

# TencentDB for MySQL 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

#### STencent Cloud

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**

Use Limits

Operation Overview

Instance Management and Maintenance

Viewing instance health status

Setting Instance Maintenance Time

Assigning Instance to Project

Adjusting Database Instance Specifications

Modifying the Instance Name

Change the Port Number

Migrating AZ

Primary/Replica Switch

**Terminating Instance** 

Instance Upgrade

Upgrading the database version of MySQL 5.5 and MySQL 5.6

Upgrading from MySQL 5.7 to MySQL 8.0

Upgrading Kernel Minor Version

Upgrading Two-Node Instances to Three-Node Instances

One-Click Upgrade to Cluster Edition

CPU Elastic Expansion

Overview

Overview

Configuring auto scale-out

Establishing Manual Scaling

Activating Event Alarm

Performance Manifestation

**Billing Overview** 

FAQs

Read-Only/Disaster Recovery Instances

Creating Read-Only Instance

Managing the RO Group of Read-Only Instance

Managing Delayed Replication of Read-Only Instance

Managing Disaster Recovery Instance

Database Proxy

Overview

Overview
Database Proxy Overview
Use Limits
Database Proxy Kernel Features
Kernel Minor Version Release Notes
Upgrading Kernel Minor Version of Database Proxy
Database Proxy Management
Enabling Database Proxy
Setting Database Proxy Access Address
Viewing and Modifying Access Policy
Adjusting Database Proxy Configuration
Switching Database Proxy Network
Viewing Database Proxy Monitoring Data
Load Rebalancing
Disabling Database Proxy
Access Mode
Transaction Split Feature
Momentary Disconnection Prevention Feature
Automatic Read/Write Separation Feature
Automatic Read/Write Separation Overview
Configuring Database Proxy Read-Write Attributes
Connection Pool Feature
Connection Pool Overview
Enabling and Disabling the Connection Pool Feature
Adaptive Load Balancing
Other Features
Hint Syntax Usage
Account Management
Creating Account
Setting Password Complexity
Resetting Password
Modifying Account Permissions
Modifying Host Addresses with Access Permissions
Modifying the Number of Connections
Deleting Account
Modifying Remarks
Database Management Center (DMC)
DMC Overview

**DMC Management** Parameter Configuration Setting Instance Parameters Managing Parameter Template **Comparison of Different Template Parameters** Performance Comparison of Parameter Templates Suggestions on Parameter Configuration Intelligent Parameter Tuning Modification Records of System Parameter Templates Network and Security Access Management Overview Authorization Policy Syntax Authorizable Resource Types **Console Examples** TencentDB Security Group Management **Network Switch** Enabling Public Network Access Through CLB **Enabling Transparent Data Encryption** Setting SSL Encryption **Enabling Public Network Address** Backup and Rollback **Backing up Databases** Configuring Transition-to-Cold Storage **Cross-Region Backup Backup Encryption Backup Retention Settings** Setting Backup Download Rules Viewing Backup Space Configuring Local Binlog Retention Policy Restoring Database from Backup File Restoring Database from Physical Backup Restoring Database from Logical Backup **Downloading Backups Rolling back Databases Deleting Backups Cloning Instances** 

Data Migration

Migrating with DTS

Importing SQL Files

Migrating Data Offline

Monitoring and Alarms

Monitoring

Alarm Policies (TCOP)

Alarm Notification (DBbrain)

Setting Event Alarms

**Operation Logs** 

**Operation Logs** 

Log Delivery

Read-Only Analysis Engine

Introduction

Feature Limits and Compatibility

Compatibility and Use Limits

Data Loading Restrictions

Description of Functions and Operators

**Description of Supported Bitwise Operations** 

Description of Supported Arithmetic Operations

Comparison Operation Support Description

Description of Supported Logical Operations

**Description of Supported String Functions** 

Description of Supported Mathematical Operation Functions

**Description of Supported Date Functions** 

Description of Supported Aggregate Functions

**JSON** Function

Description of Supported Type Conversion Functions

Data Type

Read-Only Analysis Engine Management

Enabling Read-Only Analysis Engine

**Basic Operations** 

Parameter Management

Data Loading

Loading Data to Read-Only Analysis Engine

Viewing Data Loading Status

Controlling Data Loading Speed

Canceling Loading Certain Objects into the Read-Only Analysis Engine

**Executing Queries** 

**Collecting Statistics SELECT Statement Description CTE Syntax Use Instructions** Window Function Use Instructions Executing SQL and Viewing the Execution Plan **Exclusive Features ETL Writeback Acceleration Ordered Pagination Feature** Performance Optimization **Optimizing SQL Execution With HINT** Runtime Filter User Manual Lazy Materialization Merge Join Columnar Storage Secondary Index System Variables System Tables SLOW\_LOG Performance White Paper **TPC-H** Testing Method **TPC-H Performance Results** 

#### Tag

Overview Editing Tags

# Operation Guide Use Limits

Last updated : 2024-07-22 14:31:43

### Limits on the Data Volume

TencentDB for MySQL imposes data volume restrictions on all types of MySQL instances to isolate performance issues due to limited resources. This document describes the technical impact of a single instance or table with a large data volume on MySQL.

**Instance with a high data volume**: The default storage engine for TencentDB is InnoDB. If the cache buffer can cache all data and index pages in the MySQL instance, the instance can support a large number of concurrent access requests. If the instance contains too much data, the cache and buffer will frequently swap data in/out, quickly hitting an IO bottleneck and reducing throughput. For example, if a TencentDB instance is designed to sustain up to 8,000 access requests per second, it can only support 700 when the data volume is twice the size of the cache and buffer.

**Table with a high data volume**: If a table contains too much data, MySQL will find it more difficult to manage table resources (data, indexes, etc.), reducing table processing efficiency. For example, if the size of a transaction table (InnoDB) exceeds 10 GB, the latency in update operations will soar, increasing the response time for transactions. In this case, the problem can only be solved through sharding and migration.

#### Note:

If the number of tables in a single instance exceeds one million, backup, monitoring, and upgrade may fail and database monitoring may be affected. Make sure that the number of tables in a single instance is below one million.

### Limit on the Number of Connections

The maximum number of connections to a MySQL instance is specified with the MySQL system variable max\_connections . When the actual number exceeds max\_connections , no more connections can be established.

The default number of connections to TencentDB can be viewed by clicking the instance ID to enter the **Database Management** > **Parameter Settings** page in the TencentDB for MySQL console, which can be adjusted if necessary. However, more connections mean that more system resources will be consumed; if the number of connections goes beyond what the actual system load capacity allows, the system service quality will be definitely undermined.

For more information on max\_connections , see MySQL official documentation.

### Limits on the MySQL Client Version

We recommend that you use the MySQL client and library that come with CVM to connect to TencentDB instances.

#### Notes on slow logs

For Linux CVM instances, you can use TencentDB's export tool to get slow query logs. For more information, see Restoring Database from Physical Backup.

For Windows CVM instances, slow query logs cannot be obtained directly at present. If you need them, submit a ticket for assistance.

#### Notes on the TencentDB binlog retention duration

TencentDB for MySQL binlogs can be retained for 7 (default value) to 1830 days (customizable by clicking the instance ID to enter the **Backup and Restoration** > **Auto-Backup Settings** page).

If binlogs are retained for an extended period of time or grow too fast, additional space will be required for backup, and fees will be incurred if the space exceeds the free tier of backup capacity.

#### Notes on the character set

TencentDB for MySQL uses the UTF8 character set by default.

Even though TencentDB supports changing the default character set, we recommend that you explicitly specify the encoding format for a table when creating it and specify the connection encoding during connection establishment. In this way, your application will be more portable.

For more information on the resources of MySQL character set, see MySQL official documentation.

You can modify the character set through SQL or in the TencentDB for MySQL console.

#### Modifying the character set through SQL

1. Run the following SQL statements to change the default character set for TencentDB instances:

```
SET @@global.character_set_client = utf8;
SET @@global.character_set_results = utf8;
SET @@global.character_set_connection = utf8;
SET @@global.character_set_server = utf8;
```

After the statements are executed, @@global.character\_set\_server will be automatically synced to a local file for persistence in approximately 10 minutes, while the other 3 variables will not. The configured value will stay unchanged even after migration or restart.

2. Run the following statements to change the character set encoding for the current connection:

```
SET @@session.character_set_client = utf8;
SET @@session.character_set_results = utf8;
SET @@session.character_set_connection = utf8;
```

#### Or

```
SET names utf8;
```

3. For PHP programs, you can configure the character set encoding for the current connection by using the following function:

```
bool mysqli::set_charset(string charset);
```

#### Or

```
bool mysqli_set_charset(mysqli link, string charset);
```

4. For Java programs, you can configure the character set encoding for the current connection as shown below:

```
jdbc:mysql://localhost:3306/dbname?useUnicode=true&characterEncoding=UTF-8
```

#### Modifying the character set in the TencentDB for MySQL console

1. Log in to the TencentDB for MySQL console and click an instance ID in the instance list to enter the instance details page.

2. Find the character set in **Basic Info** and click the modification icon to modify it.

Basic Info	
Instance Name	1
Instance ID	To and the second se
Status/Task	Running /
Region/AZ	North China (Beijing)/ Beijing Zone 6 Migrate to New AZ
Project	Default Project Switch to Another Project
Placement Group	Add
GTID	Enabled
Character Set	UTF8

3. In the pop-up window, select a character set and click **OK**.

Nodify Characte	er Set			
Character Set*	O UTF8	🔵 GBK		UTF8MB4
		ОК	Cancel	

#### Limits on operations

1. Do not modify the information and permissions of the existing accounts for a MySQL instance; otherwise, some cluster services may become unavailable.

2. InnoDB is recommended for creating databases and tables, so that the instances can better support a large number of concurrent access requests.

3. Do not modify or terminate the source-replica relationship; otherwise, hot backup may fail.

#### Limits on the table name

Chinese table names are not supported because they may result in failures of processes such as rollback and upgrade.

### Database Account Permission

TencentDB for MySQL no longer provides the super user permission. To modify parameters that require this permission, log in to the TencentDB for MySQL console, click the instance ID to enter the **Database Management** > **Parameter Settings** page, and modify them.

### **Network Selection**

We recommend that you use a VPC. In the VPC, you can freely define IP range segmentation, IP addresses, and routing policies. Compared with the classic network, VPC is more suitable for scenarios where custom network configurations are required. For the comparison of VPC and classic network, see <u>Overview</u>.

## **Operation Overview**

Last updated : 2025-06-03 18:23:54

This document describes operations on the instance list and management pages in the TencentDB for MySQL console.

### Instance list page

You can log in to the TencentDB for MySQL console and enter the instance list page to view instance information and manage your instances.

Feature	Description
Log In	Click <b>Log In</b> in the Operation column to log in to the database using the next-generation database tool. For more information, see DMC Overview.
Manage	Click <b>Manage</b> in the Operation column to enter the instance management page.
Upgrade Version	Two-node and three-node instances: Click <b>More</b> > <b>Upgrade Version</b> in the Operation column to upgrade the database version. For more information, see Upgrading Database Engine Versions. Single-node and Cluster Edition instances: Upgrading the database engine version is not yet supported.
Purchase Same Configuration	After the target instance is selected, click <b>More</b> > <b>Purchase Same Configuration</b> in the Operation column to quickly create a new instance with the same configuration as the selected one on the purchase page.
Adjust Configurations	Click <b>More</b> > <b>Adjust Configurations</b> in the Operation column to adjust the configuration of the database instance (scaling). Both instance upgrade and downgrade are supported. For more information, see Adjusting Database Instance Specifications. For Cluster Edition instances, you can also quickly add and delete read-only nodes by adjusting the configuration. For more information, see Adjusting Instance Configuration.
Pay-as-You-go to Monthly Subscription	Click <b>More</b> > <b>Pay-as-You-Go to Monthly Subscription</b> in the Operation column to change the billing mode of pay-as-you-go instances. For more information, see Pay-as-You-Go to Monthly Subscription.
Edit Tag	Click <b>More</b> > <b>Edit Tag</b> in the Operation column to edit and manage the tags of instance resources. For more information, see Tag Editing.
Terminate/Return	Click <b>More &gt; Terminate/Return</b> in the Operation column to return the instance by



	yourself. For more information, see Terminating an Instance.			
Security Group	Click <b>More</b> > <b>Security Group</b> in the Operation column to configure or modify the security group of the instance. For more information, see Cloud Database Security Group Management.			
Create	Click <b>Create</b> above the instance list to go to the purchase page and create an instance. For more information, see Creating a MySQL Instance.			
Renew	After the target instance is selected, click <b>Renew</b> above the instance list to manually renew the instance. For more information, see <b>Renewal Instructions</b> .			
Restart	After the target instance is selected, click <b>Restart</b> above the instance list to restart the instance. Batch restart is supported (selecting multiple instances). During the restart, the instances will not be accessible normally and existing connections will be disconnected. Please be prepared to avoid any impact. Restart will fail if there are a large number of writes and dirty pages during the restart. In this case, the instance will roll back to the status before the restart and can still be accessed. Be sure to restart the instance during off-hours so as to ensure success and reduce the impact on your business.			
Roll Back	After the target instance is selected, click <b>More</b> > <b>Roll Back</b> above the instance list to roll back a database to a specified time based on backups. For more information, see Rolling Back a Database.			
Parameter Settings	After the target instance is selected, click <b>More</b> > <b>Parameter Settings</b> above the instance list to modify the parameter values of the instance. Batch settings are supported (selecting multiple instances). For more information, see <u>Setting Instance Parameters</u> .			
Enable Auto- Renewal	After the target instance is selected, click <b>More</b> > <b>Enable Auto-Renewal</b> above the instance list to make the non-auto-renewal instance automatically renew monthly after expiration. Batch enablements are supported (selecting multiple instances). For more information, see Renewal Instructions.			
Disable Auto- Renewal	After the target instance is selected, click <b>More</b> > <b>Disable Auto-Renewal</b> above the instance list to cancel auto-renewal for instances that are already set to auto-renew. Batch cancellations are supported (selecting multiple instances).			
Quick Migration	After the required target instance is selected, click <b>Quick Migration</b> above the instance list to migrate TencentDB for MySQL two-node/three-node instances to TDSQL-C for MySQL. For more information, see Perform Quick Migration.			
Quick Check	Click <b>Quick Check</b> above the instance list to go to the TencentDB for DBbrain Console, where you can view the health status of instances under your account and perform performance optimization on instances. For more details, see MySQL Performance Optimization.			



Comparative Monitoring After 2 to 6 target instances are selected, click **Comparative Monitoring** above the instance list to view and compare the monitoring situation of multiple instances in the popup window on the right of the console.

### Instance management page

Log in to the TencentDB for MySQL console, click the target instance ID in the instance list, or click **Manage** in the **Operation** column to enter the instance management page, where you can switch to view Instance Details, Instance Monitoring, Database Management, Security Group, Backup and Restoration, Operation Log, Read-Only Instance, Database Proxy, Data Security, and Connection Check.

Feature Page	Feature Item	Description
Instance Details	Viewing instance health status	Help users find the abnormal situation of the instance in time and optimize system performance and security configurations. For more information, see Viewing Instance Health Status.
	Setting instance maintenance time	The system will aperiodically perform maintenance on instances during the maintenance time. For more information, see Setting Instance Maintenance Time.
	Assigning instances to projects	Assign instances to different projects for management. For more information, see Assigning Instances to Projects.
	Adjusting database instance specifications	<ul> <li>Quickly adjust instance specifications through the console for flexible scaling.</li> <li>For more information on single-node, two-node, and three-node instances, see Adjusting Database Instance Specifications.</li> <li>For more information on Cluster Edition instances, see Adjusting Instance Configuration.</li> </ul>
	Modifying the port number	Support modification of the port number for the instance's private network address. For more information on single-node, two-node, and three-node instances, see Modifying the Port Number. For more information on Cluster Edition instances, see Read-Write Address Management and Read-Only Address Management.
1		



	Migrate to New AZ	Migrate TencentDB for MySQL instances to another availability zone in the same region. For more information, see Migrating Availability Zones.
	Source/Replica Switch	Switch the primary and secondary databases of instances. For more information on two-node and three-node instances, see Primary/Replica Switch. For more information on Cluster Edition instances, see Primary/Replica Switch.
	CPU elastic scale-out	Automatically or manually expand the CPU upper limit of the current instance. For more information, see CPU Elastic Scale-out. Cluster Edition instances do not support this feature.
	Managing disaster recovery instances	Two-node and three-node instances support the creation of disaster recovery instances. For more information, see Managing Disaster Recovery Instances. Single-node and Cluster Edition instances do not support the creation of disaster recovery instances.
	Network switch	Modify the private network of instances. For more information on single-node, two-node, and three-node instances, see Network Switch. For more information on Cluster Edition instances, see Read-Write Address Management and Read-Only Address Management.
	Enabling public network connection address	For more information, see Enabling Public Network Connection Address.
	Methods to Modify Data Replication	Set the data replication mode of the instance according to the business characteristics to improve TencentDB availability. For operations, see Modifying the Data Replication Mode.
Instance Monito	pring	You can view the monitoring of many core metrics of the current database in dimensions such as access, load, query cache, Ttable, InnoDB, and MyISAM. For details, see Monitoring Features and Alarm Feature. Cluster Edition instances support switching nodes on the instance monitoring page to view the monitoring situation. For more information, see View Instance Monitoring.
Database Management	Database List	The DMC console allows you to create databases and tables, manage databases, import/export data, and supports SQL window. For more information, see DMC Management.
	Parameter Settings	View and modify some parameters through the console, and query parameter modification records in the console. Set instance parameters.



		Apply parameter templates. Compare different template parameters. For more information on single-node, two-node, and three-node instances with 4 cores and above, see Intelligent Parameter Tuning. Cluster Edition instances currently do not support intelligent parameter tuning.	
	Account Management	You can manage the default root account of the system or create and manage new accounts. Create accounts. Set password complexity. Reset passwords. Modify account permissions. Modify host addresses with access permissions. Delete accounts.	
Security Group		You can configure security groups for your databases. For more information, see TencentDB Security Group Management.	
Backup and Restoration	Back up a database	You can automatically or manually back up the database. For more information, see Back up a Database.	
	Configuring transition-to- cold storage	Support transition-to-cold storage for backup files to reduce backup storage costs. For more information on two-node and three-node instances, see Configuring Transition-to-Cold Storage. Single-node and Cluster Edition instances currently do not support configuring transition-to-cold storage.	
	Cross-Region Backup	Support storing backup files in another region. For more information on two-node, three-node, and Cluster Edition instances, see Cross-Region Backup. Single-node instances currently do not support cross-region backup.	
	Backup Encryption	Support encryption of physical backups and log backups (binlog). For more information on two-node and three-node instances, see Backup Encryption. Single-node and Cluster Edition instances currently do not support backup encryption.	
	Setting backup download rules	Set restrictions for backup download. For more information on two-node and three-node instances, see Setting Backup Download Rules. Single-node and Cluster Edition instances currently do not support setting backup download rules.	
	Viewing backup space	View the instance backup space and free quota. For more information, see Viewing Backup Space.	

	Setting local Binlog retention	Set the binlog retention period for the instance. For more information, see Setting Local Binlog Retention. Single-node instances do not support setting local binlog retention.	
	Restoring a database using physical backup	Two-node and three-node instances support the use of XtraBackup tool to restore MySQL physical backup files to a self-built database on another host. For more information, see Restoring a Database Using Physical Backup.	
	Restoring a database using logical backup	Two-node and three-node instances support using logical backup files for manual data restoration. For more information, see Restoring a Database Using Logical Backup.	
	Downloading a backup	Logical and physical backups of two-node and three-node instances can be downloaded. For more information, see <u>Downloading a Backup</u> . Snapshot backups of single-node and Cluster Edition instances cannot be downloaded.	
	Rolling back a database	Data corrupted can be recovered through rollback in a self-service manner. For more information, see Rolling back a Database.	
-	Deleting a backup	Only manual backup files of two-node and three-node instances can be deleted. For more information, see Delete a Backup.	
	Cloning an instance	Help users quickly roll back an instance to a newly purchased MySQL instance. For more information, see Cloning an Instance.	
Operation Log		For Cluster Edition, two-node, and three-node instances, slow logs, error logs, and rollback logs can be viewed, and logs can be shipped. For more information, see Operation Logs and Log Shipping. For single-node instances, slow logs and error logs can be viewed. For more information, see Operation Logs.	
Read-Only Instance	Creating a read-only instance	Create one or more read-only instances to support users' read-write separation and one-primary-multiple-secondary application scenarios. For more information on two-node and three-node instances, see Creating a Read-Only Instances. For Cluster Edition instances, 1-5 read-only nodes can be created. For more information, see Adding a Node for an Instance. For Cluster Edition instances, 1-15 independent read-only instances can also be created. For more information, see Creating a Read-Only Instance.	
	Managing the RO group of read-only instance	Two-node and three-node instances support managing the RO group of read-only instance. For more information, see Managing RO Group of Read-Only Instance.	



	Managing delayed replication of read-only instance	Two-node and three-node instances support managing delayed replication of read-only instances. For more information, see Managing Delayed Replication of Read-Only Instance.	
Database Proxy	/	It is used to proxy requests when the application server accesses the database. It provides advanced features such as automatic read/write separation, connection pool, and connection retention. Single-node and Cluster Edition instances do not support database proxy. For more information on two-node and three-node instances, see Database Proxy Overview.	
Data Security	Enabling Transparent Data Encryption (TDE)	Two-node and three-node instances support enabling TDE. For more information, see Enabling Transparent Data Encryption.	
	Setting SSL encryption	Two-node and three-node instances support setting SSL encryption. For more information, see Setting SSL Encryption.	
Connection Check		You can detect potential connection access issues with your TencentDB and address them using the provided solutions to ensure that your TencentDB can be accessed normally. For more information, see One-Click Connectivity Checker.	

# Instance Management and Maintenance Viewing instance health status

Last updated : 2025-06-09 14:23:41

The MySQL Instance Details page now displays the health status, exception alarms, and configurations related to availability, performance, and security of the instance. This helps users detect anomalies in a timely manner, so that they can optimize system performance and security configurations, and improve the security and stability of instances. This section describes how to view the health status of the instance and related operations on the Instance Details page.

### Prerequisites

You have created a MySQL instance..

### Viewing Instance Health Status

1. Log in to the TencentDB for MySQL console..

In the instance list, click an **instance ID** or **Manage** in the **Operation** column to enter the instance details page.
 On the right side of the page, you can view the health status of the instance, including the health status score, exception alarm items, and functional configuration deployment status.

### **Related Operations**

Page	Operation
	One-Click Connection Check
Instance Details	One-Click Diagnosis
Page	Viewing Alarms
	Viewing Configurations

#### 1. One-Click Connection Check

If you cannot access a MySQL instance via either the private or public network, the One-Click Connection Check tool can help to troubleshoot the connection issues. Click the **Quick Connection Check** following the private network address to navigate to the connection check page. For subsequent operations, please refer to One-Click Connection Check Tool.

#### 2. One-Click Diagnosis\*\*\*\*

Under the instance health status, the current instance health score is displayed. With a perfect health score of 100, you can promptly know whether the instance is running healthily. This score can be manually refreshed. Click the **Quick Chec**k button following the score to navigate to the DBbrain console's Exception Diagnosis page. This page offers real-time performance monitoring, health checks, failure diagnosis, and optimization suggestions for the user's instance. It enables users to see the real-time operation status of the database instance intuitively, locate newly emerged performance anomalies, and optimize the system based on the optimization suggestions. For instructions on viewing monitoring information, diagnostic information, etc. on the exception diagnosis page, please refer to Exception Diagnosis.

#### 3. Viewing Alarms

On the instance details page, the Exception Alarm section displays the exception alarms of the instance within the last three hours. Based on the exception alarm prompts, users can acquire detailed information about risk items of different risk levels at a certain stage of the instance. This information helps users make optimizations and adjustments accordingly. The exception alarm items are collected by the Exception Diagnosis feature of DBbrain. Click **View Alarm**. A sidebar is displayed, listing exception alarms of different levels associated with the current instance.

Click **Alarm Details** in the top-right corner to go to the Exception Alarms page of the DBbrain console. Here you can review details such as basic information about the alarmed instance, risk level, diagnostic items, and the duration. For details about the information and operations on the Exception Alarms page, please refer to Exception Alarms.

#### 4. Viewing Configurations

MySQL provides users with configuration details about database functionality from three perspectives: availability, performance, and security. Users can select a category based on their actual requirements and reference the corresponding configurations.

Click **View Configuration**. A sidebar is displayed, listing detailed configuration information of the database functionality. As shown in the following table:

Availability (The configuration item description is shown based on the availability level.)

Single Point Failure

**AZ-level Faults** 

Region-Level Failure

Highly-Availability Architecture Applying to All Scenarios

Configuration Item	Description	Configuration Level	Operation



Space	Forecasts whether the remaining space can satisfy the needs for the forthcoming 30 days.	-	Space Analysis
Momentary Disconnection Prevention	Ensures application continuity without loss during planned detrimental database instance switches and transfers, thus averting connection interruptions.	Enhanced Configuration	Enable Database Proxy View Instructions
Multi-AZ Deployment	The multi-AZ deployment feature of MySQL provides your database instance with high availability and failover support.	-	View Instructions
Multi-AZ Deployment (Database Proxy)	The Multi-AZ deployment feature of MySQL database proxy offers multi-region failover support for the proxy.	-	View Instructions
Three-Node	The three-node mode employs the backup architecture with one primary and two secondary databases, supporting strong synchronous replication method and offering financial-grade reliability and high availability.	Enhanced Configuration	View Instructions
Multiple RO Instances RO Group	A single read-only instance is subject to single point failure risks, hence, it's recommended to keep at least two read-only instances within a single RO group to ensure availability.	Mandatory Configuration	View Instructions
Multi-AZ RO Instance	Within a single RO group, RO instances can only be deployed in one AZ. To enable cross-AZ deployment, you can create multiple RO groups.	-	View Instructions
Disaster Recovery Instance	Helps users in enhancing their cross-region service continuity at a lower cost while	-	View Instructions



	simultaneously improving data reliability.		
Cross-Region Backup	Synchronizes backup files to another region, enabling users to enhance their regulatory compliance and disaster recovery capabilities.	-	View Instructions

Configuration Item	Description	Configuration Level	Operation
Space	Forecasts whether the remaining space can satisfy the needs for the forthcoming 30 days.	-	Space Analysis
Momentary Disconnection Prevention	Ensures application continuity without loss during planned detrimental database instance switches and transfers, thus averting connection interruptions.	Enhanced Configuration	Enable Database Proxy View Instructions
Multi-AZ Deployment	MySQL's Multi-AZ deployment ensures high availability and failover support for database instances.	Mandatory Configuration	View Instructions
Multi-AZ Deployment (Database Proxy)	The Multi-AZ deployment feature of MySQL database proxy offers multi-region failover support for the proxy.	Mandatory Configuration	View Instructions
Three-Node	Three-node instances utilize the architecture with one primary and two secondary databases, which can support a strong synchronous replication mode, and providing finance-grade reliability and high availability.	Enhanced Configuration	View Instructions
Multiple RO Instances RO Group	A single read-only instance is subject to single point failure risks, hence, it's recommended to keep at least two read-only instances within a single RO group to ensure availability.	Mandatory Configuration	View Instructions
Multi-AZ RO Instance	Within a single RO group, RO instances can only be deployed in one AZ. To enable cross-AZ deployment, you can create multiple RO groups.	Mandatory Configuration	View Instructions

Disaster Recovery Instance	Helps users in enhancing their cross-region service continuity at a lower cost while simultaneously improving data reliability.	-	View Instructions
Cross-Region Backup	Synchronizes backup files to another region, enabling users to enhance their regulatory compliance and disaster recovery capabilities.	-	View Instructions

Configuration Item	Description	Configuration Level	Operation
Space	Forecasts whether the remaining space can satisfy the needs for the forthcoming 30 days.	-	Space Analysis
Momentary Disconnection Prevention	Ensures application continuity without loss during planned detrimental database instance switches and transfers, thus averting connection interruptions.	Enhanced Configuration	Enable Database Proxy View Instructions
Multi-AZ Deployment	MySQL's Multi-AZ deployment ensures high availability and failover support for database instances.	Mandatory Configuration	View Instructions
Multi-AZ Deployment (Database Proxy)	The Multi-AZ deployment feature of MySQL database proxy offers multi-region failover support for the proxy.	Mandatory Configuration	View Instructions
Three-Node	Three-node instances utilize the architecture with one primary and two secondary databases, which can support a strong synchronous replication	Enhanced Configuration	View Instructions



	mode, and providing finance-grade reliability and high availability.		
Multiple RO Instances RO Group	A single read-only instance is subject to single point failure risks, hence, it's recommended to keep at least two read-only instances within a single RO group to ensure availability.	Mandatory Configuration	View Instructions
Multi-AZ RO Instance	Within a single RO group, RO instances can only be deployed in one AZ. To enable cross-AZ deployment, you can create multiple RO groups.	Mandatory Configuration	View Instructions
Disaster Recovery Instance	Helps users in enhancing their cross- region service continuity at a lower cost while simultaneously improving data reliability.	Mandatory Configuration	View Instructions
Cross-Region Backup	Synchronizes backup files to another region, enabling users to enhance their regulatory compliance and disaster recovery capabilities.	Enhanced Configuration	View Instructions

Configuration Item	Description	Configuration Level	Operation
Space	Forecasts whether the remaining space can satisfy the needs for the forthcoming 30 days.	-	Space Analysis



Momentary Disconnection Prevention	Ensures application continuity without loss during planned detrimental database instance switches and transfers, thus averting connection interruptions.	Mandatory Configuration	Enable Database Proxy View Instructions
Multi-AZ Deployment	MySQL's Multi-AZ deployment ensures high availability and failover support for database instances.	Mandatory Configuration	View Instructions
Multi-AZ Deployment (Database Proxy)	The Multi-AZ deployment feature of MySQL database proxy offers multi-region failover support for the proxy.	Mandatory Configuration	View Instructions
Three-Node	Three-node instances utilize the architecture with one primary and two secondary databases, which can support a strong synchronous replication mode, and providing finance-grade reliability and high availability.	Mandatory Configuration	View Instructions
Multiple RO Instances RO Group	A single read-only instance is subject to single point failure risks, hence, it's recommended to keep at least two read-only instances within a single RO group to ensure availability.	Mandatory Configuration	View Instructions
Multi-AZ RO Instance	Within a single RO group, RO instances can only be deployed in one AZ. To enable cross-AZ	Mandatory Configuration	View Instructions



	deployment, you can create multiple RO groups.		
Disaster Recovery Instance	Helps users in enhancing their cross- region service continuity at a lower cost while simultaneously improving data reliability.	Mandatory Configuration	View Instructions
Cross-Region Backup	Synchronizes backup files to another region, enabling users to enhance their regulatory compliance and disaster recovery capabilities.	Mandatory Configuration	View Instructions

#### Performance

Configuration Item	Description	Operation
Parallel Query	Harnesses multiple computing kernels, significantly reducing the response time for large-scale queries.	View Instructions
Asynchronously Dropping Large Tables	The DROP TABLE operation gradually clears the original database file (.ibd) in the background, mitigating the impact on performance.	View Instructions
Instant Column Addition	The Instant Column Addition feature enables swift column addition to large tables by solely modifying the data dictionary, minimizing the impact on system.	View Instructions
Instant Column Modification	The Instant Column Modification Feature enables swift column modification in large tables by solely modifying the data dictionary, minimizing the impact on system.	-
Read-Only Instance	Creates Read-Only instances to facilitate read/write segregation and master-replica application scenarios, thereby enhancing the database's read load capacity.	View Instructions
CPU Elastic Scaling	Once Auto CPU Elastic Scaling is activated, it effectively guarantees the stability of online operations.	View Instructions



Automatic	By configuring access strategies for Read/Write Separation through	View
Read/Write	database proxy, you can achieve elastic expansion of reading capability,	Instructions
Separation	thereby mitigating database pressure.	

#### Security

Configuration Item	Description	Operation
Security Group	Security group is a type of virtual firewall that regulates the inbound and outbound traffic at the instance level, serving as a crucial means for network security isolation.	View Instructions
SSL	SSL authentication is the verification process from the client to the cloud database server, which authenticates both the user and the server.	View Instructions
Password Complexity	Configures the password complexity to enhance the robustness of the database access passwords, safeguarding the security of your database.	View Instructions
Transparent Data Encryption (TDE)	TDE performs encryption and decryption operations transparently to the user, doing so in real-time, and meets the compliance requirements for static data encryption.	View Instructions
Latency Read-Only	By setting delay replication and initiating the recovery to a specific time during the delay period, efficient data rollback and rapid fault traceback can be achieved.	View Instructions
Backup Encryption	After the Backup Encryption feature is enabled, newly created physical backup files will be automatically encrypted for storage.	View Instructions
Download Restrictions on Backups	By default, backup files can be downloaded over a public or private network. Restrictions can be implemented via the download configuration settings.	View Instructions
Database Auditing	Records accesses to databases and execution of SQL statements, improving your database security level.	View Instructions

# Setting Instance Maintenance Time

Last updated : 2024-07-22 14:33:38

### Overview

Maintenance time is a very important concept for TencentDB for MySQL. To ensure the stability of your TencentDB for MySQL instance, the backend system performs maintenance operations on the instance during the maintenance period from time to time. To minimize the potential impact on your business, we recommend that you set an acceptable maintenance period for your business instance, usually during off-peak hours.

In addition, we also recommend you perform operations involving data migration during the maintenance time, such as instance specification adjustment, instance version upgrade, and instance kernel upgrade. Currently, the maintenance time can be customized for source, read-only, and disaster recovery instances.

Take the database instance specification upgrade as an example. As this operation involves data migration, after the upgrade is completed, a momentary disconnection from the database may occur. When the upgrade is initiated, you can select **During maintenance time** for **Switch Time**, so that the instance specification will be switched during the next **maintenance time** after the instance upgrade is completed. Note that when you select **During maintenance time** for **Switch Time**, the switch will not occur immediately after the database specification upgrade is completed; instead, the sync will continue till the instance goes into the next **maintenance time** when the switch will be performed. As a result, the overall time it takes to upgrade the instance may be extended.

#### Note:

Before maintenance is carried out for TencentDB for MySQL, notifications will be sent to the contacts configured in your Tencent Cloud account by SMS and email.

TencentDB for MySQL will perform a data consistency check within the maintenance time configured for the database instance to ensure the source-replica data consistency and reduce the risk of data exceptions after the instance switch. The database performance will drop slightly during the check; therefore, we recommend that you select a maintenance time during off-peak hours. If the current load of the database is high, the check will be initiated during the maintenance time.

Instance switch will cause a momentary disconnection from the database. Make sure that your business has a reconnection mechanism.

### Directions

#### Setting the maintenance time

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation column to enter the instance details page.

2. In the Maintenance Info section on the instance details page, click Modify.



3. In the pop-up window, select Maintenance Window and Maintenance Time and click OK.

Modify Maintenance Window and Time $\qquad \qquad \qquad$				
Maintenance Window	✓ Mon ✓ Tue ✓ Sun	💙 Wed 💙 Thu 💙 Fri 💙 Sat		
Maintenance Time	Start Time 03 💌	: 00 -		
	Duration 1 -	hr		
Delay Threshold for Data Consistency Check (i) 10 sec				
		Enter an integer from 1 to 10 (The new		
		configuration may take longer to take		
		effect if the threshold is lower.)		
	ОК	Cancel		

#### Performing the switch

If a task is configured to be switched during the maintenance time, but you need to switch it urgently under special circumstances, you can click **Switch Now** in the **Operation** column.

#### Note:

Immediate switch is applicable to operations involving data migration such as instance specification adjustment, instance version upgrade, and instance kernel upgrade.

For version upgrade, if it is associated with multiple instances, the switch will be performed in the order from disaster recovery instance to read-only instance to source instance.

# Assigning Instance to Project

Last updated : 2024-07-22 14:35:34

TencentDB for MySQL supports assigning instances to different projects for management. Notes for assignment are listed below:

Read-only instances and disaster recovery instances are the associated instances of the source instance and should be in the same project as the source instance.

Assigning and reassigning TencentDB instances will not affect the services provided by the instances.

You need to specify a project to which a new instance belongs when purchasing it. **Default Project** will be used if you don't specify one.

Assigned instances can be reassigned to other projects through the Assign to Project feature.

### Assigning the project on the purchase page

1. Log in to the TencentDB for MySQL purchase page.

2. When purchasing a new instance, you can directly specify the project to assign the instance to.

Project	·	G		
	C	ξ		
Tag 🛈	International Contraction of Contrac		*	×
	In the existing tags do not meet your requirements, pie	ase go to	Creat	te Tags12

### Modifying the project on the purchase page

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance details page.

2. Click **Switch to Another Project** after **Basic Info** > **Project**, select the target project in the pop-up window, and click **OK**.

Note:

Assign to Project		×
You have selected <b>1 instance(s).</b> , View D	etails 👻	
Search by project name/description		Q
Project Name	Description	
$\bigcirc$ (	DEFAULT PROJECT	
0(		
	OK Cancel	

You can also batch assign multiple instances to a project by selecting them and clicking **More** > **Assign to Project** above the instance list.

# **Adjusting Database Instance Specifications**

Last updated : 2025-05-07 17:40:17

TencentDB for MySQL supports flexible scaling that allows you to quickly adjust instance specification in the console. You can do so at any time (at the start, during rapid development, or during peak/off-peak hours), so you can get the most out of your resources and reduce unnecessary costs in real time. For more information, see Instance Adjustment Fee.

In cluster edition instances, the adjustment configuration operation also allows for adding or deleting read-only nodes, change of node availability zones, etc. Please see Adjust Instance Configuration.

### Instance Disk Space Description

To ensure business continuity, upgrade your instance specifications or purchase additional disk capacity in time before disk capacity is used up.

#### Note:

The disk space can be viewed on the instance details page of the MySQL console. You can also configure alarms to receive disk space alarm messages in time.

When the amount of data stored in an instance exceeds the current storage space of the instance, the instance will be locked and become read-only, and you will not be able to write data to it. You need to expand the capacity or delete some database tables in the console to unlock it. If your instance uses a single-node architecture, the only way to unlock it is through expansion.

To avoid the database from triggering the lock status repeatedly, a locked instance will be unlocked and allowed for reads and writes only when its remaining available capacity is more than 20% of its total capacity or over 50 GB.

### **Configuration Adjustment Description**

By default, the instance configuration is adjusted in the normal mode, which requires a data migration for the adjustment to complete after you adjust the instance configuration in the console. But if the physical machine where the instance is located has sufficient remaining resources (aka local resources), you can choose the QuickChange mode. The adjustment process is as follows:

**Normal mode**: to adjust the instance configuration, instance data needs to be migrated from the original physical machine to a new one. Because adjustment in this mode requires data migration, comparison, and verification, the

overall adjustment process will take a long time in case of a huge amount of data. Besides, an instance switch may occur after the adjustment is completed.

**QuickChange mode**: Overall configuration adjustment takes a short time. Whether it involves primary-secondary switching depends on the actual configuration adjustment page reminder. Currently, it is categorized into the following scenarios.

Disk Types	Configuration Adjustment	Influence
local disk	Only instance computing specifications changes or combination adjustments (specification upgrades/downgrades + disk expansion/reduction) are permitted	When local resources are abundant and the parameters involved in the change do not require a restart, the current modification does not involve master-slave switching or rebooting. The configuration change page will indicate: "Current changes will not involve data migration or instance restart."
		When local resources are abundant, yet the parameters involved in the changes necessitate a restart, the current modifications entail a master- slave switch and a reboot. The configuration adjustment page will indicate: "During the configuration adjustment, the instance access is not affected. After the adjustment is completed, there will be a momentary disconnection due to source/replica switch. Please ensure that your business has a reconnection mechanism."
		When local resources fail to meet the requirements for configuration changes, the current modification involves master-slave switching and restarting. The configuration change page will indicate: "During the configuration adjustment, the instance access is not affected. After the adjustment is completed, there will be a momentary disconnection due to source/replica switch. Please ensure that your business has a reconnection mechanism."
	Disk Expansion and Contraction	When only modifying the disk capacity, there is no need for master-slave switching or restarts. The configuration change page will indicate: "Current changes will not involve data migration or instance restart."

#### Note:

The ability to trigger a rapid configuration change during instance configuration change is closely related to the resource availability of the host machine. Please take the activation status of the rapid configuration change switch on the actual configuration adjustment page as the criterion.

If the Rapid Allocation Switch is activated, it means that the current host resources are sufficient to support rapid configuration changes.

If the rapid configuration change switch is not activated, it means that the host machine lacks sufficient resources to perform rapid configuration change.

When local remaining resources are sufficient to meet the conditions for rapid configuration change, the rapid configuration change will be used by default. If not to use it, turn it off on the Configuration Adjustment page. (i.e. you can choose to enable or disable the "Rapid Configuration Change" button on the Configuration Adjustment Page).

### Use Limits

A read-only instance with an exclusive VIP enabled does not support QuickChange.

If a read-only instance group (RO group) enables the "Remove Delayed RO Instances" feature, read-only instances whose delay exceeds the specified threshold will be removed from the RO group until the number of remaining ones reaches the allowed minimum number. If the number of the remaining available read-only instances in an RO group is less than or equal to the allowed minimum number, none of those instances will support QuickChange. If an RO group only has one read-only instance, the instance does not support QuickChange. When upgrading a minor version, whether a restart is involved, the prompts on the configuration change page shall prevail.

### **Configuration Adjustment Rules**

You can adjust the configuration of a TencentDB for MySQL instance and its associated read-only and disaster recovery instances only when they are in normal status (running) and are not executing any task.

You cannot cancel a configuration adjustment operation in progress.

The name, access IP, and access port of the instance remain unchanged after configuration adjustment. During configuration adjustment, you should try to avoid such operations as modifying MySQL's global parameters and user password.

Data migration may be involved in configuration adjustment. During data migration, the TencentDB for MySQL instance can be accessed normally and the business will not be affected.

Instance switchover may be needed after configuration adjustment is completed (i.e., the MySQL instance may be disconnected for seconds). It is recommended that applications be configured with auto reconnection feature and that instance switchover be conducted during the instance maintenance period. For more information, see <u>Setting Instance</u> <u>Maintenance Window</u>.

Basic single-node TencentDB for MySQL instances are unavailable for about 15 minutes in the process of configuration adjustment. We recommend you adjust instance configuration during off-peak hours.

### Instance Specification and Storage Table

#### Two-node/Three-node (local SSD disk)

Isolation Policy	CPU and Memory	Maximum IOPS	Storage Space	
	1-core 1000 MB	1200	25GB - 3000GB	
	1-core 2000 MB	2000		
	2-core 4000 MB	4000		
	4-core 8000 MB	8000	-	
	4-core 16000 MB	14000		
	8-core 16000 MB	20000		
	8-core 32000 MB	28000	250D 40000D	
Conorol	16-core 32000 MB	32000	2368 - 400068	
General	16-core 64000 MB	40000		
	16-core 96000 MB	40000	-	
	16-core 128000 MB	40000	25GB - 8000GB	
	24-core 192000 MB	60000	25GB - 6000GB	
	24-core 244000 MB	60000	25GB - 8000GB	
	32-core 256000 MB	80000		
	48-core 488000 MB	120000	25CP 12000CP	
	80-core 690000 MB	140000	2300 - 1200000	
Dedicated	2-core 16000 MB	8000	25GB - 4000GB	
	4-core 16000 MB	10000		
	4-core 24000 MB	13000		
	4-core 32000 MB	16000		
	8-core 32000 MB	32000		
	8-core 48000 MB	36000		



	8-core 64000 MB	40000	
	12-core 48000 MB	36000	
-	12-core 72000 MB	40000	
	12-core 96000 MB	48000	
	16-core 64000 MB	60000	
	16-core 96000 MB	60000	
	16-core 128000 MB	60000	
	24-core 96000 MB	72000	2568 - 800068
	24-core 144000 MB	76000	
	24-core 192000 MB	80000	
	32-core 128000 MB	80000	
	32-core 192000 MB	90000	
	32-core 256000 MB	100000	
-	48-core 192000 MB	120000	
	48-core 288000 MB	140000	
	48-core 384000 MB	140000	25GB - 10000GB
	64-core 256000 MB	150000	
	64-core 384000 MB	150000	
	64-core 512000 MB	150000	25CP 12000CP
	90-core 720000 MB	150000	2300 - 1200000

#### Note:

The storage space cap of an instance specification may vary by region as displayed on the purchase page.

### Adjusting Instance Configuration in the Console


1. Log in to the TencentDB for MySQL console, locate the desired instance in the instance list, and select **More** > **Adjust Configurations** in the **Operation** column.

2. In the pop-up window, select the desired configuration and click **Submit**.

#### Note:

When the local resources are sufficient, you can enable the QuickChange mode by turning on the **QuickChange** toggle on the configuration adjustment page in the console.

When you are only adjusting the cloud disk (storage space), which does not involve a restart, the system will provide options for the execution time if rapid configuration change is enabled.

In some cases, instance restart is not required in the QuickChange mode and the configuration adjustment will take effect immediately after you submit the request, as shown below:

When the primary instance triggers migration configuration change, the data validation sensitivity settings will be displayed, and you can set the data validation speed to high, standard, or low.

High: Suspends the data consistency comparison task when database load pressure increases, to avoid additional performance impact from the task, which will extend the duration of the data consistency comparison task.

