query record is available for download.

4. In the pop-up dialog box, click **Get Download Address** to get the download address in a VPC.
5. Log in to CVM (Linux system) under the VPC where the database resides as instructed in [Customizing Linux CVM Configurations](#) and run the `wget` command to download the file.

Note:

Download from Public Network: Enable this option in **Download Settings** on the **Database Backup** page. Then, you can directly copy the download link to a browser for download.

Download from Private Network: Access the instance in the VPC and use the `wget` command for download: `wget -O <custom name.log> '<slow log download address>'`.

The address is valid for 15 minutes. Refresh the page to get a new one after expiration.

Decompressing Backups and Logs

Last updated : 2024-01-11 15:28:38

For the sake of compression performance and ratio, backup files and log files (binlog files) of TencentDB for MariaDB are compressed with LZ4 (Extremely Fast Compression Algorithm). You can use LZ4 for decompression. This document describes how to use this tool.

Windows

Installing LZ4

Download LZ4 and install it by following the installation wizard.

Decompressing files

Right click the LZ4 file to be decompressed and select **Decode with LZ4**.

Linux

Installing LZ4

There is an LZ4 component in the yum repository of CVM. [Log in to your CVM instance](#) and run the following command to install it.

```
$ yum install lz4
```

The installation is successful if the result is returned as shown below after you execute `lz4`.

```
[root@UM_240_177_centos ~]# lz4
Incorrect parameters
Usage :
    lz4 [arg] [input] [output]

input    : a filename
           with no FILE, or when FILE is - or stdin, read standard input
Arguments :
  -1      : Fast compression (default)
  -9      : High compression
  -d      : decompression (default for .lz4 extension)
  -z      : force compression
  -f      : overwrite output without prompting
  -h/-H  : display help/long help and exit
```

Decompressing files

Execute the following command to decompress a file.

```
$ lz4 -d xxx.lz4
```


Restoring Instances from Backup Files

Last updated : 2024-01-11 15:28:38

You can view historical data by using the rollback feature of TencentDB for MariaDB. To restore your database instance locally, you can do so by following the instructions below.

Prerequisites

Preparing a server

To restore the database instance locally, ensure that the basic configuration of the server meets the following requirements:

CPU: 2 or more cores.

Memory: 4 GB or above.

Disk capacity: It must be greater than the used space of the database and leave enough temporary space for the system.

Operating system: CentOS

Preparing a database

Note:

The version of the local database must be the same as that of the TencentDB instance.

Take installation of MariaDB 10.0.10 as an example:

1. Add the yum source.

```
vi /etc/yum.repos.d/mariadb-10.0.10.repo):  
# MariaDB 10.0 CentOS repository list - created 2016-05-30 02:16 UTC  
# http://downloads.mariadb.org/mariadb/repositories/  
[mariadb]  
name = MariaDB  
# baseurl = http://yum.mariadb.org/10.0/centos7-amd64  
baseurl = http://archive.mariadb.org/mariadb-10.0.10/yum/centos6-amd64/  
gpgkey=https://yum.mariadb.org/RPM-GPG-KEY-MariaDB  
gpgcheck=0
```

2. Check whether the version of the MariaDB instance corresponding to the yum source is 10.0.10.

```
yum makecache  
yum info MariaDB-server
```

3. Install MariaDB-server.

```
yum install MariaDB-server
```

Note:

If the system prompts a conflict with a legacy version, you need to remove the previously installed package, such as `yum remove mariadb-libs`.

Installing the auxiliary tool**1. Install the MariaDB client**

```
yum install MariaDB-client
```

2. Install the LZ4 decompression software. For more information, see [Decompressing Backups and Logs] (<https://intl.cloud.tencent.com/document/product/237/2088>). LZ4 is installed in the `mysqlagent/bin` directory by default. You can also install it in the `/usr/bin` directory and import it as an environment variable.

```
yum install -y lz4
percona-xtrabackup
yum install http://www.percona.com/downloads/percona-release/redhat/0.1-3/percona-release-0.1-3.noarch.rpm
yum install percona-xtrabackup
```

Downloading a backup

In the [TencentDB for MariaDB console](#), click an instance ID to enter the instance management page, and get the backup download address on the **Backup and Restoration** tab.

Sample download command:

```
wget --content-disposition
'http://1x.2xx.0.27:8083/2/noshard1/set_1464144850_587/1464552298xxxxxxxx'
```

Restoring Databases from Backup Files (Unencrypted)

1. Enter the cold backup file download directory and decompress the file with LZ4

```
lz4 -d set_1464144850_587.1464552298.xtrabackup.lz4
```

2. Decompress the file to a temporary directory `xtrabackuptmp` with `xbstream`

```
mkdir xtrabackuptmp/
mv set_1464144850_587.1464552298.xtrabackup xtrabackuptmp/
xbstream -x < set_1464144850_587.1464552298.xtrabackup
```

After the decompression, the directories and files are as shown below:

```
[root@VM_0_2_centos xtrabackuptmp]# ll
total 6347108
-rw-r----- 1 root root      358 May 30 16:28 backup-my.cnf
-rw-r----- 1 root root 1073741824 May 30 16:31 ib_logfile0
-rw-r----- 1 root root 1073741824 May 30 16:30 ib_logfile1
-rw-r----- 1 root root 1073741824 May 30 16:31 ib_logfile2
-rw-r----- 1 root root 1073741824 May 30 16:31 ib_logfile3
-rw-r----- 1 root root 2147483648 May 30 16:31 ibdata1
drwx----- 2 root root      4096 May 30 16:28 mysql
drwx----- 2 root root      4096 May 30 16:28 performance_schema
-rw-r----- 1 root root 2149044297 May 30 16:27 set_1464144850_587.1464552298.xtrabackup
drwx----- 2 root root      4096 May 30 16:28 sysdb
drwx----- 2 root root      4096 May 30 16:28 test
-rw-r----- 1 root root        25 May 30 16:28 xtrabackup_binlog_info
-rw-r----- 1 root root        58 May 30 16:30 xtrabackup_binlog_pos_innodb
-rw-r----- 1 root root       117 May 30 16:30 xtrabackup_checkpoints
-rw-r----- 1 root root       858 May 30 16:28 xtrabackup_info
-rw-r----- 1 root root 2097152 May 30 16:30 xtrabackup_logfile
[root@VM_0_2_centos xtrabackuptmp]#
```

3. Use `innobackupex` to apply logs

```
mkdir /root/dblogs_tmp
innobackupex --apply-log --use-memory=1G --tmpdir='/root/dblogs_tmp/'
/root/xtrabackuptmp/
```

After the operation succeeds, `completed OK!` will be displayed.

```
InnoDB: Initializing buffer pool, size = 1.0G
InnoDB: Completed initialization of buffer pool
InnoDB: Setting log file ./ib_logfile101 size to 1024 MB
InnoDB: Progress in MB: 100 200 300 400 500 600 700 800 900 1000
InnoDB: Setting log file ./ib_logfile1 size to 1024 MB
InnoDB: Progress in MB: 100 200 300 400 500 600 700 800 900 1000
InnoDB: Setting log file ./ib_logfile2 size to 1024 MB
InnoDB: Progress in MB: 100 200 300 400 500 600 700 800 900 1000
InnoDB: Setting log file ./ib_logfile3 size to 1024 MB
InnoDB: Progress in MB: 100 200 300 400 500 600 700 800 900 1000
InnoDB: Renaming log file ./ib_logfile101 to ./ib_logfile0
InnoDB: New log files created, LSN=210681783
InnoDB: Highest supported file format is Barracuda.
InnoDB: 128 rollback segment(s) are active.
InnoDB: Waiting for purge to start
InnoDB: 5.6.24 started; log sequence number 210681868
xtrabackup: starting shutdown with innodb_fast_shutdown = 1
InnoDB: FTS optimize thread exiting.
InnoDB: Starting shutdown...
InnoDB: Shutdown completed; log sequence number 210683014
160530 16:31:18 completed OK!
[root@VM_0_2_centos ~]# ls
```

4. Stop the database and clear data files

```
service mysql stop
```

Clear data files (in data directories, tablespace directories, and log directories):

```
mkdir /var/lib/mysql-backup  
mv /var/lib/mysql/* /var/lib/mysql-backup
```

5. Modify the database parameter file

Modify the database parameter file `(/etc/my.cnf.d/server.cnf)` . For specific parameter values, see parameters in the extracted `backup-my.cnf` file. **Do not directly replace the parameter file with `backup-my.cnf` .**

```
[mysqld]  
skip-name-resolve  
datadir=/var/lib/mysql  
innodb_checksum_algorithm=innodb  
innodb_log_checksum_algorithm=innodb  
innodb_data_file_path=ibdata1:2G:autoextend  
innodb_log_files_in_group=4  
innodb_log_file_size=1073741824  
innodb_page_size=4096  
innodb_log_block_size=512  
innodb_undo_tablespaces=0
```

6. Use `innobackupex` to load the image

```
innobackupex --defaults-file=/etc/my.cnf --move-back /root/xtrabackuptmp/
```

After loading succeeds, `completed OK!` will be displayed.

```

160530 15:17:36 [01] Moving ./performance_schema/mutex_instances.frm to /var/lib/mysql/performance_schema/mutex_instances.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_statements_history_long.frm to /var/lib/mysql/performance_schema/events_statements_history_long.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_waits_summary_by_thread_by_event_name.frm to /var/lib/mysql/performance_schema/events_waits_summary_by_thread_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_stages_summary_by_user_by_event_name.frm to /var/lib/mysql/performance_schema/events_stages_summary_by_user_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_statements_summary_global_by_event_name.frm to /var/lib/mysql/performance_schema/events_statements_summary_global_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/users.frm to /var/lib/mysql/performance_schema/users.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_statements_history.frm to /var/lib/mysql/performance_schema/events_statements_history.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_waits_summary_by_host_by_event_name.frm to /var/lib/mysql/performance_schema/events_waits_summary_by_host_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/socket_summary_by_instance.frm to /var/lib/mysql/performance_schema/socket_summary_by_instance.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_statements_summary_by_digest.frm to /var/lib/mysql/performance_schema/events_statements_summary_by_digest.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_waits_summary_by_user_by_event_name.frm to /var/lib/mysql/performance_schema/events_waits_summary_by_user_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_waits_summary_by_account_by_event_name.frm to /var/lib/mysql/performance_schema/events_waits_summary_by_account_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/events_stages_summary_by_host_by_event_name.frm to /var/lib/mysql/performance_schema/events_stages_summary_by_host_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 [01] Moving ./performance_schema/file_summary_by_event_name.frm to /var/lib/mysql/performance_schema/file_summary_by_event_name.frm
160530 15:17:36 [01] ...done
160530 15:17:36 completed OK!

```

7. Start the database

```

chmod 777 -R /var/lib/mysql
service start mysql

```

If you fail to start the database, you need to check and fix the error, and then try again.