Standard: Suspends the data consistency comparison task when the database load pressure is significant, to prevent noticeable performance impact.

Low: Perform the data consistency comparison task even when the database load pressure is significant, to shorten the duration of the task.

# Adjusting Instance Configuration Through API

You can adjust the instance configuration using the UpgradeDBInstance API. For more information, see UpgradeDBInstance.

# FAQs

#### Will there be a primary-secondary delay issue during the configuration change?

During the configuration change of the primary instance, data comparison occurs, which might lead to primarysecondary delay.

#### Will instance configuration adjustment affect instances?

In the process of TencentDB for MySQL configuration adjustment, data migration may occur, and instances can still be accessed during the process. After the migration is completed, there is a switch which causes a short disconnection lasting for just seconds, ensure that your business has a reconnection mechanism. Basic single-node TencentDB for MySQL instances are unavailable for about 15 minutes in the process of configuration adjustment. We recommend you adjust instance configuration during off-peak hours.

#### Why can't my instance be downgraded?

It may be because the used storage capacity has reached the maximum capacity of the hard disk. To downgrade you instance, you need to clean up data first and make sure the remaining available capacity accounts for more than 20% of the total capacity or over 50 GB.

# Why is my instance in the "Waiting for switch" status for a long time after I adjust instance configuration in the console?

It may be because you select **During maintenance time** as the **Switch Time** when you adjust instance configuration in the console, so the instance will not be switched immediately after the adjustment. To switch immediately, find the target instance in the instance list and click **Switch Now** in the **Operation** column. The switch will cause a momentary disconnection. Make sure that your business has a reconnection mechanism.

#### How long does it take to upgrade instance configuration?

The time it takes depends on the instance's data volume and data replication speed. Instances can still be accessed during the upgrade, but after the upgrade is completed, there is a VIP switch which causes a short disconnection lasting for just seconds.

#### How do I view the progress of instance configuration adjustment?

You can view the progress in Task List in the console.

#### What should I do if the disk space is being used up?

If over 85% disk space is used, we recommend that you delete data no longer used or expand disk space in the console by selecting **More** > **Adjust Configurations** in the **Operation** column on the right of the instance list.

# How do I know whether the QuickChange mode is supported when I adjust instance memory or disk capacity?

On the configuration adjustment page, if the QuickChange mode is supported, the **QuickChange** toggle can be turned on or off as needed; otherwise, the toggle cannot be used.

#### Does expanding memory or disk capacity affect the kernel minor version of the instance?

If the kernel minor version of the instance is not the latest when memory or disk capacity is expanded, it will be upgraded to the latest. In this case, enabling the QuickChange mode will cause the database to restart.

#### Will the instance restart if the QuickChange mode is enabled?



The instance needs to be restarted in some cases. A prompt will appear asking if you want to restart the instance on the configuration adjustment page.

#### Note:

If the instance is running the latest kernel minor version, it won't restart if you only adjust the disk capacity in the QuickChange mode.

# How do I know whether the QuickChange mode is enabled when I upgrade instance configuration in the console?

You can check the **QuickChange** switch on the configuration adjustment page: if the switch is toggled on, the QuickChange mode is enabled; otherwise, the mode is disabled.

# How do I know whether the QuickChange mode is enabled when I use APIs to adjust instance configuration?

APIs do not support QuickChange for the time being, so instance configuration can be adjusted only by migrating data if you use APIs. APIs will support QuickChange in the future.

#### Will database parameters be modified during database configuration adjustment?

The innodb\_buffer\_pool\_size parameter will be modified according to the configuration changes.

# Will database parameters be modified during database configuration adjustment in the QuickChange mode?

It is the same with the normal mode. In the QuickChange mode, some parameters will be modified according to the configuration changes.

#### What is the difference between QuickChange mode and normal mode?

The QuickChange mode requires no data migration, so it takes less time to adjust instance configuration.

# Modifying the Instance Name

Last updated : 2025-04-02 10:16:48

The instance name is mainly used to distinguish and manage TencentDB for MySQL instances. You can modify the instance name through the console.

This document describes how to modify the instance name through the console.

# Instance ID/Name Description

Primary instances, read-only instances, disaster recovery instances, and read-only analysis engine instances support modifying the instance name.

The instance name can contain no more than 60 characters and supports only digits, uppercase and lowercase letters, Chinese characters, and special characters: - ... /()[] () +=: : @

### Modifying the Instance Name in the Instance List

1. Log in to the TencentDB for MySQL console.

2. Find the target instance in the instance list and click the editing icon next to the instance name.

3. Enter a new instance name and click OK.

# Modifying the Instance Name on the Instance Details Page

1. Log in to the TencentDB for MySQL console.

2. Find the target instance in the instance list and click **Instance ID** or **Manage** in the **Operation** column to go to the instance management page.

3. Click the editing icon next to the instance name on the instance details tab of the instance management page.

4. Enter a new instance name and click OK.

### **Related APIs**

API

Description



ModifyDBInstanceName This API is used to modify the name of a TencentDB instance.

# Change the Port Number

Last updated : 2023-12-27 10:16:47

MySQL instances support modifying the port number of their private IP addresses. This document describes the steps to change the port number through the console.

#### Note:

Only supports modifying the port number of private IP address of the MySQL instance; the port for public IP address is not supported.

# Prerequisites

#### You have created a MySQL instance..

#### Note:

Before modifying the port for the private IP address of a read-only instance, you must verify that the private network address for the read-only instance has been enabled.

P	cdb_ I Running
K	Read-Only Instance 0 seconds delayed to the source instance $\phi$
Instance ID	cdbro-:
Region/AZ	North China(Beijing) / Beijing Zone 6
Private Netw	vork Address (i) Enable

# Modifying the private address port numbers of Master/Read-only Instances

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operations column to go to the instance details page.

2. Click the edit icon following the private network address.

CC Hi	gh-availability instance	Running The current instance has a source-replica delay of
Instance ID	cdb-{	
Region/AZ	North C	hina(Beijing) / Beijing Zone 5 Migrate to New AZ
Private Network A	Address (i) One Port: 33	15 Click Connectivity Check New Architecture
Public Network A	ddress (i) Enable	
Database Proxy	Enabled	2 address(es)
3. In the displayed dialog box,	modify the private network	port number and then click <b>OK</b> .

Modify Priva	te Network Address	<
Private IP*	5	
	Available Private IP Range: 172.21.0.0/20	
Private Port*	3308	
	Port value range: 1024-65535	
Note: modifying	the private network address will affect the database service being accessed.	
	OK Cancel	

# Modifying the Private Network Address Port Number of the read-only Group

1. Log in to the TencentDB for MySQL console..

2. Locate the master instance whose read-only group's private network address needs to be modified in the instance list, then click its **instance ID** or **Manage** in the operations column to access the instance details page.

3. Navigate to the read-only instance page and click the edit icon following the private network address of the readonly group.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Security	Connection Check	Read-Only Instance (
Create										
RO Group ID:cdbrg- RO Group:ro_g	Region:N Network:	orth China(Beijing)/Beijing Zone Det	6	Private Network Address: Public Network Address: Ene	):3306 💉	De	ayed Replication Status:	Closed		Create Configura Rebalancing

4. In the displayed dialog box, modify the private network port number and then click OK.

# Migrating AZ

Last updated : 2025-06-10 09:29:39

You can migrate a TencentDB for MySQL instance to another AZ in the same region. All its attributes, configurations, and connection address will stay unchanged after the migration. The time it takes to migrate the instance is subject to the instance's data volume.

For example, you can migrate to a new AZ in the following scenarios:

If you want to modify an instance's type, but the current AZ doesn't support the new instance type, you can migrate the instance to an AZ supporting the new type.

If the current AZ has no remaining resources for scaling, you can also migrate the instance to another AZ in the same region with sufficient resources to meet your business needs.

# Prerequisites

The instance is running and uses the two-node or three-node architecture. The region where the instance is located has multiple AZs to support cross-AZ migration.

# **Billing description**

This feature is free of charge. There is no charge even for migrating an instance from a single AZ to multiple AZs.

# Feature overview

AZ migration will not result in a VIP change.

The source instance is not decoupled from the read-only instances after AZ migration and can still be synced with them.

You can choose the AZ of read-only instances.

If the target instance has a task lock on the cloud platform during the DTS task, cross-AZ migration cannot be performed.

If there are ongoing DTS tasks, after AZ migration, you need to restart such tasks.

DTS export will fail if the source instance undergoes a cross-AZ migration switch during dumper export.

### Impact

The instance will be momentarily affected when its AZ is switched; therefore, make sure that your application has an automatic reconnection mechanism.

# Use limits

During the migration of the availability zone, the system will check if the instance disk is overused. If the disk is overused, the migration of the availability zone cannot be carried out. It is recommended to retry after expanding the disk. If the disk space has exceeded the maximum storage limit of the current instance specification, it is recommended to retry after upgrading the instance specification configuration. For detailed specifications, disk limits, and related operations, please refer to Adjusting Database Instance Specifications.

RO instances do not support cross-regional migration.

The migration of the availability zone does not currently support enabling instances with the database proxy. Please turn off the database proxy before migrating across availability zones.

During the migration switch, access via the RO group is not possible (excluded).

Under the dual-node and triple-node architecture, the selection of the primary and secondary availability zones for migration is subject to regional and remaining resource restrictions. When migrating in the console, select the target availability zone, and the secondary availability zone option will automatically update.

RO instances in the same availability zone as the primary instance will by default migrate to the target availability zone with the primary instance. RO instances not in the same availability zone as the primary instance do not currently support the migration of availability zone operations.

# Migration type

Migration Type	Applicable Scenario	Supported Instance Types
Migration from one AZ to another AZ	The AZ where the instance is located is under full load or other conditions that affect the instance performance.	Source, read-only, and disaster recovery instances
Migration from one AZ to multiple AZs	You can improve the disaster recovery capability of the instance and implement cross-data center disaster recovery. Source and replica instances are located in different AZs. Multi-AZ instances can withstand higher levels of disasters than single-AZ instances. For example, the latter can tolerate server- and rack-level failures, while the former can tolerate data center-level failures.	Source, read-only, and disaster recovery instances

Migration	You want to meet specific feature requirements.	Source,
from		read-only,
multiple		and disaster
AZs to		recovery
one AZ		instances

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the **Operation** column to enter the instance details page.

2. On the Instance Details page, select Basic Info > Region/AZ and click Migrate to New AZ, or select Availability Info > Deployment Mode and click Modify Replica AZ.

3. In the pop-up window, adjust the relevant configurations and click **Submit** after confirming that everything is correct.

**New AZ**: You can change the source AZ in the drop-down list or select **Yes** for **Multi-AZ Deployment** to modify the replica AZ.

**Delay Threshold for Data Consistency Check**: This option is available only when you change the primary AZ. The threshold can be an integer between 1 and 10 seconds.

#### Note:

There may be a delay during the data consistency check. You need to set a data delay threshold. The database consistency check will be paused when the delay exceeds the set value and will be resumed when the delay drops below the threshold. If the threshold value is too small, the migration may take longer.

**Switch Time**: You can choose to switch during the maintenance time or upon migration completion. For more information, see <u>Setting Instance Maintenance Time</u>.

**Same Availability Zone RO Instance Migration**: (This setting only appears when the primary instance has a readonly instance in the same availability zone.) Select whether the same availability zone RO should follow the primary instance to migrate to the new availability zone.

#### Note:

During the process of adjusting configurations, data migration may occur within the instance. However, this will not affect the accessibility of the instance. After the migration is complete, a switch will occur, resulting in a second-level disconnection. Please ensure that your operations possess a reconnection mechanism.

# Primary/Replica Switch

Last updated : 2023-12-27 10:17:46

MySQL supports the switch of primary and secondary databases. In the event of an instance failure, the system automatically promotes the secondary database to primary to ensure system availability and data integrity. In addition to automatic switch during anomalous conditions such as failures, you can also manually switch the databases via the console.

# Background

In enterprise-level applications, the database often forms a critical component of the business system. Any database failure or downtime can have serious implications on business operations. To safeguard system availability and data integrity, high-availability solutions such as primary-replica replication are implemented. Primary-replica switch is a crucial technical method within the replication scheme, which enables a quick switch to the backup database when the primary database is faulty, averting business interruptions and data losses.



### Prerequisites

The instance architecture is set to either two-node or three-node.

The instance is in running status, with no ongoing tasks.

### Precautions

During the primary-secondary database switch, a momentary disconnection occurs. We recommend that you perform the switch during off-peak hours and make sure that your application has a reconnection mechanism. After the switch, the instance connection address remains the same. The original primary instance converts into a secondary instance, and the application automatically connects to the new primary instance (formerly the secondary instance).

### The steps are as follows:

#### Note:

Automatic fallback is not supported after the primary/secondary switch, but manual fallback can be done if needed (i.e., perform another primary/secondary switch after the first one is completed).

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operation column to access the instance details page.

2. On the Instance Details page, click Switch Primary/Secondary under Availability Info.

Availability Info	
Data Replication Mode	Async Modify Replication Mode
Deployment Mode	Multi-AZ Modify Replica AZ
Source AZ	Beijing Zone 5 Source/Replica Switch
Replica AZ	Beijing Zone 6

3. In the pop-up window, select the switch time, check "With the Primary/Secondary instance switch, there will be a brief disconnection; ensure that your application has a reconnection mechanism", and click OK. During Maintenance Period: When data synchronization between primary and secondary AZs has been completed, the system will switch the primary and secondary instances in the next maintenance period. During this period, the instance status shows as Pending Primary/Secondary Switch and operations such as configuration adjustments or instance termination cannot be performed. For operations related to setting the instance maintenance time, please see Setting Instance Maintenance Time.

**Switch Immediately**: The switch will occur immediately upon the completion of data synchronization between the primary and secondary AZs.

Source/Rep	olica Switch		
Switch Time	During maintenance time	Switch Now	Learn more about switch
	time 🖸		
	Maintenance Time: 03:00-04:00 (	modify on the "Ins	tance Details" page)
	Source/replica switch is accollasting for just seconds. Pleas reconnection mechanism.	mpanied by a disc se make sure that Cancel	connection from the database your business has a

# FAQs

#### Will the instance address change after the primary and secondary database switch?

The instance address remains the same after the switch. The original primary instance works as a secondary, with applications automatically connecting to the new primary instance (formerly the secondary).

# After a manual switch, if an instance is upgraded or expanded, will this affect or reset the information of the primary and secondary AZs?

The switch will not affect or reset the information of the primary and secondary AZs.

# **Terminating Instance**

Last updated : 2025-06-03 10:25:02

# Overview

You can return the instances in the console based on your business needs.

After a monthly subscribed instance is returned, it will be moved to the TencentDB recycle bin and retained there for seven days. During the retention period, the instance cannot be accessed, but it can be restored after renewal. After a pay-as-you-go instance is returned, it will be moved to the TencentDB recycle bin and retained there for 24 hours. During the retention period, the instance cannot be accessed, but it can be restored after renewal. When an instance is returned and its status has changed to **Isolating**, it will no longer generate fees.

# Precautions

After the instance is terminated, its data cannot be retrieved, and its backup files will also be terminated. Data recovery on the cloud is impossible. Transfer your backup files to a safe location in advance. It is suggested to configure backup retention.

After an instance is terminated, its IP address is released. For the read-only or disaster recovery instance associated with this instance:

The read-only instance will be terminated as well.

The disaster recovery instance will be disassociated and automatically upgraded to the primary instance.

After an instance is terminated, the following refund policies are applied:

The common self-service refund is returned to your TencentDB account based on the proportion of cash and free credits used for the purchase.

For promotion channel orders, 25% of the actually paid order amount will be charged as the refund fee. Such orders do not support self-service refunds. You can submit a ticket to apply for a refund.

# Directions

1. Log in to the TencentDB for MySQL console, select the required instance in the instance list, and choose **More** in the **Operation** column > **Terminate/Return** or **Terminate Instance**.

2. In the pop-up dialog box, complete the **Backup File Retention** settings, and click **Confirm Terminate**.

Backup File

Description

Retention	
Do not retain	If you select Do not retain, backups will be terminated simultaneously after instance termination and will not be saved.
Retain	If you select Retain, backups within the specified scope will be retained after the instance is terminated. The scope options are as follows: Retain all: Retain all backups generated by the instance in the current region. You can choose whether to retain the binlog files. Retain the latest one: Retain the latest 1 backup generated by the instance in the current region. By default, binlog files cannot be retained. Retain those within the specified time range: Retain backups generated by the instance in the current region and after the selected time. You can choose whether to retain the binlog files. For example: You can retain backups generated after April 28, 2025.

3. After the operation is completed, the instance is moved to the recycle bin, and its status changes to Isolating.

#### Note:

If you retain the backup file, after the instance is removed from the recycle bin, the retained backup file can be queried and downloaded in Database Backup > Backup List > Backups of Terminated Instances.

# Instance Upgrade Upgrading the database version of MySQL 5.5 and MySQL 5.6

Last updated : 2023-11-20 15:25:07

# Overview

This document describes how to upgrade the TencentDB for MySQL engine in the console.

TencentDB for MySQL supports database engine upgrade on the following versions:

From MySQL 5.5 to MySQL 5.6

From MySQL 5.6 to MySQL 5.7

From MySQL 5.7 to MySQL 8.0

#### Note:

Database engine downgrading is not supported.

Upgrading across major versions is not supported. For example, to upgrade a TencentDB for MySQL 5.5 instance to

MySQL 5.7 or later, you have to upgrade it to MySQL 5.6 first.

Currently, MySQL 5.7 cannot be upgraded to MySQL 8.0.

Single-node instances of the cloud disk edition don't support engine version upgrade.

# Upgrade Rules

The create table ... as select ... syntax is not supported.

Source-replica sync in TencentDB for MySQL 5.6 and 5.7 is implemented based on GTID. Only InnoDB is supported by default.

MyISAM tables will be converted to InnoDB tables during upgrade from MySQL 5.5 to 5.6 if the last full backup is logical cold backup. We recommend you complete the conversion first before upgrading.

During upgrade, TencentDB for MySQL will clear the slow\_log table. Save the logs before upgrading if necessary.

If an instance to be upgraded is associated with other instances (e.g., source or read-only instances), these instances will be upgraded together to ensure the data consistency.

TencentDB for MySQL upgrade involves data migration and generally takes a relatively long time. Your business will not be affected during the upgrade and can be accessed as per usual.

Instance switch will be required after version upgrade is completed (that is, the MySQL database may be disconnected for seconds). We recommend you use applications configured with auto reconnection feature and



conduct the switch during the instance maintenance time. For more information, see Setting Instance Maintenance Window.

If the number of tables in a single instance exceeds one million, upgrade may fail and database monitoring may be affected. Make sure that the number of tables in a single instance is below one million.

### Directions

1. Log in to the TencentDB for MySQL console, find the target instance in the instance list, and select **More** > **Upgrade Version** in the **Operation** column.

#### Note:

MySQL 8.0 cannot be upgraded to a later version.

2. In the pop-up window, select the target database version and click Upgrade.

As database upgrade involves data migration, a momentary disconnection from the MySQL database may occur after the instance upgrade is completed. You can set the switch time to **During maintenance time**, so that the switch will be initiated within the next maintenance time after the upgrade.

#### Note:

When you select **During maintenance time** for **Switch Time**, the switch will not occur immediately after the database specification upgrade is completed; instead, the sync will continue till the instance goes into the next **maintenance time** when the switch will be performed. As a result, the overall time it takes to upgrade the instance may be extended.

Database version upgrade instructions:						
The source instance and the associated rea at the same time. The version can be upgra	d-only instances and ded via any associat	d disaster recov ted instance.	very instances are upgra	ded to th	e target da	tabase version
The upgrade process has no impact on exis	ting business.					
<ul> <li>There is a second-level interruption during to period to reduce the impact on the service.</li> </ul>	he upgrade switch p	rocess. It is rec	commended to switch in r	naintena	nce time o	r off-peak
istance ID						
istance Name						
istance Type	Source Instance					
elay Threshold for Data Consistency Check 🛈	10	sec				
	Enter an integer fr is lower.)	om 1 to 10 (Th	e new configuration may	take lon	ger to take	effect if the thr
witch Time	During mainte	nance time	Upon upgrade comp	etion	Switch Ti	me Descriptior
	Maintenance Time	03:00-04:00(M	odify maintenance time of	on the in	stance det	ails page)
atabase Version	MySQL5.5	MySQL5.6	(Current Version)	MyS	QL5.7	
	It may take longer	to migrate data	a if there are tables witho	ut orima	rv kevs. Ct	eck Tables wit
	Primary Keys 🛂		199 10		.,,	
I have read and agreed to Database Version U	porade Rule					
I have read and agreed to Database version o	pgrade rule					

### FAQs

#### Will TencentDB for MySQL automatically back up data before upgrade?

TencentDB for MySQL adopts the daily real-time dual-server hot backup mechanism, which supports lossless restoration of data from the last 7–1830 days based on data backup and log backup (binlog).

#### Can TencentDB for MySQL be downgraded from MySQL 5.7 to MySQL 5.6?

No. Database engine downgrading is not supported. To use a MySQL 5.6 instance, you need to terminate or return the existing MySQL 5.7 instance, and then purchase a MySQL 5.6 instance.

#### Will there be a source-replica delay during upgrade?

Source instance upgrade requires data comparison and may cause a source-replica delay.

# Will the instance switch after database engine version upgrade affect my TencentDB for MySQL instance?

The upgrade won't affect your business, but the TencentDB for MySQL instance may be disconnected for a few seconds. We recommend you configure an automatic reconnection feature for your application and conduct the switch during the instance maintenance time.

# How long will it take to upgrade the database engine version of a TencentDB for MySQL instance? How do I check the upgrade progress?

The time it takes depends on the instance's data volume and the read requests to replicate data.

TencentDB for MySQL upgrade involves data migration and generally takes a relatively long time. Your business will not be affected during the upgrade and can be accessed as per usual.

#### Why is the instance always in the "Waiting for switch (after upgrade)" status?

It may be because you select **During maintenance time** for **Switch Time**, and the switch will be initiated within the next maintenance time after the upgrade.

To switch immediately, find the target instance in the instance list and click **Switch Now** in the **Operation** column. The switch will cause a momentary disconnection. Make sure that your business has a reconnection mechanism.

# Upgrading from MySQL 5.7 to MySQL 8.0

Last updated : 2025-06-10 09:31:35

TencentDB for MySQL supports direct database version upgrade through the console. This document describes how to upgrade the CDB engine from MySQL 5.7 to MySQL 8.0 in the console.

### Overview

MySQL 8.0 offers significant performance improvements and enhanced features. Its performance under high concurrency and large data volume scenarios is particularly prominent, offering greater system performance and stability.

The upgrade of the following database engine versions is supported by TencentDB for MySQL:

From MySQL 5.5 to MySQL 5.6 From MySQL 5.6 to MySQL 5.7

From MySQL 5.7 to MySQL 8.0

#### Note:

Single-node instances of the cloud disk edition do not support engine version upgrade.

The upgrade process has no impact on existing business.

### Considerations

#### Feature Limits

Limit Type	Description
Instance Limits	The database upgrade operation from MySQL 5.7 to MySQL 8.0 is only supported for local disk instances with two-node or three-node architecture.
	If a master instance has multiple read-only instances, the read-only instances will be upgraded together. If a master instance has a disaster recovery instance, the disaster recovery instance must first be released before upgrading. After the upgrade is completed, a new disaster recovery instance can be recreated.
	Before initiating the database version upgrade, ensure the instance status is running, and there are no other tasks being executed or to be executed. If there are other tasks, please wait until the task is completed.
	Before upgrading, a compatibility check will be performed on the instance. The check results



	are valid for 24 hours. If it exceeds this time, you need to recheck it. During the check, if you add a read-only instance or a disaster recovery instance, you also need to re-execute the task.
Upgrade Limits	Upgrading across major versions is not supported. For example, MySQL 5.6 cannot be directly upgraded to MySQL 8.0. You need to upgrade it to MySQL 5.7 first, and then upgrade to MySQL 8.0.
	Downgrading after upgrade is not supported. <b>Note:</b> After the upgrade is completed, backup sets of the old version cannot be used to restore the new version instance. If you need to perform cloning or rollback operations, please select the backup set or point in time generated after the instance upgrade. Backup sets of older versions can be downloaded for local restoration operations.
Database Limits	After upgrading the database instance to MySQL 8.0, the unsupported settings in the sql_mode parameter will be reset to the default values of MySQL 8.0.
	If a partition table is used in the RocksDB storage engine, upgrading to MySQL 8.0 is not supported. You need to modify it to the InnoDB engine or delete the partition table to upgrade.
	If there are tables, views, stored procedures and triggers that contain MySQL 8.0 reserved keywords in the database instance, the upgrade will fail.
	If stored procedures, triggers, views, or functions in the database contain Changes in MySQL 8.0, the upgrade will fail.
	If you are using the MyISAM or Memory engines, upgrading is not supported. Please convert to the InnoDB engine first. Note:
	If your database is already running on the InnoDB engine, but some tables are still using other engines, please use the ALTER TABLE  engine=InnoDB; command to convert those tables to the InnoDB engine before upgrading.

#### Preparation

Please be sure to carefully compare the differences between the versions before and after the upgrade. We recommend testing syntax on a new MySQL 8.0 instance-this will help avoid issues with lower-version syntax or features not being supported after the upgrade. For more information, please refer to List of functional differences between MySQL8.0 and MySQL5.7.

We recommend that you clone the original instance for upgrade testing first, and confirm that all functions are normal before upgrading.

Before upgrading, please check whether there is a successful record of full data backup in the past week. If

necessary, please perform a full data backup.

During the database upgrade, a brief disconnection of the TencentDB for MySQL service may occur. It is suggested that you carry out the upgrade during off-peak period or ensure your application has an automatic reconnection

mechanism.

Please ensure that free disk space is available before upgrading.

Before upgrading, review the Keywords and Reserved Words that you should avoid using them when creating custom functions.

To ensure stability and performance of MySQL in the new version, after the database version has been upgraded, CDB will not support the inspection or modification of MySQL 5.7 parameters that have been deprecated in MySQL 8.0. For more information, please refer to Options and Variables Removed in MySQL 8.0. Before upgrading, it is recommended that you back up modification records of relevant parameters for subsequent operations and audits.

### Directions

1. Log in to the TencentDB for MySQL console, fin the target instance in the instance list, and select **More** > **Upgrade Version** in the Operation column.

#### Note:

You can also click on the **Instance ID** to go to the instance details page and click **Upgrade Version** after the database version.

2. In the Database Version Upgrade window, complete the required settings, read and **select** the Database Version Upgrade Rule, then click on **Upgrade Check**.

Parameter	Description
Delay Threshold for Data Consistency Check	Enter an integer from 1 to 10 ( <b>The new configuration may take longer to take effect if the threshold is lower.</b> )
Switch Time	During maintenance time: Set the switch time to During maintenance time, so that the switch will be initiated within the next maintenance time after the upgrade. For more information, see Setting Instance Maintenance Time. Upon upgrade completion: Set the switch time to Upon Upgrade Completion, so that the switch will be initiated immediately after the instance data migration has been completed.
Database Version	MySQL 8.0 is selected by default.
default_collation_for_utf8mb4	Specify the default collation for a character set. For the collation of a specific selectable character set, refer to the actual page selections.

3. During the upgrade check process, the system will start a check task to verify whether the instance to be upgraded meets the criteria of **not having any disaster recovery instances and having parameter configurations that** 

meet the specifications. After the check is completed and the upgrade conditions are met, click Upgrade.

3.1 Verify whether there are any **disaster recovery instances** under the master instance. If master instance has a disaster recovery instance, the check will fail. You need to release the disaster recovery instance first, and recreate it once the upgrade is completed.

3.2 The system will automatically verify if the parameter configuration is up to the specified standards. If it doesn't meet the specification requirements, adjustments can be made according to the error message in the check box. **Note:** 

After the upgrade check task is executed successfully, please perform the upgrade within 24 hours. If it times out, check again.

4. After upgrading, return to the instance list. Select **Task List** from the left sidebar to track the progress and details of the database version upgrade task for the instance.

5. Find the Task ID of the recent database version upgrade in the task list, and you can query the task progress. By clicking on **Task Details** in its Operation column, more upgrade-related information will be displayed.

6. The database version upgrade is completed when the instance status changes from **Changing configuration** to **Running**.

#### Note:

If you select **During Maintenance Time** as the switch time for the upgrade, the instance status will display **Awaiting switch** if the next maintenance time has not been reached after the instance upgrade is completed. If an immediate switch is necessary, you can click **Switch Now** in the **Operation** column of the instance list. However, be aware that the switch will involve a momentary disconnection lasting for a few seconds. Please ensure that your business is equipped with a reconnection mechanism.

# Appendix: MySQL 8.0 vs. MySQL 5.7

MySQL 8.0 introduces over 300 new features, including invisible indexes, descending indexes, functional indexes, common table expressions, window functions, instant column addition, implicit primary keys and more. These capabilities make development more adaptable and efficient.

MySQL 8.0 introduces improvements in performance, implementing features such as Histograms, HASH JOINS, parallel queries, and JSON columns that enhance capability. These additions have served to increase both the query performance and response time of the database, enabling it to process large-scale data and complex queries more swiftly.

MySQL 8.0 supports the caching\_sha2\_password authentication plugin, thus enhancing security.

MySQL 8.0 Optimizer enhances functionality.

MySQL8.0 enhances replication performance.

MySQL 8.0 supports the creation of multi-valued indexes and optimization under derived conditions.

MySQL 8.0 supports the reading of MySQL authorization tables.

MySQL 8.0 supports resource allocation control.

### FAQs

#### Will TencentDB for MySQL automatically back up data before upgrade?

TencentDB for MySQL adopts the daily real-time dual-server hot backup mechanism, which supports lossless restoration of data from the last 7–1830 days based on data backup and log backup (binlog).

#### Can TencentDB for MySQL be downgraded from MySQL 8.0 to MySQL 5.7?

No. Database engine downgrading is not supported. To use a MySQL 5.7 instance, you need to terminate or return the existing MySQL 8.0 instance, and then purchase a MySQL 5.7 instance.

#### Will there be a source-replica delay during upgrade?

Source instance upgrade requires data comparison and may cause a source-replica delay.

# Will the instance switch after database engine version upgrade affect my TencentDB for MySQL instance?

The upgrade won't affect your business, but the TencentDB for MySQL instance may be disconnected for a few seconds. We recommend you configure an automatic reconnection feature for your application and conduct the switch during the instance maintenance time.

# How long will it take to upgrade the database engine version of a TencentDB for MySQL instance? How do I check the upgrade progress?

The time it takes depends on the instance's data volume and the read requests to replicate data.

TencentDB for MySQL upgrade involves data migration and generally takes a relatively long time. Your business will not be affected during the upgrade and can be accessed as per usual.

You can monitor the progress and details of tasks in the task list. Keep an eye on the task execution progress for insight into the upgrade status.

#### Why is the instance always in the "Waiting for switch (after upgrade)" status?

It may be because you select **During Maintenance Time** for **Switch Time**, and the switch will be initiated within the next maintenance time after the upgrade.

To switch immediately, find the target instance in the instance list and click **Switch Now** in the **Operation** column. The switch will cause a momentary disconnection. Make sure that your business has a reconnection mechanism.

# **Upgrading Kernel Minor Version**

Last updated : 2025-06-10 09:37:51

TencentDB for MySQL will progressively release new minor versions of the database kernel. By upgrading to these newer minor versions, your database instance can benefit from enhanced performance, improved security, support for new features, better compatibility, and bug fixes. This document provides guidance on how to upgrade the minor version of your database kernel through the console, along with detailed explanations.

#### Note :

For details on the TencentDB for MySQL kernel minor version, see Kernel Version Release Notes.

# **Operational Scenarios**

Upon the release of a new minor version of the database engine, or when the minor version of the database engine in your current instance becomes outdated and necessitates an upgrade to a newer version, you may manually execute the upgrade via the console.

#### Note:

In the event of a bug in the backend program or the detection of security vulnerabilities, the system will dispatch upgrade notifications through internal messages, text messages, and other means. TencentDB for MySQL will undergo upgrades and repairs during the maintenance window of your instance.

When operations that trigger instance migration in TencentDB for MySQL occur (such as scaling instance specifications, expanding or contracting disk capacity, database version upgrades, etc.), the system will not automatically upgrade your instance to the latest minor kernel version. Should you require, you may manually perform the upgrade. When the primary instance possesses RO instances, the system will automatically assess master-slave synchronization compatibility, ensuring that migrations of the primary instance do not employ a minor version newer than that of the RO instances.

### Upgrade rules

If an instance to be upgraded is associated with other instances (e.g., source instance and read-only instances), these instances will be upgraded together to ensure data consistency.

TencentDB for MySQL upgrade involves data migration. The time it takes to migrate an instance depends on the size of the instance's data. Your business will not be affected during the upgrade and can be accessed as per usual.

### Notes

Instance switch will be required after version upgrade is completed, which will cause your instance to disconnect for seconds. We recommend that you conduct the switch during the instance maintenance time. Make sure that your business has a reconnection mechanism. For more information, see Setting Instance Maintenance Window. If the number of tables in a single instance exceeds 300000, upgrade may fail and database monitoring may be affected. Make sure that the number of tables in a single instance of tables in a single instance of tables in a single instance. The kernel minor version cannot be downgraded once upgraded.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click the target **Instance ID** or **Manage** in the **Operation** column to enter the instance details page.

2. In the Configuration Info section, click Upgrade Kernel Minor Version.

3. In the pop-up window, set the configuration items and click **Upgrade**.

**Delay Threshold for Data Consistency Check**: This option is available only during source instance upgrade, which can be an integer between 1 and 10 seconds. There may be a delay during the data consistency check. You need to set a data delay threshold. The database consistency check will be paused when the delay exceeds the set value and will be resumed when the delay drops below the threshold. If the threshold value is too small, the migration may take longer.

#### Note:

As database upgrade involves data migration, a momentary disconnection from the MySQL database may occur after the upgrade is completed. We recommend that you set the switch time to **During maintenance time**, so that the switch will be initiated within the next maintenance time after the instance upgrade is completed.

# Upgrading Two-Node Instances to Three-Node Instances

Last updated : 2024-07-22 14:39:43

### Overview

This document describes how to upgrade a TencentDB for MySQL instance from two-node to three-node in the console.

#### Note:

Only TencentDB for MySQL two-node instances can be upgraded to three-node.

The upgrade will not interrupt the instance service.

You can also purchase three-node instances on the purchase page.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance details page.

2. In the Configuration Info section, click Upgrade to Three-Node next to Architecture.

Configuration Info					
Architecture	Two-Node Upgrade to Three-Node				
Configuration	General-1core1000MB MEM,50GB storage, Adjust Configurations				
Max IOPS	1200				
Database Version	MySQL5.7 Upgrade Version Upgrade Kernel Minor Version				
Used/Total	<b>(</b> )				
Billing Mode	Pay as You Go				
Creation Time	2021-08-31 20:23:28				

3. In the pop-up dialog box, specify the data replication mode and availability zones and click **Submit**.



Data Replication Mode: For more information, see Database Instance Replication.

**Multi-AZ Deployment**: Multi-AZ deployment protects database service from being interrupted in case of database failures or availability zone failures.

For a three-node instance, its replicas can be deployed in the source AZ or different AZ. We recommend that you deploy one replica in the source AZ and another replica in a different AZ.

Currently, only some source AZs support selecting different AZs as replica AZs. You can view these source AZs and the supported replica AZs on the purchase page.

Upgrade to Three-No	de X
Instance Name	dts-test-stacy-1M
Data Replication Mode	Async Replication Semi-sync Replication
	Async is used for data replication.
Source AZ	Guangzhou Zone 6
Multi-AZ Deployment 🕥	Yes No
	It may take longer to migrate data if there are tables without primary keys. Check Tables without Primary Keys 🛛
New Fees	USD /hour/server (configuration fee, tiered pricing (), Billing Details ☑) USD /GB (Public traffic cost, free of charge now. About FREE ☑ )
	Submit Cancel

4. After the payment is completed, you will be redirected to the instance list. After the status of the instance changes from **Changing configuration** to **Running**, it can be used normally.

### FAQs

How do I view the architecture information of an instance?

Log in to the TencentDB for MySQL console. In the instance list, view the architecture information in the **Configuration** column; or click an instance ID or **Manage** in the **Operation** column to enter the instance details page, where you can view the architecture information in the **Configuration Info** > **Architecture** section.

Instance ID/Name 🗡	Monitoring/Status/Task <b>T</b>	Availability Zone 🔻	Configuration <b>T</b>
	Running	Guangzhou Zone 3	Two-node 1core1000MB/50GB Network: Default-VPC Default-Subnet

#### How do I view the source and replica AZs of an instance?

Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance details page, where you can view the source and replica AZs in the **Availability Info** section.

Availability Info Mod	ify Replication Mode	Source/Replica Switch
Data Replication Mode	Async	
Deployment Mode	Single-AZ	
Replication Status	Async	
Source database AZ	Guangzhou Zone 3	
Replica database AZ_1	Guangzhou Zone 3	

# **One-Click Upgrade to Cluster Edition**

Last updated : 2024-12-24 15:16:05

This document describes how to upgrade a TencentDB for MySQL dual-node or three-node instance to a cluster instance using the console's one-click upgrade feature.

# Prerequisites

The TencentDB for MySQL dual-node or three-node instance architecture is used. The instance is running.

### Notes

After an upgrade to the cluster edition, rollback is not supported. Proceed with caution.

After an upgrade to the cluster edition, the backup files of the original instance will not be retained. If you need to keep the backup files of the original instance, please ensure to download and save it in advance.

If an upgrade to the cluster edition is required, neither the original instance nor its read-only instances can have public network access enabled.

### Directions

#### 1. Log in to the TencentDB for MySQL console.

2. In the instance list, find the instance to be upgraded and click **One-click Upgrade** in the **Billing Mode** column, or check the target instance and click **One-click Upgrade** above the instance list.

Create One-Click Upgrad	de Quick Check (	Comparative Monitoring	Restart Renew	Separat	te keywords with	" "; press Enter to separate fil	ter tags		Q Ø
- Instance ID/Name T	Monitoring/Status/Task 🔻	AZ T	Configuration <b>T</b>	Databa 🔻	Engine <b>T</b>	Private Netwo 🛈 🔻	Billing Mode 🝸 💲	Project <b>T</b>	Operation
cdb-: Renew	<b>II</b> ⊙ Running	Beijing Zone 7	Two-Node(Local Disk) General-1core1000MB/25GB	MySQL8.0	InnoDB	3306 Network: c	Monthly Subscription Expire at 2024-10-04	Default Project	Log In Manage M

3. In the pop-up window, select Upgrade to TencentDB for MySQL Cluster Edition and click Upgrade Now.

4. On the right side of the pop-up interface, click Change Configuration.

5. In the pop-up window, set the relevant configuration information for the cluster, and then click OK.

Parameter	Description

AZ	It is used to select the availability zone for read-write nodes and read-only nodes.
Disk type	It is the cloud disk by default.
Instance specification	It is used to select the type, CPU, and memory of the instance.
Disk	It is used to select the disk type and storage space. The storage space range is from 30 to 32,000 GB, with an increment of 10. The storage space should be at least 1.2 times the data file size and not be less than the disk size of the original instance.
Data protection space	It is used to set data protection space, with a range from 1 to 10 GB. To ensure the recoverability of database instances, cluster instances require the setting of some space for data protection of database instances. The space cannot store data, belongs to the system protection space, and aims to prevent instances from malfunctioning when storage is full.

6. Due to the upgrade of cluster edition alarm metrics, you need to create an alarm policy. For operations, refer to Alarm Policy.

7. After changing the configuration and selecting the alarm policy, click **Upgrade Check** at the bottom of the redirected interface.

8. Check items for the upgrade check are listed in the table below. If any item fails, make a necessary adjustment as reminded, and then re-run the upgrade check. Once all the check items pass the upgrade check, you can proceed to the next step.

Check Item	Check Instance Type	Description
Audit	Primary instance Read-only instance	Check whether the original instance has database audit enabled. If so, migration is not possible. You can back up the database audit logs, disable database audit, and then upgrade to the cluster edition again. For the disabling of database audit, see <u>Disabling Audit Service</u> .
Backup	Primary instance	Check whether the original instance has regular retention policy, transition- to-cold storage for backup files, or transition-to-cold storage for binlog files enabled. If so, migration is not possible. You need to disable them and try again. For more details and operations, see Backing up Databases and Configuring Transition-to-Cold Storage.
Backup encryption	Primary instance	Check whether the original instance has backup encryption enabled. If so, migration is not possible. You need to disable it and try again. For more details and operations, see Backup Encryption.
DR instance	Primary instance	Check whether there are disaster recovery instances under the original instance. If so, migration is not possible. You need to release the disaster

TencentDB for MySQL



		recovery instances and try again. For instance releasing operations, see Terminating Instances.
Elastic expansion	Primary instance Read-only instance	Check whether the original instance has CPU elastic expansion enabled. If so, migration is not possible. You need to disable it and try again. For more details and operations, see CPU Elastic Expansion.
Instance status	Primary instance Read-only instance	Check whether the original instance's running status is normal. If other tasks are in progress, retry after the tasks are complete and the instance is in running state. Other abnormal status can be adjusted according to the suggested solutions.
Single node	Primary instance Read-only instance	Check the architecture of the original instance. Only dual-node and three- node instances support one-click upgrade to the cluster edition. If the original instance adopts a single-node architecture, migration is not possible. Select an instance architecture that supports migration.
Remote backup	Primary instance	Check whether the original instance has cross-region backup enabled. If so, migration is not possible. You can disable cross-region backup and try again. For related operations, see Cross-Region Backup.
SSL	Primary instance Read-only instance	Check whether the original instance has SSL encryption enabled. If so, migration is not possible. You can disable SSL encryption and try again. For related operations, see Setting SSL Encryption.
Version check	Primary instance Read-only instance	Check the kernel minor version of the original instance. Migration is possible only if the following conditions are met. For kernel minor version upgrade operations, see Upgrading Kernel Minor Version. MySQL 5.7 20230601 and later versions. MySQL 8.0 20230630 and later versions.
TDE encryption	Primary instance	Check whether the original instance has TDE encryption enabled. If so, migration is not possible. Use an instance with TDE encryption disabled instead for an upgrade.

9. Select the upgrade switching time. You can choose to switch during the maintenance time or upon completion of the upgrade. For the setting of maintenance time, see Setting Maintenance Time.

10. Read and check the reminder: During the instance upgrade, data will be migrated, and instance access will not be affected. After upgrade preparation is complete, an instance switch will occur, causing a momentary disconnection. Please ensure that your business has a reconnection mechanism.

11. Click **Start Upgrade**. You can check the upgrade progress in the Task List.

12. If you choose to switch upon upgrade completion, the original instance will be immediately upgraded to a cluster edition after data consistency verification. The upgrade is complete when the instance is in running state.

13. If you choose to switch during the maintenance time, the upgrade will be performed within the maintenance time you have set. If you want to complete the upgrade immediately, you can go to the instance list, and click **Switch Now** in the operation column of the target instance. After reading the precautions in the pop-up window, click **Switch** again. The upgrade is complete when the instance is in running state.

#### Note:

When you perform instance switching operation, the instance will switch to the cluster architecture.

After switching to the cluster architecture, a new read-only IP address will be generated.

After switching to the cluster architecture, the IP address of the read-write node will remain consistent with the preupgrade instance.

# CPU Elastic Expansion Overview

Last updated : 2023-12-27 10:19:09

This section introduces the features related to CPU Elastic Scaling in MySQL.

# **Chapter Overview**

CPU Elastic Scaling		Description	Document Guidance
Feature Overview		This section introduces the background, implementation process, instructions, and operation scope of the CPU Elastic Scaling feature.	INTRODUCTION
Auto Scaling		Provides guidance on how to enable, disable, and set threshold policies for the automatic CPU Elastic Scaling feature.	Configuring Auto Scaling
Operation	Manual Scaling	Provides guidance on how to enable, disable, and set threshold policies for the manual CPU Elastic Scaling feature.	Configuring Manual Scaling
	Event Alarm	Introduces metrics related to the CPU Elastic Scaling feature and the methods to configure event alarms.	Enabling Event Alarms
Performance Test		Evaluates the performance after enabling the CPU Elastic Scaling feature by employing the performance test tool SysBench.	Performance
Billing		Introduces the billing method and unit price of the CPU Elastic Scaling feature.	Billing Description
FAQs		Updates on common issues encountered when using the CPU Elastic Scaling feature and their solutions.	FAQs

# Overview

Last updated : 2023-12-27 11:01:17

MySQL supports manual scaling of the current instance's CPU upper limit and automatic scaling based on set CPU average usage thresholds and observation periods. This provides better mitigation for sudden request performance pressure, accommodating peak business traffic and ensuring online business stability. This section provides a brief overview of the CPU Elastic Expansion feature.

# Background

In many business scenarios, peak traffic periods or sudden traffic surges can impose tremendous pressure on server loads. The CPU utilization of the application can skyrocket, and if CPU resources are not scaled in a timely manner, it could lead to system performance degradation, increased request response time, and other issues, severely affecting user experience and satisfaction. Beyond burst traffic scenarios, situations such as low cache hit rate in the database, ineffective indexes, database deadlocks, complex database query statements, a high volume of concurrent updates, coupled with insufficient hardware configurations, can also lead to CPU overloads.

#### **Traditional Solutions**



Upon detecting a CPU overload, users typically manually upgrade it through the console and wait for the upgrade to be completed. The upgrade duration, influenced by the amount of data, is uncertain. Moreover, the upgrade process may result in transient disruptions, affecting regular operations. After the business traffic become normal after the upgrade, users need to manually apply for a downgrade through the console. Thus, the traditional solution poses challenges such as the necessity for rapid response, unpredicted time frames, risk of transient disruptions, and an extensive follow-up workload.

**CPU Elastic Scaling**




The CPU Elastic Scaling feature, underpinned by the advantages of the cloud environment, allows for the dynamic assignment of CPU resources. As the database traffic increases or CPU resource usage escalates, it can automatically allocate more CPU resources, then scale in after peak periods. Users can choose whether to enable the CPU Elastic Scaling feature via the console, and dynamically configure the database's CPU resources based on the business needs and traffic volume. Consequently, elastic scaling can be achieved to handle peak pressure, ensuring superior performance, high availability, and stability of the database instance.

### Feature Type

The CPU Elastic Scaling feature can be categorized into automatic scaling and manual scaling. If automatic scaling is enabled, you need to set the thresholds for both scale-out and scale-in. For manual scaling, you need to set the number of CPU cores to be added, and manual scaling doesn't automatically scale in - it must be manually stopped.

#### Automatic scale-out

#### Auto scale-out

Upon enabling **CPU Elastic Scaling** -> **Auto Scaling**, when the average CPU utilization of the database instance within the observation window reaches the set threshold, the system will double the number of CPU cores on the basis of the original computing specification. For instance, if the original computing specification has 4 cores, it will be increased to 8 cores. Concurrently, the IOPS of the instance will also rise. For each additional core, the IOPS increases by 1,000. If the user's Tencent Cloud account balance is insufficient or there are not enough CPU resources on the host (which rarely occurs), the scaling will not proceed, and a scaling failure event will be generated later. **Note :** 

Only supports doubling the number of CPU cores based on the original computing specification, and no further increase is supported. For instance, after the original CPU cores are doubled to 8, it cannot be increased to 16. **Auto Scale-in** 



Upon enabling **CPU Elastic Scaling** -> **Auto Scaling**, during the set observation window for auto scale-in, if the CPU utilization is below the set threshold, the system will automatically reduce both the CPU cores and IOPS to the original computing specifications.

#### Note:

After enabling **CPU Elastic Scaling** -> **Auto Scaling**, the system will monitor the database instance based on the most recent configuration parameters. When the instance meets the conditions for automatic scale-in, the system will automatically scale in the instance.

For details about the operations of enabling auto-scaling, disabling auto-scaling, and setting the scale-in and scale-out thresholds, see Setting Auto Scaling.

#### Manual scaling

#### Manual Scale-out

When **CPU Elastic Scaling** -> **Manual Scaling** is enabled, the instance operates according to the number of CPU cores after scale-out.

#### Manual Scale-in

Upon enabling **CPU Elastic Scaling** -> **Manual Scaling**, when business no longer requires additional CPU resources, users must manually disable CPU Elastic Scaling, otherwise charges will continue. For details about the operations of enabling and disabling manual scaling, see <u>Setting Manual Scaling</u>.

#### Comparison between Automatic and Manual Scaling

Feature Type	Screenshot After Successful Feature Enabling	CPU Upper Limit	Scale-in Process	Event Alarm	Billing
Auto Scaling	Hadrana Configuration Agen Configuration Exails: CPU Equivation Manual expansion, vite 7014 statist superstant introduct Manual expansion, vite 7014 statist superstant introduct	Twice the original number of CPU cores	Automatically scales in according to the set scale-in threshold.	Support	Billing begins only after auto scale-out is successfully triggered, and ceases after the scale-in.
Manual Scaling	Interna Contigonition General Ellivorational Mol ANN Brance, Alex Contigonation Electric CPU Experiment ① Manual expension, with estimationes added. Molety Class	Supports 1-core unit scale-out, up to twice the original number of	Unable to automatically scale in, manual stop of scaling is required. After manual scaling	Support	Billing will continue from the enabling of manual scaling until

	CPU cores.	is stopped, the system immediately scales in to the original CPU	it is disabled.
		core count.	

### Event Alarm Explanation

For both automatic and manual scaling out, relevant event alarms can be set and configured for notifications. The following are the event metrics and descriptions related to the CPU Elastic Scaling feature.

Scaling Type	Event Metrics Description	
	CPUExpansion	Automatic CPU scale-out is successful.
Auto Scaling	CPUExpansionFailed	Automatic CPU scale-out failed.
	CPUConstraction	Automatic CPU scale-in is successful.
	CPUExpansion	Manual CPU scale-out is successful.
Manual Scaling	CPUExpansionFailed	Manual CPU scale-out failed.
	CPUConstraction	Manual CPU scale-out is stopped.

There are two causes for automatic and manual scaling failures: insufficient balance in the Tencent Cloud account of the user or inadequate CPU resources on the host (rarely occurs). If insufficient account balance leads to a scaling failure, users can recharge their account and retry. For other causes, please submit a ticket for assistance. For details about the steps to configure event alarms, please refer to Activating Event Alarms.

### Feature Impact Overview

The CPU Elastic Scaling feature is concurrently deployed on both primary and secondary nodes on two-node and three-node instances. Therefore, if an HA switch occurs after the CPU Elastic Scaling feature is enabled, the switched-to node will also reach the scaled-out CPU specifications. After automatic scale-in or stop of manual scaling, the system will restore to the original CPU specifications.

The feature should be separately enabled for read-only instances and disaster recovery instances, which will not simultaneously scale out. In other words, the read-only instances and disaster recovery instances attached to the primary instance will not simultaneously expand their CPU resources even after the CPU Elastic Scaling feature of the primary instance is triggered. The CPU Elastic Scaling feature must be enabled for these two types of instances separately.

### **Related Operations**

Configuring Auto Scaling Configuring Manual Scaling Enabling Event Alarms

## Configuring auto scale-out

Last updated : 2024-12-26 11:57:40

This section describes how to enable and disable automatic scaling.

### Prerequisites

The instance architecture uses MySQL general two-node/three-node or single-node local disk (read-only instance).

The instance specification is equal to or less than 32 cores before the CPU Elastic Scaling feature is enabled.

Your Tencent Cloud account has sufficient balance.

#### Note:

For information on billing of automatic scaling, please see Billing Overview.

### Enabling Automatic Scaling

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operations column to go to the instance details page.

2. On Instance Details > CPU Auto Scaling, click Enable.



3. In the **CPU Auto Scaling** window, complete the following settings, and after confirming the scaling fee, click **Scale Now**.

(i) Once the feature is of when the flexible ex	enabled, yo pansions tri	u will only be char ggered. <mark>Learn Mo</mark>	ged for the expar	nded CPU co	res
Elastic CPU Expansion Type	Automa	atic Expansion	Manual Expa	nsion	
Automatic Expansion (	)	Autom	atic Reductio	n	
Automatic CPU Elastic Expansion	70%	▼ Elastic Reducti	CPU on Threshold	30%	•
Dbservation Duration	1min	Observ: ▼	ation Duration	5min	•
		USD	'hour		

Parameter	Description
CPU Auto Scaling Type	Select Automatic Scaling.
CPU Auto Scaling Threshold	Configure the threshold for automatic elastic scale-out triggered by average CPU utilization. Available options are 40%, 50%, 60%, 70%, 80%, and 90%.
Observation Period	Set the observation period to 1 minute, 3 minutes, 5 minutes, 10 minutes, 15 minutes, or 30 minutes. Within this specified time period, the system will monitor whether the average CPU utilization of the instance reaches the set scaling threshold. If it does, automatic elastic scaling will be triggered. <b>Note:</b> To mitigate the impact of continuous intermittent peaks, there is a minimum protection period of 10 minutes during automatic scaling. This means, if the automatic scaling lasts more than 10 minutes, the costs are calculated based on the actual scaling duration. If the automatic scaling lasts for no more than 10 minutes, the costs are calculated based on a 10-minute duration.
CPU Automatic Elastic Scale- In Threshold	Configure the threshold for automatic elastic scale-in triggered by average CPU utilization. Available options are 30%, 20%, and 10%.
Observation Period	Set the observation period to 5 minutes, 10 minutes, 15 minutes, or 30 minutes. This specifies an interval during which the system monitors the average CPU utilization of the instance and



determines whether it reaches the set scale-in threshold. If the threshold is reached, automatic elastic scale-in is triggered.

4. After automatic scale-out is triggered successfully, the interface will display as the following:

Instance Configuration	General-2core4000MB MEM,50GB storage, Adjust Configurations
Elastic CPU Expansion 🛈	Manual expansion, with 70% elastic expansion threshold Modify Close

#### Note :

After automatic scaling is enabled, if you need to modify the automatic scaling policy, you can navigate to **Instance Details** -> **CPU Auto Scaling**, and click **Modify** for reconfiguration.

### **Disabling Automatic Scaling**

#### Note:

The billing for automatic scaling starts from the successful triggering of the scale-out operation to the moment of scale-in. If there is no scale-out operation, no extra fees will be charged.

After automatic scaling is disabled, all expanded CPU capacity will immediately scale in to the original number of CPU cores, and no further scale-out operations will be triggered based on the scale-out threshold.

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operations column to go to the instance details page.

2. Navigate to Instance Details > CPU Auto Scaling and then click Disable.

3. In the pop-up window End CPU Scaling, click OK.



# How to Determine If Auto-Scaling Is Effective After the Auto-Scaling Feature Is Enabled for CPU Elastic Scaling?

When the average CPU utilization of an instance reaches the set auto-scaling threshold, the instance will automatically scale out. You can check the effectiveness on the **Instance Details** page after instance configuration. If the **actual core count** is additionally displayed after the CPU core count , it indicates that auto-scaling is effective. For example,

If an instance is originally configured with 4 cores and 8000MB, and the configuration displays 4 cores and 8000MB, it means auto-scaling has not been triggered. If the configuration shows 4(8) cores and 8000MB, it indicates that auto-scaling to 8 cores is currently reached.

# **Establishing Manual Scaling**

Last updated : 2023-12-27 11:02:38

This section describes how to enable and disable manual scaling.

### Prerequisites

The instance architecture uses MySQL general two-node/three-node or single-node local disk (read-only instance).

The instance specification is equal to or less than 32 cores before the CPU Elastic Scaling feature is enabled.

Your Tencent Cloud account has sufficient balance.

#### Note:

For details about billing of manual scaling, please see Billing Overview.

### **Enabling Manual Scaling**

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operations column to go to the instance details page.

2. On Instance Details -> CPU Auto Scaling, click Enable.



3. In the CPU Auto Scaling window, complete the settings, confirm the scaling cost, and click **Scale Now**.





4. After manual scaling is enabled, the instance's CPU will immediately scale out according to the number of additional CPU cores. You can view the current instance's CPU on the Instance Detail -> CPU Elastic Scaling page. Before Manual Scaling:

Instance Configuration	General2(4)core4000MB MEM,50GB storage, Adjust Configurations
Elastic CPU Expansion	(i) Enable
anual Scaling:	

Elastic CPU Expansion (i) Manual expansion, with additional cores added. Modify Clos

Note :

After manual scaling is enabled, if you need to modify the number of additional CPU cores, you can access **Instance Detail** -> **CPU Elastic Scaling** and click **Edit** for reconfiguration.

### **Disabling Manual Scaling**

#### Note:

If manual scaling is not disabled, billing will continue.

After manual scaling is disabled, the expanded CPU will promptly restore to its original core count.

1. Log in to the TencentDB for MySQL console. In the instance list, click an **instance ID** or **Manage** in the operations column to go to the instance details page.

2. Navigate to Instance Details > CPU Elastically Scalable and click Disable.

3. In the pop-up window End CPU Scaling, click OK.

### **Disable CPU Expansion**

When you disable CPU expansion, the number of CPU cores will be reduced from 4 to 2 for the current instance.



# **Activating Event Alarm**

Last updated : 2023-12-27 10:27:01

MySQL supports the Event Alarm feature. When a user sets an event alarm related to CPU elastic scaling and configures the alarm notification, it will be triggered if detected: Auto-scaling (triggered, failed, scale-in) Manual scaling (enabled, failed, disabled) Event alarms will notify users via various available methods. This section describes the process of setting event alarms related to CPU elastic scaling.

### **Operation Scenarios**

Users can implement event alarms to trigger alarms and dispatch corresponding notifications when the relevant status changes (CPU performance scaling, CPU performance scaling failure, CPU performance scale-in) after the CPU Elastic Scaling feature is enabled in MySQL.

### The steps are as follows:

#### **Creating an Alarm Policy**

1. Log in to the TCOP Console,, then navigate to the Alarm Configuration > Alarm Policy > Policy Management page by clicking the left sidebar.

2. In the Alarm Policy list, click **New Policy**.

Alarm Records	Policy Management	Basic Configuration			
() If you have an	ly questions or suggestions, s	can QR code to join our community on WeChat or WeCom.			
Create Policy	Delete More 💌		Advanced Filter	Separate keywords with * *; press Enter to separate filter tags	Q¢

3. On the policy creation page, configure basic information, alarm rules, and alarm notifications.

**Policy Type**: Categorized as Primary Monitoring and Secondary Monitoring, each corresponds to different types of instances.

**Alarm Object**: You can locate the object instance to be associated by selecting the region where the object is located or searching for the instance ID.

**Trigger Conditions**: Locate the event alarm, click '**Add Event**'and add the following alarm events: CPU Performance Scaling, CPU Performance Failure, and CPU Performance Scale-in.

Metric Alarm	Event Alarm
AuditLowRisk	<b>▼</b> <u><u></u></u>
AuditMediumRisk	<b>▼</b> <u></u>
AuditHighRisk	<b>▼</b> <u>□</u>

#### Add Event

Alarm Notification: You can select a preset or customized notification template. Each alarm policy can be bound to three notification templates at most. For instructions on how to customize a notification template, see Creating Notification Template.

Select the preset template.

u have selected 1 notification template, and 2 more	can be selected.	
Search for notification template		Q
Notification Template Name	Included Operations	
✓ F	Recipient: 1	
ble	Recipient: 1	
xi	Recipient: 1	
Total items: 3	20 💌 / page 🛛 🖂 🚽 1 🛛 / 1 pag	e 🕨
	OK Cancel	

Create a new template.

Create Notif	ication Template		
Basic Info			
Template Name	Up to 60 characters	5	
Notification Type	✓ Alarm Trigger	✓ Alarm Recovery	
Notification Language	English		▼
Tag	Tag Key	▼ Tag Value	v
	+ Add () Paste		

User Notification	You can add a user only for receiving messages.						
	Recipient Object	User 🔻					
	Notification Cycle	✔ Mon ✔ Tue ✔ Wed ✔ Thu 🖌 Fri ✔ Sat ✔ Sun					
	Notification Period	00:00:00 ~ 23:59:59					
	Receiving Channel	Comparing the second se					
		Add User Notification					
API Callback		Add API Callback					
	i It supp	ports pushing to the WeCom group robot Try Now 🛂					
Ship to CLS	Enable (1)						
	Please select	ct a regior 🔻 Select a logset 🔹 Select a log topic 🔹 🗘 Create Log Topic 🗳					



#### **Associating Alarm Objects**

After creating an alarm policy, you can also associate it with other alarm objects (instances requiring CPU

performance scaling-related alarms). An alarm will be triggered when the object meets the alarm trigger condition.

1. In the alarm policy list, click the **alarm policy name** to enter the alarm policy management page.

2. On the alarm policy management page, click Add Object in the Alarm Object column.

3. In the pop-up dialog box, select the alarm objects to be associated with, and click **OK**. The alarm objects will be associated.

# **Performance Manifestation**

Last updated : 2023-12-27 10:27:56

This section describes how to use the performance testing tool SysBench to compare the performance after the CPU Elastic Scaling is enabled.

### SysBench Tool Introduction

SysBench is a modular, cross-platform, and multi-threaded performance measurement tool that evaluates primary system parameters concerning its operational performance under heavy loads. The tool circumvents the need for complex database benchmark configurations, enabling swift understanding of database performance, even without the necessity of installing a database.

### **Testing Environment**

#### Note:

The following environment is only used in the example. **Environment Utilized for MySQL Performance Testing:** Region/AZ: Beijing - Beijing Zone 7 Client: S5.8XLARGE64 (32-core 64 GB Standard S5) Client OS: CentOS 8.2 64-bit Network: Both the CVM and MySQL instances are under the same VPC subnet. **Tested MySQL Instance:** Three groups of instances: Beijing Zone 7, General 2-core 4 GB, two-node MySQL Group A: NO-CPU-Expand, the CPU Elastic Scaling feature is not enabled. Group B: AUTO-CPU-Expand, the Automatic CPU Elastic Scaling feature is enabled. Group B: Group C: MANUAL-CPU-Expand, the Manual CPU Elastic Scaling feature is enabled. Group B:

Set console parameter 'max\_prepared\_stmt\_count' to its highest value to cope with high concurrency.

### **Testing Process**

Step 1:

Create a 'sbtest' database in the system, enable the Automatic CPU Elastic Scaling feature for group B, and enable the Manual CPU Elastic Scaling feature for group C.

Step 2:

Utilize Sysbench to prepare data, including 10 tables, with each table containing a data volume of 100,000.

Step 3:

Perform a Sysbench read/write test on 10 tables, with each table containing a data volume of 100,000.

Step 4:

Proceed with Sysbench to clear data and disable the Manual CPU Elastic Scaling feature for group C.

### **Testing Description**

The Test can be affected by factors such as testing time, duration, machine used, etc. We suggest performing multiple tests.

### **Executed Commands:**

Following the above steps, conduct the tests respectively.

Note:

Please replace the 'XXX' in the following command with the internal network address, port number, username, user password, and database name of the MySQL instance. The specific parameters are explained as follows:

host: Private network address of the tested instance

port: Port number

user: Username

password: Password corresponding to the username

time: Time

threads: Number of concurrent threads

1. Prepare data: 100,000 data volumes in each table, with 10 tables in total.

```
sysbench --db-driver=mysql --mysql-host=xxx --mysql-port=xxxx --mysql-user=xxx
--mysql-password='XXXXXXXXX' --mysql-db=sbtest --table_size=100000 --tables=10
--events=0 --time=xxxx --threads={xxxx} oltp_read_write prepare
```

2. Execute high concurrency operations on the 100,000 data volume in each table, with 10 tables in total.

```
sysbench --db-driver=mysql --mysql-host=xxx --mysql-port=xxxx --mysql-user=xxx
--mysql-password='XXXXXXXXX' --mysql-db=sbtest --table_size=100000 --tables=10
--events=0 --time=xxx --threads={xxx} --percentile=95 --report-interval=1
oltp_read_write run
```

3. Clear the data.

```
sysbench --db-driver=mysql --mysql-host=xxx --mysql-port=xxxx --mysql-user=xxx
--mysql-password='XXXXXXXXX' --mysql-db=sbtest --table_size=100000 --tables=10
--events=0 --time=xxx --threads={xxx} --percentile=95 oltp_read_write cleanup
```

### **Test Results**

#### Note:

As illustrated below, in the comparative experiment results of the CPU Elastic Scaling feature of three groups of instances, we can observe:

Phase One: At 21:03, data injection was initiated. The CPU usage of both NO-CPU-Expand and AUTO-CPU-Expand exceeded 70%, but the duration did not exceed 1 minute. Therefore, the AUTO-CPU-Expand group did not trigger automatic scaling, while the MANUAL-CPU-Expand group scaled out immediately upon being operated and ran with 4 cores, with the CPU usage less than 50%.

Phase Two: At 21:07, high concurrency testing was initiated. The CPU usage of the NO-CPU-Expand group was remained over 70%, the CPU usage of the AUTO-CPU-Expand group also exceeded 70%, and it triggered automatic scaling when the CPU usage remained high for 1 minute. The scaling succeeded within seconds and reduced the usage to less than 50%. The CPU usage of the MANUAL-CPU-Expand group remained below 50%.



#### Note:

In addition to significant changes in CPU utilization, the performance of instances that have enabled the CPU Elastic Scaling feature also changes significantly. As demonstrated in the graph below, the TPS and QPS of instances with the feature enabled are higher than those for which the feature is not enabled. It should be noted, the performance in the test is affected by factors such as the testing period, duration, machine, etc. Thus, it is recommended to perform multiple tests and compare the results.



# **Billing Overview**

Last updated : 2025-03-21 14:09:31

This section introduces the billing instructions related to the CPU Elastic Scaling feature.

### Billing for Automatic Scaling

#### **Fee Calculation**

The CPU Elastic Scaling feature uses pay-as-you-go billing mode, where fees are charged on a minute basis and deducted once per hour.

Billing Formula = (Unit price x Number of added CPU cores) x Scaling duration (minutes) / 60.

#### Automatic Scaling Billing Example

1. Suppose a two-node instance in Guangzhou has a CPU with 4 cores. After Automatic Scaling is triggered, it is expanded to 8 cores. The scaling duration is 1 hour. Given that the unit price in Guangzhou is 0.08 USD/core/hour, the total cost would be calculated as follows: 0.08 (Unit price) x 4 (additional CPU cores) x 1 (duration of expansion) = 0.32 USD.

2. Suppose a two-node instance in Guangzhou has a CPU with 2 cores. After Automatic Scaling is triggered, it is expanded to 4 cores, with a scaling duration of 30 minutes. Considering that the unit price in Guangzhou is 0.08 USD/core/hour, the total cost would be calculated as follows: 0.08 (Unit price) x 2 (additional CPU cores) x 30/60 (scaling duration) = 0.08 USD.

#### Note:

**Upon enabling CPU Auto Scaling**: you will be billed according to the number of CPU cores added when the elastic scaling is successfully triggered. There will be no additional charges when no scaling is performed. To mitigate the impact of persistent intermittent peaks, there is a minimum protection time of 10 minutes during auto-scaling. That is, if the auto-scaling lasts for more than 10 minutes, the cost will be calculated based on the actual scaling time. If the auto-scaling lasts no more than 10 minutes, the cost will be calculated with the duration of 10 minutes.

### Manual Scaling Billing Description

#### **Fee Calculation**

The Manual CPU Elastic Scaling feature adopts the pay-as-you-go billing method, where fees charged on a minute basis and deducted once per hour.

Billing Formula = (Unit price x Number of added CPU cores) x Scaling duration (minutes)/60.

#### Manual Scaling Billing Example

In Beijing, a CPU elastic scaling is manually enabled for a two-node instance with a CPU with 2 cores on 10 AM, with 2 cores added. The manual CPU elastic scaling process is then disabled at 5 PM, which means the scaling lasts for 7 hours. Given that the unit price in Beijing is 0.08 USD/core/hour, the total fee would be: 0.08 (unit price) x 2 (additional CPU cores) x 7 (scaling duration) = 1.12 USD.

#### Note:

Once **Manual CPU Elastic Scaling** is started, the number of CPU cores for the instance will immediately increase based on your configurations and fees will be charged synchronously until you disable manual scaling. Therefore, the system will continuously charge since manual scaling is enabled and until it is disabled. For instructions on how to disable manual scaling, please see Disabling Manual Scaling for operation to turn off manual scaling.

### Unit Price per Core

#### Single-node Local Disk (Read-only Instance)

#### Two-node Instance

Three-node Instance

Region	Unit Price (USD/Core/Hour)
Chengdu, Chongqing	0.03
Guangzhou, Shanghai, Beijing, Nanjing	0.04
Hong Kong (China), Tokyo, Seoul, Bangkok	0.0495
Frankfurt, Sao Paulo	0.0365
Singapore, Jakarta, Silicon Valley, Virginia	0.061

Region	Unit Price (USD/Core/Hour)
Chengdu, Chongqing	0.06
Guangzhou, Shanghai, Beijing, Nanjing	0.08
Hong Kong (China), Tokyo, Seoul, Bangkok	0.099
Frankfurt, Sao Paulo	0.073



Singapore, Jakarta, Silicon Valley,	0.122
Virginia	

Region	Unit Price (USD/Core/Hour)
Chengdu, Chongqing	0.09
Guangzhou, Shanghai, Beijing, Nanjing	0.12
Hong Kong (China), Tokyo, Seoul, Bangkok	0.1485
Frankfurt, Sao Paulo	0.1095
Singapore, Jakarta, Silicon Valley, Virginia	0.183

## FAQs

Last updated : 2024-12-30 18:00:35

This document gives FAQs related to the CPU Elastic Scaling feature.

#### Why can't I find the CPU Elastic Scaling feature in the console?

The CPU Elastic Scaling feature is only available for MySQL Universal Edition with two or three-node instances, or single-node instances with local disks (read-only instance). You can verify whether an instance is a single-node (cloud disk version), dedicated instance, or cluster edition (cloud disk version) instance. The aforementioned types of instances currently do not support the CPU Elastic Scaling feature.

# During stress testing, why does the monitoring show that the CPU utilization remains at 100% even after the feature is enabled?

CPU utilization is calculated based on the actual number of CPU cores. If the service load fully utilizes all CPU performance, performance-related metrics such as TPS, QPS, and response time of the database system should be given attention.

#### Why does the CPU Elastic Scaling feature fail to be enabled?

There are two possible reasons for the failure to enable the CPU elastic scaling feature:

Your Tencent Cloud account does not have a sufficient balance.

The CPU resources on the host is insufficient (rarely occurs).

Ensure sufficient balance in your Tencent Cloud account or please submit a ticket for feedback.

#### Why am I not notified of events when the CPU is scaled-out or scaled-in?

You can troubleshoot and handle the issue as follows.

Check whether the alarm for CPU elastic scaling events has been enabled. If not, enable it by referring to Enabling Event Alarms.

Check the notification recipient list; it is possible that no relevant users have been added. You can add the notification recipient by referring to Adding Message Recipients.

#### What should I do if a HA switch occurs after the CPU Elastic Scaling feature is enabled?

The CPU Elastic Scaling feature will take effect simultaneously on both the primary and secondary nodes on two-node and three-node instances. If a HA switch occurs on a two-node or three-node instance after the CPU Elastic Scaling feature takes effect, the switched node will also have the expanded CPU specification. When automatic scaling is disabled or manual scaling is stopped, it will automatically restore to the original CPU specification. Please note that for read-only instances and disaster recovery instances, the feature need to be enabled separately. That means when the CPU Elastic Scaling feature is enabled and takes effect on the primary instance, the read-only instances and

disaster recovery instances attached to this primary instance will not expand their CPU specification simultaneously. You need to enable the CPU elastic Scaling feature for them separately.

#### How can I query the progress and details related to CPU elastic scaling tasks?

You can query the progress and details related to CPU elastic scaling tasks through the task list, with the following steps:

1. Log in to the TencentDB for MySQL console.

2. Choose Task List on the left sidebar, select the region, click **Task Type**, choose **Configure Elastic Expansion Strategy**, and then click **OK**.

3. In the filtered task list, locate the target task to inquire about its progress. Click **Task Details** in the operation column to view detailed task information.

All	Today Last 24	ours Last 7 days	Last 30 days	Select date			Ins	istance ID 🔻	Please enter Instance IE	Q Ç
Task ID	Task Typ	T Ins	tance ID	Task Progress	Task Status <b>T</b>	Task Start Time	Task End Time		Operation	
3895669	Configur Expansio	Elastic Policy(Enable)	o-8ycmh6zo	100%	Successful	2023-12-13 15:46:45	2023-12-13 15:46:5	52	Task Details	

#### Why am I continuously billed after the CPU Elastic Scaling feature is enabled?

If you choose manual scaling when enabling the feature, the scaling will continue until you manually stop it. Therefore, after enabling manual scaling, please be sure to manually stop it based on your business requirements.

It is recommended to choose automatic scaling. After the configuration, the CPU specifications will be automatically and elastically scaled based on the user-configured scale-out and scale-in thresholds. Also, you will only be billed for the number of cores scaled out when the scaling is successfully triggered. No fees are charged when there is no scaling.

# Within one hour after automatic scaling is enabled, it automatically scales in only 6 minutes after scaling was triggered. Why am I being charged for 10 minutes?

To mitigate the impact of continuous intermittent peaks, there is a minimum protection period of 10 minutes during automatic scaling. This means, if the automatic scaling lasts more than 10 minutes, the costs are calculated based on the actual scaling duration. If the automatic scaling lasts for no more than 10 minutes, the costs are calculated based on a 10-minute duration.

# Read-Only/Disaster Recovery Instances Creating Read-Only Instance

Last updated : 2024-07-30 14:59:21

### Overview

TencentDB for MySQL allows you to create one or more read-only instances, which are suitable for read/write separation and one-primary-multiple-replica application scenarios and capable of greatly enhancing the read load capacity of your database.

Database proxy is supported. After creating a read-only instance, you can purchase the database proxy service to enable the read/write separation feature. Then, you can configure the database proxy address in your application so as to automatically forward write requests to the primary instances and read requests to the read-only instances. **Note:** 

For read-only instance pricing, see Product Pricing.

You can now configure a custom and exclusive private network address (IP and port) for a read-only instance on the instance details page.

GTID	Enabled
Character Set	UTF8
Private Network Address (i)	Enable

The deployment architecture of a Cluster Edition instance consists of 1 read-write node and up to 5 read-only nodes. If you have purchased a Cluster Edition instance, see Adding a Node for an Instance and Deleting a Node for an Instance for operations on adding or deleting read-only nodes. The Cluster Edition instance also supports adding 1 to 15 independent read-only instances. You can refer to the method described in this document for creation, and should note that the independent read-only instances cannot be switched to read-write nodes.

#### Concepts

Read-only group: It consists of one or more load balancing-enabled read-only instances. If there are multiple read-only instances in one read-only group, read request volume can be evenly distributed among the instances. Read-only groups provide IPs and ports for access to databases.

Read-only instance: A single-node (with no replica) instance that supports read requests. It cannot exist independently; instead, it must be in a read-only group.

#### Architecture

The MySQL primary-replica binlog sync feature is adopted for read-only instances, which can sync the changes in the primary instances (source database) to all read-only instances. Given the single-node architecture (without a replica) of read-only instances, repeated attempts to restore a failing read-only instance will be made. Therefore, we recommend that you choose an RO group over a read-only instance for higher availability.

#### Note:

If there is only one read-only instance in the read-only group, there will be risks with single points of failure, and this group will not be included in the overall availability calculation of the TencentDB for MySQL service. A single read-only instance does not provide SLA guarantee. We recommend you purchase at least two read-only instances to ensure the availability of the read-only group.



### Feature limits

Read-only instances cannot be created for single-node instances of the cloud disk edition.

Read-only instances can be purchased only for **two-node or three-node or Cluster Edition primary instances on MySQL 5.6 or later with the InnoDB engine at a specification of 1 GB memory and 50 GB disk capacity or above**. If your primary instance is below this specification, upgrade it first.

The minimum specification of a read-only instance is 1 GB memory and 50 GB disk capacity, which must be equal to or greater than the storage capacity purchased for the primary instance.

A two-node or three-node primary instance supports creating up to 5 read-only instances, and a Cluster Edition instance supports creating up to 15 read-only instances.

A primary instance can create up to 5 read-only instances.

Backup and rollback features are not supported.

Data cannot be migrated to read-only instances.

Database creation/deletion is not supported and neither is phpMyAdmin (PMA).

Operations including account creation/deletion/authorization and account name/password modification are not supported.

### Use limits

There is no need to maintain accounts or databases for read-only instances, which are synchronized with those of the primary instance.

If the MySQL version is 5.6 but GTID is not enabled, you need to enable GTID in the console first before creating a read-only instance.

The operation takes a long time, and the instance will be disconnected for several seconds. We recommend that you do so during off-peak hours and add a reconnection mechanism in the programs that access the database.

Read-only instances only support the InnoDB engine.

Data inconsistency between multiple read-only instances may occur due to the delay in data sync between the readonly instances and the primary instance. You can check the delay in the console.

The specification of a read-only instance can be different from that of the primary instance, which makes it easier for you to upgrade the read-only instance based on the load. We recommend that you keep the same specifications of read-only instances in one RO group.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to access the instance details page.

2. Click **Add Read-Only Instance** in the **Instance Architecture Diagram** section on the **Instance Details** tab (add an independent read-only instance for the Cluster Edition) or click **Create** on the **Read-Only Instance** tab.

3. On the displayed purchase page, specify the following read-only instance configurations, confirm that everything is correct, and click **Buy Now**.

#### Note:

If you need to unify the expiration time of the primary and read-only instances, you can set the collective expiration date in the renewal management console as instructed in Renewal Management.

The newly added independent read-only instance for the Cluster Edition does not involve the operations related to specifying RO groups.

Te	encentDB MyS	QL Read-only Instance								
	Master Instance Instance Name Network Architecture Project Availability Zone Version	Instance ID         Argion       East China (Shanghai)         High-Availability Edition       Instance specifications         Shanghai Zone 4       MySQL5.6								
	Specify RO Group	Assigned by system System automatically assigns RO group	Learn about RO Group <sup>12</sup> os which are not enabled latency elimina	tion. Please retain at least one instance.						
	Region	East China (Shanghai)								
	Version	MySQL5.6								
	Architecture	Single-node high IO edition Although the single-node architecture is cost-effective, there is a single point of failure for a single read-only instance. It is strongly recommended to purchase at least two read-only instances in the service RO group that requires availability.								
	Availability Zone	Shanghai Zone 4								
	Instance specifications	4core8000MB	v							
	Hard Disk			500 I GR (Increment: 5GR)						
	Fees Configu Buy	Iration cost (Tiered Pricing ②)	Traffic Fee 0.00 USD/GB ⑦							
Paramete	er	Description								
Specify F	O Group	You can use the RO g select an existing one Assigned by system: be assigned to an ind assigned by the syste Create RO group: Cre time, all of them will b allocated by the syste Existing RO group: S	group automatically assig e. If multiple instances are p ependent RO group, and em by default. eate an RO group. If multi e assigned to this new R0 em automatically by defau pecify an existing RO gro	ned by the system, create one, or burchased at a time, each of them will their weights will be automatically ple instances are purchased at a O group, and their weights will be lt. up. If multiple read-only instances are						

purchased at a time, all of them will be assigned to the RO group. Their weights will



	be allocated as configured in the RO group. If assignment by the system is set for the RO group, the instances will be added to the group automatically according to the purchased specifications. If custom allocation is set, their weights will be zero by default. As the same private IP is shared within an RO group, if a VPC is used, the same security group settings will be shared. If an RO group is specified, it is not possible to customize any security group when instances are purchased.
RO Group Name	You can set the name of the new RO group when creating it. The name can contain up to 60 letters, digits, hyphens, underscores, and dots.
Remove Delayed RO Instances	Select whether to enable the removal policy. If this option is enabled, you need to set **Delay Threshold** and **Least RO Instances**, and the weight of the removed instances is automatically set to 0. If a read-only instance is removed when its delay exceeds the threshold, it will become inactive, its weight will be automatically set to 0, and alarms will be sent to corresponding recipients. The instance will be put back into the RO group when its delay falls below the threshold. For more information on how to configure the read-only instance removal alarm and recipients, see Alarm Policies (TCOP). No matter whether delayed read-only instance removal is enabled, a read-only instance that is removed due to instance failure will rejoin the RO group when it is repaired.
Billing Mode	Monthly subscription and pay-as-you-go billing are supported.
AZ	When creating a new RO group, you can select the same AZ as the primary instance or a different AZ. There are no substantial differences between different AZs. If you choose a different AZ, the RO group will reside across AZs and can improve data disaster recovery, but there will be a few milliseconds of network latency.

4. After the purchase is completed, you will be redirected to the instance list. When the status of the instance changes to **Running**, it can be used normally.

### FAQs

#### What are the rules for removing read-only instances?

After **Remove Delayed RO Instances** is enabled, the RO group will determine the removal rule based on the delay threshold and "Least RO Instances" value. If the delay of a read-only instance reaches the threshold, it will be removed and made inactive, with its weight set to 0 and an alarm sent to corresponding recipients. When its delay falls below the threshold, it will be returned to the RO group.

Delay Threshold: Set a delay threshold for a read-only instance. When the threshold is exceeded, the instance will be removed from the RO group.

Least RO Instances: This is the minimum number of instances that should be retained in the read-only group. When there are fewer instances in the read-only group, even if an instance exceeds the delay threshold, it will not be removed.

#### If read-only instances are terminated or returned, how will the primary instance be affected?

The termination and return of read-only instances has no impact on the primary instance.

#### Why can't I select a specific AZ when creating a read-only instance?

If you can't select an AZ, it means that there are no resources in this AZ. You can choose another AZ as displayed on the actual purchase page, which will not affect your use of the read-only instance.

#### Can I select an AZ different from that of the primary instance when creating a read-only instance?

Yes. When creating a read-only instance, you can choose to create a new RO group for it and select an AZ different from that of the primary instance. However, if you select an existing RO group, the AZ of the new read-only instance can only be the same as that of the selected existing RO group, which may not necessarily be the same as that of the primary instance.

# Managing the RO Group of Read-Only Instance

Last updated : 2024-07-22 14:52:52

### Overview

TencentDB for MySQL allows you to create one or more read-only instances to form an RO group, which is suitable for read/write separation and one-source-multiple-replica application scenarios and capable of greatly enhancing the read load capacity of your database.

An RO group is a set of read-only instances sharing the same address. You can set their weights to balance the traffic load, set the policy of removing delayed read-only instances, and perform other configurations. You can deploy an RO group as needed and send the corresponding read requests to read-only instances according to certain rules. In addition, you can implement disaster recovery by configuring multiple read-only instances in the same RO group.

### Prerequisites

A source instance must be created first before a read-only instance can be created. For more information, see Creating MySQL Instance.

### Creating an RO group

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.

2. Click Add Read-Only Instance in the Instance Architecture Diagram section on the Instance Details tab or click Create on the Read-Only Instance tab.

Instance Details

Instance Monitoring

Database Managem

Instance Details	Instance Monitoring Manage Database	Security Group	Backup and Restore	Operation Logs	Read-only Instance	Connection Check	
Basic Info			Instance Architectu	ıre Diagram 🧳			
Instance Name	kaylal_cdb238160 🎤		South China (Guangzh	iou)			
Instance ID	cdb-qak9w3yf I⊡		🧿 Guangzhou Zon	ie 4	cdb-qak9w3yf kayl	al_cdb238160 ( Running ) Current Ins	stance
Status/Task	Running /		ve 6				
Region/Availability Zone	South China (Guangzhou)/ Guangzhou Zone 4				Async Del	ay 0MB Standby Database Guangzhou Zone 4	🚉 Cold Standby
Project	Default Project Switch to another project						
Placement Group	-				+ Add	I Read-Only Instance	
GTID	Enabled						
Character Set	UTF8 🎤						
- cdh225942 ~dh≘	2306 E					Log In Restart Poll Po	ck Adjust Configurations

3. On the displayed purchase page, specify the following read-only instance configurations, confirm that everything is correct, and click **Buy Now**.

Backup and Restoration

Operation Log

Read-Only Instance

Database Proxv

Data Security

Cor 
Bead-Only In

rity Group

**Specify RO Group**: Select **Create RO group**. If multiple instances are purchased at a time, all of them will be assigned to this new RO group, and their weights will be automatically assigned by the system by default. - **Assign Read Weight**: It is assigned by the system.

**RO Group Name**: The RO group name doesn't need to be unique and can contain up to 60 letters, digits, hyphens, and underscores. - **Instance Name**: It must contain less than 60 letters, digits, or symbols (- .).

**Remove Delayed RO Instances**: Select whether to enable the removal policy. If a read-only instance is removed, its weight will be automatically set to 0.

If a read-only instance is removed when its delay exceeds the threshold, it will become inactive, its weight will be automatically set to 0, and alarms will be sent to corresponding recipients. The instance will be put back into the RO group when its delay falls below the threshold. For more information on how to configure the read-only instance removal alarm and recipients, see Alarm Policies (TCOP).

No matter whether delayed read-only instance removal is enabled, a read-only instance that is removed due to instance failure will rejoin the RO group when it is repaired.

**Delay Threshold**: Set a delay threshold for the read-only instance. When the threshold is exceeded, the instance will be removed from the RO group.

**Least RO Instances**: This is the minimum number of instances that should be retained in the RO group. When there are fewer instances in the RO group, even if an instance exceeds the delay threshold, it will not be removed.

Assign Read Weight: It is assigned by the system.

Billing Mode: Monthly subscription and pay-as-you-go billing are supported.

Region: It is the same as that of the source instance by default.

Database Version: It is the same as that of the source instance by default.

**Engine**: It is the same as that of the source instance by default.

**Architecture**: Select **Single-node**. Although the single-node architecture is cost-effective, there is a single point of failure for a single read-only instance. It is recommended to purchase at least two read-only instances in the service RO group that requires availability.

#### Data Replication Mode: Async replication

**AZ**: When creating a new RO group, you can select the same AZ as the source instance or a different AZ. There are no substantial differences between different AZs. If you choose a different AZ, the RO group will reside across AZs and can improve data disaster recovery, but there will be a few milliseconds of network latency.

For other configuration items, see Creating MySQL Instance.

4. Return to the instance list. The status of the created instance is **Delivering**. If the status changes to **Running**, the read-only instance has been successfully created.

### Configuring an RO group

On the RO group configuration page, you can configure the basic information of the group such as ID, name, delayed replication, replication delay, removal policy, delay threshold, least read-only instances, and read weight. **Note:** 

Read-only instances in an RO group can use different specifications, and their read traffic weights can be set.

Read-only instances in the same RO group can have different expiration dates and billing modes.

Once enabled, delayed replication will take effect for all read-only instances in this RO group, but won't change their replication statuses.

The replication delay option will appear only after delayed replication is enabled.

1. Log in to the TencentDB for MySQL console. In the instance list, click a source instance ID to enter the instance management page.

2. On the instance management page, click the **Read-Only Instance** tab and click **Configure** in the **RO Group** column to enter the RO group configuration page.

Instance Details	Instance Mor	itoring Manage	Database Securit	y Group Backup and	I Restore	Operation Logs	Read-only instance	Connection Check	Re	ad-only Instance Documentatic
Create										
RO Group ID:cdbrg-1ldx RO Group:ro_group_244	fcpx 4146	Region:South Chin Network: Default-V	na (Guangzhou)/Guang /PC - <mark>Default-Subnet</mark> C	zhou Zone 4 Change Subnet	Private I Public IF	P: P			Cre	eate Configuration Rebak
cdbro-1ldxfcpx cdb_ro_244146	Running	Weight1	Pay as you go	Expiry time		Type: 1c	ore1000MB MEM,70GB stora	ige,	Ма	inage Upgrade Terminate/R
RO Group ID:cdbrg-ddfo RO Group:ro_group_24	122az 4145	Region:South Chin Network: Default-V	na (Guangzhou)/Guang: /PC - Default-Subnet C	zhou Zone 4 Change Subnet	Private I Public IF	P: Enable	¢		Cre	eate Configuration Rebak
cdbro-ddfd22az cdb_ro_244145	Running	Weight1	Pay as you go	Expiry time	-	Type: 1c	ore1000MB MEM,70GB stora	ige,	Ма	.nage Upgrade Terminate/R

3. On the RO group configuration page, configure the RO group information and click **OK**.

You can set delayed replication and select to replay by flashbacked position or global transaction identifier (GTID) during the delay to efficiently roll back data and fix failures.

**Replication Delay**: You can configure the time of delayed replication between a read-only instance and its source instance. The value range is 1–259200 seconds.

**Remove Delayed RO Instances**: Select whether to enable the removal policy. If a read-only instance is removed, its weight will be automatically set to 0. If a read-only instance is removed when its delay exceeds the threshold, it will become inactive, alarms will be sent to corresponding recipients. For more information on how to configure the read-only instance removal alarm and recipients, see Alarm Policies (TCOP).

**Delay Threshold**: Set a delay threshold for the read-only instance. When the threshold is exceeded, the instance will be removed from the RO group.

**Least RO Instances**: This is the minimum number of instances that should be retained in the RO group. When there are fewer instances in the RO group, even if an instance exceeds the delay threshold, it will not be removed.

**Assign Read Weight**: The RO group supports two weight assignment methods: automatic assignment by the system and custom assignment. The weight value must be an integer between 0 and 100. Below is the list of read weights automatically set for two-node and three-node TencentDB for MySQL instances by the system:

Memory Specification (MB)	1000	2000	4000	8000	12000	16000	24000	32000	48000	6400
Weight	1	1	2	2	4	4	8	8	10	12

#### Rebalance:

If rebalance is disabled, modifying weight will take effect only for new loads but will not affect the read-only instances accessed by existing persistent connections or cause a momentary disconnection from the database.

If rebalancing is enabled, all connections to the database will be temporarily disconnected, and the loads of newly added connections will be balanced based on the set weights.


RO Group ID	cdbrg-1ldxfcpx
RO Group Name	ro_group_244146
Eliminate instances with out-of-limit delay	What is elimination of instances with out-of-limit delay 🛂
Assign Read Weight	• Assigned by system • Custom How to set weight [2] Weight is automatically removed in case of instance downtime or delay exceeding limit, and is restored automatically upon instance recovery. Weight is automatically removed when the instance is released.
Load Rebalancing	If Load Rebalancing is disabled, modifying weight only takes effect for new loads and will not affect the read-only instances accessed by the original persistent connection and cause flash disconned database.
Cancel OK	

### Terminating and deleting an RO group

#### Note:

RO groups cannot be deleted manually.

An RO group will be automatically deleted when the last read-only instance in it is eliminated.

Empty RO groups cannot be retained.

1. Log in to the TencentDB for MySQL console. In the instance list, click a source instance ID to enter the instance management page.