8. Connect to the database to check data

After starting the database, you may need to connect to the database with the original account and password to view data.

Restoring Databases from Backup Files (Encrypted)

TDE is only supported for Percona 5.7 in Hong Kong region and MySQL 8.0.24 , but it will be available to more kernel versions in the future. You can access **Data Security > Data Encryption** on the instance management page in the [TencentDB for MariaDB console] (<https://console.intl.cloud.tencent.com/mariadb>)

After data encryption is enabled, the database instances can't be restored from a backup file. It is recommended to restore them as instructed in [Rolling back Databases] (<https://intl.cloud.tencent.com/document/product/237/8719>).

Note:

To use the data encryption feature, [submit a ticket] (<https://console.intl.cloud.tencent.com/workorder/category>) to apply for it.

Rolling back Databases

Last updated : 2024-01-11 15:28:38

Rollback description

TencentDB for MariaDB can roll back data to any time point in the last 30 days based on the retention of backups and logs. With the database rollback feature, system loss can be minimized.

This rollback feature doesn't affect a production instance and can directly roll back data to a newly created pay-as-you-go instance. The new rollback instance is a standard one that can be configured as needed.

Limits


When a temp instance is being rolled back or created, some management features of the production instance are unavailable until the operation is completed.

The binlogs may be forcibly sharded during the rollback operation, and files smaller than 100 MB in size will be backed up separately.

The newly purchased instance after rollback will have the parameter information of the production instance (such as account and database parameters). Therefore, you must pay attention to account management.

Instance rollback

1. Log in to the [TencentDB for MariaDB console](#) and click an instance ID to enter the instance management page.
2. On the instance management page, select the **Backup and Restoration** > **Rollback Instance** tab and click **Create Rollback Instance**.



3. In the pop-up window, set the rollback time and click **OK**.

Rollback Settings



Create a pay-as-you-go instance (the rollback instance) based on the data of the selected instance (the data source instance) backed up at the specified point in time

Set Rollback Time *

Select time



1. The rollback time must be within the backup retention period (2022-05-19 16:09:21 - 2022-05-26 14:44:29).
2. The pay-as-you-go instance will be created based on backup data without affecting the data source instance.

OK

Cancel

4. On the instance purchase page, adjust configuration based on your needs, click **Buy Now**, and wait for instance rollback to be completed.

5. After the rollback is completed, you can view the generated rollback instance on the **Backup and Restoration > Rollback Instance** page or in the instance list.

Modifying Data Replication Mode

Last updated : 2024-01-11 15:28:38

Data Replication Mode

Data replication mode, aka data sync mode, is a mechanism for data replication of primary and replica nodes in the high availability scheme of database. Currently, TencentDB for MariaDB supports the following modes:

Async replication: an application initiates an update request such as adding, deletion, or modification. After completing the corresponding operation, the primary node (primary) responds to the application immediately and replicates data to the replica node (replica) asynchronously. Therefore, in async replication mode, an unavailable replica does not affect operations on the primary, while an unavailable primary may cause data inconsistency.

Strong sync (non-downgradable) replication: an application initiates an update request. After completing the operation, the primary replicates data to the replica immediately. After receiving the data, the replica returns a success message to the primary. Only after receiving the message from the replica will the primary respond to the application. The data is replicated synchronously from the primary to the replica. Therefore, an unavailable replica will affect operations on the primary, while an unavailable primary will not cause data inconsistency.

Note:

When you perform strong sync replication, the primary database will be hanged if it is disconnected from the replica database or the replica database fails. If there is only one primary or replica database, the high-availability scheme is unavailable, because if only one single server is used, part of data will be lost completely when a failure occurs, which does not meet the requirements for finance-level data security.

Strong sync (downgradable) replication: batch processing and writing large amounts of transaction data in the business system may cause severe delay in the replica; moreover, in strong sync (non-downgradable) mode, the remaining node will be hanged. These mechanisms designed to ensure data consistency may cause exceptions in the business system.

To address this problem, TencentDB provides a scheme where the strong sync mechanism can be downgraded to async mode. When the replica delay is longer than or equal to 15 seconds, the system will automatically downgrade the strong sync mode to async mode; if it then becomes shorter than 15 seconds, the system will automatically upgrade the async mode to strong sync mode. In this way, strong sync (downgradable) replication is an efficient scheme to ensure eventual data consistency.

Note:

Different from the open-source semi-sync mechanism of Google, strong sync replication uses thread pools instead of the worker thread mode. Besides, the downgradable scheme is better than the semi-sync mode.

Modifying Data Replication Mode

Note:


The one-primary-one-replica architecture of TencentDB for MariaDB provides only two schemes: strong sync (downgradable) replication and async replication. If data consistency is required, please purchase the three-node edition with one primary and two replicas.

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance details page.
2. Click the edit icon next to **Data Replication Mode** in the **Availability Info** section.

Note:

The modification will not affect normal operations of the instance and will take effect in up to five seconds.

Availability Info

Data Replication Mode	Strong sync (downgradable) 
Deployment Mode	Single-AZ
Primary Server AZ	Guangzhou Zone 3 Primary-Replica Switch
Replica Server AZ	Guangzhou Zone 3

Migrating Data

Importing Data to TencentDB for MariaDB

Instances with DTS

Last updated : 2024-01-11 15:28:38

The database migration tool for TencentDB for MariaDB has been upgraded, and its entry has been moved to [DTS](#). For more information on migration, see [Migration from MySQL to TencentDB for MariaDB \(MySQL/MariaDB/Percona\)](#).

Importing Data with mysqldump

Last updated : 2024-01-11 15:28:38

mysqldump is easy to use for data import but needs long downtime, so it is only suitable for small amounts of data or situations that allow relatively long downtime.

1. Use mysqldump to export data from the local database to a data file.

Note:

Do not update data during the export. This step only exports data excluding procedures, triggers, and functions.

The export account needs to have the permission of `select on *.*`.

```
mysqldump -h localIp -u userName -p --opt --default-character-set=utf8 --hex-  
blob dbName --skip-triggers > /tmp/dbName.sql
```

Parameters:

localIp: IP address of the local database server.

userName: Migration account of the local database.

dbName: Name of the database that needs to be migrated.

/tmp/dbName.sql: Name of the generated backup file.

2. Use mysqldump to export procedures, triggers, and functions.

Note:

If the database does not use procedures, triggers, or functions, you can skip this step. During the export, you need to remove the definer in order to make the data compatible with TencentDB.

```
mysqldump -h localIp -u userName -p --opt --default-character-set=utf8 --hex-  
blob dbName -R | sed -e 's/DEFINER[ ]*=[ ]*[^(]*\*/\*/' >  
/tmp/triggerProcedure.sql
```

Parameters:

localIp: IP address of the local database server.

userName: Migration account of the local database.

dbName: Name of the database that needs to be migrated.

/tmp/triggerProcedure.sql: Name of the generated backup file.

3. Upload the data file and procedure file to a CVM instance. Make sure that the CVM and TencentDB instances can be properly connected and there is sufficient available storage capacity in the CVM instance.

4. Log in to CVM and import the data file and procedure files to the target TencentDB instance. Make sure that you have the database account with corresponding permissions; otherwise, you need to generate an account in the console.

```
mysql -h xxx.xxx.xxx.xxx:xxxx -u userName -p dbName < /tmp/dbName.sql  
mysql -h xxx.xxx.xxx.xxx:xxxx -u userName -p dbName < /tmp/triggerProcedure.sql
```

Parameters:

xxx.xxx.xxx.xxx:xxx: Instance connection address. In this document, a private network address is used as an example.

userName: Migration account of the TencentDB instance.

dbName: Name of the database that needs to be imported to.

/tmp/dbName.sql: Name of the data file that needs to be imported.

/tmp/triggerProcedure.sql: Name of the procedure file that needs to be imported.

Database Sync(Legacy)

Overview

Last updated : 2024-01-11 15:28:38

Data Sync Requirements

The data sync feature supports real-time data sync between two database sources. It is suitable for various business scenarios such as local data disaster recovery, query and report analysis, business intelligence (BI) analysis, and real-time data warehousing.

Supported instance types

Source Instance Type	Target Instance Type
Public cloud TencentDB for MariaDB/MySQL	**Public cloud** TencentDB for MariaDB/MySQL
	Public cloud TencentDB for PostgreSQL
	Public cloud CKafka
	Private cloud TencentDB for MariaDB/MySQL
Public cloud TDSQL	**Public cloud** TDSQL
	Public cloud CKafka
	Private cloud TDSQL

Supported sync topologies

Database sync supports serial topology of "one-to-one one-way sync", "one-to-many one-way sync", "cascading one-way sync", and "many-to-one one-way sync", but does not support the topology of "two-way ring sync".

Supported instance versions

MariaDB 10.0, 10.1

MySQL 5.6, 5.7

PostgreSQL 9.6

All versions of CKafka

Supported sync syntax

DML operations: INSERT, UPDATE, DELETE

DDL operations: CREATE TABLE, ALTER TABLE, ADD COLUMN, DROP COLUMN, RENAME COLUMN

Database sync objects

Sync objects support exact and regex match of databases and tables. Non-table objects such as stored procedures, views, indices, or triggers are not supported currently. If you need to sync a non-table object, please create the same one in the target database.

Database sync fees

Currently, data sync is free of charge.

Database sync network

Currently, database sync is supported only between Tencent Cloud intra-region database instances or databases in Direct Connect. For cross-region data sync or disaster recovery, please [submit a ticket](#) and indicate the IDs of source and target instances and sync requirements. Only after the ticket is approved can the feature be enabled.

Data Sync Feature

Data type conversion

Sync from MariaDB/MySQL to PostgreSQL

MariaDB/MySQL and PostgreSQL have different data types. The database sync tool will convert the data types according to the following rules:

MariaDB/MySQL Data Type	Data Type After Sync to PostgreSQL
tinyint/smallint	smallint
mediumint/int	integer
bigint	bigint
real/double	double precision
float	float
bit	bit
bool/boolean	boolean
char/varchar	char/varchar
binary/varbinary	text
year/date/time/datetime/timestamp	timestamp without time zone
tinyblob/mediumblob/blob/longblob/long	bytea
mediumtext/text/longtext	text

decimal/numeric	decimal/numeric
enum	integer

Note:

The data length of columns in the types listed above stays unchanged during conversion. Types that are not listed above stay unchanged too during conversion.

Sync from MariaDB/MySQL to TDSQL

MariaDB/MySQL and TDSQL have the same data types; therefore, data type conversion is not needed.

Sync from MariaDB/MySQL to CKafka

Relevant data will be converted to JSON format. For more information, please see [Binlog Consumption Format](#).

Sync from TDSQL to CKafka

Relevant data will be converted to JSON format. For more information, please see [Binlog Consumption Format](#).

Data subscription

Data sync to CKafka supports data subscription. Before configuring this feature, you need to purchase a CKafka instance first. For more information, please see [CKafka](#). Below are configuration recommendations:

Specification

You can evaluate the CKafka instance specification based on the daily updated data volume, daily peak or average bandwidth, and desired data retention period in CKafka of the instance to be synced.