2. On the instance management page, click the Read-Only Instance tab and click Terminate/Return or

Terminate/Return & Refund in the Operation column on the right.

Instance Details	Instance Monitorin	g Database Mana	gement Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Encryption	Connection Check	Read-only Instance
Create										
RO Group IC RO Grou	Reg	gion:East China (Shang work	hai)/Shanghai Zone 3	Private Network A Public Network A	ddress: ddress: Enable		Delayed Replication	Status: Closed		Create Config Rebalancing
	Running	Weight1 Pay a	is You Go	Expiration Time:	Type: 1core1000MI	B MEM,500GB storage,				Manage Upgra Terminate/Return

3. In the pop-up window, read and agree to the **Termination Rules** and click **Terminate Now**.

### FAQs

#### Why can't I select a specific AZ when creating a read-only instance?

If you can't select an AZ, it means that there are no resources in this AZ. You can choose another AZ as displayed on the actual purchase page, which will not affect your use of the read-only instance.

#### Can I select an AZ different from that of the source instance when creating a read-only instance?

Yes. When creating a read-only instance, you can choose to create a new RO group for it and select an AZ different from that of the source instance. However, if you select an existing RO group, the AZ of the new read-only instance can only be the same as that of the selected existing RO group, which may not necessarily be the same as that of the source instance.

# Managing Delayed Replication of Read-Only Instance

Last updated : 2025-06-10 09:49:19

This document describes how to set delayed replication for read-only instances and enable/disable replication in the TencentDB for MySQL console. You can set delayed replication (i.e., delay between a read-only instance and its source instance) and select to replay by flashbacked position or global transaction identifier (GTID) during the delay to efficiently roll back data and fix failures.

Delayed Replication: you can enable and configure the replication delay between a read-only instance and its source instance on its RO group configuration page or management page.

Enabling/Disabling Replication: you can manually enable or disable data sync between a read-only instance and its source instance.

### **Delayed Replication Description**

After delayed replication is enabled for a read-only instance, it will be removed from the RO group with its weight set to 0, and a removal alarm will be triggered. At this point, the traffic will not be forwarded to the removed instance if the RO VIP is used to access the RO group. What's more, the removed instance can only be accessed via the instance's VIP.

After the delayed replication is disabled for a read-only instance (the corresponding RO group has enabled delayed-replication-read-only-instance removal), the weight of this read-only instance in the RO group will be recovered only if the delay time of this instance is less than the delay threshold of the RO group. And a restoration alarm will be triggered at the same time.

During replay by flashbacked position, you cannot restart the instance, adjust its configuration, upgrade its version, or upgrade its kernel minor version.

### **Enabling Delayed Replication**

#### Note:

**Delayed Replication** is **disabled** for a read-only instance by default. If it is enabled, the delayed replication time will be displayed.

#### Enabling on the read-only instance's RO group configuration page

1. Log in to the TencentDB for MySQL console. In the instance list, click a source instance ID to enter the instance management page.

2. On the instance management page, select the **Read-Only Instance** tab, locate the desired RO group, and click **Configuration** to enter the RO group configuration page.

3. On the RO group configuration page, enable **Delayed Replication**, set the **Replication Delay**, and click **OK**. You can set delayed replication and select to replay by flashbacked position or global transaction identifier (GTID) during the delay to efficiently roll back data and fix failures.

**Replication Delay**: You can configure the time of delayed replication between a read-only instance and its source instance. The value range is 1–259200 seconds.

**Remove Delayed RO Instances**: This option indicates whether to enable the removal policy. If a read-only instance is removed when its delay exceeds the threshold, its weight will be set to 0 automatically, and alarms will be sent (for more information on how to configure the read-only instance removal alarm and recipients, see Alarm Policies (TCOP)).

**Delay Threshold**: This sets a delay threshold for the read-only instance. When the threshold is exceeded, the instance will be removed from the read-only group.

**Least RO Instances**: This is the minimum number of instances that should be retained in the RO group. When there are fewer instances in the RO group, even if an instance exceeds the delay threshold, it will not be removed.

**Assign Read Weight**: The RO group supports two weight assignment methods: automatic assignment by the system and custom assignment. The weight value must be an integer between 0 and 100.

#### Load Rebalancing:

Modifying weight will only affect new loads if rebalancing is disabled. The operation has no impact on read-only instances accessed by existing persistent connections and does not cause momentary database disconnection. If rebalancing is enabled, all connections to the database will be temporarily disconnected, and the loads of newly added connections will be balanced according to the set weights.

#### Enabling on the read-only instance management page

1. Log in to the TencentDB for MySQL console. In the instance list, click a read-only instance ID or **Manage** in the **Operation** column to access the read-only instance details page.

2. On the read-only instance details page, click **Enable** in **Deployment Info > Delayed Replication**.

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

#### Note:

The delay ranges from 1 to 259,200 seconds.

Read-Only instances in the same RO group share the same replication delay. If one instance is modified, the rest will be modified automatically at the same time.

### Modifying Delayed Replication

#### Modifying on the read-only instance's RO group configuration page

1. Log in to the TencentDB for MySQL console. In the instance list, click a source instance ID to enter the instance management page.

2. On the instance management page, select the **Read-Only Instance** tab, locate the desired RO group, and click **Configuration** to enter the RO group configuration page.

3. On the RO group configuration page, modify **Replication Delay** and click **OK**.

#### Modifying on the read-only instance management page

1. Log in to the TencentDB for MySQL console and click a read-only instance ID in the instance list to enter the readonly instance details page.

2. On the read-only instance details page, click **Modify** in **Deployment Info > Delayed Replication**.

3. In the pop-up window, set the delay and click OK.

### **Disabling Delayed Replication**

#### Disabling on the read-only instance's RO group configuration page

1. Log in to the TencentDB for MySQL console. In the instance list, click a source instance ID to enter the instance management page.

2. On the instance management page, select the **Read-Only Instance** tab, locate the desired RO group, and click **Configuration** to enter the RO group configuration page.

3. On the RO group configuration page, click **Delayed Replication** and click **OK**.

#### Disabling on the read-only instance management page

1. Log in to the TencentDB for MySQL console and click a read-only instance ID in the instance list to enter the readonly instance details page.

2. On the read-only instance details page, click **Disable** in **Deployment Info > Delayed Replication**.

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

#### Note:

If delayed replication is disabled, the delayed replication time will be 0 seconds, that is, real-time data sync is resumed between the read-only instance and its source instance. >

### **Enabling Data Replication**

#### Note:

The **Replication Status** of a read-only instance is **Normal** by default. If you set delayed replication and delete data accidentally during the delayed replication time period, you can restore the read-only instance to the specified position or GTID of the binlog file to implement quick data restoration.

1. Log in to the TencentDB for MySQL console and click a read-only instance ID in the instance list to enter the readonly instance details page.

2. On the read-only instance details page, click **Enable** in **Deployment Info > Replication Status**.

3. In the pop-up window, click OK.

#### Note:

Once the replication is enabled, the data sync between the read-only instance and the source instance will resume. 4. You can also select **Replay by Flashbacked Position** in **Deployment Info** > **Operation**. Then, data can be restored to the specified time point or corresponding GTID. After the restoration, replication of the read-only instance will be disabled until the start mode is switched to **Replicate all**.

Time: You can select a time point between the replication end time and the current time of the source instance.

GTID: You can select all logs after the binlog not applied. If GTID is selected, data will be replicated until the specified GTID is reached.

The length of instance server\_uuid is fixed at 36 bits, and the GTID must be in

server\_uuid:transaction\_id format.

#### Note:

If the entered binlog position has already been applied on the read-only instance or is after the source instance position, replication will fail to be enabled.

When you enable replication, if there is any breakpoint in the binlog, replication will fail to be enabled.

To avoid overusing the disk space of a delayed read-only instance when replication is disabled for it, the I/O thread of the read-only instance will be paused when its available disk space drops below 5 GB.

5. During the process of replay by flashbacked position, you can click **Replaying** after **Replication Status** to query and refresh the task details.

6. After the replication is completed, click **Enable** after **Replication Status**, and the read-only instance can continue to replicate.

### **Disabling Data Replication**

Note:

Only after delayed replication is enabled can it be disabled; otherwise, the **Disable** button is grayed out.

Once replication is disabled, I/O and SQL threads will also be ended.

1. Log in to the TencentDB for MySQL console and click a read-only instance ID in the instance list to enter the readonly instance details page.

2. On the read-only instance details page, click **Disable** in **Deployment Info > Replication Status**.

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

### FAQs

#### How do I get the GTID?

We recommend you run the flush log command to get the binlog file to locate the position and GTID of the maloperation.

#### How do I view the delay?

You can view the delay between a read-only instance and its source instance on the instance details page in the console.

#### How do I view the task information of replay by flashbacked position?

You can view the task progress and details on the task list page in the console.

## Managing Disaster Recovery Instance

Last updated : 2024-06-18 16:12:14

### Overview

For applications with high requirements of service continuity, data reliability, and compliance, TencentDB for MySQL provides cross-region disaster recovery instances to help enhance your capability to deliver continued services at low costs and improve data reliability.

#### Note:

A disaster recovery instance costs the same as the source instance associated with it. For more information, see Product Pricing.

#### Features

With a separate database connection address, a disaster recovery instance can offer read access capability for various scenarios such as nearby access and data analysis with lower costs of device redundancy. Its highly available source/replica architecture helps avoid single points of failure for databases.

Data is synced over a private network with lower latency and higher stability compared to public network.

The traffic of data sync over the private network is currently free of charge during the promotion period. Start date of charging for commercial purpose is subject to further notice.

#### How it works

If a TencentDB instance is used as a disaster recovery database, this instance will be the backup of the source instance.

When any change takes place in the source instance, the log information recording the change will be copied to the disaster recovery instance and then data sync will be implemented through log replay.

If any failure occurs in the source instance, the disaster recovery instance can be activated in seconds to provide full read/write capability.

### Feature limits

Disaster recovery instances cannot be created for single-node instances of the cloud disk edition.

The primary instance can support up to 4 disaster recovery instances.

Only supports specifications of 1 GB memory, 50 GB hard disk and above, and MySQL 5.6 or later versions (MySQL 5.6 needs to submit the ticket to use this feature), and the primary instance of the InnoDB engine with high availability edition can purchase disaster recovery instances. If the primary instance is below this specification, please upgrade it first.

The minimum specification of a disaster recovery instance is 1 GB memory and 50 GB disk capacity and must be at least 1.1 times the storage capacity used by the source instance.

Disaster recovery instance does not support the following features: transferring project, rollback, SQL operation, changing character set, account management, changing port, data import, rollback log, and read-only instance.

### Creating a disaster recovery instance

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to access the details page.

2. Make sure that the GTID feature is enabled by viewing the basic information of the instance on the **Instance Details** page. Click **Add Disaster Recovery Instance** in the instance architecture diagram to enter the disaster recovery instance purchase page.

Basic Info		Instance Architecture Diagram $\phi$	
Instance Name	: /	South China	F
Instance ID	6	©	(Running) Current Instance
Status/Task	Running /		
Region/AZ	North to New York		Async Delay in 0 sec
	Migrate to New		
Project	Switch to Another Project		Add Read-Only Instance
Placement Group	Add		Add Disaster Recovery Instance
GTID	Enabled		
Character Set	UTF8 🎤		

3. On the purchase page, set basic information of the disaster recovery instance such as **Billing Mode** and **Region**. **Note:** 

The time required to complete the creation depends on the amount of data, and no operations can be performed on the source instance in the console during the creation. We recommend you do so at an appropriate time.

If the sync policy is **Sync Now**, the data will be synced immediately when the disaster recovery instance is created. Only the data of the entire instance can be synced. Make sure that the disk space is sufficient.

You need to ensure that the source instance is in the running state and none of configuration adjustment tasks, restart tasks, and other modification tasks are executed. Otherwise, the sync task may fail.

- 4. After confirming that everything is correct, click **Buy Now** and wait for the delivery of disaster recovery instance.
- 5. Return to the instance list. After the status of the instance changes to Running, it can be used normally.

### Managing a disaster recovery instance

#### View a disaster recovery instance

A disaster recovery instance can be viewed from the region where it resides. You can use the instance list to filter out all instances in a specific region.



#### View the relationship

Click the icon on the right of a disaster recovery instance or source instance to view their relationship.



#### View sync delay

View the sync delay between the source instance and the disaster recovery instance at the top of the **Instance Details** page of the disaster recovery instance.



#### **Disaster recovery instance features**

A disaster recovery instance has various features, such as instance details, instance monitoring, database management, security group, and backup and restoration.

### Promoting a disaster recovery instance to a source instance

You can promote a disaster recovery instance to source instance in the console as needed.



#### Note:

After the disaster recovery instance is promoted to source instance, it will replace the original one as the new source instance to continue your business, and its access address will change. Therefore, you need to update the new address for your business.

After the promotion, the sync link to the original source instance will be disconnected and cannot be reconnected. You must exercise caution with this operation.

1. Log in to the TencentDB for MySQL console, click the ID of the disaster recovery instance to be promoted in the instance list, and access the instance management page.

2. Click **Promote to Source Instance** in the top-right corner to promote the disaster recovery instance to source instance. After the promotion, the sync link with the original source instance will be disconnected, so that the promoted instance can get data write capability and full MySQL functionality.

	0-second delay behind the	source instance $\phi$		Login Res	tart Manual E	Backup Promot	e to Source Instanc
Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Data Encryption	Connection Cl
Basic Info		Ir	stance Architectu	re Diagram $\phi$			

# Database Proxy Overview

Last updated : 2024-02-23 10:18:00

This document introduces the related features of the TencentDB for MySQL database proxy.

### **Document Overview**

Database Pro	ху	Description	Documented Guidance
Overview	Database Proxy Overview	Introducing the use cases of the database proxy, along with its proxy features and more.	Database Proxy Overview
Overview	Use Limits	Introducing the use limits for database proxy and the versions supported by various capabilities.	Use Limits
Kornol	Kernel Minor Version Release Notes	Introducing the viewing method and the update history of the database proxy's minor kernel version.	Kernel Minor Version Release Notes
Features	Upgrading Kernel Minor Version of Database Proxy	Certain features require a higher kernel version for the database proxy. In instances where it is not sufficient, you can obtain methods on how to upgrade the kernel version from this document.	Upgrading Kernel Minor Version of Database Proxy
Managing Database	Enabling Database Proxy	Introducing the prerequisites, use limits, and directions of the database proxy.	Enabling Database Proxy
Ргоху	Setting Database Proxy Access Address	Introducing how to modify, add, or delete database proxy access addresses after activating the database proxy.	Setting Database Proxy Access Address
	Viewing and Modifying Access Policy	Introducing how to view and modify the access policy of the database proxy access address after activating the database proxy.	Viewing and Modifying Access Policy
	Adjusting Database Proxy Configuration	Introducing the operations involved in adjusting the proxy configuration (node specifications/quantity, availability zone) after activating the database proxy.	Adjusting Database Proxy Configuration



	Switching Database Proxy Network	Introducing the procedure of altering the network connection address of the database proxy after activating the database proxy.	Switching Database Proxy Network
	Viewing Database Proxy Monitoring Data	Introducing how to view each database proxy node's connections, requests, CPU usage, memory usage, and more after activating the database proxy.	Viewing Database Proxy Monitoring Data
	Load Rebalancing	Introducing the manual operation of Cloud Load Balancer when there is an imbalance of load after activating the database proxy on each proxy node.	Load Rebalancing
	Disabling Database Proxy	Introducing the procedure for disabling the database proxy.	Disabling Database Proxy
Transaction S	plit Feature	Introducing the transaction split feature of the database proxy.	Transaction Split Feature
Momentary D Prevention Fe	sconnection ature	Introducing the momentary disconnection prevention feature of the database proxy.	Momentary Disconnection Prevention Feature
Automatic Read/Write	Automatic Read/Write Separation Overview	Introducing the automatic read/write separation feature of the database proxy.	Automatic Read/Write Separation Overview
Separation Feature	Configuring Database Proxy Read-Write Attributes	Introducing the procedures for configuring the read/write attributes of the database proxy connection address after activating the database proxy.	Configuring Database Proxy Read-Write Attributes
Connection	Connection Pool Overview	Introducing the connection pool feature of database proxy.	Connection Pool Overview
Pool Feature	Enabling and Disabling the Connection Pool Feature	Introducing how to enable and disable the connection pool feature after activating the database proxy.	Enabling and Disabling the Connection Pool Feature
Other Features	Hint Syntax Usage	Introducing the Hint syntax usage in the database proxy.	Hint Syntax Usage

# Overview Database Proxy Overview

Last updated : 2024-07-22 14:56:11

This document describes the new version of the TencentDB for MySQL database proxy.

Database proxy is a network proxy service between the TencentDB service and the application service. It is used to proxy all requests when the application service accesses the database.

The database proxy access address is independent of the original database access address. Requests arriving at the proxy address are all relayed through the proxy cluster to access the source and replica nodes of the database.

Read/Write requests are separated, so that read requests are forwarded to read-only instances, which lowers the load of the source database. This implements high availability, high performance, and Ops support.

TencentDB for MySQL database proxy also provides advanced features such as automatic read/write separation,

transaction split, connection pool, cross-AZ read-only instance mounting, and enabling multiple proxy addresses.

### Billing

Currently, database proxy is in beta test free of charge.

### Use cases

Businesses with low performance when there is a large number of non-persistent connections.

Businesses that use multiple read-only instances with read/write separation is implemented manually on applications, resulting in higher maintenance costs and risks.

High instance loading caused by too many connections.

High source instance load due to a large number of requests in the transaction.

Businesses that need to be assigned different loads through the access address.

Nearby access in cross-AZ deployment for reduced latency.

### Read-write attribute

Each database proxy address can have its own read-write attribute.

Read/Write: Supports read/write separation to implement linear business scaling.

This attribute contains at least one source instance and one read-only instance, and write requests are sent to the

source instance only. It supports read/write separation features such as transaction split and connection pool as well as policies such as **Remove Delayed RO Instances**, **Least RO Instances**, and **Failover**.

Read-Only: Supports read-only businesses like reports.

This attribute contains at least one read-only instance, and the source instance is not involved in routing. It supports features such as transaction split and connection pool as well as policies such as **Remove Delayed RO Instances** and **Least RO Instances**.

### Features

#### **High stability**

The database proxy is deployed in the cluster architecture, with multiple nodes ensuring smooth failover.

#### High availability

The database proxy adopts cross-AZ deployment to improve its availability.

#### Strong isolation

The database proxy provides the proxy service for the current instance with independent resources. The resources of each proxy are isolated and not shared.

#### **Ultra-high performance**

Each proxy can process up to 100,000 requests per second.

#### Convenient and fast scaling

You can dynamically add 1-60 proxy nodes, with only 6 nodes supported during the beta test.

#### Comprehensive performance monitoring

Performance metrics are monitored at the second level, such as the number of read/write requests, CPU, and memory. The number of proxies can be adjusted according to the monitoring data as described in Viewing Database Proxy Monitoring Data and business planning.

#### Hot reload

If the source instance is switched, has configuration adjustment, or has read-only instances added or removed, the database proxy can dynamically hot load the configuration without causing network disconnections or restarts.

#### Automatic read/write separation

Enabling the read/write attribute of the database proxy address can effectively reduce the read load of the source instance. Read-only instances can be added to horizontally scale the database cluster and automatically implement read/write separation, which eliminates the complexity of manually separating read and write requests of the business. This is especially suitable for scenarios with a high read load.

For example, only one proxy connection address needs to be configured in the application (when Read-Write

Attribute is set to **Read/Write Separation**), and this address will automatically implement read/write separation and send read requests to read-only instances and write requests to the source instance. Even if you add or remove read-only instances, you don't need to adjust the application settings.



#### **Connection pool**

This feature can mitigate excessively high instance loads caused by too many connections or frequent new connections in non-persistent connection businesses.

#### **Transaction split**

This feature separates reads and writes in one transaction to different instances for execution and forwards read requests to read-only instances to reduce the load of the source instance.

### Feature page

Instance Details Instance	Monitoring Datab	ase Management	Security Group	Backup and Rest	oration	Operation Log	Read-Only Instance	Database Proxy	Data Security	Connection
Overview Access Policy	Performance M	onitoring								
Basic Info			Disable Dat	abase Proxy 🗘	Proxy	Node				
Status/Task	Running				Node	e ID	Connections	AZ	Status	
Region	North China region(Be	ijing)			proxy	/node-	1	Beijing Zone 3	Running	
Proxy Version	1.3.4 Upgrade Kernel	Minor Version			proxy	mode-	1	Beijing Zone 6	Running	
Node Quantity	2 Adjust Configuratio	ins						,-,-,-		
Specification	2-core 4000 MB memo	ory								
Connection Persistence Timeout	5 sec 🎤									
Connection Address + A	dd Access Address									
Private Net Read-Writ	e Connection	Network	Remarks 0	peration						
172.2 For Read/Write Port:3306 For Separation	e Disabled		/ R C	Details Idjust Configurat Iebalance Close						

## Use Limits

Last updated : 2024-11-11 17:01:00

This document describes the use limits for the TencentDB for MySQL database proxy.

When using the proxy connection address, if you don't enable transaction split, transaction requests will be routed to the source instance.

The database proxy supports cross-AZ configuration. The number of selectable AZs depends on how many AZs are available in the current region. You can select up to three AZs. If only one AZ can be selected, there is only one available AZ in the region.

You can create multiple database proxy addresses, the number of which is the same as that of proxy nodes. When a proxy connection address is used to implement read/write separation, the consistency of non-transactional reads is not guaranteed. If your business requires read consistency, you can encapsulate it into transactions or use the hint syntax.

When a proxy connection address is used, show processlist will merge the results of all nodes before returning them.

For the PREPARE statement, the database proxy will first send PREPARE to all nodes. When a subsequent EXECUTE request comes in, it will determine the execution route according to the prepared statement type. For example, if a write statement is prepared, it will send the statement to the source database during execution, and if a non-transactional read statement is prepared, it will send the statement to a read-only instance.

After a business connection arrives at the database proxy, the proxy will connect to the source instance and all configured read-only instances. The proxy itself does not have a limit on the maximum number of connections, which is mainly subject to the maximum number of connections of the backend database instance. The smallest value of this parameter of the source and read-only instances will affect the business performance.

After the database proxy is enabled, when a read-only instance is added or restarted, only new connection requests will be routed to it. You can view the performance metrics of each proxy node through the overview or performance monitoring as described in Viewing Database Proxy Monitoring Data. If you find that the numbers of connections on the nodes are unbalanced, you can distribute the connections through rebalancing.

Versions supported for different database proxy capabilities:

Cross-AZ configuration: Proxy v1.3.1 or later.

Connection persistence timeout: Proxy v1.2.1 or later.

Connection pool: Proxy v1.3.12 or later.

Transaction split: Proxy v1.3.1 or later.

Connection address adding: Proxy v1.3.1 or later.

#### Note:

For more information on how to upgrade the proxy, see Upgrading Kernel Minor Version of Database Proxy.

# Database Proxy Kernel Features Kernel Minor Version Release Notes

Last updated : 2024-07-22 14:57:27

This document describes the kernel version updates of the TencentDB for MySQL database proxy.

### Viewing the database proxy version

You can view the version in **Overview** > **Proxy Version** on the **Database Proxy** tab.

Instance Details Instance	Monitoring Database Management	Security Group	Backup and Re	storation Operati	on Log Read-Only Insta	nce Database Proxy	Data Security	Conr
Overview Access Policy	Performance Monitoring							
Basic Info		Disable Database P	Proxy Ø	Proxy Node				
Status/Task	Running			Node ID	Connections	AZ	Status	
Region	South China(Guangzhou)				di 1		Running	
Proxy Version	1.3.4 Upgrade Kernel Minor Version				JL 0		Rupping	
Node Quantity	2 Adjust Configurations				iii v		Kunning	
Specification	2-core 4000 MB memory							
Connection Persistence Timeout	5 sec 🧨							

### Version upgrade description

#### Note:

For more information on the kernel version updates of the TencentDB for MySQL database proxy, see Kernel Version Release Notes > Database Kernel Version Release Notes.

If the MySQL kernel version requirements are not met, you can upgrade the kernel version of your database first as instructed in Upgrading Kernel Minor Version.

# Upgrading Kernel Minor Version of Database Proxy

Last updated : 2024-07-22 14:57:50

This document describes how to manually upgrade the kernel minor version of the database proxy in the console.

### Prerequisite

You have enabled the database proxy. For more information, see Enabling Database Proxy.

### Notes

There will be a momentary disconnection when you upgrade the kernel minor version of the database proxy. Therefore, upgrade the version during off-peak hours and make sure that your application has an automatic reconnection mechanism.

If the upgrade fails, the reason may be that the kernel minor version of the database is earlier than 20211031. You need to upgrade the kernel minor version of the database before upgrading the kernel version of the database proxy. For more information, see Upgrading Kernel Minor Version

### Directions

1. Log in to the TencentDB for MySQL console. Select the region at the top and click the target instance ID to enter the instance management page.

2. On the instance management page, select the Database Proxy tab.

3. In Overview > Basic Info > Proxy Version on the Database Proxy tab, click Upgrade Kernel Minor Version.

Basic Info	
Status/Task	Running
Region/AZ	
Proxy Version	1.0.1 Upgrade Kernel Minor Version
Node Quantity	2
Read/Write Separation	Enable
Connection Pool	To use this feature, you need to upgrade both the proxy and the source instance kernel to the latest version.

4. In the pop-up window, check the target version and select the upgrade switch time. After confirming that everything is correct, click **OK**.

Switch Time:

During maintenance time: The upgrade will be performed during the maintenance time, which can be modified on the instance details page.

Upon upgrade completion: The upgrade will be performed immediately after the upgrade operation is confirmed.

#### Note:

There will be a momentary disconnection during upgrade. Therefore, make sure that your business has a reconnection mechanism.

Except for abnormal nodes, all nodes are upgraded at the same time by default during the kernel minor version upgrade.



Upgrade Kern	el Minor Version			×	<
Current Version	1.0.1				
Target Version	1.1.2 To use the features of the latest printstance kernel to the latest version For the differences between the latest version document ≧	oroxy kernel, you ion as well. kernel minor vers	need to upgrade th ions, please see re	e source ference	
Switch Time	During maintenance time         Switch Time Description ☑         Maintenance Time 03:00-04:00 (lpage)         ✓       A flash disconnection will on business has a reconnection method	Upon upgrad Modify maintenar ccur during the up chanism.	e completion nce time on the inst ograde. Make sure	ance details that your	
Nodes to Upgrade	e, View Details 🔺				
Node ID	Specification		Status		
	2-core 4000 MB	8 memory	Running		
	2-core 4000 MB	8 memory	Running		
	ОК	Cancel			

# Database Proxy Management Enabling Database Proxy

Last updated : 2025-03-21 11:47:39

This document describes how to enable database proxy in the TencentDB for MySQL console. Database proxy is a network proxy service between the TencentDB service and the application service. It is used to proxy all requests when the application service accesses the database. It provides advanced features such as automatic read/write separation, connection pool, and connection persistence and boasts high availability, high performance, Ops support, and ease of use.

### Prerequisite

The instance is running on the two-node or three-node architecture.

### Note

Database proxy currently is supported in the following regions:

Beijing (except Zones 1, 2, and 4), Shanghai (except Zone 1), Guangzhou (except Zones 1 and 2), Shanghai Finance (except Zones 1 and 2), Beijing Finance, Chengdu, Chongqing, Nanjing, and Hong Kong (China) (except Zone 1). Tokyo (except Zone 1), Bangkok (except Zone 1), Virginia (except Zone 1), Silicon Valley (except Zone 1), Seoul (except Zone 1), and Singapore (except Zones 1 and 2).

Database proxy currently is supported on the following versions: two-node and three-node MySQL 5.7 (with kernel minor version 20211030 or later) and three-node MySQL 8.0 (with kernel minor version 20211202 or later). If you upgrade the kernel minor version of the source instance, the associated read-only and disaster recovery instances will be upgraded at the same time. For more information, see Upgrading Kernel Minor Version.

### Directions

Log in to the TencentDB for MySQL console. In the instance list, select the source instance for which to enable database proxy and click its ID or Manage in the Operation column to enter the instance management page.
 On the instance management page, select the Database Proxy tab and click Enable Now.

-0
Database proxy is not enabled yet.
Enable Now
Database proxy helps handle requests from the application to the database and provides advanced features such as automatic read/write separation, connection pooling, persistent connections, etc. It is easy to use, improves database and performance, and makes database OPS easier.
Database proxy is now in beta and available for free.

3. In the pop-up window, configure the following items and click **OK**.



Enable Database I	Ргоху	×
Database pr	roxy is free-of-charge in beta, after which a commercial version will be released.	
Select Network	ч • Ф	
	If the existing networks do not meet your requirements, go to Create VPCs 🗹 or Create Subnets 🖸 In the current network environment, only devices in the "can access this database instance.	
Proxy Specification	2-core 4000 MB memory 🔍	
	AZ Node Quantity Chengdu Zone 2	
	It's recommended to set the number of proxy nodes to 1/8 (rounded up to the nearest integer) of the sum of the CPU cores per node of the source instance and the CPU cores of all its read-only instances. For example, if the source instance uses 4 CPU cores per node and its read-only instances use 8 CPU cores in total, then the recommended number of proxy nodes is $(4+8)/8 \approx 2$ . If the recommended number of proxy nodes you calculated exceeds the maximum purchasable quantity, please choose a higher proxy node specification.	
Security Group	Selected 1 item	
	Preview Rules Instruction II To access through the database proxy, you need to configure security group policies and open the private port (3306). For more information, see MySQL Security Groups II	
Remarks	Enter remarks.	
	OK Cancel	
ameter	Description	_

Network	Select the network of the database proxy, which can only be a VPC.
Proxy Specification	Select 2-core 4000 MB memory, 4-core 8000 MB memory, or 8-core 16000 MB memory.



AZ and Node Quantity	1. Select the database proxy AZ. You can click <b>Add AZ</b> to add more AZs. The number of selectable AZs depends on how many AZs are available in the current region. You can select up to three AZs. 2. Select the number of nodes. We recommend that you set the quantity to 1/8 (rounded up) of the total number of CPU cores on the source and read-only instances; for example, if the source instance has 4 CPU cores, and the read-only instance has 8 CPU cores, then the recommended node quantity will be $(4 + 8) / 8 \approx 2$ . <b>Note:</b> If the selected proxy and the source instance are not in the same AZ and you access the instance through the proxy, the write performance may be reduced. If the recommended number of proxy nodes you calculated exceeds the maximum purchasable quantity, choose a higher proxy node specification.
Security group	The security group of the source instance is selected by default. You can also select another existing security group or create a new one as needed. <b>Note:</b> To access through the database proxy, you need to configure security group policies and open the private port (3306). For more information, see TencentDB Security Group Management.
Remarks	(Optional) Enter the remarks of the database proxy service to be enabled.

4. After successfully enabling the service, you can manage proxy nodes and view their basic information on the database proxy page. You can also modify the access address, network type, and remarks of the database proxy, view and adjust the connection configuration, and perform rebalance in the **Connection Address** section.

#### Note:

You can view **Connections** in the proxy node list or view the performance monitoring data of each proxy node to check whether the numbers of connections on the nodes are unbalanced, and if so, you can distribute the connections by clicking **Rebalance**.

Rebalance will cause proxy nodes to restart, and the service will become unavailable momentarily during the restart. We recommend that you restart the service during off-peak hours. Make sure that your business has a reconnection mechanism.

Overview A	Access Policy	Performance M	Monitoring		
Basic Info				Disable	Database Proxy 🗳
Status/Task	F	Running			
Region					
Proxy Version	1	1.3.4 Upgrade Kern	el Minor Version		
Node Quantity	2	2 Adjust Configurat	ions		
Specification	2	2-core 4000 MB men	nory		
Connection Persist	tence Timeout	5 sec 🧨			
Connection Ad	lduare to the		0		
Connection Ad	aaress + Add /	Access Address (1/2)	0		
Private Net	Read-Write	Connection	Network	Remarks	Operation
1 Г <u>–</u> Port:3306 Г <u>–</u>	Read/Write Separation	Disabled	с-	/	Details Adjust Configura Rebalance

## Setting Database Proxy Access Address

Last updated : 2024-07-22 15:16:55

This document describes how to set the database proxy connection address in the TencentDB for MySQL console. The database proxy connection address is independent of the original database connection address. Requests proxied at the proxy address are all relayed through the proxy cluster to the source and replica nodes of the database, so that read and write requests are separated, with reads being forwarded to read-only instances. In this way, the load of the source database is lowered.

After database proxy is enabled for TencentDB for MySQL, a database proxy connection address will be added by default, and you can also add, modify, or delete the connection address for database proxy.

### Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

### Modifying Database Proxy Connection Address

Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation column of the instance with the proxy enabled to enter the instance management page.
 On the instance management page, select the Database Proxy tab and click the

icon next to Private Network Access Address on the Connection Address tab.

us/Task Running ion/AZ South China (Guangzhou) / Guangzhou Zone 4 ion/AZ 5.7 le Quantity 2 d/Write Separation Enabled Interction Address atabase Proxy Address Network Remarks i 3306 / i Opfault-VPC - Default-Subnet	Basic Info		
ion/AZ South China (Guangzhou) / Guangzhou Zone 4 y Version 5.7 le Quantity 2 d/Write Separation Enabled Interction Address atabase Proxy Address Network Remarks 1 3306 ✓ To Default-VPC - Default-Subnet	Status/Task	Running	
y Version 5.7 le Quantity 2 d/Write Separation Enabled meetion Addresss atabase Proxy Address Network Remarks ■ :3306 ▲ Ti Default-VPC - Default-Subnet	Region/AZ	South China (Guangzhou) / Guangzhou Zone 4	
le Quantity 2 d/Write Separation Enabled meetion Address atabase Proxy Address Network Remarks ■ 3306 ▲ Ta Default-VPC - Default-Subnet -	Proxy Version	5.7	
d/Write Separation Enabled	Node Quantity	2	
atabase Proxy Address Network Remarks	Read/Write Separation	Enabled	
atabase Proxy Address     Network     Remarks       3306 / To     Default-VPC - Default-Subnet	Connection Address		
3306 V To Default-VPC - Default-Subnet	Database Proxy Addre	ss Network	Remarks
	:3306 🧪 🗗	Default-VPC - Default-Subnet	

3. In the pop-up dialog box, modify the proxy address and click **OK**.

#### Note:

Modifying the proxy address affects the database business being accessed. We recommend you modify the address during off-peak hours. Make sure that your business has a reconnection mechanism.

Private IP*	172.16.0.2
	Available Private IP Range: 172.16.0.0/20
Private Port*	3306
	Port value range: 1024-65535
Note: modifyin	g the private network address will affect the database service being accessed.

### Adding Database Proxy Connection Address

#### Note:

The number of connection addresses is the same as that of database proxy nodes.

An connection address will be created by default when the database proxy is enabled.

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation

column of the instance with the proxy enabled to enter the instance management page.



2. On the instance management page, select the **Database Proxy** tab, and click **Add Access Address** next to **Connection Address**.

Connection Add	dress + Add Ac	cess Address (1/2)	D		
Private Net	Read-Write	Connection	Network	Remarks	Operation
172.2 17 🖻 Port:3306 🖻	Read/Write Separation	Disabled		/	Details Adjust Configurat Rebalance

3. In the **Create Connection** window, set the following configuration items and click **OK**.

#### Step 1. Configure a network

Parameter	Description
Network	Database proxy only supports VPC. You can choose to automatically assign or specify an address.
Security group	The security group of the source instance is selected by default. You can also select another or more existing security groups or create a new one as needed. <b>Note :</b> To access through the database proxy, you need to configure security group policies and open the private port (3306). For more information, see TencentDB Security Group Management.
Remarks	Optional. You can add remarks for the new database proxy connection address.

#### Step 2: Configure a policy

Parameter	Description
Read-Write Attribute	Select the read-write attribute of the proxy access address, which can be Read/Write Separation or Read-Only.
Remove Delayed RO Instances	Set the Remove Delayed RO Instances policy. After this option is enabled, you can set Delay Threshold and Least RO Instances. The system will try removing or restoring a failed read-only instance no matter whether this option is enabled. Delay Threshold: Enter an integer greater than or equal to 1 (in seconds). Least RO Instances: It is subject to the number of read-only instances owned by the source instance. If it is set to `0`, when all read-only nodes are removed, all read requests will be routed to the source instance until at least one of the removed read-only instances rejoins the database proxy to continue processing the read requests.
Connection Pool Status	The connection pool feature mainly mitigates the instance load caused by frequent new connections in non-persistent connection business scenarios. After this option is enabled,



	you can select the supported connection pool type, which currently can only be the session-level connection pool by default.
Transaction Split	You can set whether to enable this feature. After it is enabled, reads and writes in one transaction will be separated to different instances for execution, and read requests will be forwarded to read-only instances to reduce the load of the source instance.
Assign Read Weight	You can select Assigned by system or Custom. If multiple AZs are configured when the database proxy is enabled, you can separately assign the weights of proxy nodes in different AZs.
Failover (with Read-Write Attribute being Read/Write Separation)	You can set whether to enable this feature. After it is enabled, if database proxy fails, the database proxy address will route requests to the source instance.
Apply to Newly Added RO Instances	You can set whether to enable this feature. After it is enabled, a newly purchased read- only instance will be automatically added to the database proxy. If Assign Read Weight is set to Assigned by system, newly purchased read-only instances will be assigned with the default weight based on their specification. If Assign Read Weight is set to Custom, when newly purchased read-only instances are added, their weights will be 0 by default, which can be modified by clicking Adjust Configurations in Connection Address on the Database Proxy tab.

### **Deleting Database Proxy Connection Address**

#### Note:

If your database proxy has multiple connection addresses, you can delete unnecessary ones while keeping the last one.

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column of the instance with the proxy enabled to enter the instance management page.

2. On the instance management page, find the target address in **Database Proxy** > **Connection Address**, and click **Disable**.



Connection Address + Add Access Address					
Private Net	Read-Write	Connection	Network	Remarks	Operation
172. 7 F	Read/Write Separation	Disabled	,	/	Details Adjust Configurat Rebalance Close
172 17 Г Ротt:3306 Г	Read/Write Separation	Disabled	,	-1	Details Adjust Configurat Rebalance Close

3. In the pop-up window, click OK.

## Viewing and Modifying Access Policy

Last updated : 2024-07-22 15:17:17

After database proxy is enabled for TencentDB for MySQL, a database proxy connection address will be added by default. You can also add connection addresses later to implement different business logics. The number of connection addresses that can be created is the same as that of database proxy nodes. You can view and modify the access policy of a connection address in the console.

### Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

### Viewing the access policy

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Access Policy**.

Instance Details Instance Monitor	ing Database Mana	agement Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Security	Connectio
Overview Access Policy Pe	rformance Monitoring							
A proxyaddr-		Diroxynode- Beijing Zone 3	Weight 1	M cdb- o cd Beijing Zone 6	lb225942 ( <b>Running</b> )			
Remove Delayed RO Instances	Disa Settings bled		Weld Harr	+ Add Read-Only In	stance			
Failover	Enabled	b proxynode-						
Apply to Newly Added RO Instances	Enabled	Beijing Zone 6						
Read-Write Attribute	Read/Wri te Separatio n	+ Add Proxy Node						
Private IP Attributes	172.21.							
Private Port	3306 🗖							
Assign Read Weight	Nearby Access							
Read/Write Separation	Disabled							

### Modifying the access policy



1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Access Policy**, find the target access policy, and click **Settings**.

#### Note:

You can also find the target access address in **Database Proxy** > **Overview** > **Connection Address**. Then, click **Adjust Configurations** in the **Operation** column.

3. In the pop-up window, modify the policy configuration and click **OK**.

	Adjust Configurations							
	Read-Write Attribute	Read/Write Separation     Read-Only						
	Remove Delayed RO Instances	Learn More 🗹						
		Note that this setting only applies to delayed KO instances. Failed KO instances are always removed directly and added backed after they're recovered.						
	Connection Pool Status	Learn about Connection Pool 🗹						
	Connection Pool Type	Session-Level Connection Pool						
	Transaction Split							
	Assign Read Weight	O Assigned by system Custom						
		Beijing Zone 3 (i) Beijing Zone 6 (i)						
		Instance ID/Na Type Enable Weight Status AZ						
		cdb- cdb225942Source InsPlease select ▼ nningBeijing Zo						
	Failover	If database provy fails, the database provy address will route requests to the source instance						
	Apply to Newly Added 80 Instances	Outerequests to the source instance.						
	Apply to newly hadea no instances	If you purchase a new non-delayed read-only instance, it will be automatically added to the database proxy.						
	_	OK Cancel						
Par	ameter	Description						
Rea	ad-Write Attribute	Modify the read-write attribute of the proxy access address, which can be Read/Write Separation or Read-Only.						



Remove Delayed RO Instances	Set the Remove Delayed RO Instances policy. After this option is enabled, you can set Delay Threshold and Least RO Instances. The system will try removing or restoring a failed read-only instance no matter whether this option is enabled. Delay Threshold: Enter an integer greater than or equal to 1 (in seconds). Least RO Instances: It is subject to the number of read-only instances owned by the source instance. If it is set to `0`, when all read-only nodes are removed, all read requests will be routed to the source instance until at least one of the removed read-only instances rejoins the database proxy to continue processing the read requests.
Connection Pool Status	The connection pool feature mainly mitigates the instance load caused by frequent new connections in non-persistent connection business scenarios. After this option is enabled, you can select the supported connection pool type, which currently can only be the session-level connection pool by default.
Transaction Split	You can set whether to enable this feature. After it is enabled, reads and writes in one transaction will be separated to different instances for execution, and read requests will be forwarded to read-only instances to reduce the load of the source instance.
Assign Read Weight	You can select Assigned by system or Custom. If multiple AZs are configured when the database proxy is enabled, you can separately assign the weights of proxy nodes in different AZs.
Failover (with Read-Write Attribute being Read/Write Separation)	You can set whether to enable this feature. After it is enabled, if database proxy fails, the database proxy address will route requests to the source instance.
Apply to Newly Added RO Instances	You can set whether to enable this feature. After it is enabled, a newly purchased read-only instance will be automatically added to the database proxy. If Assign Read Weight is set to Assigned by system, newly purchased read-only instances will be assigned with the default weight based on their specification. If Assign Read Weight is set to Custom, when newly purchased read-only instances are added, their weights will be 0 by default, which can be modified by clicking Adjust Configurations in Connection Address on the Database Proxy tab.

## Adjusting Database Proxy Configuration

Last updated : 2024-07-22 15:17:32

This document describes how to adjust the database proxy configuration in the console, including the proxy specification, AZ, and node quantity.

### Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

### Notes

The database proxy will be automatically upgraded to the latest version during configuration adjustment if its version is old.

If the selected proxy and the source instance are not in the same AZ and you access the instance through the proxy, the write performance will be reduced.

If the recommended number of proxy nodes you calculated exceeds the maximum purchasable quantity, choose a higher proxy node specification.

If resources are sufficient in the region, you can select up to three AZs and need to retain at least one AZ (the one in the first row).

### Impact description

Changing different configuration items has different impacts, some of which will not cause a momentary disconnection, while some will. The specific change items and their impacts are as detailed below:

#### Scenario 1: Changing the proxy specification but not AZ or node quantity

Proxy Specification	AZ	Node Quantity	Load Balancing Mode	Switch Time	Impact
Upgrade or downgrade	Unchanged	Unchanged	Automatic	During maintenance time	There will be a momentary disconnection lasting seconds. Make sure that your business has a reconnection mechanism.
				Upon upgrade	



		completion
	Manual	During maintenance time
		Upon upgrade completion

#### Scenario 2: Changing the node quantity but not proxy specification or AZ

Proxy Specification	AZ	Node Quantity	Load Balancing Mode	Switch Time	Impact
Unchanged Unchanged				During maintenance time	There will be a momentary disconnection lasting seconds.
	Increased	Automatic	Upon upgrade completion	Make sure that your business has a reconnection mechanism.	
		Manual	-	There will be no momentary disconnections.	
	Degraged	Automatic	During maintenance time	There will be a momentary disconnection lasting seconds.	
			Upon upgrade completion		
		Decreased	Manual	During maintenance time	Make sure that your business has a reconnection mechanism.
				Upon upgrade completion	

#### Scenario 3: Changing the AZ and node quantity but not proxy specification

Proxy	AZ	Node	Load	Switch Time	Impact
Specification		Quantity	Balancing		
			Mode		
-----------	-----------	-----------	-----------	-------------------------------	---
Unchanged			Automatic	During maintenance time	There will be a momentary disconnection lasting seconds.
		Increased		Upon upgrade completion	Make sure that your business has a reconnection mechanism.
			Manual	-	There will be no momentary disconnections.
	Increased	Decreased	Automotia	During maintenance time	There will be a momentary disconnection lasting seconds. Make sure that your business
			Automatic	Upon upgrade completion	has a reconnection mechanism.
			Manual	During maintenance time	
				Upon upgrade completion	
	Decreased			During maintenance time	
		Increased	Automatic	Upon upgrade completion	
		nicieaseu	Manual	During maintenance time	
			ivialiual	Upon upgrade completion	
		Decreased	Automatic	During maintenance time	



			Upon upgrade completion
		Manual	During maintenance time
		Manual	Upon upgrade completion
		Automotic	During maintenance time
	Increased	Automatic	Upon upgrade completion
		Manual	During maintenance time
Changed			Upon upgrade completion
Changed		Automotio	During maintenance time
	Decreased	Αυτοπατις	Upon upgrade completion
	Decleased	Manual	During maintenance time
		manual	Upon upgrade completion

## Directions

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target

instance to enter the instance management page.