For example, if an instance updates 10 million rows every day and each row is 2 KB, the daily write peak volume is 20,000 rows/second, 3 replicas are desired in CKafka, and data needs to be retained in CKafka for 3 days.

For this instance, the daily updated data volume is 20 GB and will be about 30 GB after being converted to JSON format, the retention period is 3 days, 3 replicas need about 270 GB of storage capacity, the peak write bandwidth of the database is about 39 MB/s, and the throughput for CKafka to sustain 3 replicas is about 117 MB/s; therefore, the

Standard Edition of CKafka is sufficient to meet the requirements.

Network

It is recommended to deploy the target instance in the same VPC as the source instance.

Topic selection

A sync task needs one topic, and multiple sync tasks need multiple different topics; otherwise, data may get disorganized.

Parameter selection

The following settings are for your reference only. Select appropriate values based on your actual needs.

Number of partitions: 1

Number of replicas: 3

cleanup.policy: deletion

min.isync.replicas: 2

unclean.leader.election.enable: false

Data consistency

Due to network disconnection or other reasons, messages produced by a database instance to CKafka may be duplicate. In this case, when the consumer replays data, idempotency can be used to clear duplicates.

Note:

In the idempotency scheme, eventual data consistency will be ensured in case messages are executed repeatedly. Specifically, when there are records with primary key conflict, the system will first delete the duplicate data before performing the `insert` operation, forcibly insert the data that cannot be matched by the `update` operation, and directly skip the data that cannot be found during the `delete` operation. Idempotency can be used to achieve data consumption consistency. To use this scheme, there must be a primary key or a non-NULL unique index in the table.

Database Sync Tool IP Range

Last updated : 2024-01-11 15:28:38

The externally exposed IP of the sync linkage from a public cloud TencentDB instance to a private cloud TencentDB instance is usually an IP of the VPC subnet where the Direct Connect gateway resides. If there is a firewall in the private cloud, you can open the IP range of this subnet to the internet.

Connectivity Test

Last updated : 2024-01-11 15:28:38

The connectivity test feature is mainly used to check the connectivity between the sync tool and target database.

Below are the two main check items:

Telnet: whether the network can be properly connected.

MySQL Connect: whether the database can be properly connected.

Test failures

Telnet test failure

Possible causes: incorrectly entered target instance IP, traffic blocked by firewall, or incorrect iptable configuration.

Solution: please check the IP and configure the firewall to allow it.

MySQL Connect test failure

Possible causes: special configuration on the target instance; for example, SSL connection is enabled, the account is configured with HOST, or the `bind_address` configuration is listened on.

Solution: please check the IP and configure the firewall to allow it.

Security Management

Notes on Information Security

Last updated : 2024-01-11 15:28:38

The following statement is hereby made for this document.

1. This document is intended to provide an overview of Tencent Cloud's security measures for TencentDB products and services, which are subject to change. If you have any mandatory requirement, you are recommended to enter into a service level agreement (SLA) with Tencent Cloud. Tencent Cloud makes no guarantees or warranties, express or implied, about the content of this document.
2. This document only involves "part of" the technical security features among the wide range of security features.
3. This document is not intended as a reference document for national or industry-specific information security standards or requirements.
4. This document has been adapted for readability. In the event of any ambiguity or inaccuracy, please refer to Item 1.
5. Tencent Cloud reserves the right to interpret this document.

1. Overview

TencentDB has passed and meets the security requirements of the following certifications:

ISO22301 Certification

ISO27001 Certification

ISO20000 Certification

ISO9001 Certification

Trusted Cloud Service Certification

Cybersecurity Classified Protection Certification (Level 3)

STAR Certification

Some features of TencentDB are designed based on the following standards:

GBT 20273-2006 Information Security Technology - Security Techniques Requirement for Database Management System (Level 2 or Above)

JRT 0072-2012 Testing and Evaluation Guide for Classified Protection of Information System of Financial Industry (Level 4)

2. Tencent Cloud TencentDB Service Security Protection (OPS Security Description)

2.1. Overview

Management and technical security requirements of TencentDB comply with China's Cybersecurity Classified Protection (Level 3). Some of the product features meet the standards of Classified Protection of Information System of Financial Industry (Level 4).

2.2. Internal personnel and system authentication

To improve the security of database server system and ensure the security of various OPS activities, Tencent Cloud has implemented a series of security reinforcement measures, including but not limited to:

Tencent Cloud carries out identification and authentication for users who log in to the operating system and database system, and guarantees the uniqueness of usernames.

Usernames and passwords must be configured as required. A password must contain at least 8 characters of 3 types and should be changed regularly.

The login failure processing mechanism can be enabled to take actions such as ending session, limiting the number of unauthorized login attempts, and automatically exiting in case of login failures.

Access to the system during remote management is under monitoring by Tencent Enterprise IT, and internal risk management and audit are provided, with all sensitive operations encrypted.

Two-factor authentication (dynamic token and password) is required for database server admins when they log in to the OPS system.

2.3. Internal personnel and system access control

For TencentDB management systems and admins, a discretionary access control scheme is implemented, including but not limited to:

Internal OPS staff and systems are controlled based on Tencent Cloud security policies (audit requirements are met). The granularity of a subject is down to user level, and that of an object to database table level.

Strict code management and access control are implemented.

High-risk systems can only be accessed over Tencent private network (development network), which is physically isolated from the internet.

2.4. Internal security audit

A comprehensive security audit and risk management mechanism is provided: audit features include but are not limited to audit for database operations, management system operations, file operations, external device operations, unauthorized external connections, IP address changes, and services and processes.

The audit range covers each operating system user and database user in the server, with crucial security-related system events audited, such as Tencent Cloud admin behaviors, exceptional system resource usage, and use of

important system commands. Audit records contain information like event date, time, type, subject ID, object ID, and result, and can be stored for over a year in a location with a higher level of security in order to avoid unexpected deletion, modification, or overwriting.

Database security audit: all operations on the database servers and databases will be audited by the database security audit system.

Management system operation audit: Tencent Cloud keeps detailed logs of all operations in both internal and external management systems for effective risk traceability.

Routine risk assessment: Tencent Cloud security team performs security assessment on database OPS management on a regular basis.

2.5. Internal intrusion prevention

Tencent Cloud takes multi-dimensional approaches to intrusion prevention for database servers:

The intrusion detection system can defend against intrusions into database servers.

Vulnerability scanning is deployed, and system security inspection is performed periodically.

The device security management system is deployed, and the patch distributing module is enabled to update systems with patches timely.

The operating system is installed on a minimal installation basis, with only necessary components and applications installed and unwanted services disabled.

Reinforcement is implemented on other security configurations based on system type.

2.6. Backup and restore

TencentDB provides data backup and restore features by default.

2.7. Secure reuse of objects

For returned or replaced devices, Tencent Cloud will clear the residual information promptly, so that the storage capacity (memory and disk) where the previous user's sensitive information such as authentication information, files, directories, and database records is stored will be released in time or completely cleared before the devices are reassigned to other users.

2.8. Non-repudiation

Tencent Cloud's internal OPS personnel are required to go through a two-factor authentication and non-repudiation process when logging in to the system. All the personnel involved have signed a NDA.

CAM

Overview

Last updated : 2024-01-11 15:28:38

If you use multiple Tencent Cloud services such as TencentDB, CVM, and VPC which are managed by different users sharing your Tencent Cloud account key, the following problems may exist:

Your password is shared by multiple users, leading to high risk of compromise.

You cannot limit the access permission of other users, which is easy to pose a security risk due to faulty operations.

This is exactly why CAM has been developed. For a detailed description of CAM, please see [CAM Overview](#).

After connecting to CAM, you can allow different users to manage different services through sub-accounts so as to avoid the above problems. By default, a sub-account doesn't have permission to use a TencentDB instance or related resources. Therefore, you need to create a policy to grant the required permission to the sub-account.

A policy is a syntax rule used to define and describe one or more permissions. It can authorize or deny the use of the designated resources by a user or user group. For more information on CAM policy, please see [Policy Syntax](#). For more information on how to use a CAM policy, please see [Policy](#).

If you do not need to manage the access permission to TencentDB resources for sub-accounts, you can skip this chapter. This will not affect your understanding and usage of other parts in the documentation.

Getting Started

A CAM policy must authorize or deny the use of one or more TencentDB operations. At the same time, it must specify the resources that can be used for the operations (which can be all resources or partial resources for certain operations). A policy can also include the conditions set for the manipulated resources.

Note:

You are recommended to manage TencentDB resources and authorize TencentDB operations through CAM policies. Although the experience stays the same for existing users who are granted permission by project, it is not recommended to continue managing resources and authorizing operations in a project-based manner. Effectiveness conditions cannot be set in TencentDB for the time being.

Policy Structure

Last updated : 2024-01-11 15:28:38

Policy Syntax

CAM policy configuration example:

```
{
  "version": "2.0",
  "statement": [
    {
      "effect": "effect",
      "action": ["action"],
      "resource": ["resource"],
      "condition": {"key": {"value"}}
    }
  ]
}
```

version is required. Currently, only "2.0" is allowed. (This value actually represents the version of TencentCloud APIs acceptable to CAM.)

statement describes the details of one or more permissions. This element contains a permission or permission set of other elements such as effect, action, resource, and condition. One policy has only one statement.

action describes the allowed or denied action. An action entered here is a string prefixed with "mariadb:" and suffixed with an [TencentDB for MariaDB API](#). This element is required.

resource describes the details of authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product. For more information on how to specify a resource, please see the documentation for the product whose resources you are writing a statement for. This element is required.

condition describes the condition for the policy to take effect. A condition consists of operator, action key, and action value. A condition value may contain information such as time and IP address. Some services allow you to specify additional values in a condition. This element is required.

effect describes whether the result produced by the statement is "allowed" (allow) or "denied" (deny). This element is required.

Note:

The API keyword in CAM of TencentDB for MariaDB is "mariadb".

Operations in TencentDB

In a TencentDB policy statement, you can specify any API operation from any service that supports TencentDB. APIs prefixed with "mariadb:" should be used for TencentDB, such as `mariadb:`

`mariadb:CloseDBExtranetAccess` (disabling public network access).

To specify multiple operations in a single statement, separate them with commas as shown below:

```
"action":["mariadb:action1","mariadb:action2"]
```

You can also specify multiple operations using a wildcard. For example, you can specify all operations beginning with "Describe" in name as shown below:

```
"action":["mariadb:Describe*"]
```

If you want to specify all operations in TencentDB, use a wildcard as shown below:

```
"action":["mariadb:*"]
```

TencentDB Resources

Each CAM policy statement has its own resources.

Resources are generally in the following format:

```
qcs:project_id:service_type:region:account:resource
```

project_id describes the project information, which is only used to enable compatibility with legacy CAM logic and can be left empty.

service_type describes the product abbreviation such as TencentDB for MariaDB.

region describes the region information, such as ap-guangzhou. For more information, please see [Regions](#).

account is the root account of the resource owner, such as uin/65xxx763.

resource describes detailed resource information of each product, such as instance/instance_id1 or instance/*.

For example:

You can specify a resource for a specific instance (tdsql-k05xdcta) in a statement as shown below:

```
"resource":["qcs::mariadb:ap-guangzhou:uin/65xxx763:instance/tdsql-k05xdcta"]
```

You can also use the wildcard "*" to specify it for all instances that belong to a specific account as shown below:

```
"resource":["qcs::mariadb:ap-guangzhou:uin/65xxx763:instance/*"]
```


If you want to specify all resources or a specific API operation does not support resource-level permission control, you can use the wildcard "*" in the "Resource" element as shown below:

```
"resource": ["*"]
```

To specify multiple resources in a single command, separate them with commas. Below is an example where two resources are specified:

```
"resource": ["resource1", "resource2"]
```

The table below describes the resources that can be used by TencentDB and the corresponding resource description methods.

In the table, words prefixed with \$ are placeholders.

`region` is region.

`account` is account ID.

Resource	Resource Description Method in Authorization Policy
Instance	<code>qcs::mariadb:\$region:\$account:instance/\$instanceId</code>

Resource-level Permissions Supported

Last updated : 2024-01-11 15:28:38

The API keyword in CAM of TencentDB for MariaDB is "mariadb".

Resource-level permission can be used to specify which resources a user can manipulate. TencentDB supports certain resource-level permission. This means that for some TencentDB operations, you can control the time when a user is allowed to perform operations (based on mandatory conditions) or to use specified resources. The following table describes the types of resources that can be authorized in TencentDB.

Types of resources that can be authorized in CAM:

Resource Type	Resource Description Method in Authorization Policy
TencentDB instance-related	<code>qcs::mariadb:\$region:\$account:instance/*</code> <code>qcs::mariadb:\$region:\$account:instance/*</code>

The table below lists the TencentDB API operations which currently support resource-level permission control as well as the resources and condition keys supported by each operation. When specifying a resource path, you can use the "*" wildcard in the path.

Any TencentDB API operation not listed in the table does not support resource-level permission. For such an operation, you can still authorize a user to perform it, but you must specify "*" as the resource element in the policy statement.

The following operations support resource-level permission control

Operation Name	API Name	Effective in Console After Configuration
Recovering dedicated instance	ActiveDedicatedDBInstance	Yes
Binding security group	AssociateSecurityGroups	Yes
Checking IP status	CheckIpStatus	Yes
Cloning account	CloneAccount	Yes

Disabling public network address	CloseDBExtranetAccess	Yes
Copying permission	CopyAccountPrivileges	Yes
Creating account	CreateAccount	Yes
Creating parameter template	CreateConfigTemplate	Yes
Creating instance	CreateDBInstance	Yes
Rolling back instances	CreateTmpInstances	Yes
Deleting account	DeleteAccount	Yes
Deleting parameter template	DeleteConfigTemplate	Yes
Deleting temp instance	DeleteTmpInstance	Yes
Getting the list of permissions	DescribeAccountPrivileges	Yes
Querying the list of accounts	DescribeAccounts	Yes
Querying audit logs	DescribeAuditLogs	Yes
Querying audit rule details	DescribeAuditRuleDetail	Yes
Querying list of audit rules	DescribeAuditRules	Yes
Querying audit policies	DescribeAuditStrategies	Yes
Getting custom backup time	DescribeBackupTime	Yes
Querying price for batch instance renewal	DescribeBatchRenewalPrice	Yes
Querying binlog time	DescribeBinlogTime	Yes
Querying parameter configuration history	DescribeConfigHistories	Yes
Querying parameter template	DescribeConfigTemplate	Yes
Querying the list of parameter templates	DescribeConfigTemplates	Yes
Querying instance objects	DescribeDatabaseObjects	Yes
Querying instance database names	DescribeDatabases	Yes
Querying the column information of instance table	DescribeDatabaseTable	Yes
Querying monitoring information details	DescribeDBDetailMetrics	Yes

Querying the key information of instance	DescribeDBEncryptAttributes	Yes
Querying instance details	DescribeDBInstanceDetail	Yes
Querying the HA information of instance	DescribeDBInstanceHAInfo	Yes
Querying instance specification	DescribeDBInstanceSpecs	Yes
Querying the list of instances	DescribeDBInstances	Yes
Getting the list of logs	DescribeDBLogFiles	Yes
Querying monitoring information	DescribeDBMetrics	Yes
Viewing database parameters	DescribeDBParameters	Yes
Viewing instance performance data	DescribeDBPerformance	Yes
Viewing instance performance data details	DescribeDBPerformanceDetails	Yes
Viewing instance resource usage	DescribeDBResourceUsage	Yes
Viewing instance resource usage details	DescribeDBResourceUsageDetails	Yes
Querying the security group information of instance	DescribeDBSecurityGroups	Yes
Getting slow query log details	DescribeDBSlowLogAnalysis	Yes
Querying the list of slow logs	DescribeDBSlowLogs	Yes
Querying instance sync mode	DescribeDBSyncMode	Yes
Querying temp instances	DescribeDBTmpInstances	Yes
Querying default parameter template	DescribeDefaultConfigTemplate	Yes
Querying dedicated cluster specification	DescribeFenceDBInstanceSpecs	Yes
Querying flow status	DescribeFlow	Yes
Querying the proxy configuration of instance	DescribeInstanceProxyConfig	Yes
Querying the SSL status of instance	DescribeInstanceSSLAttributes	Yes
Querying latest DBA check result	DescribeLatestCloudDBAReport	Yes
Viewing backup log settings	DescribeLogFileRetentionPeriod	Yes

Querying order information	DescribeOrders	Yes
Querying price	DescribePrice	Yes
Querying projects	DescribeProjects	Yes
Querying the security group information of project	DescribeProjectSecurityGroups	Yes
Querying the renewal price of instance	DescribeRenewalPrice	Yes
Querying purchasable AZs	DescribeSaleInfo	Yes
Getting SQL logs	DescribeSqlLogs	Yes
Querying the list of sync tasks	DescribeSyncTasks	Yes
Querying the upgrade price of instance	DescribeUpgradePrice	Yes
Querying the information of user tasks	DescribeUserTasks	Yes
Unbinding security groups from Tencent Cloud resources in batches	DisassociateSecurityGroups	Yes
Setting permissions	GrantAccountPrivileges	Yes
Initializing instances	InitDBInstances	Yes
Isolating dedicated instance	IsolateDedicatedDBInstance	Yes
Modifying account remarks	ModifyAccountDescription	Yes
Setting auto-renewal in batches	ModifyAutoRenewFlag	Yes
Setting custom backup time	ModifyBackupTime	Yes
Modifying parameter template	ModifyConfigTemplate	Yes
Modifying instance encryption information	ModifyDBEncryptAttributes	Yes
Renaming instance	ModifyDBInstanceName	Yes
Modifying security groups bound to TencentDB instance	ModifyDBInstanceSecurityGroups	Yes
Modifying instance project	ModifyDBInstancesProject	Yes
Modifying database parameters	ModifyDBParameters	Yes
Modifying instance sync mode	ModifyDBSyncMode	Yes

Modifying instance network	ModifyInstanceNetwork	Yes
Modifying remarks	ModifyInstanceRemark	Yes
Modifying SSL information	ModifyInstanceSSLAttributes	Yes
Modifying instance VIP	ModifyInstanceVip	Yes
Modifying backup log settings	ModifyLogFileRetentionPeriod	Yes
Enabling public network address	OpenDBExtranetAccess	Yes
Renewing instance	RenewDBInstance	Yes
Resetting account password	ResetAccountPassword	Yes
Enabling smart DBA	StartSmartDBA	Yes
Switching instance HA	SwitchDBInstanceHA	Yes
Replacing original instance with temp instance	SwitchRollbackInstance	Yes
Terminating dedicated instance	TerminateDedicatedDBInstance	Yes
Scaling up instance	UpgradeDBInstance	Yes
Upgrading dedicated instance	UpgradeDedicatedDBInstance	Yes

Console Examples

Last updated : 2024-01-11 15:28:38

Sample CAM Policies for TencentDB

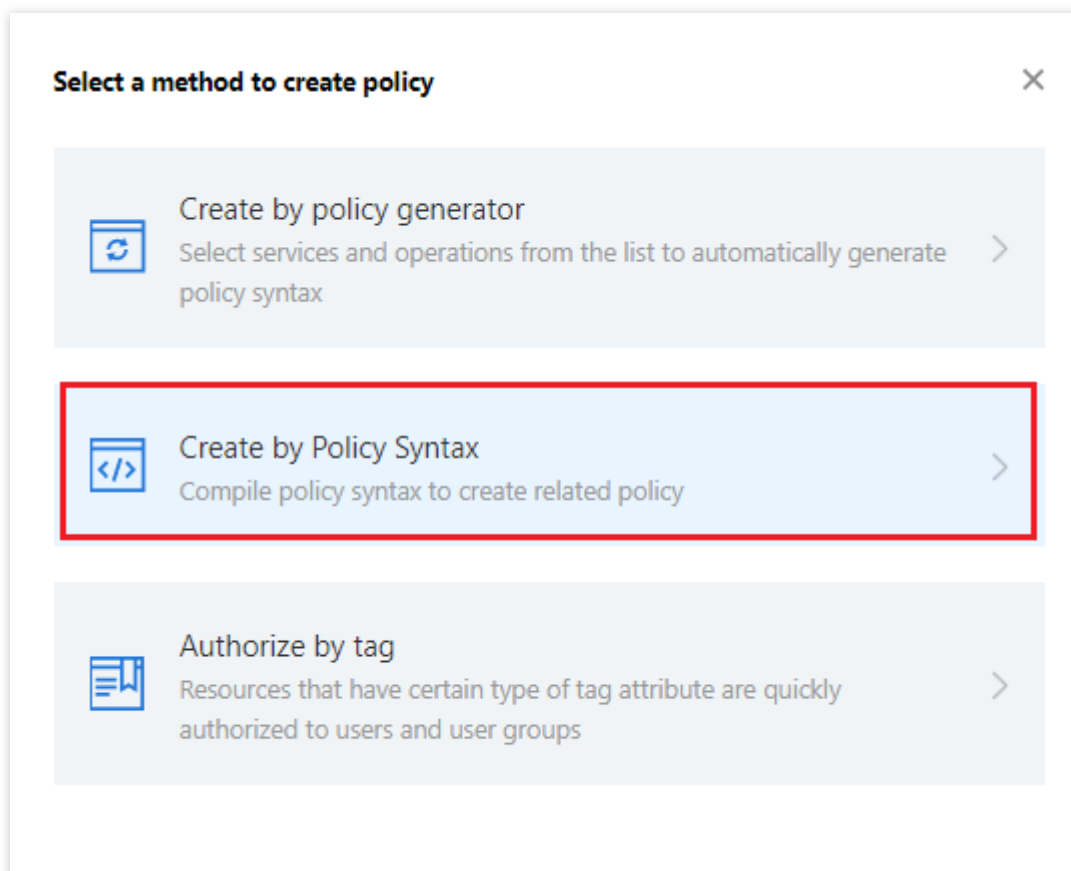
You can grant a user the permission to view and use specific resources in the TencentDB console by using a CAM policy. The sample below shows how to allow a user to use certain policies in the console. Currently, you can configure TencentDB for MariaDB to support the CAM feature only by using the **creation by policy syntax** method.