2. On the instance management page, select the **Database Proxy** tab.

3. On the **Overview** tab on the **Database Proxy** tab, select **Basic Info** > **Node Quantity** and click **Adjust Configurations**.

Instance Det	tails Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Encryption	Conne
Overview	Read/Write Separation	Performance Monitoring							
							Adjust Configuratio	ns Disable Da	tabase Pro
Basic Info				Proxy	/ Node				
Status/Task	Running			Re	ebalance (j				

4. In the pop-up window, modify the proxy specification, AZ, and node quantity of the database proxy as needed and click **OK**.

Adjust Datab	ase Proxy Configurations	>
The proxy w adjustment i	ill be automatically upgraded to the latest version during the configuration fit is the ealier version	
Proxy	4-core 8000 MB memory	
Node Quantity	2 (1 to 4)	
	To ensure the high availability of proxy please purchase at least two proxy podes	
	It's recommended to get the number of provide the 1/16 (rounded up to the	-
	It's recommended to set the number of proxy nodes to 1/16 (rounded up to the	
	hearest integer) of the sum of the CPU cores per hode of the source instance and	1
	the CPU cores of all its read-only instances. For example, if the source instance	
	uses 8 CPU cores per node and its read-only instances use 16 CPU cores in total	Ι,
	then the recommended number of proxy nodes is $(8+10)/10 \approx 2$ .	_
	numerecommended number of proxy nodes you calculated exceeds the maximum	
Switch Time	During maintenance time Upon upgrade completion	
	Switch Time Description 🗾	
	Maintenance Time 03:00-04:00 (Modify maintenance time on the instance details	
	Dage)	
	A flash disconnection will occur upon adjustment completion. Make sure tha	t
	your business has a reconnection mechanism.	
	OK Cancel	

5. In **Basic Info** on the **Database Proxy** tab, after the task status changes from **Upgrading** to **Running**, the configuration adjustment is completed.

#### Note:

If you select **Upon upgrade completion**, after the configuration is adjusted, the system will automatically switch to the new configuration.

If you select **During maintenance time**, after the configuration is adjusted, the system will switch the configuration during the specified maintenance time period.

If you select **During maintenance time** but need to switch the configuration earlier due to your business

requirements, after the configuration is adjusted, you can go to **Database Proxy** > **Overview** > **Basic Info** > **Status/Task** and click **Complete Now** after **Waiting for switch (after upgrade)**.



If you select manual rebalance after the configuration adjustment, perform the operation in **Overview** > **Connection Address** on the **Database Proxy** tab after the configuration is adjusted.

# Switching Database Proxy Network

Last updated : 2024-07-22 15:17:44

This document describes how to switch the database proxy network in the TencentDB for MySQL console.

## Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

#### Notes

Switching the network may cause the change of the instance database proxy IP. The old IP will be retained for 24 hours by default and up to 168 hours. Then, it will become invalid. Therefore, modify the IP on the client promptly. If **Valid Hours of Old IP** is set to 0 hours, the IP is released immediately after the network is changed.

You can only select a VPC in the region of the MySQL instance, but you can choose a subnet in any AZ and view its IP range.

When you add multiple access addresses for the database proxy, you can set and modify the network for each of them.

## Directions

1. Log in to the TencentDB for MySQL console. Select a region on the top, and click the instance ID to enter the instance management page.

- 2. On the instance management page, select the Database Proxy tab.
- 3. On the Database Proxy tab, click

under **Overview** > **Connection Address** > **Network**.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoratio	n Operation Log	Read-Only Instance	Database P
Overview Rea	d/Write Separation	Performance Monitoring					
							Adjust Conf
Basic Info				Pro	xy Node		
Status/Task	Running / Waiting for	switch (after upgrade) Comple	te Now		Rebalance (j)		
Region/AZ				Ν	ode ID	Connections	Specificatio
Proxy Version	1.1.2 Upgrade Kern	el Minor Version				0	4-core 8000
Node Quantity	2						
Read/Write Separation	n Enable					0	4-core 8000
Connection Pool	Enable						
Connection Addre	255						
Database Proxy Ad	ldress Netwo	rk	Remarks				
ī / īī		2	- /				

4. In the pop-up dialog box, select a new network and click **OK**.

<ol> <li>When the network</li> <li>the change by definition</li> <li>If "Valid Hours of</li> <li>You can only se</li> </ol>	ork is changed, the ault. Please modify of Old IP" are set to elect a VPC and sub	IP of the instance will be r the client program accord 0 hours, the IP is released onet in the same region as	eplaced with a ingly in time. I immediately the instance.	a new one. The old one will become invalid 24 hours follo after the network is changed.
Select Network				
	•	:4	- ¢	CIDR 253 subnet IPs in total, with 246 available
If the existing network	s do not meet your	requirements, go to Creat	e Subnets 🛂	, ···,
In the current network	environment, only	CVMs in the "	" can access	this database instance.
Valid Hours of Old IP	24	hr Range: 0-168 hours		
🔾 Auto-Assign IP				
O Specify IP				

#### Select Auto-Assign IP or Specify IP.

5. After successfully changing the network, you can view the new network under **Connection Address**.

# Viewing Database Proxy Monitoring Data

Last updated : 2024-07-22 15:17:58

This document describes how to view database proxy node monitoring data in the TencentDB for MySQL console.

## Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

#### Supported Monitoring Metrics

Metric Name	Unit	Description
Current connections	Count	Current node connections
Requests	Count/sec	Node access requests
Read requests	Count/sec	Read operation requests
Write requests	Count/sec	Write operation requests
CPU utilization	%	CPU utilization
Memory utilization	%	Memory utilization
Memory usage	MB	Used memory

#### Directions

#### Option 1:

1. Log in to the TencentDB for MySQL console. In the instance list, click the ID or **Manage** in the **Operation** column of the source instance with proxy enabled to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Performance Monitoring** page, and click a proxy node name to view its monitoring data. You can switch between different nodes by clicking their names.

#### Note:

The default granularity of monitoring within four hours is 5 seconds.



Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	g Read-Only Instance	Database Proxy	Data Security	Connection
Overview Acces	ss Policy Perform	ance Monitoring							
Proxy Node prox	kynode-		Last	hour Last 24 hours Last 7 day	s Last 30 days	2023-01-11 09:36 ~ 2023-	01-11 10:35 📋	(Granularity: 5 se	econds) (j
Current Cor	nnection: « Curren	t Connections (ProxyCu	urrentConnections, Ur	nit: pcs)					

#### **Option 2:**

1. Log in to the TencentDB for MySQL console. In the instance list, click the ID or **Manage** in the **Operation** column of the source instance with proxy enabled to enter the instance management page.

2. On the instance management page, select the **Database Proxy** > **Overview**, and click icon

after the target node ID in the **Proxy Node** column, then you can view the performance monitoring data of the node.

Instance Details Instan	ce Monitoring Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Security	Connection
Overview Access Poli	cy Performance Monitoring							
Basic Info		Disable Datab	base Proxy 🗘 Prox	y Node				
Status/Task	Running		Nod	le ID	Connections	AZ	Status	
Region	North China region(Beijing)		prox	ynode <b>II</b>	1	Beijing Zone 3	Running	
Proxy Version	1.3.4 Upgrade Kernel Minor Version			umada li	1	Paiiing Zono 6	Punning	
Node Quantity	2 Adjust Configurations		prov	linoue III	1	beijing Zone o	Kunning	

The displayed page after redirection is as follows.

← uniteresting		proxynode-			Configure Alarm Policy
Instance Details Ir	nstance Monitoring Database Man	1 hour	Time granularity: 5 sec	▼ 🗘 Disable ▼ 🚥 ✓ Show lege	nds
Overview Access	s Policy Performance Monitoring				
Basic Info		cpu use rate(%) (j)	<b>↓</b> C2 ····	current connections(count) $({\rm \hat{i}})$	<b>≜</b> ⊡ •
		0.8	11:20 0.70	1.2	11:30 1.0
Status/Task	Running	0.6	na na hanna na han na sana na na hana na ha	0.6	
Region	North China region(Beijing)	0.2		0.3	
Proxy Version	1.3.4 Upgrade Kernel Minor Ve	10:38 10:45 10:53 11:00	11:08 11:15 11:23	10:38 10:45 10:53 11:00	11:08 11:15 11:23
Node Quantity	2 Adjust Configurations	proxy   proxynode-1 c Max: 0.70 Min: 0.1	0 Avg: 0.48	proxy   proxynode- : Max: 1.00 Min: 1.0	00 Avg: 1.00

# Load Rebalancing

Last updated : 2024-07-22 15:18:09

After the database proxy is enabled, you can view **Connections** in the proxy node list or the performance monitoring data of each proxy node to check whether the numbers of connections on the nodes are unbalanced. If there is a large number of persistent connections, increasing the number of the database proxy nodes may also cause load imbalance among them, and if so, you can distribute the connections by clicking **Rebalance**. This document describes how to manually rebalance the node in the console.

#### Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

#### Directions

1. Log in to the TencentDB for MySQL console. Select the region at the top of the page and click the target instance ID to enter the instance management page.

 On the instance management page, select Database Proxy > Overview, find the target access address in Connection Address, and click Rebalance in the Operation column.

Connection Address + Add Access Address							
Private Net	Read-Write	Connection	Network	Remarks	Operation		
1720.7 🖻 Port:3306 🖻	Read/Write Separation	Disabled	,	/	Details Adjust Configurat Rebalance Close		

3. In the pop-up window, click **OK**.

#### Note:

Rebalancing will cause the disconnection of the session to the address, during which the service will be unavailable momentarily. We recommend that you restart the service during off-peak hours. Make sure that your business has a reconnection mechanism.

## **Disabling Database Proxy**

Last updated : 2024-07-22 15:18:21

This document describes how to disable database proxy in the TencentDB for MySQL console.

#### Note:

Before disabling the database proxy, make sure that your client does not access the database through the proxy address, as the generated data will be deleted and cannot be recovered after the database proxy is disabled.

#### Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

#### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, select the desired instance, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.

2, On the instance management page, select **Database Proxy** > **Overview** > **Basic Info** > **Disable Database Proxy**.

Overview Read	/Write Separation Performance Monitoring					
						Disable Database I
Basic Info			Proxy Node In beta, 1	he specification and quantity of pre	oxy nodes cannot be modified once the	proxy is enabled.
Status/Task	Running		Rebalance			
Region/AZ	South China (Guangzhou) / Guangzhou Zone 4		Node ID	Connections	Specification	Status
Proxy Version	5.7		proxynode-	2	2-core 4000 MB memory	Running
Node Quantity	2		proxynode	2	2-core 4000 MB memory	Running
Read/Write Separation	Enabled					
Connection Addres	55					
Database Proxy Add	Iress Network	Remarks				
:3306 🧪	Default-VPC - Default-Subnet					

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

# Access Mode

Last updated : 2024-08-29 14:08:08

The access mode is mainly used to control the connection path between an application or the client and the database proxy, which includes **Balanced Distribution** and **Nearby Access**. Here, we introduce rules, advantages, and disadvantages of these two access modes above for the database proxy of TencentDB for MySQL.

## Prerequisites

You have enabled the database proxy.

The two-node and cluster edition database architecture is required.

#### **Balanced Distribution**



**Rule**: In the Balanced Distribution access mode, the client application connects to all available database proxy nodes. If the database proxy nodes are spread across multiple availability zones, application connections to the database are evenly distributed among all availability zone nodes.

Advantage: Traffic is evenly distributed, resolving the issue of excessive load on a single node.

**Disadvantages**: If there are proxy nodes in different availability zones, high request latency is prone to occur due to extended access paths.

#### **Nearby Access**



**Rule**: In the Nearby Access mode, the application will connect to the proxy node in the same availability zone or the one with the shortest access path. Even if the database proxy nodes are spread across multiple availability zones, the application will still connect the proxy node with the shortest access path.

Advantage: Low access latency and high speed.

**Disadvantage**: If your client application changes to an availability zone with fewer database proxy nodes than the original zone, it can lead to excessive load on the nodes and performance delays.

## **Changing Access Mode**

The access mode can be changed. For details, see Viewing and Changing the Access Policy.

# **Transaction Split Feature**

Last updated : 2024-07-22 15:19:24

The TencentDB for MySQL database proxy provides the transaction split feature. This feature separates read and write operations in one transaction to different instances for execution and forwards read requests to read-only instances, thereby lowering the load of the source instance.

## Background

By default, the TencentDB for MySQL database proxy forwards all requests in a transaction to the source instance to ensure the transaction correctness. In some frameworks, however, all requests are encapsulated to transactions that are not committed automatically, causing an excessive load of the source instance. In this case, you can use the transaction split feature.

The transaction split feature is disabled by default. You can enable it by adjusting the configuration of the database proxy address.

## Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

## Directions

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Access Policy**, find the target access policy, and click **Settings**.

#### Note:

You can also find the target access address in **Database Proxy** > **Overview** > **Connection Address**. Then, click **Adjust Configurations** in the **Operation** column.

3. In the pop-up window, toggle on **Transaction Split** and click **OK**.



Transaction Split	
Assign Read Weight	O Assigned by system Custom

# Momentary Disconnection Prevention Feature

Last updated : 2024-12-30 16:27:05

This document describes the momentary disconnection prevention feature of the TencentDB for MySQL database proxy.

## Background

During the instance Ops, you may need to make some adjustments, such as configuration modification, planned HA switch, and planned restart. Theses Ops may cause the problems, such as session interruption, momentary disconnection and failed new connections. The TencentDB for MySQL database proxy provides the momentary disconnection prevention feature that enables lossless application continuity to prevent disconnection and transaction interruption.

## How It Works

The momentary disconnection prevention feature implements MySQL's session track mechanism. When the lossy behavior is perceived, the database proxy will disconnect the client from the source node before the switch, and connect to the source node after the switch. Then the session-related system variables, user variables, and character set encoding information will be transferred to the new backend connection through the session track mechanism, so as to realize the lossless switching on the application side.



## Notes

If the statement uses the temp tables associated with each session, the connection cannot be recovered, and an error will be reported.

To use the momentary disconnection prevention feature, you need to upgrade the database proxy kernel version to v1.3.1 or later.

The disconnection prevention feature will stop the transactions over 3 seconds.

When the connection is switched, if the database proxy is receiving the result message from the database, only part of the data will be transferred due to the source/replica switch, and the momentary disconnection prevention feature will be disabled.

If there is a prepared statement in use during the momentary disconnection prevention, the momentary disconnection prevention feature cannot maintain the connection.

## Performance Testing

The performance test of the momentary disconnection prevention feature for Tencent DB for MySQL database proxy is described in the following.

#### **Test environment**

Region/AZ: Beijing - Beijing Zone 7 Client: S5.8XLARGE64 (32-core 64 GB Standard S5) Client operating system: CentOS 8.2 64-bit Network: Both the CVM and TencentDB for MySQL instances are in the same VPC subnet. The information of the tested TencentDB for MySQL instance is as follows: Storage type: Local SSD disk Instance type: General Parameter template: High-performance template

#### Test tool

SysBench, as the tool for the performance test, is a modular, cross-platform, and multi-threaded benchmark tool for evaluating OS parameters that are important for a system running a database under intensive load. The idea of this benchmark suite is to quickly get an impression about system performance without setting up complex database benchmarks or even without installing a database at all.

#### Test method

In different Ops scenarios, you can analyze the ratio of momentary disconnections before and after the operation to test whether a database proxy provides momentary disconnection prevention for the high-availability MySQL instance.

## **Test Results**

In the following Ops scenarios, the high-availability MySQL instance maintains a 100% connection keep-alive rate by the momentary disconnection prevention capability of the database proxy.

Ops scenario	Keep-alive rate
Performing source-replica switch	100%
Upgrading kernel minor version	100%
Adjusting the instance specifications	100%

# Automatic Read/Write Separation Feature Automatic Read/Write Separation Overview

Last updated : 2024-07-22 15:20:14

This document describes the automatic read/write separation feature of the database proxy service in TencentDB for MySQL and its strengths and routing rules.

## Automatic Read/Write Separation

Currently, businesses of many users in the production environment have problems such as more reads and less writes and unpredictable business loads. In application scenarios with a large number of read requests, a single instance may not be able to withstand the pressure of read requests, which even may affect the businesses. To implement the auto scaling of read capabilities and mitigate the pressure on the database, you can create one or multiple read-only instances and use them to sustain high numbers of database reads. However, this solution requires that businesses can be transformed to support read/write separation, and the code robustness determines the quality of business read/write separation, which imposes high technical requirements and has low flexibility and scalability. Therefore, after creating a read-only instance, you can purchase a database proxy, configure access address policy,

and configure the database proxy address in your application so as to automatically forward write requests to the source instance and read requests to the read-only instance. In addition, this method provides natural solutions to other business challenges as detailed below:

Scenarios where the load is unpredictable or fluctuates irregularly with obvious peaks and troughs In internet business scenarios, business load and access pressure are often unpredictable and unstable, and there will be frequent great fluctuations. If the business uses a large number of non-persistent connections to access the database, it is easy to generate many new connections. In other words, the number of connections between the database and the application is likely to fluctuate as the business access pressure changes frequently, and this is often difficult to predict.

Connection management for the dedicated database proxy allows you to efficiently reuse database connections to appropriately scale applications that handle unpredictable workloads. First, this feature allows multiple application connections to share the same database connection to effectively use database resources. Second, it allows you to adjust the number of open database connections to maintain predictable database performance. Finally, it allows you to delete unusable application requests to guarantee the overall application performance and availability.

Scenarios where the application is frequently connected to and disconnected from the database Applications built based on technologies such as serverless, PHP, or Ruby on Rails may frequently open and close database connections to process application requests. The dedicated database proxy can help you maintain a database connection pool to prevent unnecessary pressure on data computing and the memory used to establish new connections.

#### Scenarios where the database access connection is idle for a long time and is not released

SaaS applications and traditional ecommerce applications may make database connections idle to minimize the response time for user reconnection. You can use the dedicated database proxy to retain idle connections and establish database connections as needed instead of excessively increasing the threshold or providing database services with higher specifications to support most idle connections.

#### Scenarios where you want to improve the smoothness and stability of database PaaS service failover

With the dedicated database proxy, you can build applications that can tolerate active and passive database failures in an imperceptible manner with no need to write complex failure processing code. The dedicated database proxy will automatically route read traffic to new database instances while retaining the application connections.



## Advantages

Read/write requests are automatically separated with a unified access address.

Native linkage support improves the performance and reduces the maintenance costs.

You can flexibly set weights and thresholds.

Failover is supported, so that even if the database proxy fails, requests can access the source database normally. When a source instance is switched, or its configuration is changed, or a read-only instance is added/removed, the database proxy can dynamically hot reload the configuration without causing network disconnections or restarts.

## Read/Write Separation Routing Rules

#### Sending to the source instance

DDL statements such as CREATE , ALTER , DROP , and RENAME . DML statements such as INSERT , UPDATE , and DELETE . SELECT FOR UPDATE statement. Statements related to temp tables. Certain system function calls (such as last\_insert\_id()) and all custom function calls. Statements related to LOCK . Statements after transaction is enabled (including set autocommit=0). Stored procedures. Multiple statements concatenated by ";". KILL (SQL statement, not command). All queries and changes of user variables.

#### Sending to the read-only instance

Non-transactional read ( SELECT ) statements.

#### Sending to all instances

show processlist statement. All changes of system variables (SET command). USE command.

# Configuring Database Proxy Read-Write Attributes

Last updated : 2024-07-22 15:21:09

After creating a read-only instance, you can purchase the database proxy service and configure the connection address policy. Then, you can configure the database proxy address in your application so as to automatically forward write requests to the source instance and read requests to the read-only instance. This document describes how to enable read/write separation in the console.

## Prerequisites

You have enabled the database proxy. For more information, see Enabling Database Proxy.

## Directions

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Overview**, find the target access address in **Connection Address**, and click **Adjust Configurations** in the **Operation** column.

Adjust Configurations		×		
Read-Write Attribute	Read/Write Separation     Read-Only			
Remove Delayed RO Instances	Learn More 🗹			
	Note that this setting only applies to delayed RO instances. Failed RO instances are always removed directly and added backed after they're recovered.			
Connection Pool Status	Learn about Connection Pool 🗹			
Transaction Split 🚯				
Assign Read Weight	O Assigned by system O Custom			
	Beijing Zone 3 (i) Beijing Zone 6 (i)			
	Instance ID/Na Type Enable Weight Status AZ			
	cdb- cdb225942Source InsPlease select nningBeijing Zo			
Failover	If database proxy fails, the database proxy address will route requests to the source instance.			
Apply to Newly Added RO Instances	If you purchase a new non-delayed read-only instance, it will be automatically added to the database proxy.			
	OK Cancel			

3. In the pop-up window, set **Read-Write Attribute**, assign the read weight, and click **OK**.

#### Note:

The weight here pertains to the allocation strategy for read requests (non-transactional).

Adjust Configurations		×	
Read-Write Attribute	Read/Write Separation     Read-Only		
Remove Delayed RO Instances	Learn More 🗹		
	Note that this setting only applies to delayed KO instances. Failed KO instances are always removed directly and added backed after they're recovered.		
Connection Pool Status	Learn about Connection Pool 🖸		
Transaction Split 🛈			
Assign Read Weight	O Assigned by system Custom		
	Beijing Zone 3 (i) Beijing Zone 6 (i)		
	Instance ID/Na Type Enable Weight Status AZ		
	cdb-     Source Ins     Please select ▼ nning     Beijing Zo		
Failover			
	If database proxy fails, the database proxy address will route requests to the source instance.		
Apply to Newly Added RO Instances	If you purchase a new non-delayed read-only instance, it will be automatically added to the		
	database proxy.		
	OK Cancel		

# Connection Pool Feature Connection Pool Overview

Last updated : 2024-11-11 17:03:01

TencentDB for MySQL database proxy supports the connection pool feature, currently offering session-level and transaction-level connection pools, which can effectively solve the issue of high instance loads due to frequent establishments of new connections by short connection services. This document introduces the features of session-level connection pool.

## Prerequisites

You have enabled the database proxy.

The database kernel minor version is MySQL 5.7 20211030 or later. The database kernel minor version is MySQL 8.0 20230630 or later.

#### Note:

The database proxy connection pool capability has been identified to cause deadlock issues during the 'change user' operation in earlier minor versions of the MySQL kernel. This issue has been rectified in the MySQL 8.0 20230630 minor kernel version. If the minor kernel version of your primary instance has not been updated to MySQL 8.0 20230630, there is a certain level of risk involved. It is recommended to disable the database proxy connection pool capability. If you want to continue utilizing the database proxy connection pool capability, it is advised to upgrade the minor kernel version of your primary instance to MySQL 8.0 20230630 or later. For instructions on upgrading, refer to upgrading kernel minor version.

## Background

Session-level connection pool



The session-level connection pool is applicable to non-persistent connection scenarios.

The session-level connection pool is used to reduce the instance load caused by frequent establishments of new nonpersistent connections. If a client connection is closed, the system will determine whether the current connection is idle, and if so, the system will put it into the proxy connection pool and retain it for a short period of time. (It is 5 seconds defaulted by the system. After the connection pool feature is enabled, manual configuration is supported. For more details, refer to enabling the connection pool feature.)

When the client initiates a new connection, if the connection pool has an available connection, the connection can be used directly to reduce the database connection overheads; otherwise, a new connection will be established as usual. **Note:** 

The session-level connection pool does not reduce the number of concurrent connections to the database, but it can reduce the overheads of the main thread of TencentDB for MySQL by lowering the speed for connecting the application to the database, so as to better process business requests. However, idle connections in the connection pool will occupy your connection quantity temporarily.

The session-level connection pool cannot be used to solve the problem of connection heap caused by a large number of slow SQL queries, which must first be solved by yourself.

#### Notes

Currently, the connection pool feature does not support setting different permissions for different IPs under the same account, as that may cause a permission error during connection reuse. For example, if <code>mt@test123</code> has

database\_a permissions, while mt@test456 does not, a permission error may occur when you enable the connection pool.

The connection pool feature is for the database proxy instead of the client. You don't need to use the connection pool of the database proxy if your client already supports it.

# Enabling and Disabling the Connection Pool Feature

Last updated : 2024-11-11 17:05:22

This document describes how to enable or disable the connection pool feature.

## Prerequisites

#### You have enabled the database proxy.

The database kernel minor version is MySQL 5.7 20211030 or later.

The database kernel minor version is MySQL 8.0 20230630 or later.

#### Note:

The database proxy connection pool capability has been identified to cause deadlock issues during the 'change user' operation in earlier minor versions of the MySQL kernel. This issue has been rectified in the MySQL 8.0 20230630 minor kernel version. If the minor kernel version of your primary instance has not been updated to MySQL 8.0 20230630, there is a certain level of risk involved. It is recommended to disable the database proxy connection pool capability. If you want to continue utilizing the database proxy connection pool capability, it is advised to upgrade the minor kernel version of your primary instance to MySQL 8.0 20230630 or higher. For instructions on upgrading, please refer to upgrading kernel minor version.

## Enabling the connection pool feature

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Access Policy**, find the target access policy, and click **Settings**.

#### Note:

You can also find the target access address in **Database Proxy** > **Overview** > **Connection Address**. Then, click **Adjust Configurations** in the **Operation** column.

3. In the pop-up window, toggle on **Connection Pool Status.** Once it is enabled, select **Session-Level Connection Pool** or **Transaction-Level Connection Pool** (database proxy kernel version 1.4.1 or later needed), and click **OK**.

	oxy is nee-on-charge in beta, and which a commercial version will be released.
Network	
Proxy	2-core 4000 MB memory v
Specification	
Node Quantity	2 (1 to 4) To ensure the high availability of proxy, please purchase at least two proxy nodes.
	It's recommended to set the number of proxy nodes to 1/8 (rounded up to the neare
	integer) of the sum of the CPU cores per node of the source instance and the CPU
	cores of all its read-only instances. For example, if the source instance uses 4 CPU
	cores per node and its read-only instances use 8 CPU cores in total, then the
	recommended number of proxy nodes is $(4+8)/8 \approx 2$ .
	If the recommended number of proxy nodes you calculated exceeds the maximum
	purchasable quantity, please choose a higher proxy node specification.
Type Connection	5 con (Panae: 1.200)
Persistence	sec (Range: 1-300)
Timeout	
Security Group	
	Selected 1 item
	×
	Preview Rules Instruction 🛛
	To access through the database proxy, you need to configure security group policies
	and open the private port (3308). For more information, see MySQL Security Group
	2
Remarks	Enter remarks.

4. After the connection pool feature is enabled, the default connection retention threshold is set to 5 seconds. It is supported to modify this threshold under the **Basic Info** section, with the modification range being from 1 to 300 seconds. If the connection pool feature is disabled, the connection retention threshold will not be displayed.

Basic Info			Disable Database Proxy 🗘
Status/Task	Running	Specifications Used	2core4000MB, 1 CU
Region	North China(Beijing)	Billing Mode	Pay as You Go
Proxy Version	1.3.13 Stable Upgrade Kernel Minor Version		
Node Quantity	1 Adjust Specification Configuration		
Specification	2-core 4000 MB memory		
Connection Persistence Timeout	5 sec 🎤		

## Disabling the connection pool feature

1. Log in to the TencentDB for MySQL console, select the region at the top of the page, and click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Access Policy**, find the target access policy, and click **Settings**.

#### Note:

You can also find the target access address in Database Proxy > Overview > Connection Address. Then, click

#### Adjust Configurations in the Operation column.

3. In the pop-up window, toggle off **Connection Pool Status** and click **OK**.

# Adaptive Load Balancing

Last updated : 2024-08-30 13:06:59

The TencentDB for MySQL database proxy provides the adaptive load balancing feature, which can route requests to instances with lower loads based on backend load conditions to achieve load balancing of database instances, ensuring the response speed of database instances.

## Background

In database usage scenarios, when multiple read-only instances are used, a database proxy is usually required for load balancing, allowing multiple read-only instances to share the request load, thereby reducing the load on the primary instance. The current TencentDB for MySQL database proxy supports load balancing by assigning read weights to each read-only instance. This is a "static" load balancing capability, ensuring that requests are distributed to the read-only instances based on the set read weights.

If different services access the instance simultaneously from different database proxy connection addresses, it might lead to load imbalance, which may not be solved by the "static" load balancing capability.

For example, a user has 1 primary instance and 2 read-only instances (RO1 and RO2), where RO2 is used by two database proxy connection addresses. When the load on RO2 increases, the "static" load balancing capability results in a fixed number of requests being assigned to it. Therefore, even if RO2's response slows down, it will continue to be allocated requests, potentially slowing down the overall response. Consequently, the TencentDB for MySQL database proxy provides an adaptive load balancing feature. When it is found that the request of RO2 slows down and the load is too high, the database proxy will reduce the requests assigned to RO2 and assign them to other read-only instances, thus achieving a "dynamic" load balancing and ensuring the overall response speed of the database instances.



## Prerequisites

The instance is running on the two-node or three-node architecture.

You have enabled the database proxy.

The kernel minor version of the database proxy is 1.3.8 or later. If this requirement is not met, see Upgrading Kernel

Minor Version of Database Proxy.

Database kernel version requirements are as follows. If they are not met, see Upgrading Kernel Minor Version of Database.

The MySQL 5.7 kernel minor version should be 20211030 or later.

The MySQL 8.0 kernel minor version should be 20211202 or later.

#### Overview

If your instance has a high overall load and a certain read-only instance is fully utilized while others are not, enabling the adaptive load balancing proxy will prioritize sending requests to instances that can still respond, making full use of all instances' resources. Note that if the instance is fully utilized due to unreasonable slow SQL, the adaptive load balancing may result in high loads on all instances, and the slow SQL issue will still need to be addressed. If all your instances have the same specifications and all instances have the same configured weights, and you need to evenly distribute the traffic, you can directly enable the adaptive load balancing, which can meet almost all scenarios' requirements.

In the following situations, it is not recommended to enable the adaptive load balancing feature:

1. If there is only one database proxy connection address, it is recommended to use the "static" load balancing capability instead of enabling the adaptive load balancing feature.

2. After the adaptive load balancing feature is enabled, and under low load conditions, requests may not be routed based on the set read weights. If you want to route requests in strict accordance with the set read weights, it is not recommended to enable the adaptive load balancing feature.

3. For situations where slow SQL causes a single read-only instance to be fully utilized, load balancing cannot alleviate the issue. In this case, it is recommended to optimize the SQL or increase the instance specifications instead of enabling the adaptive load balancing feature.

## Precautions

Changes to load balancing policies only take effect for new connections after the configuration is changed. If there are existing connections, you may need to perform the Load Rebalancing operation to make them take effect for the old connections.

"Dynamic" load balancing can only take effect on traffic that goes through the database proxy. If there is traffic that directly connects to the database without going through the database proxy, the load balancing policy effect may be unstable.

## Enabling/disabling Adaptive Load Balancing

#### Note:

The following steps are only for enabling/disabling the adaptive load balancing. For detailed information about the specific settings and instructions under the adjusting configurations, including the rules and impacts of adjusting the read weights, see Modifying Access Policy.

1. Log in to the MySQL Console, select a region at the top, and then click the ID of the target instance to enter the instance management page.

2. On the instance management page, select **Database Proxy** > **Overview**, find the target access address under the Connection Address and click **Adjust Configurations** in the **Operation** column.



Connection Address + Add Access Address (1/2)					
Private Network	Read/Write Attrib	Connection Pool	Network	Remarks	Operation
10 🚺 / 🗗 🎤 Port:3306 🔂	Read/Write Separation	Disabled	net 🍂	/	Details Adjust Configurati Load Rebalancing

3. On the Adjust Configurations page, click the button after **Adaptive Load Balancing** to enable or disable it, and then click **OK**.

Adjust Configurations	
Read/Write Attribute	Read/Write Separation     Read-Only
Access Mode	◯ Balanced allocation O Nearby Access How to Select Access Mode
Remove Delayed RO Instances	Learn More Z Note that this setting only applies to delayed RO instances. Failed RO instances are always removed directly and added backed after they're recovered.
Adaptive Load Balancing	What Is Adaptive Load Balancing 🗹
Connection Pool Status	Learn about Connection Pool
Transaction Split (j)	
Assign Read Weight	O Assigned by system Custom
	Beijing Zone 6 (i) Beijing Zone 7 (i)
	Instance ID/Na Type Enable (i) Weight (i) Status AZ
	Cluster Running Beijing Zon
	Read-write node dbn- 2(auto-assig v lunning Beijing Zon
	Read-only node dbn-(
Failover (j)	If the read-only instance fails, the database proxy will be routed to the primary instance.
Apply to Newly Added RO Instances	If you purchase a new non-delayed read-only instance, it will be automatically added to the database proxy.
	OK Cancel

# Other Features Hint Syntax Usage

Last updated : 2024-07-22 15:21:57

This document describes how to use the hint syntax on the database proxy.

The hint syntax can be used to forcibly execute SQL requests on the specified instance. A hint has the highest routing priority. For example, it is not subject to consistency and transaction constraints. You need to reasonably evaluate whether it is required in your business scenario before using it.

#### Note:

When using the MySQL command line tool to connect and use the HINT statement, you need to add the -c option in the command; otherwise, the hint will be filtered by the tool.

The database proxy kernel version greater than or equal to version 1.1.3 supports **PREPARE** statement when the hint syntax is used through the database proxy.

Currently, three types of hints are supported:

Assign to the source instance for execution:

```
/* to master */
or
/*FORCE_MASTER*/
```

Assign to a read-only instance for execution:

```
/* to slave */
or
/*FORCE_SLAVE*/
```

Specify an instance for execution:

/\* to server server\_name\*/

<code>server\_name can be a short ID, such as /\* to server test\_ro\_1 \*/</code> .
# Account Management Creating Account

Last updated : 2024-07-22 15:27:14

### Overview

In addition to the default root account created by the system, you can also create other database accounts in the TencentDB for MySQL console as needed.

### Directions

1. Log in to the TencentDB for MySQL 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 **Database Management**> Account Management and click Create Account.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy
Database List	Parameter Settings	Account Managemen	nt				
Create Account	Export Account List						Enter account nam
Account Name	Host		Maximum Conn	ections Rema	rks	Operatio	'n
-						Modify Permissi More ▼	ons Clone Accou
						Reset Password	
2 in total						10 🔻 / page 🛛 🛏 🔺	1 / 1 page

3. In the pop-up window, enter the account name, host, and password, confirm the password, and click OK.

#### Account Name:

In MySQL 5.5 and 5.6, the database account name can contain 1–16 letters, digits, or underscores. It must begin with a letter and end with a letter or digit.

In MySQL 5.7 and 8.0, the database account name can contain 1–32 letters, digits, or underscores. It must begin with a letter and end with a letter or digit.

**Host**: Specify a host address to access the database. It can be an IP and contain (indicating no limit on the IP range). Multiple hosts should be separated by line break, space, semicolon, comma, or vertical bar.

Example 1: Enter 👔 to indicate no limit on the IP range, that is, clients at all IP addresses are allowed to use this account to connect to the database.

Example 2: Enter 10.5.10.% to indicate that clients with an IP range within 10.5.10.% are allowed to use this account to connect to the database.

**Password**: The password must contain 8–64 characters in at least two of the following character types: letters, digits, and special symbols  $()_{+-\&amp}; = !@#$%^{()}$ . You can also set the password complexity to improve the database security as instructed in Setting Password Complexity.

**Maximum Connections**: The number of connections for this account, which must be no more than 10,240. If you don't enter a value, no limit will be imposed (subject to the maximum number of connections).

4. After successful creation, the database account can be managed in the database account list of the current instance.

### **Related APIs**

API Name	Description
CreateAccounts	Creates a TencentDB account

# Setting Password Complexity

Last updated : 2024-07-22 15:27:34

TencentDB for MySQL allows you to set the password complexity to improve the strength of database access passwords and optimize the database security.

### Prerequisites

The database version is:

MySQL 5.6, with minor version 20201231 or later. MySQL 5.7, with minor version 20201231 or later. MySQL 8.0, with minor version 20201230 or later. The instance must be on a two-node/three-node architecture.

### **Use Limits**

For any password set during account creation or reset in the TencentDB for MySQL console, it must at least meet the following initial account password requirements:

It can contain 8 to 64 characters

It must contain at least three of the following character types: uppercase letters, lowercase letters, digits, and special symbols.

Special characters are  $\_+-&=!@#$%^*()$ .

### **Enabling Password Complexity**

#### Note:

After the password complexity feature is enabled, any password set during account creation or reset must follow the password complexity policy.

#### Enabling when creating an instance on the purchase page

- 1. Log in to the TencentDB for MySQL purchase page.
- 2. Configure parameters as needed. Select Enable after the Password Complexity parameter.

Password Complexity			
O Enable Close			
Min Number of Uppercase or Lowercase Letters	_	1	+
Min Number of Digitals	-	1	+
Min Number of Symbols	-	1	+
Min Number of Password Characters	-	8	+

#### 3. Set the following parameters.

Parameter	Description
Min Number of Uppercase or Lowercase Letters	You can select a number between 1 (default value) and 16.
Min Number of Digitals	You can select a number between 1 (default value) and 16.
Min Number of Symbols	You can select a number between 1 (default value) and 16.
Min Number of Password Characters	You can select a number between 8 (default value) and 64. The minimum value must be greater than the sum of the above three parameters.

#### Enabling for existing instances in the console

1. Log in to the TencentDB for MySQL 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 **Database Management** > Account Management and click

Password Complexity (disabled by default).

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restora	tion Operation Log
Database List	Parameter Settings	Account Management	t		
Create Account	Export Account List	Password Complexity: [Dis	able] Use [	)ynamic Credentials (i)	

3. Select **Enable** in the **Password Complexity** pop-up window, set the following parameters, and click **OK**.

Close Close Min Number of Uppercase or - 1 +
Min Number of Uppercase or - 1 +
Lowercase Letters
Min Number of Digitals - 1 +
Min Number of Symbols - 1 +
Min Number of Password - 10 +

ОК	Close
Parameter	Description
Min Number of Uppercase or Lowercase Letters	You can select a number between 1 (default value) and 16.
Min Number of Digitals	You can select a number between 1 (default value) and 16.
Min Number of Symbols	You can select a number between 1 (default value) and 16.
Min Number of Password Characters	You can select a number between 8 (default value) and 64. The minimum value must be greater than the sum of the above three parameters.

### **Disabling Password Complexity**

1. Log in to the TencentDB for MySQL 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 **Database Management**> **Account Management** and click **Password Complexity**.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restor	ration Operation
Database List	Parameter Settings	Account Management	t		
Create Account	Export Account List	Password Complexity: [Ena	able] Use Dyn	amic Credentials 🛈	

3. Select **Disable** in the **Password Complexity** pop-up window and click **OK**.

Password Complexity				
◯ Enable	O Close			
		ОК	Close	

### **Relevant Documentation**

Creating Account Resetting Password

## **Resetting Password**

Last updated : 2024-07-22 15:27:49

### Overview

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

#### Note:

For TencentDB for MySQL, the password resetting feature has been connected to CAM; therefore, we recommend that you exercise tighter control over the permission to this API (password resetting) or sensitive resources (i.e., TencentDB for MySQL instances) by granting such permission only to appropriate personnel. For data security, we recommend that you regularly reset the password at least once every three months.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation column to access the instance management page.

2. On the instance management page, select **Database Management** > **Account Management**, find the account for which to reset the password, and click **Reset Password** or select **More** > **Reset Password**.

Database List	Parameter Settings	Account Management			
Create Account	Export Account List	Use Dynamic Credentials (i)			Enter account name
Account Name		Host	Maximum Connections	Remarks	Operation
					Modify Permissions Clone Account More 👻
					Reset Password
					Reset Password Reset Permissions

3. In the pop-up window, enter a new password, confirm the password, and click OK.

#### Note:

The password should be a combination of 8–64 characters containing at least two of the three types: letters, digits, and symbols (\_+-&=!@#\$%^\*()).

### **Related APIs**

API Name	Description
ModifyAccountPassword	Modifies the password of a TencentDB account

# **Modifying Account Permissions**

Last updated : 2023-12-05 17:32:58

### **Operational Scenario**

You have the capability to manage existing database accounts through the MySQL Cloud Database Console, granting them either global privileges or object-level privileges. Concurrently, you also possess the ability to revoke these authorizations.

### Account Authorization Description

Authorization	Description
ALTER	The authority to alter the structure of tables within the database, encompassing the modification of table fields, the addition and deletion of table fields, as well as the creation and removal of table indices.
ALTER ROUTINE	Modify the permissions of routines such as stored procedures and functions.
CREATE	The authority to create new objects such as databases, tables, views, stored procedures, and functions.
CREATE ROUTINE	The authority to create routine objects such as stored procedures and functions.
CREATE TEMPORARY TABLES	The authorization to create temporary tables.
CREATE USER	This authority allows for the creation of new users and the allocation of corresponding permissions to them.
CREATE VIEW	The authority to create view objects.
DELETE	The authority to delete data from a specified table.
DROP	The authority to delete objects such as databases, tables, views, stored procedures, and functions.
EVENT	The authority to create, modify, and eradicate events.
EXECUTE	This authority permits users to execute pre-existing objects such as stored procedures and functions.



INDEX	The authority to create and eradicate indices.
INSERT	The authority to insert or write new data into the table.
LOCK TABLES	Securing the requisite permissions for executing read and write operations on the designated data table.
PROCESS	The authority to view the information of all currently active threads and processes.
REFERENCES	The authority to create or remove foreign key constraints within the current database.
RELOAD	The authority to reload certain system configuration files, as well as to restart or shut down the database.
REPLICATION CLIENT	Authority to observe and administer the status and parameters of the replication process.
REPLICATION SLAVE	The authority to replicate data from the Master server.
SELECT	The authority to query data from the designated table.
SHOW DATABASES	Display the permissions of all databases in the current database server.
SHOW VIEW	Displaying the permissions of the view.
TRIGGER	Authority to create and administer triggers.
UPDATE	Permission to update data in a specified table.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click the instance ID or Manage in the Operation column to access the instance management page.

2. Select the **Database Management** > **Account Management** tab, find the account for which to modify the permissions, and click **Modify Permissions**.

nstance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	on Operation Log	Read-Only Instance	Database Proxy
Database List	Parameter Settings	Account Managemen	t				
Create Account	Export Account List						Enter account nam
Account Name	Host		Maximum Conne	ections Re	marks	Operatio	n
						<u>Modify Permissio</u> More ▼	Clone Accou
						Reset Password	
2 in total						10 🔻 / page 🛛 🖌	1 / 1 page

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

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

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

et Database Permissions				
Global Privileges	DROP		INSERT	
+ Object Level Privilege	SELECT		TRIGGER	
	CREATE TEMPORARY TABLE	S	CREATE VIEW	
	DELETE		EVENT	
	REFERENCES		SHOW VIEW	
	UPDATE		ALTER	
	EXECUTE		LOCK TABLES	
	ALTER ROUTINE		CREATE	
	INDEX		PROCESS	
	All			

### **Related APIs**

API Name	Description
ModifyAccountPrivileges	Modifies the permissions of a TencentDB instance account

# Modifying Host Addresses with Access Permissions

Last updated : 2024-07-22 15:28:14

### **Operation Scenario**

By modifying the host address authorized by the database account in the TencentDB for MySQL console, you can control the access to the database to improve the access security.

### Directions

1. Log in to the TencentDB for MySQL console.

2. In the instance list, select the instance to be modified and click the instance name or **Manage** in the **Operation** column to enter the instance management page.

3. Select **Database Management** > **Account Management**, find the account for which to modify the host, and select **More** > **Modify Host**.



4. In the pop-up dialog box for modifying host, enter the new host address and click **OK**.

#### Note:

The host address may come in the format of an IP address. To allow all the clients to access the database using the database account, enter '%'.

Modify Host	
Instance Name	kaylal_cdb238160
Account Name	mysql.sys
Server	localhost
New Host*	Enter the host.
	IP format, supporting %
	OK Cancel

# Modifying the Number of Connections

Last updated : 2025-04-02 10:20:11

TencentDB for MySQL supports modifying the number of connections (maximum number of connections from each account to TencenDB instances). This can prevent a single account from exhausting all connections to TencentDB for MySQL.

This document describes how to modify the number of connections for a TencentDB account through the console.

### Prerequisites

A TencentDB account has been created. See Creating Account for details.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

2. On the instance management page, select **Database Management** > **Account Management**, find the account requiring modification of the number of connections, and choose **More** > **Modify Maximum Connections**.

#### Note:

The root account does not support modifying the number of connections.

3. Enter the new connection limit in the pop-up and click OK.

#### Note:

The entered number of connections should be less than or equal to 10240.

### **Related APIs**

API	Description
ModifyAccountMaxUserConnections	This API is used to modify the maximum number of connections of one or more TencentDB accounts.

# **Deleting Account**

Last updated : 2024-07-22 15:28:27

### **Operation Scenarios**

In order to disable an existing database account, it must be deleted in the TencentDB for MySQL console.

#### Note:

A database account cannot be recovered once deleted.

In order to avoid accidental deletion from interrupting normal use by your business, you need to make sure that the database account to be deleted is no longer used by any applications.

### Directions

1. Log in to the TencentDB for MySQL Console. In the instance list, click an instance ID or **Manage** in the "Operation" column to enter the instance management page.