Note:

The API keyword of TencentDB for MariaDB in CAM is `mariadb`.

Step 1. Create a custom policy

1. Enter the [Policy Syntax](#) configuration page and click **Create Custom Policy**.
2. Click **Create by Policy Syntax** in the pop-up window.



3. Select **Blank Template** and click **Next**.

✓

Select policy template

>

2

Edit Policy

Template Type: All Templates

Search policy name

Select template type

All Templates (282 items in total)

Blank Template

Custom

AdministratorAccess

System

This policy allows you to manage all users under your account and their permissions, financial information and cloud assets.

QCloudResourceFullAccess

System

This policy allows you to manage all cloud assets in your account.

QCloudFinanceFullAccess

System

This policy allows you to manage all financial items in your account, such as payment and billing.

QcloudNARMSReadOnlyAccess

System

QcloudNARMSReadOnlyAccess

QcloudAccessForAegisRole

System

Aegis' access to cloud resources

QcloudAccessForBAASRole

System

QcloudAccessForBAASRole

QcloudAccessForBKRole

System

BlueKing\'s access to cloud resources

QcloudAccessForCDNRole

System

CDN permissions (including but not limited to): CLS (add/delete/modify/query CLS logsets/log topics, and search...)

QcloudAccessForCFSRole

System

CFS permissions(including but not limited to): KMS (create data key, encrypt/decrypt data using data key)

QcloudAccessForCloudAuditCARole

System

QcloudAccessForCloudStudioRole

System

Next

4. Enter the corresponding policy syntax.

✓ Select policy template

>

2 Edit Policy

Policy Name *

policygen-20200622144345

Description

Policy content

1

2

3

4

```
"version": "2.0",  
"statement": []
```

Previous

Done

Step 2. Associate the sub-account/collaborator and verify

After the policy is created, associate it with a user/group. After the association is completed, use another browser (or server) to verify whether the sub-account/collaborator can work normally. If the policy syntax is written correctly, you can observe the following:

You can access the target products/resources and use all features as expected.

You will be prompted with "You have no permission for this operation" when accessing other unauthorized products or resources.

To avoid mutual impact of multiple policies, we recommend that you associate only one policy with a sub-account at a time.

The change to account access permission will take effect within 1 minute.

Appendix. Commonly Used Policy Syntax

Policy for authorizing the use of all features in all TencentDB instances

To grant a user permission to create and manage TencentDB instances, implement the policy named

`QcloudMariaDBFullAccess` for the user.

The policy syntax is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "mariadb:*"
      ],
      "resource": "*",
      "effect": "allow"
    }
  ]
}
```

Policy for authorizing the query of all TencentDB instances

To grant a user permission to view TencentDB instances but not create, delete, or modify them, implement the policy named `QcloudMariaDBInnerReadOnlyAccess` for the user.

The policy syntax is as follows:

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "mariadb:Describe*"
      ],
      "resource": "*",
      "effect": "allow"
    }
  ]
}
```

The above policy achieves its goal by allowing the user to separately authorize the use of all operations beginning with "Describe" in TencentDB with the CAM policy.

Note:

Because not all functional APIs are now supported, a limited number of operations may be excluded from CAM, which is normal.

Policy for granting a user permission to manipulate TencentDB instances in one specific region

To grant a user permissions to manipulate TencentDB instances in a specific region, you can associate the following policy with the user. For example, the policy below allows the user to manipulate the TencentDB instances in Guangzhou.

```
{
  "version": "2.0",
  "statement": [
    {
      "action": "mariadb:*",
      "resource": "qcs::mariadb:ap-guangzhou:*",
      "effect": "allow"
    }
  ]
}
```

Policy for granting a user permission to manipulate TencentDB instances in multiple specific regions

To grant a user the permission to manipulate TencentDB instances in a specific region, associate the following policy with the user. For example, the policy below allows the user to manipulate the TencentDB instances in Guangzhou and Chengdu.

```
{
  "version": "2.0",
  "statement": [
    {
      "action": "mariadb:*",
      "resource": "qcs::mariadb:ap-guangzhou:*","qcs::mariadb:ap-chengdu:*",
      "effect": "allow"
    }
  ]
}
```

Policy for granting a user permission to manipulate one specific TencentDB instance

To grant a user the permission to manipulate a specific database, associate the following policy with the user. For example, the policy below allows the user to manipulate the TencentDB instance "tdsql-xxx" in Guangzhou.

```
{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "mariadb:*"
      ],
      "resource": "qcs::mariadb:ap-guangzhou::instance/tdsql-xxx",
      "effect": "allow"
    }
  ]
}
```

Policy for granting a user permission to manipulate multiple TencentDB instances

To grant a user the permission to manipulate TencentDB instances in batches, associate the following policy with the user. For example, the policy below allows the user to manipulate the TencentDB instances "tdsql-xxx" and "tdsql-yyy" in Guangzhou and "tdsql-zzz" in Beijing.

```
{
  "version": "2.0",
  "statement": [
    {
      "action": "mariadb:*",
      "resource": ["qcs::mariadb:ap-guangzhou::instance/tdsql-xxx",
        "qcs::mariadb:ap-guangzhou::instance/tdsql-yyy", "qcs::mariadb:ap-
        beijing::instance/tdsql-zzz"],
      "effect": "allow"
    }
  ]
}
```

Policy for granting a user different permissions to manipulate multiple TencentDB instances

To grant a user the permission to manipulate TencentDB instances in batches, associate the following policy with the user. For example, the policy below allows the user to manipulate the TencentDB instances "tdsql-xxx" and "tdsql-yyy" in Guangzhou and "tdsql-zzz" in Beijing.

```
{
  "version": "2.0",
  "statement": [
    {
```

```

        "action": "mariadb:Describe*", "mariadb:Create*",
        "resource": ["qcs::mariadb:ap-guangzhou::instance/tdsql-xxx",
"qcs::mariadb:ap-guangzhou::instance/tdsql-yyy", "qcs::mariadb:ap-
beijing::instance/tdsql-zzz"],
        "effect": "allow"
    }
]
}

```

Denying a user permission to create TencentDB accounts

To deny a user permission to create TencentDB accounts, configure `"effect": "deny"` as shown below.

```

{
  "version": "2.0",
  "statement": [
    {
      "action": "mariadb:CreateAccount",
      "resource": "*",
      "effect": "deny"
    }
  ]
}

```

Other custom policies

If preset policies cannot meet your requirements, you can create custom policies as shown below:

```

{
  "version": "2.0",
  "statement": [
    {
      "action": [
        "Action"
      ],
      "resource": "Resource",
      "effect": "Effect"
    }
  ]
}

```

Replace "Action" with the operation to be allowed or denied.

Replace "Resource" with the resources that you want to authorize the user to manipulate.

Replace "Effect" with "Allow" or "Deny".

CAM-enabled Operations

Last updated : 2024-01-11 15:28:38

The following operations support resource-level permission control

Operation Name	API Name	Effective in Console After Configuration
Recovering dedicated instance	ActiveDedicatedDBInstance	Yes
Binding security group	AssociateSecurityGroups	Yes
Checking IP status	CheckIpStatus	Yes
Cloning account	CloneAccount	Yes
Disabling public network address	CloseDBExtranetAccess	Yes
Copying permission	CopyAccountPrivileges	Yes
Creating account	CreateAccount	Yes
Creating parameter template	CreateConfigTemplate	Yes
Creating instance	CreateDBInstance	Yes
Rolling back instances	CreateTmpInstances	Yes
Deleting account	DeleteAccount	Yes
Deleting parameter template	DeleteConfigTemplate	Yes
Deleting temp instance	DeleteTmpInstance	Yes
Getting list of permissions	DescribeAccountPrivileges	Yes
Querying the list of accounts	DescribeAccounts	Yes
Querying audit logs	DescribeAuditLogs	Yes
Querying audit rule details	DescribeAuditRuleDetail	Yes
Querying the list of audit rules	DescribeAuditRules	Yes
Querying audit policies	DescribeAuditStrategies	Yes
Getting custom backup time	DescribeBackupTime	Yes

Querying price for batch instance renewal	DescribeBatchRenewalPrice	Yes
Querying binlog time	DescribeBinlogTime	Yes
Querying parameter configuration history	DescribeConfigHistories	Yes
Querying parameter template	DescribeConfigTemplate	Yes
Querying the list of parameter templates	DescribeConfigTemplates	Yes
Querying instance objects	DescribeDatabaseObjects	Yes
Querying instance database names	DescribeDatabases	Yes
Querying the column information of instance table	DescribeDatabaseTable	Yes
Querying monitoring information details	DescribeDBDetailMetrics	Yes
Querying key information of instance	DescribeDBEncryptAttributes	Yes
Querying instance details	DescribeDBInstanceDetail	Yes
Querying the HA information of instance	DescribeDBInstanceHAInfo	Yes
Querying instance specification	DescribeDBInstanceSpecs	Yes
Querying the list of instances	DescribeDBInstances	Yes
Getting the list of logs	DescribeDBLogFiles	Yes
Querying monitoring information	DescribeDBMetrics	Yes
Viewing database parameters	DescribeDBParameters	Yes
Viewing instance performance data	DescribeDBPerformance	Yes
Viewing instance performance data details	DescribeDBPerformanceDetails	Yes
Viewing instance resource usage	DescribeDBResourceUsage	Yes
Viewing instance resource usage details	DescribeDBResourceUsageDetails	Yes
Querying the security group information of instance	DescribeDBSecurityGroups	Yes
Getting slow query log details	DescribeDBSlowLogAnalysis	Yes
Querying the list of slow logs	DescribeDBSlowLogs	Yes

Querying instance sync mode	DescribeDBSyncMode	Yes
Querying temp instances	DescribeDBTmpInstances	Yes
Querying default parameter template	DescribeDefaultConfigTemplate	Yes
Querying dedicated cluster specification	DescribeFenceDBInstanceSpecs	Yes
Querying flow status	DescribeFlow	Yes
Querying the proxy configuration of instance	DescribeInstanceProxyConfig	Yes
Querying the SSL status of instance	DescribeInstanceSSLAttributes	Yes
Querying latest DBA check result	DescribeLatestCloudDBAReport	Yes
Viewing backup log settings	DescribeLogFileRetentionPeriod	Yes
Querying order information	DescribeOrders	Yes
Querying price	DescribePrice	Yes
Querying projects	DescribeProjects	Yes
Querying the security group information of project	DescribeProjectSecurityGroups	Yes
Querying the renewal price of instance	DescribeRenewalPrice	Yes
Querying purchasable AZs	DescribeSaleInfo	Yes
Getting SQL logs	DescribeSqlLogs	Yes
Querying the list of sync tasks	DescribeSyncTasks	Yes
Querying the upgrade price of instance	DescribeUpgradePrice	Yes
Querying the information of user tasks	DescribeUserTasks	Yes
Unbinding security groups from Tencent Cloud resources in batches	DisassociateSecurityGroups	Yes
Setting permissions	GrantAccountPrivileges	Yes
Initializing instances	InitDBInstances	Yes
Isolating dedicated instance	IsolateDedicatedDBInstance	Yes
Modifying account remarks	ModifyAccountDescription	Yes

Setting auto-renewal in batches	ModifyAutoRenewFlag	Yes
Setting custom backup time	ModifyBackupTime	Yes
Modifying parameter template	ModifyConfigTemplate	Yes
Modifying instance encryption information	ModifyDBEncryptAttributes	Yes
Renaming instance	ModifyDBInstanceName	Yes
Modifying security groups bound to TencentDB instance	ModifyDBInstanceSecurityGroups	Yes
Modifying instance project	ModifyDBInstancesProject	Yes
Modifying database parameters	ModifyDBParameters	Yes
Modifying instance sync mode	ModifyDBSyncMode	Yes
Modifying instance network	ModifyInstanceNetwork	Yes
Modifying remarks	ModifyInstanceRemark	Yes
Modifying SSL information	ModifyInstanceSSLAttributes	Yes
Modifying instance VIP	ModifyInstanceVip	Yes
Modifying backup log settings	ModifyLogFileRetentionPeriod	Yes
Enabling public network address	OpenDBExtranetAccess	Yes
Renewing instance	RenewDBInstance	Yes
Resetting account password	ResetAccountPassword	Yes
Enabling smart DBA	StartSmartDBA	Yes
Switching instance HA	SwitchDBInstanceHA	Yes
Replacing original instance with temp instance	SwitchRollbackInstance	Yes
Terminating dedicated instance	TerminateDedicatedDBInstance	Yes
Scaling up instance	UpgradeDBInstance	Yes
Upgrading dedicated instance	UpgradeDedicatedDBInstance	Yes

Security Group Configuration

Last updated : 2024-01-11 15:28:38

Security group serves as a stateful virtual firewall with filtering feature for configuring network access control for one or more TencentDB instances. It is an important network security isolation tool provided by Tencent Cloud.

Instances in VPC with the same network security isolation demands in one region can be put into the same security group, which is a logical group (not supported for instances in the classic network currently). TencentDB and CVM share the security group list and are matched with each other within the security group based on rules. Rules not supported by TencentDB will not take effect.

Note:

TencentDB security group currently only supports network access control for VPCs but not the classic network or the public network.

TencentDB Security Group Management

1. Log in to the [TencentDB for MariaDB console](#), click an instance ID in the instance list, and enter the instance management page.
2. Select **Data Security** > **Security Group** to manage TencentDB security groups.

Note:

TencentDB shares the security group rules of CVM. You can match or adjust the rule priority as needed on the TencentDB security group management page.

You cannot create or delete security group rules on the TencentDB security group management page. To do so, see [Virtual Private Cloud > Security Groups](#).

Security Group Policy

Security group policies are divided into "allowing" and "rejecting" traffic. You can configure security group rules to allow or reject inbound traffic of instances deployed in VPC.

Default Policy of a TencentDB Security Group

Currently, if you select VPC as the network type when purchasing a TencentDB instance, there is no need to associate a security group. In this case, the default policy is to "open all IPs and ports to Internet".

Security Group Templates

You can create a custom security group, or create a security group from a template. You can control the inbound and outbound packets of CVMs by configuring security group rules.

All ports opened: the access to TencentDB from all IP addresses is allowed, which comes with certain security risks.

Please use this template with caution.

Security Group Rules

Security group rules are used to control the inbound and outbound traffic of instances associated with the security group (filtered based on the rules from top to bottom). By default, a new security group rejects all traffic (All Drop). You can modify security group rules at any time, and the new rules take effect immediately.

Each security group rule involves the following items:

Protocol and port: as TencentDB only provides access over fixed ports, security group rules configured with other ports won't take effect for TencentDB. For example, if the TencentDB instance uses port 3306 for access, you can configure `TCP:3306` or `ALL` in the security group rule.

Authorization type: access based on address ranges (CIDR/IP).

Source (inbound rules) or target (outbound rules): choose one of the following options:

Specify a single IP in CIDR notation.

Specify an IP address range in CIDR notation, such as 203.0.113.0/24.

Policy: allow or reject access requests

Security Group Priority

You can set security group priority in the TencentDB console, and the smaller the number, the higher the priority. If an instance is associated with multiple security groups, the priority is used as a basis for evaluating the security rules for this instance.

In addition, if the last policy in multiple security groups associated with an instance is **ALL Traffic Denied**, then the last policy **ALL Traffic Denied** of all security groups except the one with the lowest priority will not take effect.

Security Group Restrictions

Security groups are applicable to TencentDB instances in [VPC](#).

Each user can configure a maximum of 50 security groups for each project in a region.

A maximum of 100 inbound or outbound rules can be configured for a security group. As TencentDB does not have active outbound traffic, outbound rules are not applicable to TencentDB.

A TencentDB instance can be associated with multiple security groups, and a security group can be associated with multiple TencentDB instances. No limit is imposed on the number.

Note:

We do not recommend associating too many instances with a security group, although no limit is imposed on the number of instances.

Feature	Quantity
Security group	50/region
Access policy	100 (inbound/outbound)
Number of security groups associated with an instance	No limit
Number of instances associated with a security group	No limit

Creating, Managing, and Deleting Security Group Rules

TencentDB shares the security group rules of CVM. You can match or adjust the rule priority as needed on the TencentDB security group management page.

To create, manage, or delete security group rules, please do so on the [security group management page](#) in the console. For more information, see [Virtual Private Cloud > Security Groups](#).

Transparent Data Encryption (TDE)

Last updated : 2024-01-11 15:28:38

Overview

TencentDB for MariaDB comes with the transparent data encryption (TDE) feature. Transparent encryption means that the data encryption and decryption are transparent to users. TDE supports real-time I/O encryption and decryption of data files. It encrypts data before it is written to disk, and decrypts data when it is read into memory from disk, which meets the compliance requirements of static data encryption.

This document describes how to enable data encryption and encrypt/decrypt data in the console.

Limits

The TDE feature is currently supported only for MySQL 8.0.24 or later and Percona 5.7.

Note:

To use TDE feature, [submit a ticket](#) to apply for it.

[KMS](#) must be activated in advance or as prompted when TDE is enabled.

KMS key permissions must be granted in advance or as prompted when TDE is enabled.

Notes

After KMS is activated, KMS fees may be incurred as detailed in [Purchase Method](#).

TDE can't be disabled once enabled.

If disaster recovery read-only instances are created, TDE cannot be enabled.

After TDE is enabled, disaster recovery read-only instances cannot be created.

After TDE is enabled, the database instances can't be restored from a backup file. You recommend that you restore them as instructed in [Rolling back Databases](#).

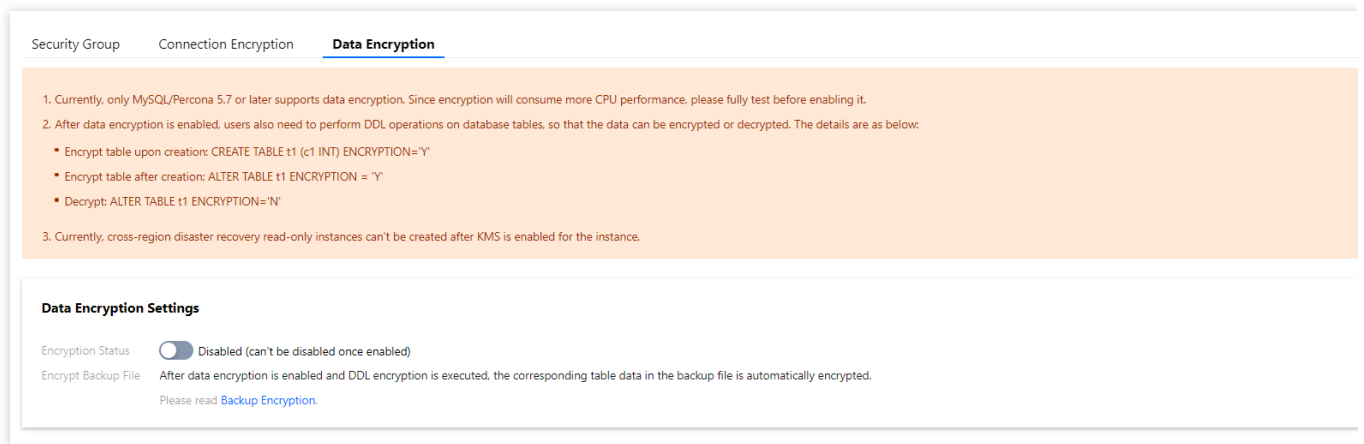
TDE enhances the security of static data while compromising the read-write performance of encrypted databases.

Therefore, use it based on your actual needs.

After TDE is enabled, more CPU resources will be consumed, and about 5% of the performance will be compromised.

Directions

1. Log in to the [TencentDB for MariaDB console](#). Click an instance ID or **Manage** in the **Operation** column to enter the instance management page.
2. On the instance management page, select **Data Security** > **Data Encryption** and toggle on **Encryption Status**.



3. In the pop-up dialog box, activate the KMS, grant the KMS key permissions, select a key, and click **OK**.

Settings

Notes

1. The data encryption feature may incur the Key Management Service (KMS) service fee. For details, see [KMS Purchase Guide](#)
2. After the feature is enabled, DDL operation is required to encrypt or decrypt the tablespace data.
3. If you want to disable this feature, you need to decrypt all the tablespaces, and submit a ticket after disabling the KMS service.
4. You can't create disaster recovery read-only instance for the instance after data encryption is enabled.

KMS Service Disabled [Enable KMS](#) 

KMS Key Authorization Authorized

Select Key ☐ Use auto-generated key

☐ I have read and agreed to the supplementary statement about data encryption in [TencentDB Service Level Agreement](#)

☐ I have read and agreed to [Disclaimer of Tencent Cloud for Information Security During Data Encryption](#)

OK

Cancel

4. After data encryption is enabled, you must perform DDL operations on the database table to encrypt or decrypt data. The detailed steps are as follows:

Encrypt a new table:

```
CREATE TABLE t1 (c1 INT) ENCRYPTION='Y'
```

Encrypt an existing table:

```
ALTER TABLE t1 ENCRYPTION='Y'
```

Decrypt a table:


```
ALTER TABLE t1 ENCRYPTION='N'
```

Monitoring and Alarms

Monitoring Feature

Last updated : 2024-01-11 15:28:38

Performance Monitoring

To make it easier for you to view and stay up to date with how instances work, TencentDB for MariaDB provides a wide variety of performance monitoring metrics. You can log in to the [TencentDB for MariaDB Console](#) and view them on the **System Monitoring** tab on the instance management page.