Select the Database Management > Manage Account tab, find the account to be deleted, and select More > Delete Account.

nstance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operatio	on Log Read-Only Ins	stance D	atabase Pro
Database List	Parameter Settings	Account Managemen	ıt					
Create Account	Export Account List						Er	iter account na
Account Name	Host		Maximum Conn	ections Rema	arks		Operation	
						Modi Mor	ify Permission e 🔻	s Clone Acco
						Reset Password Modify Maximum Connection	sword	
					[	Delete Account Modify Remarks	rmission	s Clone Acco
3 in total						10 👻 / page	k 4 1	/ 1 pag

3. In the pop-up dialog box, click OK to delete the account.

Delete Account		×
You've selected 1 account , View Details $\bullet$		
ОК	Cancel	

### **Related APIs**

API Name	Description
DeleteAccounts	Deletes a TencentDB account

# **Modifying Remarks**

Last updated : 2025-04-02 10:18:37

You can allocate and manage account resources based on account remarks. This document describes how to modify the remarks of a TencentDB account through the console.

### Prerequisites

A TencentDB account has been created. See Creating Account for details.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

2. On the instance management page, select **Database Management** > **Account Management**, find the account that needs to modify remarks, and select **More** > **Modify Remarks**.

#### Note:

The root account does not support modifying the remarks.

3. Enter new remarks in the pop-up window and click OK.

#### Note:

You can enter up to 255 characters for the remarks.

### **Related APIs**

API	Description
ModifyAccountDescription	This API is used to modify the remarks of a TencentDB account.

# Database Management Center (DMC) DMC Overview

Last updated : 2024-07-22 15:31:37

### Overview

Database Management Center (DMC) is a one-stop Tencent Cloud database management tool. Services supported include database/table-level operations, real-time monitoring, instance session management, SQL window, and data management.

DMC is currently supported for TencentDB for MySQL, TencentDB for MariaDB, TDSQL-C for MySQL, TencentDB for Redis, and TDSQL for MySQL.

#### Features

#### Data management

It supports database operations such as creating and deleting databases/tables, modifying table structures, and maintaining table data.

#### Database instance session management

It displays multidimensional instance session information to help you easily view and perform relevant operations.

#### **Database monitoring**

It supports multidimensional monitoring of database status, InnoDB row operations, database connections, and traffic.

#### SQL window

It provides a SQL compiler and can save frequently used SQL templates and custom SQL statements. It can also display and export SQL statement execution results.

### Benefits

#### Ease of use

DMC aims to provide easy-to-use database management services. You can use the convenient operation interface to perform database operations such as creating databases/tables.

#### Visualization

You do not need to separately install drivers or applications to use DMC. Instead, you just need to log in to the DMC console to manage, monitor, and manipulate databases in multiple dimensions through the visualization management page.

#### Real-Time database status monitoring

DMC can collect database instance data in real time such as CPU utilization, number of connections, and storage capacity usage, so as to intuitively display database instance running conditions.

#### Real-Time database performance monitoring

DMC supports database performance monitoring at the second level, covering multiple monitoring items related to MySQL status, InnoDB row operations, threads, and network to help you monitor fluctuations in each performance item in real time. In addition, it supports the categorized calculation of sessions and allows you to kill instance sessions, helping you quickly locate causes of exceptions and implement performance optimization.

# **DMC** Management

Last updated : 2024-12-27 22:04:06

This document describes the features of DMC, such as database/table creation, database management, instance monitoring, instance session management, and visualized table data editing.

### Database/Table Creation

1. Log in to DMC and select **Create** > **Create Database** > **Create Database** or **Create** > **Create Table** on the navigation bar at the top.

P		e v	Database	Management	Instance Monitoring	Instance Session	SQL Window New	Data Management 🗸	
infor	mation_schema (System dat	abase)	φ	Homepage	Database Manageme	nt			
III Tab	Fuzzy match table name	Q	+	Create Database	e				
ō	No table four	d		Database Name			Characte	r Set	Collation
🗐 View				information_schen	na (System database)		utf8 (UTF	-8 Unicode)	utf8_general_ci (Default)
LB Stored procedure									
fx Func									

2. In the pop-up dialog box, configure the new database or table.

#### Note:

For more information on the character set and collation, see MySQL official documentation.

Database creation dialog box:

Database Name *	Please enter database name	
Character Set *	utf8 (UTF-8 Unicode)	•
Collation *	utf8_general_ci (Default)	•

Table creation dialog box:

omepage I	Database Management	Create Table
Database: info	rmation_schema	
Basic Info	Column Info Inde	x Foreign Key Partition
Basic Info		
Table name *	Enter a table name	
Remarks		
Storago opgino	InnoDR	
otorage engine	initobb	
Character Set	Please select	•
Check Rules	Please select	▼
More Options		
Row Format	Please select	▼
Average Row Len	gth Please enter an inte	ger value
Min Rows	0	
Max Rows	0	

### Database Management

🕗 Tencent Cloud

Log in to DMC, select **Database Management** on the navigation bar at the top, and create, edit, or delete databases on the displayed page.

P	DMC Create v	Database	e Management	Instance Monitoring	Instance Session	SQL Window New	Data Management $\checkmark$	
infor	mation_schema (System database)	τ¢	Homepage	Database Management	t Create Table			
III Tak	Fuzzy match table name	<b>}</b> +	Create Databas					
Ĭ	No table found		Database Name			Character	Set	Collation
🗐 View			information_schen	na (System database)		utf8 (UTF-	8 Unicode)	utf8_general_ci (Default)
[]号 Stored procedure								
f.								

### **Instance Session Management**

Log in to DMC, select **Instance Session** on the navigation bar at the top, and enter the instance session management page. You can view the details of all database sessions in the instance from four dimensions: session overview, users, access sources, and databases.

DMC allows you to kill sessions, facilitating your session management.

Homep	age Instance Se	ession									
All	Active Sleep	Kill Selected Sessio	Kill Session by Cor	dition Filter 🔻					Server time: 01-21 16:35:25	Auto-refresh: 28	e F
	s	ession ID 💲	Duration ↓	User 🛊	Source \$	Databas	se ‡	CMD \$	SQL ‡	Status \$	÷
•									select ID,USER,TIM ()		1
						· · · · · · · · · · · · · · · · · · ·			1		
Session	Overview		By User(1)			By Access Source(1)			By Database(1)		
No.			User			Source			Database Name		Total
1			test			10.59.195.177					1
2											
3											
4	DMC System Sessio	ons 1									

### SQL Window

Log in to DMC, select **SQL Window** on the navigation bar at the top, or click **SQL Operation** on the **Operation** menu on the left sidebar to access the SQL window, which supports the following features:

- Run SQL commands and view results
- Optimize SQL statement formats
- View SQL command execution plans
- Save commonly used SQL statements
- Use SQL templates
- Export SQL statement execution results

Execute	Format Optimization	Execution Plan	Save	information_schema (Syste 🔻	This is a system database. Please proc	ceed with caution.		
1								My Template Sys
								Enter name to search
								concat
							•	
					×			
Info								
Clear								
No execution	record							

### Data Management

Log in to DMC, select **Data Management** > **Data Importing** or **Data Exporting** on the navigation bar at the top, and you can import data into or export data from a database.

Create ~	Databa	se Management	Instance Monitoring	Instance Session	SQL Window New	Data Management A
stem database)	τ¢	Homepage	Instance Session	SQL		Data Importing
) name Q	+					Data Exporting
able found		Execute	Format Optimization	Execution Plan	Save informati	bn_ochema (Syste = This is a sj
		1				

### Visualized Table Data Editing

DMC for MySQL supports inserting, deleting, and updating data. You can click a table in the left **Table** list to insert, delete, and update its data in batches in the right pane, and then click **OK** in the **Quick Operation** pane to preview the SQL statements and implement the modification.

# Parameter Configuration Setting Instance Parameters

Last updated : 2025-03-20 16:46:05

You can view and modify certain parameters and query parameter modification logs in TencentDB for MySQL console.

#### Note:

You can modify the parameters of both source and read-only instances in the console in the same way as detailed below.

### **Use Limits**

To ensure instance stability, only some parameters can be modified in the console. These parameters are displayed on the **Parameter Settings** page.

If the modified parameter requires instance restart to take effect, the system will ask you if you want to restart. We recommend that you do so during off-peak hours and ensure that your application has a reconnection mechanism. If you want to return to the default formula, clear the entered parameters and apply.

### Modifying Parameters in Parameter List

#### **Batch modifying parameters**

1. Log in to the TencentDB for MySQL console. In the instance list, click an **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

- 2. Select Database Management > Parameter Settings and click Batch Modify Parameters.
- 3. UI for the source instance:



Database List	Parameter Settings	Account	Management				
Batch Modify Parar	neters Default Temp	NEW Cust	tom Template	Import Parameters	Export Parameters	Save as Template	
Parameter Name		Instanc	Default Value (i)	Current	t Value	Acce	eptable Va
automatic_sp_privileg	ges(i)	No	ON	ON		[ON	OFF]
avoid_temporal_upgr	ade(j)	No	OFF	OFF		[ON	OFF]
		Vec	3000	3000		[1_65]	5351

Ir	stance Details	Instance Monitoring	Databa	ise Management	Security Group	Backup and Restoration	Operation Log			
	Database List	Parameter Settin	ngs A	ccount Managem	ent					
	Batch Modify Para	ameters Default Te	mplate	Custom Template	Import Paramet	ters Export Parameters	Save as Template	NEW Intelligent Parameter Tuning		Recent Modificati
									Ent	er parameter name
	Parameter Name			Instance	Default Value	Current Value		Acceptable Values		
	auto_increment_inc	rement <sup>(i)</sup>		No	1	1		[1-65535]		
	auto_increment_offs	et(i)		No	1	1		[1-65535]		

4. Locate the desired parameters, and modify their values in the **Current Value** column. After confirming that everything is correct, click **Confirm Modification**.



Database List Paran	neter Settings Acco	ount Management				
Batch Modify Parameters	NEW Default Template	Custom Template	Import Parameters	Export Parameters	Save as Template	
Parameter Name	Instan	c Default Value	(i) Current	t Value	Acce	eptable Val
automatic_sp_privileges(i)	No	ON	ON	• ~ X	[ON	OFF]
avoid_temporal_upgrade(i)	No	OFF	OFF	•	[ON	OFF]
back_log(i)	Yes	3000	3000	~ ×	[1-655	535]

5. In the pop-up window, select an option in the Execution Mode and click OK.

#### Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately. If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see <u>Setting Instance Maintenance Window</u>.

#### Modifying one parameter

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. On the **Database Management** > **Parameter Settings** tab, locate the desired parameter in the parameter list and click

in the **Current Value** column.

Database List Para	neter Settings	Account M	anagement				
Batch Modify Parameters	NEW Default Template	Custor	m Template	Import Parameters	Export Parameters	Save as T	emplate
Parameter Name	In	istanc	Default Value(	i) Curre	nt Value		Acceptable Values
automatic_sp_privileges	No	)	ON	ON			[ON   OFF]
avoid_temporal_upgrade(i)	No	D	OFF	OFF	Click to modify the parameter	er value	[ON   OFF]
back_log	Ye	s	3000	3000			[1-65535]

3. Modify the value within the restrictions stated in the Acceptable Values column and click

to save the modification. You can click

#### $\times$

to cancel the operation.

Parameter Name	Restart	Default Value	Current Value	Acceptable Values
auto_increment_increment③	No	1		[1-65535]
auto_increment_offset	No	1	1	[1-65535]

4. In the pop-up window, select an option in the **Execution Mode** and click **OK**.

#### Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately. If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see <u>Setting Instance Maintenance Window</u>.

### Modifying Parameters by Importing Parameter Template

#### Option 1. Importing a parameter template on the Parameter Settings page

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. Select Database Management > Parameter Settings and click Custom Template. If you haven't configured a commonly used custom template yet, you can select Custom Template on the left sidebar in the TencentDB for MySQL console, click Create Template to configure a parameter template, and then import it from the custom template as described in this step.

Database List	Paran	neter Settings	Ac	count Management				
			NEW					
Batch Modify Parame	eters	Default Template		Custom Template	Import Parameters	Export Parameters	Save as Template	

3. In the pop-up window, select a parameter template and click Import and Overwrite Original Parameters.



port from P	arameter Template		
As the instand MySQL5.7 ca	ce version is <b>MySQL5.7</b> , only the parameter te in be selected.	mplate whose version is	
Select Parameter		<b>*</b>	
Template*			
	Import and Overwrite Original Parameters	Cancel	

4. After confirming that everything is correct, click **Confirm Modification**.

Instance Details	Instance Monitoring	Database Mana	agement	Security Group	Backup and	Restoration	Operation Log	Read-Only Ins
Database List	Parameter Settings	Account N	Managemer	ıt				
Confirm Modificat	ion Cancel							Enter parameter na
Parameter Name		Instanc	Default Va	alue(i)	Current Value			Acceptable Val
auto_increment_inc	rement(j)	No	1		1			[1-65535]
auto_increment_off	set(j)	No	1		1			[1-65535]
automatic_sp_privil	eges(j)	No	ON		ON	~		[ON   OFF]

5. In the pop-up window, select an option in the Execution Mode and click OK.

#### Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately. If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see <u>Setting Instance Maintenance Window</u>.

#### Option 2. Modifying parameters by importing a parameter configuration file

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. Select Database Management > Parameter Settings and click Import Parameters.



Instance Detai	ls Instance	Monitoring	Database	Management	Security Group	Backu	p and Restoration	Operation Log	Read-Only Instance
Database Li	st <b>Para</b>	neter Settings	Acc	count Managemer	ıt				
Batch Modify	Parameters	Default Tem	NEW plate	Custom Template	Import Parame	eters	Export Parameters	Save as Templa	ate

3. Click Select File to locate the desired parameter file and click Import and Overwrite Original Parameters.

Import fr	rom Local Files		×
Select a N	MySQL .cnf file (less than 1 MB)	Select File	
	Import and Overwrite Original Parameters	Cancel	

- 4. After confirming that everything is correct, click **Confirm Modification**.
- 5. In the pop-up window, select the **execution mode** and click **OK**.

#### Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately. If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see <u>Setting Instance Maintenance Window</u>.

#### Option 3. Importing a parameter template on the "Parameter Template" page

For more information, see Managing Parameter Template > Applying a Parameter Template to a Database.

### Restoring to Default Template

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. Select Database Management > Parameter Settings, click Default Template, select High-Stability

Template or High-Performance Template, and click Import and Overwrite Original Parameters.

Import from Par	mport from Parameter Template						
As high-perform parameters may	nance templates give priority to performance, y affect instance stability in some cases.	some of the ter	nplate				
Select Parameter Template*	High-Stability Template (Hot) High-Stability Template (Hot) High-Performance Template	•					
In	nport and Overwrite Original Parameters	Cancel					

3. Click **Confirm Modification** to redirect to the parameter modification confirmation window.

,				
<ul> <li>Notes</li> <li>Adjust Parameter (innodb_buffer_pool_size):</li> <li>innodb_buffer_pool_size automatically becomes an integer multiple of innodb_buffer_pool_chunk_size * innodb_buffer_pool_instances.</li> <li>innodb_buffer_pool_chunk_size is 128 MB.</li> </ul>				
u'll modify 32 parameters, of whi	ch 6 parameters will cause <b>instance</b>	restart (will take 50 sec). Are you su	Instance Restart	
arameter Name	Guilent value	New value	Instance Restart()	
			Vez	
character_set_server	utf8	LATIN1	Yes	
character_set_server nnodb_buffer_pool_instanc	utf8 1	LATIN1 {MIN(DBInitMemory/2000,16)}	Yes No instance restart required because the updated value calculated by the formula is th same as the current value.	
character_set_server nnodb_buffer_pool_instanc es	utf8 1	LATIN1 {MIN(DBInitMemory/2000,16)}	Yes No instance restart required because the updated value calculated by the formula is th same as the current value.	
character_set_server	utf8 1 536070012 execution During maintenanc	LATIN1 {MIN(DBInitMemory/2000,16)} (DDInitMemory/2000,16)} e time Learn More	Yes No instance restart required because the updated value calculated by the formula is th same as the current value.	

4. In the pop-up window, select an option in **Execution Mode**, read and indicate your consent to the **Restart Rules**, and click **OK**.

#### Note:

If you select **Immediate execution**, the parameter modification task will be executed and take effect immediately.

If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see Setting Instance Maintenance Window.

### Parameter Formula

You can use a formula to set the instance parameters. To do so, set the parameters related to the instance specification as a formula, and when the instance specification is changed, the parameter values in the formula will be dynamically changed accordingly and still take effect after the specification change. In this way, the instance is always in the optimal state for running business smoothly.

Taking the {DBinitMemory\\\*786432}value of the parameter innodb\_buffer\_pool\_sizeas anexample, when the DBinitMemoryin the instance specification is changed, the parameter configuration heredoesn't need to be modified, and the value of innodb\_buffer\_pool\_sizewill be changed automatically.

innodb_buffer_pool_size()	Yes	{DBInitMemory*786432}	{DBInitMemory*786432}	[[DBInitMemory*524288]-{DBInitMemory*94371
innodb_change_buffer_max_size③	No	25	Notes     The formula is as follows:     (DBInitMemory*x). where the "x"	[0.50]
innodb_change_buffering()	No	ALL	variables are positive integers	[NONE   INSERTS   DELETES   CHANGES   PL

#### Expression syntax is supported as follows:

Supported Type	Description	Sample
Variable	DBinitMemory: The memory size of the instance specification, which is measured in megabytes (MB) and is an integer value. For instance, if the memory size of the instance specification is 1,000 MB, then the value of DBinitMemory is 1024. DBInitCpu: The number of CPU cores in the instance specification, which is an integer value. For example, if the instance specification includes 8 cores, then the value of DBInitCpu is 8.	{DBinitMemory * 786432} = memory size (DBinitMemory) * percentage (75% by default) * 1024 * 1024 (unit conversion)
Operator	Formula syntax: It should be enclosed in braces ( {} ). Division operator (/): It divides the dividend by the divisor and returns an integer quotient. If the calculation result is a decimal number, only the integer part will be retained. Decimal numbers are not supported; for example, {MIN(DBInitMemory/4+500,1000000)} instead of {MIN(DBInitMemory0.25+500,1000000)} is supported. Multiplication operator (*): It multiplies two numbers and returns an integer product. If the calculation result is a decimal number, only the integer part will be retained. Decimal number calculation is not supported.	-

Function	MAX(): It returns the greatest value in an integer or parameter formula	{MAX(DBInitCpu/2,4)}
	list. MIN(): It returns the smallest value in an integer or parameter	
	formula list.	

#### Parameters that support parameter formulas

#### Note:

TencentDB for MySQL continuously optimizes parameter settings. The following lists only certain parameters that support parameter formula. You can learn more about parameter formulas in the parameter template in the console.

Parameter	Description	Default Formula
thread_pool_size	The number of thread groups in the thread pool. The default value means that the number of thread groups is the same as the number of CPU cores.	{MIN(DBInitCpu,64)}
table_open_cache_instances	The number of partitions where MySQL caches table handles.	{MIN(DBInitMemory/1000,16)}
table_open_cache	The size of the table descriptor, which can reduce the file open/close times.	{MAX(DBInitMemory*512/1000,2048)}
table_definition_cache	The number of opened table cache instances.	{MAX(DBInitMemory*512/1000,2048)}
max_connections	Max connections.	{MIN(DBInitMemory/4+500,100000)}
join_buffer_size	The minimum size of the buffer used for normal index scans, range index scans, and table joins that perform full-table scans.	{MIN(DBInitMemory*128,262144)}
innodb_write_io_threads	The number of I/O threads in InnoDB used for write operations.	{MAX(DBInitCpu/2,4)}
innodb_read_io_threads	The number of I/O threads in InnoDB used for read operations.	{MAX(DBInitCpu/2,4)}
innodb_buffer_pool_instances	The number of partitions in the InnoDB buffer pool.	{MIN(DBInitMemory/2000,16)}
innodb_buffer_pool_size	The size of the buffer pool in bytes, i.e., the memory zone where	{DBInitMemory*786432}
InnoDB caches tables and index data.

### Exporting Parameter Configuration as File

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. Select **Database Management** > **Parameter Settings** and click **Export Parameters** to export the parameter configuration file.



### Exporting Parameter Configuration as Template

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

 Select Database Management > Parameter Settings and click Save as Template to save the existing parameter configuration as a parameter template.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log Rea	d-Only Instan
Database List	Parameter Setting	Account Manageme	ent			
Batch Modify Par	ameters Default Ter	nplate Custom Templat	e Import Param	eters Export Parameters	Save as Template	

# Modifying Parameters During Custom Time Window

Before you confirm the parameter modification, the **Modify Parameters** dialog box will pop up for you to select a custom time window for the modification to take effect.

#### Note:

If you select **During maintenance time**, the parameter modification task will be executed and take effect during this time. For more information, see Setting Instance Maintenance Window.

Parameter Name	Current	/alue New	Value	Instance Restart(i)	
back log	3000	3001		Yes	
		0001		100	
ecution Mode	Immediate execution	During maintenance time	Learn More 🛂		

### Canceling Parameter Modification Task

If a parameter modification task (to be executed **during maintenance window**) has been submitted but you want to cancel it, you can select **Task List** on the left sidebar in the console, locate the task, and click **Cancel** in the **Operation** column. You can cancel a task only before it is executed. The task status should be **Waiting for execution**.

# Viewing Parameter Modification Log

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. On the **Database Management > Parameter Settings** tab, click **Recent Modifications** on the right.

Instance Details	Instance Monitoring	Database Ma	nagement Se	ecurity Group	Backup and Resto	re Operation Log	Read-only Instance	Data Encryption	Connection Check			
Database List	Parameter Settin	igs Manage	Account									
Batch Modify Par	ameters () Import	Parameters 🕄	Export Paramet	ters Impor	t from Template 🕄	Save as Template				Enter parameter name	Q	Recent Modifica
Parameter Name			Restart	Default Val	ie(j)	Current Value	Accep	table Values				
auto_increment_ind	crement(i)		No	1		2	[1-6553	5]				

3. You can view the recent parameter modification records here.

### Subsequent Operations



You can use templates to manage database parameters in batches. For more information, see Managing Parameter

Template.

For suggestions on the configuration of key parameters, see Suggestions on Parameter Settings.

# Managing Parameter Template

Last updated : 2024-07-22 15:41:39

Besides the system parameter templates provided by TencentDB for MySQL, you can create custom parameter templates in order to configure parameters in batches.

You can use the parameter template to configure and manage the parameters of a database engine. A template is like a container of the values of database engine parameters, which can be applied to one or more TencentDB instances. You can log in to the TencentDB for MySQL console and click **Parameter Templates** on the left sidebar to create, view, and manage parameter templates which support the following features:

Support default parameter templates.

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

Import the MySQL configuration file my.conf to generate templates.

Save parameter configurations as templates.

Import parameters from templates to apply to one or more instances.

#### Note:

If the parameters in the template are updated, the instance parameters are not updated unless they are manually reapplied to the instances.

You can apply the parameter changes to single or multiple instances by importing a template.

### Creating a Parameter Template

You can create a parameter template, modify the parameter values, and apply the template to instances.

1. Log in to the TencentDB for MySQL console, select **Parameter Templates** on the left sidebar, and click **Create Template**.

Create Template		
Template ID / Name	Database Version <b>T</b>	Template Description
	MySQL 5.6	33 3
	MySQL 5.7	44fo

2. In the pop-up dialog box, specify the following configurations, and click Create and Set Parameters.

**Template Name**: enter a unique template name.

Database Version: select a database version.



Template Description: enter a brief description of the parameter template.

1 Create Templa	ate > 2 Set Template Parameters	
Template Name *	Please enter template name	
Database Version *	Please select 💌	
Template Description	Describe the template	

3. After the creation is completed, you can modify, import, and export parameters on the template details page.

### Applying a Parameter Template to Instances

1. Log in to the TencentDB for MySQL console and select Parameter Templates on the left sidebar.

2. In the parameter template list, locate the desired template and click Apply to Instance in the Operation column.

Template ID/Name	Database Ver 🔻	Template Description	Template	Operation
Televis -	MySQL 5.7	44fo	Custom Template	View Details Apply to Instance Delete Export Compare

3. On the displayed page, specify the execution mode and instances, make sure that all parameter values are correct, and click **Submit**.

**Execution Mode**: **Immediate execution** is selected by default. If you select **During maintenance window**, the parameter modification task will be executed and take effect during the instance maintenance period.

MySQL Instance: select one or more instances that need to apply the parameter template in the specified region.

Parameter Comparison: view the changed parameter values of the selected instance.

#### Note:

Before applying a parameter template to multiple instances, please make sure the parameters can be applied to those instances.



Template ID / Name					
Database Version	MySQL 5.6				
Execution Mode	Immediate execution During maintenance window Learn More 🛛				
Region					
MySQL Instance	Available Instance		Sele	cted (0) Instances	
	Filter by instance ID / name	Q,	In	stance ID / Name Region	Restart Required (i)
	Instance ID / Name			Please select an instance	to apply template parameters from the left
	cdb-fnjiqnzi (Parameter error)				
			↔		
	Support for holding shift key down for multiple selection		Con	pare Again Remove All Instances	
Back Submit					

4. (Optional) If a parameter modification or batch modification task has been submitted but you want to cancel it, you can select **Task List** on the left sidebar in the console, locate the task, and click **Cancel** in the **Operation** column. You can cancel a task only before it is executed. The task state should be "Waiting for execution".

Task ID	Task Type 🔻	Instance ID	Task Progress	Task Status 🔻	Task Start Time	Task End Time	Operation
-	Parameter settings		0%	Waiting for execution			Task Details Cancel

### Copying a Parameter Template

To include most of the custom parameters and values of an existing parameter template in a new template, you can copy the existing template.

### Method 1. Copying an existing parameter template

1. Log in to the TencentDB for MySQL console, select **Parameter Templates** on the left sidebar, click the parameter template name or **View Details** in the **Operation** column in the template list, and enter the parameter template details page.

- 2. Select **More** > **Save as Template** at the top.
- 3. In the pop-up dialog box, specify the following configurations:

Template Name: enter a unique template name.

Template Description: enter a brief description of the parameter template.

4. After confirming that everything is correct, click **Save**.

### Method 2. Saving parameters of an instance as a parameter template

1. Log in to the TencentDB for MySQL console, select Instance List on the left sidebar, and click an instance

ID/name in the instance list to access the instance management page.

2. Select Database Management > Parameter Settings and click Save as Template.

3. In the pop-up dialog box, specify the following configurations:

**Template Name**: enter a unique template name.

Template Description: enter a brief description of the parameter template.

4. After confirming that everything is correct, click **Create and Save**.

## Modifying Parameter Values in a Parameter Template

1. Log in to the TencentDB for MySQL console, select **Parameter Templates** on the left sidebar, click the parameter template name or **View Details** in the **Operation** column in the template list, and enter the parameter template details page.

2. Click Batch Modify Parameters or click

### 1

in the Current Value column to modify the parameter value.

### Note:

If you click **Import Parameters**, you need to upload a local parameter configuration file. To avoid importing failures, the configuration file should be in the same format as the configuration file of the MySQL database server, or you can use the file template of the exported parameters.

Batch Modify Parameters	Import Parameters Ap	oly to Instances More 🔻		
Parameter Name	Restart Required 🛈	Default Value (i)	Current Value	Acceptable Valu
automatic_sp_privileges (j)	No	ON	on 🖋	[ON   OFF]
auto_increment_increment (j)	No	1	1	[1-65535]
auto_increment_offset (j)	No	1	1	[1-65535]

### Importing a Parameter Template

1. Log in to the TencentDB for MySQL console, select **Parameter Templates** on the left sidebar, click the parameter template name or **View Details** in the **Operation** column in the template list, and enter the parameter template details page.

2. Click Import Parameters.

#### Note:

If you click **Import Parameters**, you need to upload a local parameter configuration file. To avoid importing failures, the configuration file should be in the same format as the configuration file of the MySQL database server, or you can use the file template of the exported parameters.

3. In the pop-up dialog box, select a file to upload and click Import and Overwrite Original Parameters.

### Exporting a Parameter Template

#### Method 1

1. Log in to the TencentDB for MySQL console and select Parameter Templates on the left sidebar.

2. In the parameter template list, locate the desired template and click **Export** in the **Operation** column.

#### Method 2

1. Log in to the TencentDB for MySQL console, select **Parameter Templates** on the left sidebar, click the parameter template name or **View Details** in the **Operation** column in the template list, and enter the parameter template details page.

2. Select **More** > **Export Parameters** at the top.

### Deleting a Parameter Template

If a parameter template is created redundantly or no longer needed, it can be easily deleted.

1. Log in to the TencentDB for MySQL console and select **Parameter Templates** on the left sidebar.

2. In the parameter template list, locate the desired template and click **Delete** in the **Operation** column.

3. Click **OK** in the pop-up dialog box.

### See Also

For suggestions for configuration of key parameters, see Suggestions on Parameter Settings.

# **Comparison of Different Template Parameters**

Last updated : 2024-07-22 15:41:51

You can compare custom parameter templates with other custom templates or default templates to view their differences in the TencentDB for MySQL console.

### Directions

1. Log in to the TencentDB for MySQL console, select **Parameter Template** on the left sidebar, select the custom templates to be compared, and click **Compare**.

Custom Template	Default Template			
Create Template			Separate keywords with " "; press E	nter to separate filter tags
Template ID/Name	Database Version <b>T</b>	Template Description	Template Type	Operation
<b>5</b>	MySQL 5.7		Custom Template	View Details Apply to Instance Delete Expo Compare
-	MySQL 5.6		Custom Template	View Details Apply to Instance Delete Expo

2. In the pop-up dialog box, select the target templates to be compared through **Select Parameter Template**.

Parameter Comparison						
Select Template *	Please select	Ŧ				
Only preview	[Default]High-Stability Template (Hot)					
	[Default]High-Performance Template					

3. View the comparison result.



Parameter Comparison	
Select Template * [Default]High-Performance Template * Only preview changed parameters	
Parameter Name	y-test2
auto_increment_increment (j)	+ 12
back_log (j)	<del>3000</del> 210
binlog_cache_size (j)	<del>2097152</del> 32768
binlog_checksum (j)	CRC32 NONE
binlog_row_image (j)	FULE MINIMAL
eq_range_index_dive_limit (j)	<del>200</del> 10
innodb_autoinc_lock_mode (j)	2 1
innodb_disable_sort_file_cache (j)	OFF ON

Total items: 35

OK

# Performance Comparison of Parameter Templates

Last updated : 2024-07-22 15:42:05

# Test Tool

Sysbench 1.0.20 is the tool used to test the database benchmark performance.

### Tool installation

Run the following code to install Sysbench 1.0.20:

```
git clone https://github.com/akopytov/sysbench.git
git checkout 1.0.20
yum install gcc gcc-c++ autoconf automake make libtool bzr mysql-devel git
mysql
cd sysbench
./autogen.sh
./configure
make -j
make install
```

### Note:

The installation directions above apply to performance stress testing on a CentOS CVM instance. For directions on installing the tool on other operating systems, see the official Sysbench documentation.

### **Test Environment**

Туре	Description
Test instance specification	Three common specifications, namely, 4-core CPU and 8 GB memory, 8-core CPU and 32 GB memory, and 16-core CPU and 128 GB memory
Client configuration	64-core CPU and 128 GB memory
Client private network bandwidth	23 Gbps
Test data volume	Database instance memory * 1.2

Test database instance versions

**Tencent Cloud** 

5.6 20210630, 5.7 20210630, and 8.0 20210330

Note on client specification: High-spec client machines are used so as to ensure that the database instance

performance can be measured through stress testing on a single client. For low-spec clients, we recommend you use multiple clients for concurrent stress testing and aggregate the results.

Note on network latency: When performing the test, make sure that clients and database instances are in the same AZ to prevent the testing result from being affected by network factors.

### **Test Method**

### **Test data preparations**

```
sysbench --db-driver=mysql --mysql-host=xxxx --mysql-port=xxxx --mysql-
user=xxxx --mysql-password=xxxx --mysql-db=sbtest --table_size=xxxx --
tables=xxxx --events=0 --time=600 --threads=xxxx --percentile=95 --report-
interval=1 oltp_read_write prepare
```

### Command for performance stress testing

```
sysbench --db-driver=mysql --mysql-host=xxxx --mysql-port=xxxx --mysql-
user=xxxx --mysql-password=xxxx --mysql-db=sbtest --table_size=xxxx --
tables=xxxx --events=0 --time=600 --threads=xxxx --percentile=95 --report-
interval=1 oltp_read_write run
```

#### Descriptions of stress testing parameters:

oltp\_read\_write indicates to implement the OLTP test by calling the

/usr/share/sysbench/oltp\_read\_write.lua script.

--tables=xxxx indicates the number of tables in this test.

--table\_size=xxxx indicates the number of table rows in this test.

--threads=xxxx indicates the number of concurrent connections of the client in this test.

--report-interval=1 indicates that the test result is output once every second.

--percentile=95 indicates the sampling rate, which is 95% by default.

--time=600 indicates the execution time of this test, which is 600 seconds.

#### Scenario model

The test cases in this document all use the Lua script of sysbench.

For the common configurations, performance testing is conducted for different parameter templates. The test results are as follows:



### Test result

#### v5.6 20210630

CPU (Core)	Memory (GB)	Threads	Test Duration	Template	SysBench QPS	SysBench TPS	avg_lat
				Default template (disused)	34428.69	1721.43	18.59 ms
4	8	32	10 min	High-performance parameter template	35917.50	1795.87	17.82 ms
			High-stability parameter template	34834.04	1741.70	18.37 ms	
			Default template (disused)	61210.19	3060.51	20.91 ms	
8	32	64	10 min	High-performance parameter template	67719.55	3385.98	18.90 ms
				High-stability parameter template	64910.09	3245.50	19.72 ms
			28 10 min	Default template (disused)	106965.44	5348.27	23.93 ms
16	128	128 128		High-performance parameter template	127955.48	6397.77	20.00 ms
				High-stability parameter template	119509.02	5975.45	21.41 ms

#### v5.7 20210630

CPU (Core)	Memory (GB)	Threads	Test Duration	Template	SysBench QPS	SysBench TPS	avg_lat
4	8	32	10 min	Default template (disused)	34428.69	1721.43	18.59 ms
				High-performance parameter template		1795.87	17.82 ms
				High-stability parameter template	34834.04	1741.70	18.37 ms

8 32		64	10 min	Default template (disused)	fault template 61210.19 sused)		20.91 ms
	32			High-performance parameter template	67719.55	3385.98	18.90 ms
				High-stability parameter template	64910.09	3245.50	19.72 ms
16		28 128	10 min	Default template (disused)	106965.44	5348.27	23.93 ms
	128			High-performance parameter template	127955.48	6397.77	20.00 ms
				High-stability parameter template	119509.02	5975.45	21.41 ms

#### v8.0 20210330

CPU (Core)	Memory (GB)	Threads	Test Duration	Template	SysBench QPS	SysBench TPS	avg_lat
				Default template (disused)	32594.79	1629.74	19.63 ms
4	8 32	32	10 min	High-performance parameter template	33383.77	1669.19	19.17 ms
			High-stability parameter template	32071.90	1603.60	19.95 ms	
		Default template (disused)	65718.22	3285.91	19.47 ms		
8	32	64	64 10 min	High-performance parameter template	70195.37	3509.77	18.23 ms
				High-stability parameter template	60704.69	3035.23	21.08 ms
16	128	128	10 min	Default template (disused)	132023.66	6601.18	19.38 ms
			_	High-performance parameter template	151021.67	7551.08	16.95 ms



	High-stability parameter template	132391.01	6619.55	19.33 ms
--	-----------------------------------	-----------	---------	-------------

# Suggestions on Parameter Configuration

Last updated : 2024-07-22 15:42:17

Parameters in TencentDB for MySQL have been optimized on the basis of official default values in MySQL. We recommend that you configure the following parameters for your TencentDB for MySQL instance after purchase based on your business scenarios.

#### character\_set\_server

Default value: UTF8

Restart required: Yes

Description: Configure the default character set for the MySQL server. TencentDB for MySQL offers four character sets: LATIN1, UTF8, GBK, and UTF8MB4. Among them, LATIN1 supports English characters, with one character having one byte; UTF8 is generally used for international encoding that contains all characters used by all countries, with one character having three bytes; in GBK, every character has two bytes; and UTF8MB4 (a superset of UTF8) is completely backward compatible and supports emojis, with one character having four bytes. Recommendation: After purchasing an instance, you need to select the appropriate character set based on the data format required in your business to ensure that the client and the server use the same character set, preventing garbled text and unnecessary restarts.

#### lower\_case\_table\_names

Default value: 0

#### Restart required: Yes

Description: When creating a database or table, you can set whether storage and query operations are case-sensitive. This parameter can be set to 0 (case-sensitive) or 1 (case-insensitive). The default value is 0. Recommendation: TencentDB for MySQL is case-sensitive by default. We recommend that you configure this

parameter based on your business needs and usage habits.

### sql\_mode

Default values:

```
NO_ENGINE_SUBSTITUTION (v5.6); ONLY_FULL_GROUP_BY, STRICT_TRANS_TABLES,
NO_ZERO_IN_DATE, NO_ZERO_DATE, ERROR_FOR_DIVISION_BY_ZERO, NO_AUTO_CREATE_USER,
NO_ENGINE_SUBSTITUTION (v5.7)
```

#### Restart required: No

Description: TencentDB for MySQL can operate in different SQL modes, which define the SQL syntax and data check that it should support.

The default value of this parameter on v5.6 is NO\_ENGINE\_SUBSTITUTION, indicating that if the used storage engine is disabled or not compiled, an error will be reported.

On v5.7 and v8.0, the default values are ONLY\_FULL\_GROUP\_BY, STRICT\_TRANS\_TABLES ,

NO\_ZERO\_IN\_DATE, NO\_ZERO\_DATE, ERROR\_FOR\_DIVISION\_BY\_ZERO, NO\_AUTO\_CREATE\_USER, and NO\_ENGINE\_SUBSTITUTION.

#### Here:

ONLY\_FULL\_GROUP\_BY means that in a GROUP BY operation, the column in SELECT or the HAVING or ORDER BY subquery must be a function column that appears in or relies on GROUP BY.

STRICT\_TRANS\_TABLES enables strict SQL mode. NO\_ZERO\_IN\_DATE controls whether the server permits dates in which the year part is nonzero but the month or day part is zero. The effect of NO\_ZERO\_IN\_DATE depends on whether strict SQL mode is enabled.

NO\_ZERO\_DATE controls whether the server permits a zero date as valid. Its effect depends on whether strict SQL mode is enabled.

ERROR\_FOR\_DIVISION\_BY\_ZERO means that in strict SQL mode, if data is divided by zero during the INSERT or UPDATE process, an error rather than a warning will be generated, while in non-strict SQL mode, NULL will be returned.

NO\_AUTO\_CREATE\_USER prohibits the GRANT statement from creating a user with an empty password.

NO\_ENGINE\_SUBSTITUTION means that if the storage engine is disabled or not compiled, an error will be

#### reported.

Recommendation: As different SQL modes support different SQL syntax, we recommend that you configure them based on your business needs and development habits.

### long\_query\_time

Default value: 10

Restart required: No

Description: Used to define the time threshold for slow queries, with the default value as 10s. If a query execution takes 10s or longer, the execution details will be recorded in the slow log for future analysis.

Recommendation: As business scenarios and performance sensitivity may vary, we recommend that you set the value in consideration of future performance analysis.

# Intelligent Parameter Tuning

Last updated : 2023-11-15 16:03:25

This document describes how to implement intelligent parameter tuning in the TencentDB for MySQL console.

# Background

The concept of "deep learning" has been gaining popularity among general users. As relevant technologies are getting more mature, the TencentDB team is thinking about how to use deep learning to improve the operating efficiency of databases. The first thing that comes to mind is database parameter tuning. Due to the huge differences in business systems, it is impossible to perform targeted tuning in a fine-grained manner like optimizing SQL, which is a challenge for DBAs. They often have to draw upon their experience to build a set of better performing parameter templates. The ability to tune database parameters is more like an exclusive skill of expert DBAs.

Between 2019 and 2021, the TencentDB team published two papers entitled "Automatic Database Tuning Using Deep Reinforcement Learning" and "An Online Cloud Database Hybrid Tuning System for Personalized Requirements" respectively, and applied for international patents for the theory. Now, the theory is developed into a usable system based on the papers, which improves the database performance by tuning parameters in real-world use cases.

#### Why database parameter tuning?

High numbers of parameters: For example, MySQL has hundreds of configuration items, so tuning is difficult.
High labor costs: A full-time DBA is required, and expert experience is needed, which incur high labor costs.
Poor tool universality: Existing tools have limited features and involve time-consuming operations, producing a poor effect.

**New requirements in the cloud**: Some users don't have a full-time Ops team, making it impossible to implement parameter tuning.

### Prerequisites

You have a running TencentDB for MySQL instance.

### **Use Limits**

Scenario-based intelligent tuning can be performed only three times per instance per month. This limit is reset on the first day of each month.

Al-based intelligent analysis (coming soon) can be performed only once per instance per month. This limit is reset on the first day of each month.

Intelligent parameter tuning can be performed only for instances with at least four CPU cores.

The intelligent parameter tuning task list keeps only the last 15 tuning results.

When an instance is terminated/returned or expires, any ongoing intelligent parameter tuning task in it will be automatically aborted and deleted.

Only one tuning task can be executed per instance at any time.

The intelligent parameter tuning feature currently is supported only in Beijing, Shanghai, and Guangzhou regions and will be available in more regions in the future.

### Directions

### Previously purchased MySQL instance

1. Log in to the TencentDB for MySQL console and select a region at the top. 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 **Database Management** > **Parameter Settings** > **Intelligent Parameter Tuning**.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log R	Read-Only Instance Database
Database List	Parameter Settings	Account Management	:			
		NEW				
Batch Modify Par	ameters Default Tem	Custom Template	Import Parame	ters Export Parameters	Save as Template	Intelligent Parameter Tuning

3. In the Intelligent Parameter Tuning pop-up window, select Scenario-Based Intelligent Tuning or Al-Based Intelligent Analysis as the tuning method and click Analyze.



If you select **Scenario-Based Intelligent Tuning**, follow the steps below:

Scenario-based intelligent tuning: An efficient and targeted intelligent analysis based on your selected scenario.

Select a business scenario in the Scenario drop-down list, which can be Order Transactions, OLTP

#### Performance Test, or Pressure Test.

After selecting a scenario, you can customize the business ratio in the scenario for more accurate system analysis.

Then, click **Analyze**.

Order Transactions (TPCC)

Customizable items: Order processing (high), Payment (high), Order query (low), Logistics (low), Warehousing (low)

Data Reading Mode: Cache (default), Disk + Cache

Concurrency: Low, Middle, High (default)

OLTP Performance Test (Sysbench)

Customizable items: Reading data (high), Writing data (none by default)

Data Reading Mode: Cache (default), Disk + Cache

Concurrency: Low, Middle, High (default)

Pressure Test (myslap)

Concurrency: Low, Middle, High (default)

cenario	Order Transactions 🔻		
rder processina	Order Transactions		0
	L OLTP Performance Test	Middle	High
ayment	Pressure Test Zero Low	Middle	High
rder query	Zero Low	Middle	High
ogistics	Zero Low	Middle	High
'arehousing		Middle	High
ata Reading Mode	Cache Disk + Cache	Middle	ngn
oncurrency	Low Middle High		

If you select AI-Based Intelligent Analysis (coming soon), follow the steps below:

**Al-Based Intelligent Analysis**: The database business type is determined through in-depth analysis of database operating metrics, the performance of different parameters in specific scenarios is analyzed with deep learning algorithms, and then parameter suggestions are presented.

After selecting AI-Based Intelligent Analysis, click Analyze.

Note:

The AI-based intelligent analysis is under improvement and will be available soon.

This is a time-consuming intelligent analysis based on deep learning algorithm and big data analytics. We recommend you perform it during off-peak hours.

4. The parameter tuning task starts when the analysis starts. You can select Intelligent Parameter Tuning > View
 Task on the Parameter Settings page to view the task details.

Intelligent Paran	neter Tuning		View Ta
	Scenario-Based Intelligent Tuning		Al-Based Intelligent Analysis
	An efficient and targeted intelligent analysis based on your selected scenario		coming soon

5. After the parameter tuning task is completed, you can click **View Results** in the **Operation** column in **Intelligent Parameter Tuning** > **View Task**.

Task List				:
Intelligent Analysis T	Scenario Type	Task Status	Completion Time	Operation
Sc .	Pressure Test	Successful	2022-06-13 11:28:19	View Results
Sc .	Pressure Test 🛈	Successful	2022-06-13 11:24:51	View Results

6. After confirming the parameter tuning recommendations, click **Apply to Instance**.

Farameter running Suggest	on			
Instance Specs Before Analysis	General-4core8000MBMEM, 200	)GBStorage Space, MySQL8.0		
Current Specs	General-4core8000MBMEM, 200	)GBStorage Space, MySQL8.0		
QPS Increase	69%			
RT Decrease	49%			
TPS Increase	69%			
Concurrency	High			
Parameter Name	Instance Restart	Original Value	Recommended Value	
binlog_order_commits	No	ON	OFF	
eq_range_index_dive_limit	No	200	200	
innodb_adaptive_flushing	No	ON	ON	
innodb_adaptive_flushing_lwm	No	10	70	

7. Confirm the parameter change in the pop-up window, select the **Execution Mode**, read and indicate your consent to the **Restart Rules**, and click **OK**.

Execution mode:

Immediate execution: The change will be applied to the instance immediately after the confirmation.

During maintenance time: The change will be applied to the instance during the maintenance time, which can be modified on the **Instance Details** page.

### Newly purchased MySQL instance

When purchasing a TencentDB for MySQL instance, after selecting a parameter template, you can choose whether to enable **Scenario-Based Intelligent Tuning**. If this feature is enabled, the system will perform secondary adjustment according to the selected business scenario, which can be **Order Transactions**, **OLTP Performance Test**, or **Pressure Test**.

You can view the change results in Parameter Settings > Intelligent Parameter Tuning > View Task.

Parameter Template	[Default]High-Stability Template (Hot)	v	G		
	Scenario-Based Intelligent Tuning	Order 1	ransactions	×	0
	If "Scenario-Based Intelligent Tuning" is selected	ed, instan	ce parameters will be tu	ined	after the parameter template is applied. Learn More $\slash$

# Modification Records of System Parameter Templates

Last updated : 2024-12-30 17:48:56

TencentDB for MySQL provides system parameter templates for batch parameter settings. The parameters in system parameter templates may be optimized and updated with MySQL iteration. This document describes the update history of the parameters in system parameter templates.

#### Note:

If the parameters in the template are updated, the instance parameters are not updated unless they are manually reapplied to the instances. You can apply the parameter changes to single or multiple instances by importing a template. For directions on how to use system parameter templates, see Managing Parameter Template.

### December 2024

Parameter name	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
innodb_stats_persistent_sample_pages	√	✓	✓	Involved template: high- stability/high-performance Involved engine: InnoDB <b>Before change</b> Default Value: 20 <b>After change</b> Default Value: 200

### November 2024

Parameter name	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
explicit_defaults_for_timestamp	✓	✓	✓	Involved template: high- stability/high-performance Involved engine: InnoDB and RocksDB <b>Before change</b> Whether to restart to take effect: Yes



				After change Whether to restart to take effect: No
innodb_page_cleaners	-	✓	✓	Involved template: high- stability/high-performance Involved engine: InnoDB <b>Before change</b> Default value: 4 <b>After change</b> Default value: 8

# September 2024

Parameter Name	MySQL 5.7	MySQL 8.0	Change Description
innodb_table_drop_mode	✓	✓	Involved template: high-stability Involved engine: InnoDB <b>Before change</b> Default value: SYNC_DROP <b>After change</b> Default value: ASYNC_DROP

### March 2024

Parameter Name	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
slave_parallel_workers	s	✓	✓	Involved template: High-performance/high- stability Involved engines: InnoDB/RocksDB <b>Before change</b> Parameter range: [0   1   2   4   8   16   32   64] Default value: 0 <b>After change</b> Parameter range: [0-64] Default value: {MIN(DBInitCpu*2,64)}

# January 2024



Parameter	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
thread_pool_oversubscribe	✓	<b>√</b>	✓	Changes in default value of thread_pool_oversubscribe parameter in high-performance parameter templates for MySQL 5.6, 5.7, 8.0. Previous default value: 10 Updated default value: 16

### December 2023

Parameter	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
thread_pool_oversubscribe	\$	\$	\$	The default value for the parameter 'thread_pool_oversubscribe' has been modified. Previous default value: 3 Updated default value: 10

## August 2023

Parameter	MySQL 5.7	MySQL 8.0	Change Description
innodb_fast_ahi_cleanup_for_drop_table	\$	✓	The default value for the parameter 'innodb_fast_ahi_cleanup_for_drop_table' has been altered. Previous default value: OFF Updated default value: ON

### June 2023

Parameter	MySQL 5.5	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
time_zone	1	1	1	5	Modification of default value for the 'time_zone' parameter



		Previous default value: SYSTEM
		Updated default value: +8:00

# May 2023

Parameter	MySQL 5.7	MySQL 8.0	Change Description
innodb_async_truncate_work_enabled	\$	-	Asynchronous table deletion function switch parameter modification for version 5.7 The parameter has been changed from innodb_async_truncate_work_enabled to innodb_table_drop_mode. (May 15, 2023)
innodb_parallel_ddl	-	1	Parameter deprecation (May 5, 2023)

### March 2023

Parameter	MySQL 5.7	MySQL 8.0	Change Description
innodb_temp_data_file_path	1	\$	Modification of default value for innodb_temp_data_file_path Parameter Prior default value : ibtmp1:12M:autoextend Altered default value : ibtmp1:12M:autoextend:max:1048576M

## January 2023

Parameter	MySQL 5.7	MySQL 8.0	Change Description
max_execution_time	✓	1	Dynamic setting is supported for <pre>max_execution_time Restart required: No Default value: 0 Value range: 0-4294967295 Unit: ms</pre>



## August 2022

Parameter	MySQL 5.7	MySQL 8.0	Change Description
innodb_buffer_pool_size	√	✓	Dynamic setting is supported for innodb_buffer_pool_size Restart required: No Default value: {DBInitMemory * 786432} Value range: {DBInitMemory * 524288} - {DBInitMemory * 943718} DBinitMemory: An integer of instance memory size

# July 2022

Parameter	MySQL 5.7	MySQL 8.0	Change Description
innodb_temp_data_file_path	<b>√</b>	✓	Parameter modification is supported for innodb_temp_data_file_path (temp tablespace size) and the parameter properties are as follows: Restart required: Yes Default value: ibtmp1:12M:autoextend Value range of ibtmp : 12 MB - 1024 MB. After you select autoextend, the maximum value can be set to 2097152 MB

### March 2022

Parameter	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
innodb_open_files	1	1	1	Removed
innodb_stats_sample_pages	-	1	1	Removed
wait_timeout	1	1	1	New value range: 1-31536000
thread_cache_size	1	1	1	New value range: 1-16384

### December 2021

Parameter	MySQL 5.6	MySQL 5.7	MySQL 8.0	Change Description
binlog_row_image	\$	✓	1	Default value: FULL . For the early created instances, the default value is MINIMAL . Manual modification is supported.

### November 2020

Parameter	MySQL 8.0	Change Description
iinnodb_flush_log_at_trx_commit	1	Parameter added
sync_binlog	1	Parameter added
local_infile	1	Parameter added
innodb_log_file_size	1	Parameter added
cdb_recycle_bin_enabled	1	Parameter added
binlog_format	1	New value range: row
innodb_autoinc_lock_mode	1	New default value: 2
table_open_cache	1	New default value: 2000
slave_pending_jobs_size_max	1	New default value: 1073741824
time_zone	✓	New value range: [SYSTEM -12:00 -11:00 -10:00 -09:00 -08:00  -05:00 -04:00 -03:00 -02:00 -01:00 +00:00 +01:00 +02:00 +03:0 +05:30 +06:00 +06:30 +07:00 +08:00 +09:00 +10:00 +11:0
max_connections	✓	New value range: 1-100000
slave_rows_search_algorithms	✓	New default value: TABLE_SCAN, INDEX_SCAN, HASH_SCAN
innodb_open_files	✓	New default value: 10240
slave_parallel_type	✓	New value range: LOGICAL_CLOCK TABLE DATABASE



# August 2020

Parameter	MySQL 5.6	MySQL 5.7	Change Description
log_warnings	1	1	Parameter added
innodb_flush_log_at_trx_commit	1	1	Parameter added
sync_binlog	1	1	Parameter added
local_infile	1	1	Parameter added
innodb_log_file_size	1	1	Parameter added
binlog_format	1	1	New range value: row
innodb_autoinc_lock_mode	1	1	New default value: 2
innodb_open_files	1	1	New value range: 1–102400
table_open_cache	1	1	New default value: 2000
slave_pending_jobs_size_max	1	1	New default value: 1 GB
time_zone	1	1	New value range: [SYSTEM -12:00 -11:00 -10:0 -05:00 -04:00 -03:00 -02:00 -01:00 +00:00 +01:0 +05:30 +06:00 +06:30 +07:00 +08:00 +09:00 +1
max_connections	1	1	New value range: 1-100000
cdb_more_gtid_feature_supported	-	1	All kernel features supported
cdb_more_gtid_feature_supported	✓	-	New default value: off
slave_parallel_workers	-	1	All kernel features supported
tls_version	-	1	Removed
slave_rows_search_algorithms	1	1	New default value: TABLE_SCAN, INDEX_SCAN
innodb_open_files	1	✓	New default value: 10240

## August 2020



Parameter	MySQL 5.5	Change Description
innodb_autoinc_lock_mode	✓	New default value: TABLE_SCAN, INDEX_SCAN, HASH_SCAN
innodb_open_files	✓	New default value: 10240

# Network and Security Access Management Overview

Last updated : 2024-07-22 15:54:11

### Issues

If you have multiple users managing different Tencent Cloud services such as CVM, VPC, and TencentDB, and they all share your Tencent Cloud account access key, you may face the following problems:

Your key will be easily compromised because it is shared by several users.

You cannot restrict the access from other users and your service will be vulnerable to the security risks caused by their maloperations.

### Solution

You can avoid the problems above by allowing different users to manage different services through sub-accounts. By default, a sub-account does not have permissions to use Tencent Cloud services or resources. Therefore, you need to create a policy to grant different permissions to the sub-accounts.

Cloud Access Management (CAM) is a web-based Tencent Cloud service that helps you securely manage and control access permissions to your Tencent Cloud resources. Using CAM, you can create, manage, and terminate users (groups), and control the Tencent Cloud resources that can be used by the specified user through identity and policy management.

When using CAM, you can associate a policy with a user or user group to allow or forbid them to use specified resources to complete specified tasks. For more information on CAM policies, see Element Reference. You can skip this section if you do not need to manage permissions to TencentDB resources for sub-accounts. This will not affect your understanding and use of the other sections of the document.

### **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:

We recommend that you manage TencentDB resources and authorize TencentDB operations through CAM policies. Although the user experience does not change for existing users who are granted permissions by project, we do not recommend that you continue to manage resources and authorize operations in a project-based manner.

Effectiveness conditions cannot be set in TencentDB for the time being.

Task	Link
Basic policy structure	Policy syntax
Operation definition in a policy	Operations in TencentDB
Resource definition in a policy	Resources in TencentDB
Resource-level permission supported by TencentDB	List of APIs supporting authorization at resource level
Console examples	Console examples

# Authorization Policy Syntax

Last updated : 2024-07-22 15:54:25

# Policy syntax

### CAM policy:

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

version is required. Currently, only the value "2.0" is allowed.

**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.

 $effect \ is \ required. \ It \ describes \ whether \ the \ declaration \ result \ is \ \ \texttt{allow} \ \ or \ explicit \ \ \ deny \ .$ 

**action** is required. It specifies whether to allow or deny the operation. The operation can be an API (prefixed with "cdb:").

**resource** is required. It describes the details of authorization. A resource is described in a six-segment format. Detailed resource definitions vary by product.

**condition** is required. It 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.

# Operations in TencentDB

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

```
cdb:CreateAccounts .
```

To specify multiple operations in a single statement, separate them by comma.

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

You can also specify multiple operations by using a wildcard. For example, you can specify all operations beginning with "Describe" in the name.

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

To specify all operations in TencentDB, use the \* wildcard.

```
"action":["cdb:*"]
```

## Resources in TencentDB

Each CAM policy statement has its own applicable 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 cdb .

region describes the region information, such as ap-guangzhou .

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

**resource** describes detailed resource information of each product, such as <code>instanceId/instance\_id1</code> or <code>instanceId/\*</code>.

For example, you can specify a resource for a specific instance (cdb-k05xdcta) in a statement.

"resource":[ "qcs::cdb:ap-guangzhou:uin/65xxx763:instanceId/cdb-k05xdcta"]

You can also use the wildcard "\*" to specify it for all instances that belong to a specific account.

"resource":[ "qcs::cdb:ap-guangzhou:uin/65xxx763:instanceId/\*"]

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.

"resource": ["\*"]

To specify multiple resources in one policy, separate them by comma. In the following example, 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, where words prefixed with s are placeholders, region refers to a region, and account refers to a naccount ID.

Resource	Resource Description Method in Authorization Policy	
Instance	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>	
VPC	<pre>qcs::vpc:\$region:\$account:vpc/\$vpcId</pre>	
Security group	<pre>qcs::cvm:\$region:\$account:sg/\$sgId</pre>	

## Authorizable Resource Types

Last updated : 2024-07-22 15:54:38

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

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

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. You can use the 💉 wildcard in a resource path when defining it.

## List of APIs supporting resource-level authorization

API	Resource Path
AddTimeWindow	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
AssociateSecurityGroups	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
CloseWanService	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
CreateAccounts	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
CreateBackup	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
CreateDBImportJob	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DeleteAccounts	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>

DeleteBackup	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
	qcs::cdp:\$region:\$account:instanceid/\$instanceid
DeleteTime/Mindow	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Delete I Imewindow	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Deparibe A coount Drivilageo	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeAccountenvileges	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Depariba	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Describeaccounts	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DepariboBackupConfig	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribebackupConlig	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribePackupDatabases	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeBackupDatabases	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeBackupDownloadDbTableCode	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeRackups	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Describebackups	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeBackunTables	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Describedackup rabies	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeBinlogs	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Describediniogs	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeDatabases	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeDatabases	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeDBImportBecords	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Described Dimport records	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Described Dinstance Onarset	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeDBInstanceConfig	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Described Binstanee Oolning	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeDBInstanceCTID	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
שפאושכשטווזגמוונפט וש	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DescribeDBInstanceRebootTime	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>

🔗 Tencent Cloud

	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeDBSwitchRecords	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DeseribeDDCesurityCroups	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeDBSecurityGroups	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Descuibelesters - Deven Descuide	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeinstanceParamRecords	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeinstanceParams	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeRoGroups	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeRollbackRange I ime	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeSlowLogs	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeSupportedPrivileges	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeTables	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeTimeWindow	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeDatabasesForInstances	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DescribeMonitorData	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
Describe l'ableColumns	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
DropDatabase I ables	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
InitDBInstances	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>

IsolateDBInstance	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
ModifyAccountDescription	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModifyAccountDescription	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
ModifyAccountPaceword	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
wounyAccountr assword	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
ModifyAppountPrivilages	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
woonyAccountentineges	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Madifu AutoDonouvElan	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModilyAutoRenewFlag	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
MadifuDaaluurCanfar	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
моопуваскирсопіід	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Madife De alevertata	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
моатуваскиріпто	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Mall'( DDI ale a Alexa	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModifyDBInstanceName	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Madifu DDInatana Draigat	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModilyDBInstanceProject	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
MadifuDDInatanaaCaaurituCrauna	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModilyDBinstanceSecurityGroups	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Madifu DD lugate regel (in) (new)	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModilyDBInstancevipvport	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
Madifulationan Davara	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
MoonymstanceParam	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
MadifuDDInatanaaMadaa	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModifyDBInstanceModes	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
ModifyTimoWindow	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
ModifyTimevindow	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
MadifuDrataatMada	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
woullyProtectiviode	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
OfflineDBInstances	<pre>qcs::cdb:\$region:\$account:instanceId/*</pre>
	<pre>qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>



OpenDBInstanceGTID	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
OpenWanService	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
ReleaseIsolatedDBInstances	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
RestartDBInstances	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
StartBatchRollback	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
SubmitBatchOperation	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
SwitchDrInstanceToMaster	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
SwitchForUpgrade	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
DisassociateSecurityGroups	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
UpgradeDBInstance	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>
UpgradeDBInstanceEngineVersion	<pre>qcs::cdb:\$region:\$account:instanceId/* qcs::cdb:\$region:\$account:instanceId/\$instanceId</pre>

## List of APIs not supporting resource-level authorization

For TencentDB API operations that don't support resource-level authorization, you can still authorize a user to perform them, but you must specify \* as the resource element in the policy statement.

API	Description
CreateDBInstance	Creates a monthly subscribed TencentDB instance
CreateDBInstanceHour	Creates a pay-as-you-go TencentDB instance

CreateParamTemplate	Creates a parameter template
DeleteParamTemplate	Deletes a monitoring template item
DescribeProjectSecurityGroups	Queries the security group information of a project
DescribeDefaultParams	Queries the list of default configurable parameters
DescribeParamTemplateInfo	Queries the details of a parameter template
DescribeParamTemplates	Queries the list of parameter templates
DescribeAsyncRequestInfo	Queries the execution result of an async task
DescribeTasks	Queries the list of tasks for a TencentDB instance
DescribeUploadedFiles	Queries the list of imported SQL files
ModifyParamTemplate	Modifies a parameter template
RenewDBInstance	Renews a TencentDB instance
StopDBImportJob	Stops a data import task
DescribleRoMinScale	Queries the minimum specification supported by a read-only instance
DescribeRequestResult	Queries the details of a task
DescribeRoMinScale	Queries the minimum specification for read-only instance purchase or upgrade

## **Console Examples**

Last updated : 2024-07-22 15:54:53

This document provides examples about how to grant a user permissions to view and use specific resources in the TencentDB console by using a CAM policy.

## Full access policy for TencentDB

To grant a user permissions to create and manage TencentDB instances, you can implement the

QcloudCDBFullAccess policy for the user.

Log in to the CAM console, select **Policies** on the left sidebar, and search QcloudCDBFullAccess in the upper right corner.

O Associate users or user groups with policies to grant permissions.			
Create Custom Policy Delete			QcloudCDBFullAccess
Policy Name	Description	Service Type 🔻	Action
QcloudCDBFullAccess	Full read-write access to TencentDB, including permissions for Tencent	Cloud Database	Associate

The policy syntax is as follows:

```
{
    "statement": [
        {
             "action": [
                 "cdb:*"
             ],
             "effect": "allow",
            "resource": "*"
        },
        {
             "action": [
                 "vpc:*"
             ],
             "effect": "allow",
             "resource": "*"
        },
        {
```



```
"action": [
           "cvm:*"
        ],
        "effect": "allow",
        "resource": "qcs::cvm:::sg/*"
    },
    {
        "action": [
           "cos:*"
        ],
        "effect": "allow",
        "resource": "*"
    },
    {
        "action": [
            "monitor:*",
            "cam:ListUsersForGroup",
            "cam:ListGroups",
            "cam:GetGroup"
        ],
        "effect": "allow",
        "resource": "*"
    },
    {
        "action": [
            "kms:CreateKey",
            "kms:GenerateDataKey",
            "kms:Decrypt",
            "kms:ListKey"
        ],
        "effect": "allow",
        "resource": "*"
    }
],
"version": "2.0"
```

The above policy achieves its goal by allowing the user to separately authorize the use of TencentDB, VPC, security group, COS, KMS, and all resources available in the monitor with the CAM policy.

## Read-only permission policy for TencentDB

To grant a user permission to view TencentDB instances but not create, delete, or modify them, you can implement the policy named <code>QcloudCDBInnerReadOnlyAccess</code> for the user.

}

#### Note:

We recommend that you configure the read-only policy for TencentDB.

Log in to the CAM console, select Policies on the left sidebar, click Service Type in the policy list and select

TencentDB for MySQL in the drop-down list, and then you can see this policy in the results.

The policy syntax is as follows:

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

## Read-only permission policy for TencentDB-related resources

To grant a user permissions to view TencentDB instances and related resources (VPC, security groups, COS, and Cloud Monitor) but not create, delete, or modify them, you can implement the <code>QcloudCDBReadOnlyAccess</code> policy for the user.

Log in to the CAM console, select **Policies** on the left sidebar, click **Service Type** in the policy list and select **TencentDB for MySQL** in the drop-down list, and then you can see this policy in the results.

The policy syntax is as follows:

```
{
    "statement": [
        {
            "action": [
               "cdb:Describe*"
        ],
            "effect": "allow",
            "resource": "*"
        },
        {
            "action": [
            "monitor:Describe*",
            "monitor:Get*",
            "cam:ListUsersForGroup",
        }
}
```



```
"cam:ListGroups",
"cam:GetGroup"
],
"effect": "allow",
"resource": "*"
}
],
"version": "2.0"
}
```

The above policy achieves its goal by allowing the user to separately authorize the use of the following operations with the CAM policy.

All operations in TencentDB that begin with "Describe".

All operations in VPC that begin with "Describe", "Inquiry", or "Get".

All operations in security groups that begin with "DescribeSecurityGroup".

All operations in COS that begin with "List", "Get", and "Head" as well as the "OptionsObject" operation. All operations in the monitor.

# Policy for granting a user permissions to use APIs not at the resource level

To grant a user permissions to use only APIs not at the resource level, you can implement the

QcloudCDBProjectToUser policy for the user.

Log in to the CAM console, select **Policies** on the left sidebar, click **Service Type** in the policy list and select **TencentDB for MySQL** in the drop-down list, and then you can see this policy in the results.

The policy syntax is as follows:

```
{
    "version": "2.0",
    "statement": [
        {
            "action": [
                "cdb:BalanceRoGroupLoad",
                "cdb:CreateBatchOperation",
                "cdb:CreateDBInstance",
                "cdb:CreateDBInstanceHour",
                "cdb:CreateDBInstanceHour",
                "cdb:CreateParamTemplate",
                "cdb:DeleteBatchJobFiles",
                "cdb:DeleteMonitorTemplate",
                "cdb:DeleteParamTemplate",
                "cdb:DeleteP
```



```
"cdb:DescribeBatchJobFileContent",
                "cdb:DescribeBatchJobFiles",
                "cdb:DescribeBatchJobInfo",
                "cdb:DescribeProjectSecurityGroups",
                "cdb:DescribeDefaultParams",
                "cdb:DescribeMonitorTemplate",
                "cdb:DescribeParamTemplateInfo",
                "cdb:DescribeParamTemplates",
                "cdb:DescribeRequestResult",
                "cdb:DescribeRoGroupInfo",
                "cdb:DescribeRoMinScale",
                "cdb:DescribeTasks",
                "cdb:DescribeUploadedFiles",
                "cdb:ModifyMonitorTemplate",
                "cdb:ModifyParamTemplate",
                "cdb:ModifyRoGroupInfo",
                "cdb:ModifyRoGroupVipVport",
                "cdb:StopDBImportJob",
                "cdb:UploadSqlFiles"
            ],
            "effect": "allow",
            "resource": "*"
        }
   ]
}
```

# Policy for granting a user permissions to manipulate a specific TencentDB instance

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

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

# Policy for granting a user permissions to manipulate TencentDB instances in batches

To grant a user permissions 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 "cdb-xxx" and "cdb-yyy" in Guangzhou region and "cdb-zzz" in Beijing region.

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

# Policy for granting a user permissions to manipulate TencentDB instances in a specific region

To grant a user permissions to manipulate TencentDB instances in a specific region, associate the following policy with the user. This policy authorizes the user to operate TencentDB instances in Guangzhou region.

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

## Custom policy

If preset policies cannot meet your requirements, you can create custom policies as shown below. If permissions are granted by resources, for a TencentDB API operation that does not support authorization at the resource level, you can still authorize a user to perform it, but you must specify \* as the resource element in the policy statement. The syntax of custom policies is as follows:

```
{
    "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".

## **TencentDB Security Group Management**

Last updated : 2024-07-22 15:55:06

## Overview

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 with the same network security isolation demands in one region can be put into the same security group, which is a logical group. TencentDB and CVM share the security group list and are matched with each other within the security group based on rules. For specific rules and limitations, see Security Group Overview.

#### Note:

TencentDB for MySQL security groups currently only support network access control for VPCs and public networks but not the classic network.

As TencentDB doesn't have any active outbound traffic, outbound rules don't apply to it.

Security groups are supported for source, read-only, and disaster recovery TencentDB for MySQL instances.

## Configuring a security group for TencentDB

### Step 1. Create a security group

#### 1. Log in to the CVM console.

2. Select **Security Group** on the left sidebar, select a region, and click **Create**.

3. In the pop-up dialog window, configure the following items, and click OK.

Template: Select a template based on the service to be deployed on the TencentDB instance in the security group,

which simplifies the security group rule configuration, as shown in the table below.

Template	Description	Description
Open all ports	All ports are opened to the public and private networks. This may pose security issues.	-
Open ports 22, 80, 443, and 3389 and the ICMP protocol	Ports 22, 80, 443, and 3389 and the ICMP protocol are opened to the public network. All ports are opened to the private network.	This template doesn't take effect for TencentDB.
Custom	You can create a security group and then add custom rules. For more information, see Add a security group rule.	-

Name: Custom name of the security group.

Project: Select a project for easier management. By default, DEFAULT PROJECT is selected.

**Remarks**: A short description of the security group for easier management.

### Step 2. Add a security group rule

1. On the Security Group page, click **Modify Rule** in the **Operation** column on the row of the security group for which to configure a rule.

2. On the security group rule page, click **Inbound Rules** > **Add Rule**.

3. In the pop-up window, set the rule.

**Type: Custom** is selected by default. You can also choose another system rule template. MySQL(3306) is recommended.

**Source** or **Target**: Traffic source (inbound rules) or target (outbound rules). You need to specify one of the following options:

Source or Target	Description
A single IPv4 address or an IPv4 range	In CIDR notation, such as 203.0.113.0, 203.0.113.0/24 or 0.0.0.0/0, where 0.0.0.0/0 indicates all IPv4 addresses will be matched.
A single IPv6 address or an IPv6 range	In CIDR notation, such as FF05::B5, FF05:B5::/60, ::/0 or 0::0/0, where ::/0 or 0::0/0 indicates all IPv6 addresses will be matched.
ID of referenced security group. You can reference the ID of: Current security group Other security group	Current security group indicates the CVM associated with the security group. Other security group indicates the ID of another security group under the same project in the same region.
Reference an IP address object or IP address group object in a parameter template.	-

**Protocol Port**: Enter the protocol type and port range or reference a protocol/port or protocol/port group in a parameter template.

#### Note:

To connect to a TencentDB for MySQL instance, you must open its port. You can log in to the TencentDB for MySQL console, click an instance ID to enter the instance details page.



Private Network Address	3306 🖍 🗗 One-Click Connectivity Check
Public Network Address	58295 Close To One-Click Connectivity Check

TencentDB for MySQL uses private network port 3306 by default and supports customizing the port. If the default port is changed, the new port should be opened in the security group.

The TencentDB for MySQL public port is automatically assigned by the system and cannot be customized. After the public network access is enabled, it will be controlled by the ACL of the security group. When configuring the security policy, you need to open the private port 3306.

The security group rules displayed on the **Security Group** page in the TencentDB for MySQL console take effect for private and public (if enabled) network addresses of the TencentDB for MySQL instance.

Policy: Allow or Reject. Allow is selected by default.

**Allow**: Traffic to this port is allowed.

Reject: Data packets will be discarded without any response.

**Remarks**: A short description of the rule for easier management.

4. Click **Complete**.

#### Use case

Scenario: You have created a TencentDB for MySQL instance and want to access it from a CVM instance.

Solution: When adding security group rules, select MySQL(3306) in Type to open port 3306.

You can also set **Source** to all or specific IPs (IP ranges) as needed to allow them to access TencentDB for MySQL from a CVM instance.

| Inbound or Outbound | Type | Source | Protocol and Port | Policy |

| Inbound | MySQL(3306) | All IPs: 0.0.0.0/0

Specific IPs: Specify IPs or IP ranges | TCP:3306 | Allow |

### Step 3. Configure a security group

A security group is an instance-level firewall provided by Tencent Cloud for controlling inbound traffic of TencentDB. You can associate a security group with an instance when purchasing it or later in the console.

#### Note:

Currently, security groups can be configured only for TencentDB for MySQL instances in VPC.

1. Log in to the TencentDB for MySQL console, click an instance ID in the instance list, and enter the instance management page.

2. On the Security Group tab, click Configure Security Group.

3. In the pop-up dialog box, select the security group to be bound and click OK.

## Importing security group rules

- 1. On the Security Group page, click the ID/name of the target security group.
- 2. On the inbound rule or outbound rule tab, click **Import Rule**.
- 3. In the pop-up window, select an edited inbound/outbound rule template file and click Import.

### Note:

As existing rules will be overwritten after importing, we recommend that you export the existing rules before importing new ones.

## Cloning a security group

On the Security Group page, locate the target security group and click More > Clone in the Operation column.
 In the pop-up window, select the target region and project and click OK. If the new security group needs to be associated with a CVM instance, do so by managing the CVM instances in the security group.

## Deleting a security group

1. On the Security Group page, locate the security group to be deleted and click **More** > **Delete** in the **Operation** column.

2. In the pop-up window, click **OK**. If the current security group is associated with a CVM instance, it must be disassociated first before being deleted.

## **Network Switch**

Last updated : 2024-07-22 15:55:22

## Overview

Tencent Cloud supports classic network and VPC as described in Overview, which are capable of offering a diversity of smooth services. On this basis, we provide more flexible services as shown below to help you manage network connectivity with ease.

#### Network switch

Switch from classic network to VPC: A single TencentDB source instance can be switched from classic network to VPC.

Switch from VPC A to VPC B: A single TencentDB source instance can be switched from VPC A to VPC B.

#### Custom IP and port

Custom source instance IP and port: The IP and port of a source instance can be customized on the instance details page in the console.

Custom read-only instance IP and port: The IP and port of a read-only instance can be customized on the instance details page in the console.

## Notes

After the switch from classic network to VPC, only clients in the same VPC can interconnect with each other. You can configure a VPC IP range to keep the VPC IP the same as the classic network IP.

The repossession time for the old IP is 24 hours by default and can be up to 168 hours. If "Valid Hours of Old IP" is set to 0 hours, the IP is released immediately after the network is changed.

The switch from classic network to VPC is irreversible. After the switch to a VPC, the TencentDB instance cannot communicate with Tencent Cloud services in another VPC or classic network.

The network switch of a primary instance doesn't apply to its associated read-only instances or disaster recovery instances, you need to manually switch the network for these instances.

## Subnet Description

A subnet is a logical network space in a VPC. You can create subnets in different AZs in the same VPC, which communicate with each other over the private network by default.

After you select a network, the subnet IPs in the AZ of the selected instance are displayed by default. You can also select subnet IPs in other AZs in the region of the instance. Business connections adopt nearby access, so the network latency will not be increased.

## Directions

### Switching the network

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance details page.

2. Click Switch to VPC or Change Network after Network in the Basic Info section.

3. In the pop-up window, select a VPC and a subnet and click OK.

#### Note:

If there is no IP address specified, one will be automatically assigned by the system.

You can only select a VPC in the region of the instance, but you can choose a subnet in any AZ and view its IP range. We recommend that you select a VPC in the region where the CVM instance resides; otherwise, the CVM instance will not be able to access TencentDB for MySQL over the private network, unless a peering connection or a CCN instance is created between the two VPCs.

#### Switch from classic network to VPC

Change Network					
<ol> <li>If the network is changed, so is the instance IP. The old IP will be released 24 hours after the change by default. Please modify the client program in time.</li> <li>If "Time to Release Old IP" is set to 0 hours, the IP is released immediately after the network is changed.</li> <li>The VPC and subnet should be in the AZ and region of the instance.</li> </ol>					
To change the network, please go to the console Create Subnet 🛂					
In the current network environment, on in the "Default-VPC VPC" can access this database instance.					
Time to Release Old IP 24 hours Available range is 0-168 hours					
O Auto-assign IP					
OK Cancel					

#### Switch from VPC A to VPC B

Select Network				
	•		τ¢	CIDR 253 subnet IPs in total, with 247 available
If the existing networks	do not meet your	requirements, go to Creat	e Subnets 🛂	
In the current network	environment, only	CVMs in the	can access	this database instance.
Valid Hours of Old IP	24	hr Range: 0-168 hours		
O Auto-Assign IP				
O Specify IP				
		o	K Ca	incel

4. Return to the instance details page where you can view the network of the instance.

### Switching the RO group network

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation column to enter the instance details page.

2. In the **RO Group Info** of a read-only instance, click **Change Subnet** or **Switch to VPC** according to the network switch type (classic network to VPC or VPC to VPC).

3. In the pop-up window, select a VPC and a subnet and click **OK**.

#### Note:

If there is no IP address specified, one will be automatically assigned by the system.

You can only select a VPC in the region of the instance, but you can choose a subnet in any AZ and view its IP range. We recommend you select a VPC in the region where the CVM instance resides; otherwise, the CVM instance will not be able to access TencentDB for MySQL over the private network, unless a peering connection or a CCN instance is created between the two VPCs.

### Customizing the IP and port

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance details page.

2. Click





after Private Network Address and Private Port in the Basic Info section.

#### Note:

Modifying the private network address and port affects the database business being accessed.

3. In the pop-up window, specify the IP and port and click **OK**.

# Enabling Public Network Access Through CLB

Last updated : 2025-04-27 14:20:34

TencentDB for MySQL supports both private and public network addresses, with the former enabled by default for you to access your instance over the private network and the latter enabled or disabled as needed. To access your database instance from a Linux or Windows CVM instance over the public network, you can enable the public network address. You can also enable public network access through CLB, but you must configure security group rules in this case.

This document describes how to enable public network access through CLB and connect to an instance through MySQL Workbench.

## Prerequisites

You have applied to use the real server feature. Currently, this feature is in beta testing. To experience this feature, submit a ticket for application.

#### Note:

Enabling the public network services through CLB is exclusively applicable when the CLB instances and the MySQL instances belong to the same VPC network. It is currently not supported for instances across different VPC networks.

## Step 1. Purchase a CLB instance

#### Note:

If you already have a CLB instance in the same region as TencentDB for MySQL, skip this step. Go to the CLB purchase page, select the configuration, and click **Buy Now**. **Note:** 

Region: You need to select the region where the TencentDB for MySQL instance is.

## Step 2. Configure the CLB instance

The following describes how to configure the CLB instance in the same VPC as the database instance and in a different VPC respectively.

## Scenario: Deploying the CLB instance in the same VPC as the TencentDB for MySQL instance

1. Enable cross-VPC access so that the CLB instance can be bound to another private IP.

a. Log in to the CLB console, select the region, and click the target instance ID in the instance list to enter the instance management page.

- b. On the **Basic Info** page, click **Configure** in the **Real Server** section.
- c. In the pop-up window, click **Submit**.
- 2. Configure a public network listener port.

a. Log in to the CLB console, select the region, and click the target instance ID in the instance list to enter the instance management page.

b. On the instance management page, select the Listener Management tab and click Create below TCP/UDP/TCP

#### SSL Listener.

c. In the pop-up window, complete the settings and click Submit.

## Step 3. Bind a TencentDB for MySQL instance

1. After creating the listener, click it in Listener Management and click Bind on the right.

2. In the pop-up window, select **Other Private IPs** as the object type, enter the IP address and port of the TencentDB for MySQL instance, and click **OK**.

#### Note:

The login account must be a standard account (bill-by-IP). If binding fails, submit a ticket for assistance.

## Step 4. Configure the TencentDB for MySQL security group

1. Log in to the TencentDB for MySQL console and select a region. 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 **Security Group** tab, click **Configure Security Group**, configure the security group rule to open all ports, and confirm that the security group allows access from public IPs. For more information on configuration, see TencentDB Security Group Management.

# Step 5. Connect to the instance through the MySQL Workbench client

1. Download MySQL Workbench from MySQL Community Downloads and install it.

1. Go to the download page and click MySQL Workbench.

- 2. Click **Download** after **Windows (x86, 64-bit), MSI Installer**.
- 3. Click No thanks, just start my download.

4. After MySQL Workbench is installed, open it and click the plus sign after **MySQL Connections** to add the information of the target instance.

5. In the pop-up window, configure the following items and click **OK**.

Parameter	Description
Connection Name	Name the connection.
Connection Method	Select Standard (TCP/IP).
Hostname	Enter the address of the CLB instance. You can view the VIP information in the basic information on the CLB instance details page.
Port	Enter the port of the CLB instance. You can view the TCP port number in listener management on the CLB instance details page.
Username	Enter the account name of the target MySQL instance, i.e., the account created in Database Management > Account Management on the instance management page.
Store in Vault	Enter the account password of the target MySQL instance in the Password field and save it.

6. Return to the MySQL Workbench homepage and click the just configured instance to connect to it.

7. The UI after successful connection is as follows:

## **Enabling Transparent Data Encryption**

Last updated : 2024-12-24 16:15:10

## Overview

TencentDB for MySQL 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.

## Key Management Instructions

TencentDB for MySQL does not provide the keys and certificates required for encryption. The keys used for transparent data encryption are generated and managed by the KMS. The relevant explanations regarding the keys are as follows.

The TDE feature incurs no additional charges; however, the Key Management System will generate extra costs. Please refer to the Billing Overview for more details.

The Key Management System (Postpaid Version) will cease operations on December 30, 2024. From this date forward, the Key Management System will no longer support the pay-as-you-go billing model, exclusively supporting a prepaid billing approach.

For existing users of the Key Management System (Postpaid Version), an account in arrears will be unable to get keys from KMS, potentially hindering tasks such as migration and upgrades from proceeding as planned.

For users who have newly purchased the Key Management System (Prepaid Edition), when their account is in arrears, the Key Management System, having been prepaid for a certain period, will not affect the retrieval of KMS keys, nor will it impact tasks such as migration and upgrading during this period. Please be mindful of the renewal time for your KMS keys. Failure to renew them upon expiration will also affect the use of Transparent Data Encryption features. To manage your KMS keys, please visit the Key Management System Console.

The regions supported by the TencentDB for MySQL instances and the Key Management System differ. When creating a key, if there is no corresponding region in China available on the Key Management System, you may opt to create it in the Guangzhou region. Conversely, if there is no corresponding overseas region available, you may choose to create it in the Hong Kong region.

After TDE is enabled, if an account (UIN) has not previously created any encrypted tables, the corresponding key information may not be displayed in the key list. Conversely, if an account (UIN) has created encrypted tables, the corresponding key information will be visible. For instructions on creating encrypted tables, please refer to the Frequently Asked Questions.

## Prerequisites

The instance architecture must be either General or Dedicated two-node/three-node.

The database version must be MySQL 5.7 or 8.0.

You have activated Key Management Service (KMS). If not, you can enable it as instructed during the TDE activation process.

You have granted KMS key permissions. If not, you can grant permissions as instructed during the TDE activation process.

Your account needs the <code>QcloudAccessForMySQLRole</code> permission. To do so, you can follow the instructions provided during the TDE activation process.

## Use Limits

Once the authorization is revoked, MySQL databases will be inaccessible upon restart.

TDE can't be disabled once enabled.

Once TDE is enabled, you need to decrypt data before you can restore it to a local database.

TDE enhances the security of static data while compromising the read-write performance of encrypted databases. Therefore, use it based on your actual needs.

If the source instance is associated with a read-only or disaster recovery instance, you only need to enable TDE for the source instance, which will then be automatically enabled for its associated instances.

When utilizing the TDE feature, please ensure that the KMS key is in a normal operational state. Failure to do so may result in an inability to get keys from KMS, potentially hindering tasks such as migration and upgrades from proceeding as expected.

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

After TDE is enabled, authenticated applications and users can access the data transparently.

After TDE is enabled, the efficiency of backup compression may decrease.

## Directions

## Enabling TDE

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or Manage in the Operation column to enter the instance management page.

2. On the **Data Security** tab, toggle on **Encryption Status**.

## Note:

An instance with TDE enabled cannot be restored from a physical backup to a self-created database on another server.

TDE can't be disabled once enabled.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Security	Conne
Data Encryption	SSL								
Instances with data If your account has If the instance is asso instance is also enab After data encryption Encrypt table upon o Encrypt table after o Decrypt: ALTER TABL	encryption enabled canno overdue payment, you can ociated with a read-only inst led at the same time; n is enabled, users must perf reation: CREATE TABLE 11 (c' reation: ALTER TABLE 11 ENC E 11 ENCRYPTION='N';	ot be restored from a physical nnot get keys from KMS, which ance or disaster recovery instance form DDL operations on the MyS form DDL operations on the MyS 1 INT) ENCRYPTION='Y'; RYPTION='Y';	backup to a self-buil n may cause migratic re, you only need to e GQL table to encrypt o	t database on another server. n, upgrade, and other tasks to nable the data encryption feature r decrypt data. The detailed step:	fail. View Details of the source instance are as follows:	ce, and the data encryption	of the read-only instan	ce and disaster recov	ery
Data Encryption	Settings								

Encryption Status Not enabled (cannot be disabled once enabled)

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

If you select Use key auto-generated by Tencent Cloud, the key will be auto-generated by Tencent Cloud.

Set Data Encryption		×
Notes		
1 Deleting secret ke	ey or key material will cause the unavailability of primary key	
2 After the authoriz unavailable after r	ation is revoked, TencentDB for MySQL database will be restart.	
3 The operation acc	count must have the QcloudAccessForMySQLRole permission	
KMS Service	Enabled	
KMS Key Authorization	Not authorized	
	To set data encryption, please click here to authorize.	
Select Key	O Use key auto-generated by Tencent Cloud	
	O Use existing custom key	
	Encrypt Cancel	

If you select **Use existing custom key**, you can select a key created by yourself.

#### Note:

If there are no custom keys, click **go to create** to create keys in the KMS console. For more information, see Creating a Key.

Set Data Encryption		×
Notes 1 Deleting secret ke 2 After the authoriz unavailable after 3 The operation acc	ey or key material will cause the unavailability of primary key ation is revoked, TencentDB for MySQL database will be restart. count must have the QcloudAccessForMySQLRole permission	
KMS Service KMS Key Authorization Select Key	Enabled Not authorized To set data encryption, please click here to authorize. Use key auto-generated by Tencent Cloud Use existing custom key	
	South China(Guangzhou)          No data for this region         If you need to use other custom keys, go to create         Encrypt       Cancel	~

### **Encrypting a table**

Once you enable TDE, you can encrypt a table of a MySQL instance by running the example DDL statements on the table.

To encrypt a table upon creation, run the following statement:

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

To encrypt an existing table, run the following statement:

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

#### Decrypting a table

Once you enable TDE, you can decrypt a table of a MySQL instance by running the example DDL statement on the table.

To decrypt an encrypted table, run the following statement:

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

## **Frequently Asked Questions**

#### Why does the key list lack key information after enabling TDE?

#### **Issue Phenomenon**

Data Encryption Settings						
Encryption Status Chabled (cannot be disabled once er	Encryption Status C Enabled (cannot be disabled once enabled)					
Kay List Data Encryption Architecture						
Key ID/Name	Status	Creation Time	Key Usage	Key Source		

#### Normal Key List After Enabling TDE Encryption

Data Encryption Settings						
Encryption Status C Enabled (canno	t be disabled once enabled)					
Key List Data Encryption Archite	cture					
Key ID/Name	Status	Creation Time	Key Usage	Key Sourc		
KMS-CDB	Enabled	2023-05-12 15:24:54	Symmetric Encryption and Decryption	KMS		

#### **Recommended Actions**

Firstly, verify the encryption status of KMS to ensure it is functioning correctly. Check whether the account is in arrears, and ascertain if there are any overdue payments for the Cloud Database MySQL instance and the Key Management System. Please ensure all the aforementioned conditions are met before attempting again.
 Should this be your inaugural endeavor with TDE encryption, and no encrypted tables have previously been established under your account (UIN), the key list will be devoid of key information. Please refer to the commands below to create an encrypted table within your data instance and attempt again.