Slave Monitoring

TencentDB for MariaDB supports monitoring slave performance. On the **System Monitoring** tab, you can switch to slave data as prompted.

Available Monitoring Metrics

Metric Name	Unit	Remarks
CPU Utilization	%	TencentDB for MariaDB adopts flexible CPU limiting, allowing your instance to use extra device CPU resources when they are idle. In this case, CPU utilization may exceed 100%.
Buffer Cache Hit Ratio	%	The percentage of data a SELECT or preprocessing query directly gets from the memory. You are recommended to keep it above 90%; otherwise, you need to increase the memory specification.
Freeable Memory (this metric will be deleted)	GB	Actual memory capacity available in Innodb_buffer. As the database uses the LRU scheduling scheme, normally this value tends to be zero. However, when large transactions are processed, it may be negative, i.e., the used database memory exceeds the assigned capacity.
Storage Space Utilization	%	Percentage of current storage capacity used by instanced data, logs, temporary data, and system files out of the purchased disk capacity. Generally, this value should be smaller than 80% of

		the purchased disk capacity; otherwise, the disk capacity should be expanded.
Free Storage Space	GB	Absolute value of remaining available disk capacity of current instance. Generally, this value should be greater than 20% of the purchased disk capacity; otherwise, the disk capacity should be expanded.
Binary Log Disk Usage	GB	Storage capacity temporarily stored in database on data disk. Generally, this value varies by percentage of written data. Please note that this value is not the log storage capacity in the backup system.
DB Connections	-	Total number of connections from client to database server.
Active Connections	-	Total number of active connections from client to database server.
IO Utilization	%	TencentDB for MariaDB adopts a flexible I/O limiting policy, allowing instances to use extra device I/O resources when they are idle. In this case, I/O utilization may exceed 100%.
SQL Throughput	Statements/second	Total number of DDL, DML, and DCL statements.
SQL Error Throughput	Statements/second	Total number of execution errors of DDL, DML, and DCL statements. If this value is high, please check the service logs as soon as possible.
SQL Success Throughput	Statements/second	Total number of successful executions of DDL, DML, and DCL statements.
Slow Query	Statements/second	Number of SQL statements whose execution time exceeds the set value of <code>long_query_time</code> . For more information, please see the performance optimization page.
DML Latency 1ms - 5ms	Statements/second	Overall statistics of SQL statement execution time.
DML Latency 5ms - 20ms	Statements/second	Overall statistics of SQL statement execution time.
DML Latency 20ms - 30ms	Statements/second	Overall statistics of SQL statement execution time.
DML Latency more than 30ms	Statements/second	Overall statistics of SQL statement execution time.
DML Throughput	Statements/second	Total number of DML statements.

SELECT	Statements/second	Total number of SELECT statements. If this value is high, you can enable read/write separation.
UPDATE	Statements/second	Total number of UPDATE statements.
INSERT	Statements/second	Total number of INSERT statements.
REPLACE	Statements/second	Total number of REPLACE statements.
REPLACE_SELECT	Statements/second	Total number of REPLACE_SELECT statements.
DELETE	Statements/second	Total number of DELETE statements.
Master Switch	-	Number of occurrences of switching from the master to the slave.
Replica Lag	Seconds	Data delay between the slave and master. With strong sync, the master will return a transaction response only when data is written to the binlog of the slave. In this case, the data has not been fully written to the disk, so delay will still occur.
Innodb Buffer Pool Read	Requests/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Buffer Pool Read Requests	Requests/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Buffer Pool Read Ahead	Requests/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Rows Deleted	Rows/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Rows Insert	Rows/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Rows Read	Rows/second	This metric is used to analyze current performance of the InnoDB storage engine.
Innodb Rows Updated	Rows/second	This metric is used to analyze current performance of the InnoDB storage engine.

Alarming Feature

Last updated : 2024-01-11 15:28:38

TencentDB for MariaDB sends alerts for critical alarms via SMS, email and the Message Center for key performance metrics. You can set alarms on the **Alarm Policy** page in the [Cloud Monitor Console](#).

Killing Threads

Last updated : 2024-01-11 15:28:38

When you use the database to run certain statements, deadlock may occur with no response due to large data volume. At this time, you need to kill the query statement that is consuming resources by running the following command:

```
KILL [CONNECTION | QUERY] thread_id
```

Every connection to MySQL runs in an independent thread. You can check which threads are running with the `SHOW PROCESSLIST` statement and kill a thread with the `KILL thread_id` statement.

```
mysql> show processlist;
+-----+-----+-----+-----+-----+-----+-----+
| Id      | User      | Host                | db    | Command | Time | State |
| Info                                | Rows_sent | Rows_examined |
+-----+-----+-----+-----+-----+-----+-----+
| 924107 | sutest    | 10.0.0.8:38314      | NULL  | Query   | 0    | starting |
show processlist | 0 | 0 |
| 924114 | sutest    | 10.0.0.8:38318      | test  | Query   | 264  | User sleep |
select sleep(20000) | 0 | 0 |
+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> kill 924114;
Query OK, 0 rows affected (0.00 sec)
```

If your business has many threads and you cannot accurately determine which transactions have not been committed, you can use an SQL statement similar to the following content to query thread ID (for example):

```
SELECT
  it.trx_id AS trx_id,
  it.trx_state AS trx_state,
  it.trx_started AS trx_started,
  it.trx_mysql_thread_id AS trx_mysql_thread_id,
  CURRENT_TIMESTAMP - it.trx_started AS RUN_TIME,
  pl.user AS USER,
  pl.host AS HOST,
  pl.db AS db,
  pl.time AS trx_run_time,
  pl.INFO as INFO
```

```
FROM
  information_schema.INNODB_TRX it,
  information_schema.processlist pl
WHERE
  pl.id=it.trx_mysql_thread_id
ORDER BY RUN_TIME DESC LIMIT 10;
```

If your business has many threads and you cannot accurately determine which transactions are in the status of lock wait, you can use an SQL statement similar to the following content to query thread ID (for example):

```
SELECT
  r.trx_id waiting_trx_id,
  r.trx_mysql_thread_id waiting_thread,
  TIMESTAMPDIFF( SECOND, r.trx_wait_started, CURRENT_TIMESTAMP ) wait_time,
  r.trx_query waiting_query,
  l.lock_table waiting_table_lock,
  b.trx_id blocking_trx_id,
  b.trx_mysql_thread_id blocking_thread,
  SUBSTRING( p. HOST, 1, INSTR(p. HOST, ':') - 1 ) blocking_host,
  SUBSTRING(p. HOST, INSTR(p. HOST, ':') + 1) blocking_port,
  IF (p.COMMAND = 'Sleep', p.TIME, 0) idel_in_trx,
  b.trx_query blocking_query FROM information_schema.INNODB_LOCK_WAITS w
INNER JOIN information_schema.INNODB_TRX b ON b.trx_id = w.blocking_trx_id
INNER JOIN information_schema.INNODB_TRX r ON r.trx_id = w.requesting_trx_id
INNER JOIN information_schema.INNODB_LOCKS l ON w.requested_lock_id = l.lock_id
LEFT JOIN information_schema. PROCESSLIST p ON p.ID = b.trx_mysql_thread_id
ORDER BY wait_time DESC;
```

Risk warning: after a large transaction is killed, it needs to be rolled back, which may lead to long wait time if the data volume is high. In this case, you can click master/slave switch in the console to promote the slave to the master for quick business restoration. **However, please note that when using an async or strong sync (downgradable) replication scheme, due to the delay of master/slave data sync, some data may get lost or disordered. Please perform master/slave switch with caution.**

Parameter Templates and Settings

Last updated : 2024-01-11 15:28:38

Parameter Template Overview

You can use parameters in a parameter template to manage configuration of a database engine. A database parameter set is just like a container of engine configuration values which can be applied to one or more database instances.

If you have already created a database parameter template and want to include most of its custom parameters and values in the new template, simply copy the parameter template.

If you want to use your own database parameter template, you only need to create it, modify the desired parameters, and configure your database instance to use it. It should be noted that all database instances that have applied a parameter template will not get all parameter updates of the template. If you want to apply new parameters to a batch of database instances, you can apply them by importing a template during batch parameter settings.

Managing Parameter Templates

A parameter template supports the following features, you can log in to the [MariaDB Console](#) and click an instance ID to enter its management page and configure the template features:

Specify the default parameter template.

Create templates by modifying the default parameters to generate custom parameter optimization schemes.

Generate templates by importing parameters from configuration file `my.conf` .

Save parameter settings as templates.

Import parameters from templates when setting parameters for one or multiple instances.

Parameter Description

Parameters are initialized by default when an instance is created.

You can modify instance parameters through entries such as parameter template and parameter configuration.

Parameters of different instances are isolated from and will not affect one another.

To avoid faulty operations, only common parameters can be set. If you need other parameters, please [submit a ticket](#) for application and specify the "instance ID and parameter names to be added".

Parameter Name	Restart Required	Default Value	Current Value (Configurable)	Valid Values
----------------	------------------	---------------	------------------------------	--------------

	(0: No; 1: Yes)			
auto_increment_increment	0	1	1	[1-65535]
auto_increment_offset	0	1	1	[1-65535]
autocommit	0	ON	ON	[ON, OFF]
character_set_server	0	utf8	utf8	[utf8, latin1, gbk, utf8mb4]
connect_timeout	0	10	10	[1-3600]
default_week_format	0	0	0	[0-7]
delay_key_write	0	ON	ON	[ON, OFF, ALL]
delayed_insert_limit	0	100	100	[1-4294967295]
delayed_insert_timeout	0	300	300	[1-3600]
delayed_queue_size	0	1000	1000	[1-4294967295]
div_precision_increment	0	4	4	[0-30]
group_concat_max_len	0	1024	1024	[4-18446744073709547520]
innodb_concurrency_tickets	0	5000	5000	[100-10000]
innodb_large_prefix	0	OFF	OFF	[OFF, ON]
innodb_lock_wait_timeout	0	50	50	[1-1073741824]
innodb_max_dirty_pages_pct	0	10	70	[10-90]
innodb_old_blocks_pct	0	37	37	[5-95]
innodb_old_blocks_time	0	1000	1000	[0-1000]
innodb_purge_batch_size	0	300	300	[1-1024]
innodb_read_ahead_threshold	0	56	56	[0-64]
innodb_stats_method	0	nulls_equal	nulls_equal	[nulls_equal, nulls_unequal, nulls_ignored]
innodb_stats_on_metadata	0	OFF	OFF	[ON, OFF]

innodb_stats_sample_pages	0	8	8	[1-4294967296]
innodb_strict_mode	0	OFF	OFF	[ON, OFF]
innodb_table_locks	0	ON	ON	[ON, OFF]
innodb_thread_concurrency	0	0	0	[0-128]
innodb_thread_sleep_delay	0	10000	10000	[1-3600000]
interactive_timeout	0	28800	28800	[10-86400]
key_cache_age_threshold	0	300	300	[100-4294967295]
key_cache_block_size	0	1024	1024	[512-16384]
key_cache_division_limit	0	100	100	[1-100]
lock_wait_timeout	0	5	5	[1-31536000]
log_queries_not_using_indexes	0	OFF	OFF	[ON, OFF]
long_query_time	0	1.000000	1.000000	[0.5-10]
low_priority_updates	0	OFF	OFF	[OFF, ON]
max_allowed_packet	0	134217728	1073741824	[16384-1073741824]
max_connect_errors	0	2000	2000	[1-4096]
max_connections	0	4096	4096	[1-32768]
myisam_sort_buffer_size	0	4194304	4194304	[262144-16777216]
net_buffer_length	0	16384	16384	[4096, 8192, 16384, 32768, 65536, 1048576]
net_read_timeout	0	30	30	[1-3153600]
net_retry_count	0	10	10	[1-4294967295]
net_write_timeout	0	60	60	[1-3153600]
query_alloc_block_size	0	8192	8192	[1024-16384]
query_cache_limit	0	1048576	1048576	[1-1048576]
query_cache_size	0	0	0	[0-104857600]
query_cache_type	0	OFF	OFF	[OFF, ON, DEMAND]

query_prealloc_size	0	8192	8192	[8192-1048576]
slow_launch_time	0	2	2	[1-1024]
sort_buffer_size	0	2097152	2097152	[32768-1073741824]
sqlasyn	0	ON	ON	[ON, OFF]
sqlasyntimeout	0	10	30	[10-100]
table_definition_cache	0	400	400	[400-2048]
table_open_cache	0	1024	10240	[400-524288]
tmp_table_size	0	33554432	33554432	[262144-67108864]
tx_isolation	0	REPEATABLE- READ	REPEATABLE- READ	[REPEATABLE-READ, SERIALIZABLE, READ- COMMITTED, READ- UNCOMMITTED]
wait_timeout	0	28800	28800	[60-259200]

TencentDB for MariaDB System Variable Description

[Index of Categories](#)

[Index of All Variables](#)

Database Audit

Enabling Database Audit

Last updated : 2024-01-11 15:28:38

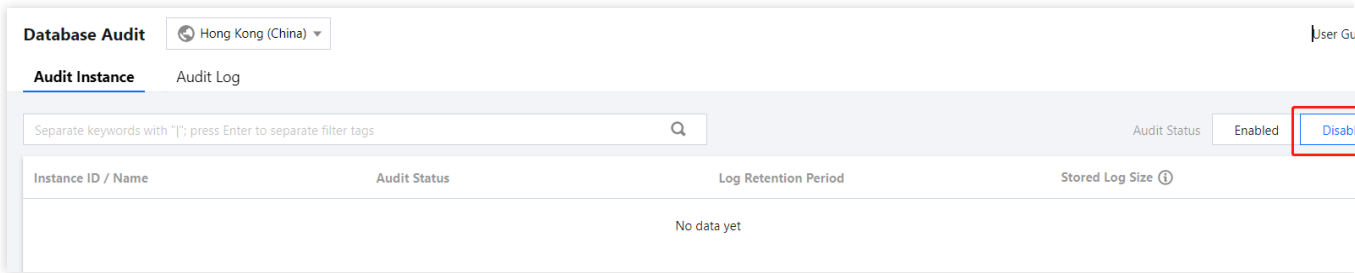
TencentDB for MariaDB has database audit capability, which can record accesses to databases and executions of SQL statements to help you manage risks and improve the database security.

Note:

The database audit feature is currently in the beta testing phase. To use this feature, [submit a ticket](#).

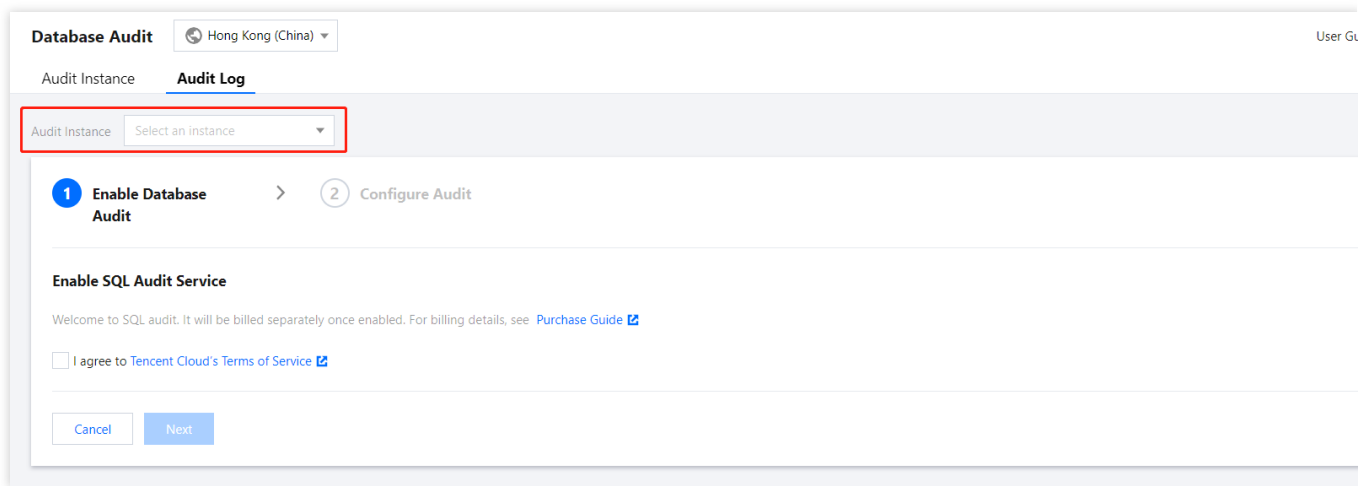
Enabling SQL Audit

1. Log in to the [TencentDB for MariaDB console](#), select **Database Audit** on the left sidebar, select a region at the top, click the **Audit Instance** tab, and click **Disabled** to filter audit-disabled instances.



Note:

Alternatively, in **Audit Instance** on the **Audit Log** tab, directly search for audit-disabled instances and then enable audit for them.



2. On the **Audit Instance** tab, click the ID of the target instance to enter the enablement page, indicate your consent to the agreement, and click **Next**.

3. On the **Configure SQL Audit** page, select the audit log retention period and click **Enable**.

Note:

You can select 7 days, 30 days, 3 months, 6 months, 1 year, 3 years, or 5 years as the audit log retention period. You can also modify it in the console after enabling audit. For more information, see [Modifying Log Retention Period](#).

In order to meet the security compliance requirements for the retention period of SQL logs, we recommend you select 180 days or above.

Viewing Audit Log

After enabling audit, you can view SQL audit logs on the **Audit Log** tab. For more information, see [Viewing Audit Logs](#).


Viewing Audit Logs

Last updated : 2024-01-11 15:28:37

Viewing Logs

1. Log in to the [TencentDB for MariaDB console](#), select **Database Audit** on the left sidebar, select a region at the top, and click the **Audit Log** tab.
2. In the audit instance section on the **Audit Log** tab, select a database instance with audit enabled to view its SQL audit logs. Or, on the **Audit Instance** tab, click an instance ID to enter the **Audit Log** tab and view audit logs.

Tool list



Click the time box and select a time period to view the audit results in the selected time period.

Note:

You can select any time period with data for search. Up to the first 60,000 eligible records can be displayed.

You can search by key tag to view audit results. Common key tags include SQL command, client IP, database name, database account, execution time, affected rows, and returned rows.

When entering multiple key tags in the text box for search, you can separate them by pressing **Enter**.

You can filter IP addresses by using the wildcard ". For example, if you enter "client IP: 10.0.0.0", IP addresses that start with "10.0.0.0" will be searched.

Log list

The **Returned Rows** field represents the specific number of rows returned by executing the SQL command, which is mainly used to determine the impact of `SELECT` commands.

The **Affected Rows** field represents the specific number of rows modified by executing the SQL command, which is mainly used to determine the impact of rewrite commands.

Audit Log Download

You can click the following icon on the **Audit Log** tab in the TencentDB for MariaDB console to obtain and view the complete SQL audit log.

Audit Instance

2022-07-20 21:23:11 ~ 2022-07-20 22:23:10

Resource Usage

Config

Separate keywords with "; press Enter to separate filter tags

Search Tips

Time	Client IP	Database Name	User Account	SQL Details	Returned Rows	Affecte...	Executi.
No data yet							

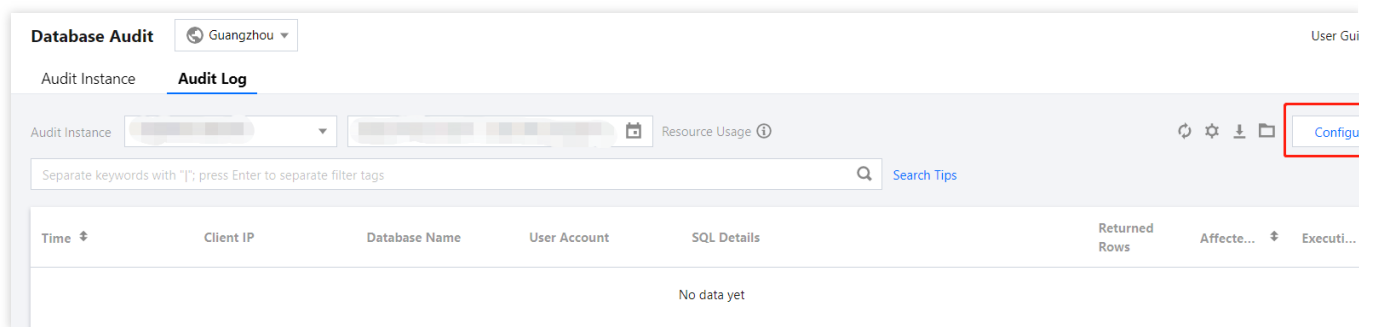
Modifying Log Retention Period

Last updated : 2024-01-11 15:28:38

This document describes how to modify the log retention period after the database audit service is activated.

Directions

1. Log in to the [TencentDB for MariaDB console](#), select **Database Audit** on the left sidebar, select a region at the top, and click the **Audit Log** tab.
2. In the top-right corner of the **Audit Log** tab, click **Configure**.



3. In the pop-up window, modify the log retention period and click **Submit**.