```
CREATE TABLE `user_test` (
 `id` bigint(20) NOT NULL AUTO_INCREMENT,
 `userId` int(11) NOT NULL,
 `age` int(11) NOT NULL,
 `name` varchar(64) DEFAULT NULL,
 `ins_date` varchar(10) DEFAULT NULL,
 PRIMARY KEY (`id`),
 KEY `idx_ins_date` (`ins_date`),
 KEY `idx_userId` (`userId`)
) ENGINE=InnoDB DEFAULT CHARSET=utf8 ENCRYPTION='Y';
```

#### Why is the key unusable after enabling TDE?

Please verify that your KMS encryption status is functioning normally, and ensure that your account is not in arrears. Additionally, confirm that both your TencentDB for MySQL instance and Key Management System are not in a state of expiration without payment. Given that the Key Management System (postpaid version) can no longer be created, and new purchases of the Key Management System are limited to the prepaid version only, users who are currently utilizing the Key Management System (postpaid version) will find that their keys become unusable in the event of an abnormal KMS encryption status or account delinquency. Therefore, it is advised to recharge your account and attempt again.

## Setting SSL Encryption

Last updated : 2025-03-28 14:11:09

## SSL encryption overview

Secure Sockets Layer (SSL) authentication is a process that authenticates the connection from the user client to the TencentDB server. After SSL encryption is enabled, you can get a CA certificate and upload it to the server. Then, when the client accesses the database, the SSL protocol will be activated to establish an SSL secure channel between the client and the server. This implements encrypted data transfer, prevents data from being intercepted, tampered with, and eavesdropped during transfer, and ultimately ensures the data security for both the client and the server.

The SSL protocol needs to be established based on reliable TCP and has the advantage of being independent from application layer protocols; therefore, high-level application layer protocols such as HTTP, FTP, and TELNET can be transparently established on it. It completes encryption algorithm processing, communication key negotiation, and server authentication before communication is made over application layer protocols. After that, all data transferred over application layer protocols will be encrypted to ensure communication privacy.

## Background

When you connect to a database in an unencrypted manner, all information transferred over the network will be in plaintext and may be eavesdropped, tampered with, and impersonated by malicious users. The SSL protocols are designed to address these risks and can bring the following benefits theoretically:

Information is encrypted and cannot be eavesdropped by a third party.

There is a verification mechanism for immediate tampering detection by both parties in the communication. Identity certificates will be used to authenticate the identity.

TencentDB for MySQL supports enhancing the connection security by enabling SSL encryption as well as downloading and installing SSL CA certificates to the required application services.

#### Note:

SSL encryption protects the traffic between the database and the server rather than the data itself. Encrypting the network connection at the transport layer can improve the security and integrity of the communication data, but will increase the response time of the network connection.

## Prerequisites

The instance version is MySQL 5.6/5.7/8.0. The instance architecture is two-node/three-node. The instance engine is InnoDB/RocksDB.

## Supported Versions

MySQL uses OpenSSL to implement secure connections. TencentDB for MySQL supports Transport Layer Security Protocol (TLS) versions 1.0, 1.1, 1.2, and 1.3, depending on the MySQL version. The following table shows the MySQL versions that support TLS.

MySQL Version	TLS1.0	TLS1.1	TLS1.2	TLS1.3
MySQL 5.7	Not Supported	Not Supported	Supported	Supported
MySQL 8.0	Supported	Supported	Supported	Supported

## Enabling SSL encryption

1. Log in to the TencentDB for MySQL 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 **Data Security** > **SSL** tab.

3. This feature is disabled by default. Toggle it on and click **OK**.

The window to enable SSL encryption in the source instance is as follows:

#### Note:

During the process of enabling SSL encryption, your database instance will be restarted to load the SSL certificate. Make sure that your business has a reconnection mechanism.

The window to enable SSL encryption in a read-only instance is as follows:

#### Note:

The configuration of the SSL encryption feature in a read-only instance will be synced to other read-only instances in its read-only group.

4. Click **Download** to download the SSL CA certificate. The certificate has a validity period of 20 years.

The downloaded file is a compressed package (TencentDB-CA-Chain.zip), which contains the following three files: .p7b file: It is used to import the CA certificate into Windows.

jks file: A storage file for truststore and keystore in Java with a unified password of tencentdb, used for importing the CA certificate chain in Java programs.

.pem file: It is used to import the CA certificate into other systems or applications.

## Configuring an SSL CA certificate

After enabling SSL encryption, you need to configure an SSL CA certificate when using a client to connect to TencentDB. The following takes Navicat as an example to describe how to install an SSL CA certificate. For other applications or clients, refer to their respective documentation.

#### Note:

Every time SSL encryption is enabled or disabled for TencentDB for MySQL, a new certificate will be generated.

- 1. Open Navicat.
- 2. Right-click the target database and select Edit Connection.
- 3. Select the **SSL** tab, select the path of the CA certificate in .pem format, complete the settings, and click **OK**. **Note**:

If the error message connection is being used is displayed, it may be because the previous session has not been disconnected. In this case, restart Navicat and try again.

4. Double-click the target database to test whether the connection is normal.

## **Disabling SSL encryption**

1. Log in to the TencentDB for MySQL 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 Data Security > SSL tab.

3. Toggle the switch off and click **OK** in the pop-up window.

#### Note:

During the process of disabling SSL encryption, your database instance will be restarted to uninstall the SSL certificate. Make sure that your business has a reconnection mechanism.

# Use the MySQL command-line client to enable the SSL-encrypted instance

If your database version is different, the connection command parameters of the MySQL client will be different. You can use the following command to query the database version you are using, and then refer to the subsequent steps to connect to the instance.

```
SELECT VERSION();
Query Result Example:
+----+
| VERSION() |
+----+
| 8.0.30-txsql |
+----+
1 row in set (0.00 sec)
```

1. Download the SSL CA certificate through the TencentDB for MySQL console. For operations, refer to Enable SSL Encryption.

2. Using the MySQL command-line client to enable the SSL-encrypted instance through the command.

For the MariaDB client, connect to the instances using the following command.

```
mysql -h <IP address> --ssl-ca=<CA certificate> --ssl -P <Port number> -u <
username > -p
```

When the client database version is MySQL 5.6, use the following command to connect to the instance.

```
mysql -P <Port number> -h <IP address> -u < username > -p< password > --ssl-
ca<CA certificate>
```

When the client database version is MySQL 5.7/8.0, use the following command to connect to the instance.

```
mysql -h <IP address> --ssl-ca=<CA certificate> --ssl-mode=REQUIRED -P <Port
number> -u < username > -p
```

To use other SSL modes, such as VERIFY\_CA or VERIFY\_IDENTITY, connect to the instance using the following command.

```
mysql -h <IP address> --ssl-ca=<CA certificate> --ssl-mode=VERIFY_CA -P <Port
number> -u <username> -p
```

#### Note:

The --ssl-mode parameter indicates the SSL mode. Generally, REQUIRED and VERIFY\_CA modes are recommended. These modes require the MySQL client to use the SSL/TLS protocol to connect to the MySQL server and to verify the SSL/TLS certificate of the MySQL server. The VERIFY\_IDENTITY mode, in addition to verifying the SSL/TLS certificate of the MySQL server, requires that the hostname used by the client matches the identifier in the server certificate; otherwise, the MySQL client will refuse to connect to the MySQL server. 3. Enter the password for the corresponding username as prompted by the system.



## Sample code for connecting to an SSL-enabled instance from common programs

#### PHP

```
$conn = mysqli_init();
mysqli_ssl_set($conn,NULL,NULL, "<path of the downloaded certificate>", NULL,
NULL);
mysqli_real_connect($conn, '<database access address>', '<database access
username>', '<database access password>', '<the specified database to be
accessed>', <access port>, MYSQLI_CLIENT_SSL);
if (mysqli_connect_errno($conn)) {
die('Failed to connect to MySQL: '.mysqli_connect_error());
}
```

PHP (Using PDO)

```
$options = array(
   PDO::MYSQL_ATTR_SSL_CA => '<path of the downloaded certificate>'
);
$db = new PDO('mysql:host=<database access address>;port=<access port>;dbname=
<the specified database to be accessed>', '<database access username>',
'<database access password>', $options);
```

#### Java (MySQL Connector for Java)

```
# generate truststore and keystore in code
String importCert = " -import "+
  " -alias mysqlServerCACert "+
  " -file " + ssl_ca +
  " -keystore truststore "+
  " -trustcacerts " +
  " -storepass password -noprompt ";
String genKey = " -genkey -keyalg rsa " +
  " -alias mysqlClientCertificate -keystore keystore " +
  " -storepass password123 -keypass password " +
  " -dname CN=MS ";
sun.security.tools.keytool.Main.main(importCert.trim().split("\\\s+"));
sun.security.tools.keytool.Main.main(genKey.trim().split("\\\s+"));
# use the generated keystore and truststore
System.setProperty("javax.net.ssl.keyStore", "<path of the downloaded
certificate>");
System.setProperty("javax.net.ssl.keyStorePassword","tencentdb");
```
```
System.setProperty("javax.net.ssl.trustStore", "<path of the downloaded
certificate>");
System.setProperty("javax.net.ssl.trustStorePassword", "tencentdb");
url = String.format("jdbc:mysql://%s/%s?serverTimezone=UTC&useSSL=true",
'<database access address>', '<the specified database to be accessed>');
properties.setProperty("user", '<database access username>');
properties.setProperty("password", '<database access password>');
conn = DriverManager.getConnection(url, properties);
```

.NET (MySqlConnector)

```
var builder = new MySqlConnectionStringBuilder
{
   Server = "<database access address>",
   UserID = "<database access username>",
   Password = "<database access password>",
   Database = "<the specified database to be accessed>",
   SslMode = MySqlSslMode.VerifyCA,
   SslCa = "<downloaded certificate>",
};
using (var connection = new MySqlConnection(builder.ConnectionString))
{
   connection.Open();
}
```

Python (MySQLConnector Python)

Python (PyMySQL)

Django (PyMySQL)



```
DATABASES = {
  'default': {
    'ENGINE': 'django.db.backends.mysql',
    'NAME': '<the specified database to be accessed>',
    'USER': '<database access username>',
    'PASSWORD': '<database access password>',
    'HOST': '<database access address>',
    'PORT': '<access port>',
    'OPTIONS': {
        'ssl': {'ca': '<path of the downloaded certificate>'}
    }
}
```

#### Node.js

```
var fs = require('fs');
var mysql = require('mysql');
const serverCa = [fs.readFileSync("<path of the downloaded certificate>",
"utf8")];
var conn=mysql.createConnection({
  host:"<database access address>",
  user:"<database access username>",
  password:"<database access password>",
  database:"<the specified database to be accessed>",
  port:<access port>,
  ssl: {
      rejectUnauthorized: true,
      ca: serverCa
  }
});
conn.connect(function(err) {
if (err) throw err;
});
```

#### Golang

```
rootCertPool := x509.NewCertPool()
pem, _ := ioutil.ReadFile("<path of the downloaded certificate>")
if ok := rootCertPool.AppendCertsFromPEM(pem); !ok {
    log.Fatal("Failed to append PEM.")
}
mysql.RegisterTLSConfig("custom", &tls.Config{RootCAs: rootCertPool})
var connectionString string
connectionString = fmt.Sprintf("%s:%s@tcp(%s:<access port>)/%s?
allowNativePasswords=true&tls=custom", "<database access username>" , "<database</pre>
```

```
access password>", "<database access address>", '<the specified database to be
accessed>')
db, _ := sql.Open("mysql", connectionString)
```

#### Ruby

# **Enabling Public Network Address**

Last updated : 2025-06-10 10:25:01

TencentDB for MySQL supports enabling the public network connection address. Once enabled, it facilitates daily testing and management, enhancing the convenience of development and use. This document describes the public network addresses and how to manually enable or disable them in the TencentDB for MySQL console. After enabling the public network addresses, you can use the system-assigned domain name and port to access TencentDB for MySQL instances from the public network. For related operations, see Connecting to Database via Public Network Address.

#### Note:

Failures of public network access will not be counted in the overall availability calculation of the TencentDB for MySQL service.

Public network access to your TencentDB for MySQL instance poses a risk to its security, hence it is advisable to employ it solely for administrative purposes, testing, or auxiliary database management. The availability of the Service Level Agreement (SLA) cannot be ensured in this context. For business access, please select private network access. It is recommended that you use public domain names rather than IP addresses for access. Because the adjustment in database instance specifications, the reactivation of public network, and the network upgrade may lead to changes in the public IP addresses. Accessing via public domain names ensures minimal impact on your operations, without the need of application modification.

# Explanation of New Public Network Architecture Upgrade

#### 1. New Architecture Release Time

To enhance the security and reliability of databases' public network links, in **May 2024, Beijing Time (UTC+8)**, TencentDB for MySQL released a new public network architecture, which adopted Cloud Load Balancer (CLB) as its underlying structure.

#### 2. Comparison of New and Old Architectures

Item	Old Public Network Architecture	New Public Network Architecture (CLB)
Architecture differences	The old public network architecture adopts single-point deployment, resulting in slow recovery and a lack of high availability in case of a single point of failure.	The new public network architecture can extend the service capability of an application system through traffic distribution, improving the availability of the application system by eliminating single points of failure.



Whether product integration involved	No.	Yes. After the public network is enabled, the system will automatically create a free, simple CLB instance in the same region in the CLB console to provide public network capabilities.

#### 3. Precautions

Currently, after the public network is enabled for TencentDB for MySQL instances, they adopt the CLB architecture. The system will automatically create a free, simple CLB instance in the same region in the CLB console to provide public network capabilities. Note the policies of the CLB architecture (as shown in the table below). If you have higher performance requirements, you can also directly purchase CLB.

Classification	Number of Concurrent Connections	Number of New Connections	Packet Volume	Inbound Bandwidth	Outbound Bandwidth
CLB	2000	200/s	Unlimited	20 Mbps	20 Mbps

#### Note:

The CLB instance automatically created upon enabling the public network address can be tried for free.

After disabling the public network address, the corresponding CLB instance on the CLB console will be automatically deleted.

Activating the Public network capabilities for TencentDB for MySQL instances is complimentary. However, the activation process includes a verification of the account's financial status. Should the account be in arrears (This means that the account balance is less than 0), activation will not be possible. To enable Public network functionalities, please ensure your account is not in a state of indebtedness. Upon successful activation, the CLB instance automatically created as a result of enabling the Public network will incur no charges. Starting mid-May 2024, the CLB health check source IP will be in the 100.64.0.0/10 subnet. Once the public network is enabled, if your simple CLB instance shows an abnormal health status, you can **allow the 100.64.0.0/10 subnet** in the security group settings for your TencentDB for MySQL instance to resolve the issue of failed health checks causing the CLB instance to display an abnormal health status. Please refer to Manage Cloud Database Security Groups for operations.

You need to configure monitoring and alarms for the above-mentioned simple CLB instances, to monitor public network connections through metrics (such as the number of new public network connections established, number of public network connections, public network outbound bandwidth, and public network inbound bandwidth) after

enabling the public network address. Refer to Alarm Policy for operation steps, and select the policy type as shown in the following figure.

# **Use Limitations**

Currently, the primary instance regions of Guangzhou, Shanghai, Beijing, Chengdu, Chongqing, Nanjing, Hong Kong (China), Singapore, Seoul, Tokyo, Silicon Valley, Virginia, and Frankfurt support enabling the public network. For regions where read-only instances can enable the public network, refer to the console. If your instance's region does not support enabling but requires the public network, see Enable Public Network Access through CLB to access TencentDB for MySQL instance via the public network.

# Enabling the Public Network Address

1. Log in to the TencentDB for MySQL console. In the instance list, click an **Instance ID** or **Manage** in the **Operation** column to access the instance details page.

2. Under Public Network Address on the instance details page, click Enable.

3. In the pop-up window, click OK.

#### Note:

After successful enabling, you may access TencentDB for MySQL via the public network by using the domain name and port assigned by the system. Note that it may take approximately 5 minutes for the change to take effect.

# **Disabling the Public Network Address**

#### Note:

After disabling, you will no longer be able to access TencentDB for MySQL via the public network by using the domain name and port assigned by the system. Furthermore, the domain will be released. We cannot guarantee that the domain will remain unchanged if reapplied. Ensure that your system does not use the public network access address to avoid unnecessary losses.

When the public network is disabled, the Cloud Load Balancer (CLB) instances that support this public network capability will be synchronously taken offline within 3 to 5 minutes.

1. Log in to MySQL Console, in the Instance List, click **Instance ID** or **Manage** in the **Operation** column to enter the Instance Details page.

2. On the Instance Details page, click **Disable** behind the **Public Network Address**.

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

4. Return to the Instance Details page, the **Public Network Address** will no longer display the public network address, indicating that the public network feature is turned off.

# Backup and Rollback Backing up Databases

Last updated : 2025-06-10 09:53:45

To avoid data loss or corruption, you can back up a database automatically or manually.

# **Backup Overview**

#### Backup modes

TencentDB for MySQL single-node (cloud disk), two-node (local disk), and three-node (local disk) instances support **automatic backup** and **manual backup** of databases.

#### Backup types

#### TencentDB for MySQL two-node and three-node instances support two backup types:

Physical backup: A full copy of physical data (supported for automatic backup).

**Logical backup**: The backup of SQL statements (supported for both manual and automatic backups). **Note:** 

To restore a database from a physical backup, you need to use xbstream to decompress the package first. For more information, see <u>Restoring Database from Physical Backup</u>.

If the number of tables in a single instance exceeds one million, backup may fail and database monitoring may be affected. Make sure that the number of tables in a single instance is below one million.

As the data of tables created by the MEMORY storage engine is stored in the memory, physical backups cannot be created for such tables. To avoid data loss, we recommend that you replace them with InnoDB tables.

If there are a high number of tables without a primary key in the instance, backup may fail, and the high availability of the instance may be affected. You need to create a primary key or secondary index for such tables in time.

Physical Backup Advantages	Logical Backup Disadvantages
High backup speed. Streaming backup and compression are supported. High success rate. Simple and efficient restoration. Faster backup-based coupling operations such as adding real-only instances and disaster recovery instances. 1/8 of average time needed for creating a logical backup.	Long time needed to restore as it takes time to run SQL statements and build indexes. Low backup speed, especially when there are massive amounts of data. Possible increase in source-replica delay due to the pressure on instances during backup. Possible loss of precision information of floating points. Potential backup failures due to wrong views and other problems.



Ten times faster than logical backups during	
import.	

Slower backup-based coupling operations such as adding read-only instances and disaster recovery instances.

#### TencentDB for MySQL single-node instances of cloud disk edition support snapshot backup.

**Snapshot backup** backs up data by creating snapshots for disks at the storage layer and is supported for both automatic and manual backups.

Snapshot Backup Advantages	Snapshot Backup Disadvantages
High backup speed. Small size.	Cannot be downloaded.

#### **Backup objects**

Data Backup	Log Backup
TencentDB for MySQL two-node and three-node instances: Automatic backup supports full physical backup. Manual backup supports full physical backup, full logical backup, and single-database/table logical backup. Both automatic and manual backups can be compressed and downloaded. Single-node instances of cloud disk edition: Automatic backup supports full snapshot backup. Manual backup supports full snapshot backup. Both automatic and manual backups cannot be downloaded.	TencentDB for MySQL single-node (cloud disk edition), two-node, and three-node instances support binlog backup: Log files occupy the instance's backup space. Log files can be downloaded but cannot be compressed. Retention periods can be set for log files.

### Note

Since February 26, 2019, the automatic backup feature of TencentDB for MySQL only supports physical backup (default type) and no longer provides logical backup. automatically. Existing automatic logical backups will be switched to physical backups

This will not affect your business access, but may have impact on your habit of automatic backup. If you need logical backups, you can use the manual backup option in the TencentDB for MySQL console or call the CreateBackup API to generate logical backups.

Manual backups are logical cold backups by default. If full backups before configuration adjustment are chosen as logical cold backups, the configuration adjustment may take longer. To reduce adjustment time, it is recommended to choose physical cold backups for full backups before configuration adjustment.

Instance backup files occupy backup space. We recommend that you plan the usage of backup space appropriately. Usage of backup space that exceeds the free tier will incur fees. For more information, see Backup Space Billing.

We recommend that you back up your database during off-peak hours.

To avoid situations where the required backup files are deleted after the retention period lapses, you need to download them to the local file system in a timely manner.

Do not perform DDL operations during the backup process to avoid backup failure due to table locking.

TencentDB for MySQL single-node instances cannot be backed up.

MySQL read-only instances do not support database backups.

# Backing up MySQL Data Automatically

#### Configuring automatic backup

1. Log in to the TencentDB for MySQL console, click an **Instance ID** on the instance list page to enter the management page, and select **Backup and Restoration** > **Auto-Backup Settings**.

2. Select backup parameters in the pop-up window as detailed below and click **OK**:

#### Note:

The rollback feature as described in Rolling back Databases relies on the backup cycle and retention days of data backups and log backups (binlog). Rollback will be affected if you reduce the automatic backup frequency and retention period. You can select the parameters as needed.

For example, if the backup cycle is set as Monday and Thursday and the retention period is set as seven days, you can roll a database back to any point of time in the past seven days, which is the actual retention days of data backups and log backups.

Automatic backups cannot be deleted manually. You can set the retention period for automatic backups, and the backups will be deleted automatically when they expire.

Increasing the retention period of data and log backups may cause additional backup space fees.

Shortening the retention period of log backups may affect the data rollback cycle of the instance.

In **Auto-Backup Settings**, you can enable periodic archive for data backups. The settings with periodic archive not enabled are non-archive backup settings. The following describes the parameters for **non-archive backup** and **archive backup**.

#### Non-archive backup settings

Parameter	Description
Backup Start Time	The default backup start time is automatically assigned by the system. You can set a backup start time range as needed. We recommend that you set it to off-peak hours. This is just the start time of the backup process and does not indicate the end time. For example, if the backup start time is set to 02:00-06:00 AM, the system will initiate a backup

	at a point in time during this period of time, which depends on the backend backup policy and backup system conditions.
Data Backup Retention Period	For TencentDB for MySQL two-node and three-node instances, data backup files can be retained for 7 (default) to 1,830 days and will be automatically deleted upon expiration. For TencentDB for MySQL single-node instances of cloud disk edition, data backup files can be retained for 7 (default) to 30 days and will be automatically deleted upon expiration.
Backup Cycle	Configuration rules: By week: 7 days from Monday to Sunday are selected by default, and you can customize the backup time; however, to ensure your data security, we recommend that you configure to back up at least twice a week. By month: To protect your data security, the interval between any two adjacent backup dates in a month cannot exceed 2 days; for example, if you choose to back up on the 1st day of a month, the next backup date cannot be the 5th day (with the 2nd, 3rd, and 4th days skipped).
Transition- to-Cold Storage (optional)	Select the target data backup transition-to-cold storage policy and specify the number of days: Standard storage days: Specify the number of days after which to transition a generated data backup to standard storage. Archive storage days: Specify the number of days after which to transition a generated data backup to archive storage. For more information on transition-to-cold storage, see Configuring Transition-to-Cold Storage. Note that the archive storage type isn't available yet.
Log Backup Retention Period	For TencentDB for MySQL two-node and three-node instances, log backup files can be retained for 7 (default) to 1,830 days and will be automatically deleted upon expiration. For TencentDB for MySQL single-node instances of cloud disk edition, log backup files can be retained for 7 (default) to 30 days and will be automatically deleted upon expiration.
Transition- to-Cold Storage (optional)	Select the target binlog backup transition-to-cold storage policy and specify the number of days: Standard storage days: Specify the number of days after which to transition a generated binlog backup to standard storage. Archive storage days: Specify the number of days after which to transition a generated binlog backup to archive storage. For more information on transition-to-cold storage, see Configuring Transition-to-Cold Storage. Note that the archive storage type isn't available yet.

#### Archive backup settings

#### Note:

You cannot configure archive backup for single-node instances of cloud disk edition.

Archive backup retention period can only be longer than non-archive backup retention period.

Parameter	Description

Backup Start Time	The default backup start time is automatically assigned by the system. You can set a backup start time range as needed. We recommend that you set it to off-peak hours. This is just the start time of the backup process and does not indicate the end time. For example, if the backup start time is set to 02:00-06:00 AM, the system will initiate a backup at a point in time during this period of time, which depends on the backend backup policy and backup system conditions.
Data Backup Retention Period	For TencentDB for MySQL two-node and three-node instances, data backup files can be retained for 7 (default) to 1,830 days and will be automatically deleted upon expiration.
Backup Cycle	Configuration rules: By week: 7 days from Monday to Sunday are selected by default, and you can customize the backup time; however, to ensure your data security, we recommend that you configure to back up at least twice a week. By month: To protect your data security, the interval between any two adjacent backup dates in a month cannot exceed 2 days; for example, if you choose to back up on the 1st day of a month, the next backup date cannot be the 5th day (with the 2nd, 3rd, and 4th days skipped).
Archive Backup Retention Period	Data backup files can be retained for 90 to 3,650 days (1,080 days by default) and will be automatically deleted upon expiration.
Archive Backup Retention Policy	You can set the number of backups by month, quarter, or year.
Start Time	The time to start archive backup.
Transition- to-Cold Storage (optional)	Select the target data backup transition-to-cold storage policy and specify the number of days: Standard storage days: Specify the number of days after which to transition a generated data backup to standard storage. Archive storage days: Specify the number of days after which to transition a generated data backup to archive storage. For more information on transition-to-cold storage, see Configuring Transition-to-Cold Storage. Note that the archive storage type isn't available yet.
Log Backup Retention Period	Log backup files can be retained for 7 (default) to 1,830 days and will be automatically deleted upon expiration.
Transition- to-Cold	Select the target binlog backup transition-to-cold storage policy and specify the number of days:



Storage	Standard storage days: Specify the number of days after which to transition a generated binlog
(optional)	backup to standard storage.
	Archive storage days: Specify the number of days after which to transition a generated binlog
	backup to archive storage.
	For more information on transition-to-cold storage, see Configuring Transition-to-Cold Storage.
	Note that the archive storage type isn't available yet.
	For more information on transition-to-cold storage, see Configuring Transition-to-Cold Storage. Note that the archive storage type isn't available yet.

#### Viewing the retention plan

#### Note:

Single-node (cloud disk) and cluster edition instances currently do not support the viewing the retention plan feature. After you select an archive backup retention policy in the backup settings, you can click **View Retention Plan** to preview it.

Blue dates are for non-archive backups.

Red dates are for archive backups.

You can click Non-archive Backup or Archive Backup to hide corresponding dates for easier preview.

The backup plan preview is currently for backups for the year to come and is for reference only.

Example 1: The backup cycle is Monday, Wednesday, Friday, and Sunday, with one backup retained for each month starting from January 11, 2022.

Example 2: The backup cycle is Monday, Wednesday, and Friday, with three backups retained for each quarter starting from January 11, 2022.

Example 3: Show only archive backups.

# Backing up MySQL Data Manually

The manual backup feature allows you to initiate a backup task manually.

#### Note:

For TencentDB for MySQL two-node and three-node instances, manual backup supports full physical backup, full logical backup, and single-database/table logical backup.

For TencentDB for MySQL two-node and three-node instances, manual backups can be manually deleted from the backup list in the console. You can delete manual backups that are no longer in use to free up space. Manual backups can be retained as long as they are not deleted until the database instance is deactivated.

For TencentDB for MySQL single-node instances of cloud disk edition, manual backup supports full snapshot backup. For TencentDB for MySQL single-node instances of cloud disk edition, manual backups cannot be deleted.

When the instance is performing daily automatic backup, no manual backup tasks can be initiated.



Manual backups will enter a backup queue with the highest priority. The specific backup start time is subject to the start time field of the task in a backup list.

Dual Node, Three Node instance operation steps

Single node (cloud disk), cluster edition (cloud disk) instance operation steps

1. Log in to the TencentDB for MySQL console, click an **Instance ID** on the instance list page to enter the management page, and select **Backup and Restoration** > **Manual Backup**.

2. In the pop-up Backup Settings dialog box, select the backup mode and object, enter the alias, and click OK.

#### Note:

If it is a logical backup of a single database or table, check the database or data table you want to back up in **Select database & table** on the left, which then will be added to the list on the right. If you do not have a database or table yet, create one first.

Manual backups are logical cold backups by default. If full backups before configuration adjustment are chosen as logical cold backups, the configuration adjustment may take longer. To reduce adjustment time, it is recommended to choose physical cold backups for full backups before configuration adjustment.

 Log in to the TencentDB for MySQL console, click an Instance ID on the instance list page to enter the management page, and select Backup and Restoration > Manual Backup.
 Enter an alias, and click OK.

### FAQs

#### 1. Can I download or restore backup files that exceed the retention period?

Expired backup sets will be automatically deleted and cannot be downloaded or restored.

We recommend that you configure a reasonable backup retention period based on your business needs or download the backup files in the TencentDB for MySQL console. Note that the backup files of single-node instances of cloud disk edition cannot be downloaded.

You can also manually back up instance data in the console. Manual backups will be retained permanently.

#### Note:

Manual backups will also take up the backup space. We recommend that you plan the usage of the backup space appropriately to reduce costs.

#### 2. Can I delete backups manually?

You can't delete automatic backups manually, but you can set the retention period for them so that they are deleted automatically upon expiration.

For two-node and three-node instances, manual backups can be manually deleted from the backup list in the TencentDB for MySQL console. Manual backups are retained permanently as long as they are not deleted. For singlenode instances of cloud disk edition, manual backups cannot be deleted.

#### 3. Can I disable data and log backups?

No. However, you can reduce the backup frequency and delete unnecessary manual backups in the TencentDB for MySQL console to lower the space usage. Note that the manual backups of single-node instances of cloud disk edition cannot be deleted.

#### 4. How do I reduce the backup space costs?

Delete manual backups that are no longer used. You can log in to the TencentDB for MySQL console, click an instance ID/name to access the instance management page, and delete manual backups on the **Backup and Restore** tab. Note that the manual backups of single-node instances of cloud disk edition cannot be deleted. Reduce the frequency of automatic data backup for non-core businesses. You can adjust the backup cycle and retention period in the console, and the frequency should be at least twice a week.

#### Note:

The rollback feature relies on the backup cycle and retention days of data backups and log backups (binlog). Rollback will be affected if you reduce the automatic backup frequency and retention period. You can select the parameters as needed. For more information, see Rolling back Databases.

Shorten the retention period of data and log backups for non-core businesses. A retention period of seven days can meet the needs in most cases.

Configure the transition-to-cold storage policy to transition backup files and reduce the storage costs. For more information, see Configuring Transition-to-Cold Storage.

Business Scenario	Recommended Backup Retention Period
Core businesses	7–3,650 days. We recommend that you enable archive backup to periodically retain backup files in the long term.
Non-core, non- data businesses	7 days
Archive businesses	7 days. We recommend that you manually back up data based on your actual business needs and delete the backups promptly after use.
Testing businesses	7 days. We recommend that you manually back up data based on your actual business needs and delete the backups promptly after use.

# Configuring Transition-to-Cold Storage

Last updated : 2025-06-06 10:33:51

TencentDB for MySQL supports transition-to-cold storage for backup files to reduce the backup storage costs. This document describes how to use this feature.

#### Note:

The cluster edition of the instance currently does not support the backup archiving feature.

The free space is not applicable to transitioned backups.

The archive storage type isn't available yet.

### Overview

**Transition-to-cold storage**: The feature automatically transitions the storage types of generated data and binlog backup files according to the configured transition-to-cold storage rules to reduce the storage costs.

After you configure the backup retention period in **Auto-Backup Settings**, backup files will be stored according to the backup policy. Within the configured retention period, you can transition the backup files generated N or X days (customizable) ago to standard or archive storage respectively through the **transition-to-cold storage** feature.

#### Backup file storage types

Standard storage: It is suitable for business scenarios with frequent access and use. It can be used for rollback and cloning, uses no free space, and only supports downloads over the private network.

Archive storage: Suitable for business scenarios with lower access frequency, Archive Storage mandates a minimum retention period of 90 days. If the storage duration is less than 90 days, fees will be calculated based on the 90-day minimum. Downloading is not supported.

# Feature overview

You can configure to transition backup files generated a certain number of days ago to standard storage. The number can range from 30 days to the configured number of backup retention days.

You can configure to transition backup files generated a certain number of days ago to archive storage. The number can range from 90 days to the configured number of backup retention days.

The backup storage type can only be transitioned from hot to cold (general storage > standard storage > archive storage) but not vice versa.

All the configured transition-to-cold storage policies will be performed at 00:00 (GMT+08:00) on the next day. As lifecycle tasks are executed asynchronously, the execution for backups after the rule is configured will be finished by

24:00 on the next day in most cases. The time is calculated based on the modification time of the transition-to-cold storage configuration.

Transition-to-cold storage can be configured for data backups and binlog backups separately.

Transition-to-cold storage can be used together with periodic archive.

# Billing overview

Transitioned backups are stored at lower prices to reduce the storage costs. For billing details, see Backup Space Billing.

# Transition-to-cold storage policy

1. When the backup retention period is  $\leq$  31 days, cold storage for backups cannot be configured. To enable this feature, please set the backup retention period to more than 31 days.

2. When 31 days < backup retention period  $\leq$  180 days, the number of days for archive storage cannot be configured.

3. The number of days for standard storage cannot exceed or equal the number of days for archive storage.

- 4. Both the standard storage duration and archive storage duration must be less than the backup retention period.
- 5. A transition-to-cold storage rule only takes effect after it is selected.

# Configuring transition-to-cold storage

1. Log in to the TencentDB for MySQL console, click an **Instance ID** on the instance list page to enter the management page, and select **Backup and Restoration** > **Auto-Backup Settings**.

2. In the **Backup Settings** pop-up window, configure the data backup retention period and cycle or configure the log backup retention period, select the target transition-to-cold storage policy, and specify the number of days.

3. Read and select **Backup Space Billing Notes** and click **OK**.

# Examples of transition-to-cold storage

#### Example 1: Non-archive backup + transition-to-cold storage configurationBackground

Company A has a TencentDB for MySQL instance in Guangzhou region. The data backup retention period is 200 days, and the non-archive backup cycle is every Monday, Wednesday, Friday, and Sunday. On December 31, 2022, the following transition-to-cold storage policy is configured: after a backup file is generated for 90 or 180 days, it will be

transitioned to standard or archive storage respectively.

#### Transition expectations

The transition-to-cold storage policy takes effect on January 1, 2023.

Existing backup files: Starting from January 1, 2023, existing backup files generated for more than 90 or 180 days are transitioned to standard or archive storage respectively.

New backup files: Starting from January 1, 2023, newly generated backup files are transitioned to standard or archive storage respectively after 90 or 180 days.

#### Example 2: Archive backup + transition-to-cold storage configurationBackground

Company B has a TencentDB for MySQL instance in Guangzhou region. The data backup retention period is 365 days, the non-archive backup cycle is every Monday, Wednesday, Friday, and Sunday, the archive backup retention period is 365 days, and one backup file is retained periodically per month. On January 31, 2023, the following transition-to-cold storage policy is configured: after a backup file is generated for 90 or 180 days, it will be transitioned to standard or archive storage respectively.

#### Transition expectations

The transition-to-cold storage policy takes effect on February 1, 2023.

Existing backup files: Starting from February 1, 2023, existing backup files generated for more than 90 or 180 days are transitioned to standard or archive storage respectively.

New backup files: Starting from February 1, 2023, newly generated backup files are transitioned to standard or archive storage respectively after 90 or 180 days.

In the backup retention schedule shown in the figure below, each newly generated backup file will be transitioned according to the time configured in the transition-to-cold storage policy, and backups can be transitioned only from hot to cold.

# **Cross-Region Backup**

Last updated : 2025-01-09 17:41:26

This document describes the cross-region backup feature of TencentDB for MySQL. This feature allows you to store backup files in another region, which enhances the regulatory compliance and disaster recovery capabilities and improves the data reliability.

# Background

Data is an important part of enterprise operations. Although the information technology brings convenience, electronic data and stored information are very vulnerable to damage or loss. Any incident, such as natural disaster, system failure, maloperation, or virus, can cause business interruptions or even disastrous losses. Therefore, ensuring the security and integrity of core data is a top priority of every enterprise.

The cross-region backup feature of TencentDB for MySQL can be used to store backup files in another region so as to minimize data corruptions caused by natural disasters or system failures. This feature ensures the high availability, security, and recoverability of data and implement various features, such as remote backup and restoration, remote disaster recovery, long-term data archive, and regulatory compliance.



# Notes on Cross-Region Backup

Cross-region backup feature is not supported currently for single-node cloud disk edition instances.

A cross-region backup can be restored to the instance regions or backup region.

Once enabled, cross-region backup doesn't affect the local default backup, and both coexist.

Cross-region backup will be triggered after the local automatic backup is completed, that is, the automatic backup will be dumped to the storage device for cross-region backup.

The retention period of cross-region backup only affects the backup cycle of the cross-region storage.

Backups and binlogs stored in a remote region cannot use the free storage space, and cross-region backups will use the space in the backup region of the source instance.

Enabling/disabling cross-region backup or changing the region won't affect existing backups.

After cross-region backup is enabled, the last valid backup and the binlogs generated from the last valid backup time to the current time will be synced to the target region.

If you enable cross-region backup again, remote backups generated before the enablement cannot be used to restore data to the specified time point.

# Billing

For more information on cross-region backup billing, see Backup Space Billing.

# Differences Between Cross-Region Backup and Local Default Backup

Comparison Item	Cross-Region Backup	Local Default Backup
Enabled by default	No. It needs to be enabled manually.	Yes.
Backup storage region	Up to two specified remote regions.	Instance region.
Backup restoration	The data can be restored to: Original instance New instance in the target region	The data can be restored to: Original instance New instance in the current region
Uses the free storage space	No.	Yes.

# Supported Regions

The regions currently supporting this feature are listed in the following table. This feature will be successively supported in other cities in the future.

Source region	Regions for cross-region backup
Beijing, Shanghai, Guangzhou, Chengdu, Nanjing, Chongqing, Hong Kong (China), Singapore	<ul> <li>Beijing, Shanghai, Guangzhou, Chengdu, Nanjing, Chongqing, Hong Kong (China), Singapore</li> <li>Note :</li> <li>Supports backup to regions other than the source region itself. The details are subject to the regions actually supported by the console.</li> </ul>
Jakarta, Bangkok, Seoul, Tokyo, Silicon Valley, Virginia, Frankfurt, Sao Paulo	Singapore

# **Enabling Cross-Region Backup**

1. Log in to the TencentDB for MySQL 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 **Backup and Restoration** > **Cross-Region Backup**.

3. In the **Cross-Region Backup** window, complete the configuration and click **OK** to enable cross-region backup.

Cross-Region Backup: It is disabled by default.

Back up Binlog: When cross-region backup is enabled, it will be enabled automatically and can be disabled separately.

Backup Region: Select one or two regions other than the source instance region.

Backup Retention Period: 7 days by default. Value range: 3–1830 days. Backup sets will be deleted automatically upon expiration.

4. After cross-region backup is completed, the backup will be synced to the target region and can be queried in the backup list of the source instance.

For cross-region backup files, all selected backup regions are displayed in the **Backup Region** column.

The backup list has the following fields:

Field	Description
File Name	You can set a filename when creating a backup, which cannot be modified after the backup is created.
Backup Time	The backup start time.
Task Start TimeTask End Time	The time points when the backup task starts and ends respectively.

Backup Size	The size of the current backup file.
Туре	The backup file type. The cross-region backup type for two-node and three-node instances is physical cold backup, while the cross-region backup type for Cluster Edition instances is snapshot backup.
Backup Mode	Automatic backup
Backup Method	Full backup
Backup Region	It displays all regions where the backup file is stored.
Status	It displays the backup task status.
Operation	DownloadClone

# **Disabling Cross-Region Backup**

You can disable the cross-region backup feature in the console. After it is disabled, cross-region backup files won't be deleted immediately; instead, they will be deleted automatically upon expiration.

1. Log in to the TencentDB for MySQL 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 **Backup and Restoration** > **Cross-Region Backup**.

3. In the Cross-Region Backup window, disable the feature and click OK.

# FAQs

#### Why are fees still incurred after I disable cross-region backup?

If your data volume exceeds the free storage space, the excessive data will be billed. After cross-region backup is disabled, no more new backups will be generated, but existing backups won't be deleted immediately. They will be retained within the retention period and use the storage space before expiration, so they will still incur fees before expiration. You can set the backup retention period to automatically clear all expired backup files, and then no more cross-region backup fees will be incurred.

# **Backup Encryption**

Last updated : 2025-06-10 09:57:37

TencentDB for MySQL supports backup encryption for physical backups and log backups (binlog). To use an encrypted backup, you need to download it and its encryption key for decryption. This document describes how to enable or disable the backup encryption feature and download a key.

# Prerequisite

The MySQL instance architecture is two-node/three-node.

### Note

After backup encryption is enabled, the previous backup will not be encrypted, but the new backup files will be automatically encrypted for storage.

You cannot modify the backup encryption key.

After backup encryption is enabled, you don't need to manually decrypt a backup in the console, as the backend will decrypt it automatically before relevant operations such as cloning. However, if you download an encrypted backup, you need to download its key for decryption. For more information, see <u>Restoring Database from Physical Backup</u>.

# Enabling backup encryption

1. Log in to the TencentDB for MySQL console. In the instance list, click an **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

2. On the instance management page, select **Backup and Restoration** and click **Backup Encryption**.

3. In the pop-up dialog box, select a backup encryption type and click **OK**.

#### Note:

After enabling the corresponding backup encryption, the physical backup time may be extended by 20%, and the log backup time may be extended by 30%.

4. After the backup encryption is enabled, the backup files generated based on the automatic backup settings will be automatically encrypted.



5. You can also click **Manual backup**. On the manual backup setting page, select the configuration, click **OK** to

encrypt the newly added physical backup.

Parameter	Description
Select backup mode	Select Physical cold backup.
Object	It is the instance by default.
Backup Encryption	It is enabled by default.
Backup Name	It can contain up to 60 letters, digits, or symbols (/()[]+=:@).

# Disabling backup encryption

#### Note:

After backup encryption is disabled, the previous backup will not be decrypted, and the new backup files will not be encrypted for storage.

1. Log in to the TencentDB for MySQL console. In the instance list, click the ID of the target instance or **Manage** in the **Operation** column to enter the instance management page.

2. On the instance management page, select **Backup and Restoration** and click **Backup Encryption**.

3. In the pop-up window, disable the corresponding backup encryption type, and click **OK**.

# Downloading a backup key

1. Log in to the TencentDB for MySQL console. In the instance list, click the ID of the target instance or **Manage** in the **Operation** column to enter the instance management page.

2. On the instance management page, select the **Backup and Restoration** > **Data Backup List** tab, find the target backup, and click **Download Key** in the **Operation** column.

3. In the pop-up window, select the file path where to save the key and click **Download**.

# **Backup Retention Settings**

Last updated : 2025-06-03 10:21:45

TencentDB for MySQL supports completing backup retention settings before or when an instance is terminated. This prevents backups from being deleted immediately in case of an accidental instance termination. Users still possess the backups of this instance for a specified period after the instance is terminated. This document describes the operations and details related to backup retention settings.

### Prerequisites

The instance architecture is two-node or three-node. Backup files have been generated for instances.

# **Billing Instructions**

For pricing details, see Backup Space Billing.

# Directions

#### **Backup Retention Settings**

1. Log in to the TencentDB for MySQL console. In the instance list, click **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

2. On the instance management page, select Backup Recovery, and click Backup Retention Settings.

3. In the pop-up window, complete the following configurations, and click OK.

Parameter	Description
Backup Files After Instance Termination	<ul><li>Whether to retain backup files after instance termination. You can select the Retain or Do not retain option.</li><li>Do not retain: If you select Do not retain, backups will be terminated simultaneously after instance termination and will not be saved. Data cannot be restored on the cloud.</li><li>Retain: If you select Retain, backups within the specified scope will be retained after the instance is terminated. These backups are billed.</li></ul>
Retention	This parameter is displayed if you choose to retain backup files after instance termination.

l	Policy	The options are as follows:
l		Retain all: Retain all backups generated by the instance in the current region. You can
l		choose whether to retain the binlog files.
l		Retain the latest one: Retain the latest 1 backup generated by the instance in the current
l		region. By default, binlog files cannot be retained.
		Retain those within the specified time range: Retain backups generated by the instance in the current region and after the selected time. You can choose whether to retain the binlog files. For example, you can retain backups generated after April 28, 2025.
	Binlog Files	When "Retain all" or "Retain those within the specified time range" is selected for Retention Policy, this parameter is displayed, and you can choose whether to retain binlog files.

#### **Querying and Downloading Backups of Terminated Instances**

For instances whose backup files are retained after termination, their backups can be queried and downloaded through the database backup page after the instances are terminated and removed.

1. Log in to the TencentDB for MySQL console and choose Database Backup on the left sidebar.

2. On the **Database Backup** page, choose **Backup List** > **Backups of Terminated Instances**. You can query the retained backup information in the list.

Filter from the list: You can filter backups by **instance decommissioning time** to quickly obtain the corresponding backups.

Filter by instance ID: You can accurately filter target backups of terminated instances by instance ID.

Refresh: You can refresh the list.

Set list fields: You can click the

button to set the fields you want to display in the list.

Download: You can download the list (containing no more than 10,000 data entries).

3. To download a backup file, find the corresponding instance in the list, click the number in the Number of Files column, and view and download the corresponding backup file in the right pop-up window.

Click Download. In the pop-up window, set the **download conditions**, then click **Generate Download Address**.
 After successful generation, you can copy the link and use the wget command to download files.
 Note:

Every download address is valid for 12 hours. If it expires, please refresh the page to obtain a new one. wget command format: wget -c '<Backup file download address>' -O Custom filename.xb

# Setting Backup Download Rules

Last updated : 2025-03-21 11:57:11

By default, you can download backup files of TencentDB for MySQL instances over public or private network. To limit the download, you can adjust backup download settings.

Note:

Backup download settings are supported in the following regions:

Guangzhou, Shanghai, Beijing, Shenzhen, Chengdu, Chongqing, Nanjing, Hong Kong (China), Beijing Finance, Shanghai Finance, Shenzhen Finance, Singapore, Silicon Valley, Frankfurt, Seoul, Bangkok, and Tokyo.

# Setting backup download rules

1. Log in to the TencentDB for MySQL console, select **Database Backup** on the left sidebar, and select a region at the top.

2. On the **Download Settings** tab, view the backup download settings and click **Edit** to modify them.

#### Note:

Download over public network is enabled by default and when it is enabled, download over private network is also allowed.



3. On the displayed page, set the download rules and click **OK**.

Download over public network:

Enabled: You cannot set any download rule.

Disabled: You can set the download rules for private network by allowing or blocking specific IPs and VPCs. Set the download rules:

If you don't specify any value, the condition won't take effect.

You should separate the values of an IP condition with commas.

You can enter IPs or IP ranges as the values of an IP condition.

If no IP and VPC requirements are set, there will be no limit on download over private network.



Download Settings			
Download over Public Network	Enable	Close	
Download Conditions	Field	Operator	Value (j)
	IP	IN 🔻	Enter IPs and separate them with commas
	VPC	IN	Please select
OK Cancel			

4. After the configuration is completed, return to the **Download Settings** tab to view the rules that take effect.

Download	Settings Edi	t	
	Field	Operator	Value
	IP	IN	/24
	VPC	IN	

### Authorizing sub-accounts to set backup download rules

By default, sub-accounts do not have the permission to set backup download rules for TencentDB for MySQL instances. Therefore, you need to create CAM policies to grant specific sub-accounts the permission.

Cloud Access Management (CAM) is a web-based Tencent Cloud service that helps you securely manage and control access permissions to your Tencent Cloud resources. Using CAM, you can create, manage, and terminate users (groups), and control the Tencent Cloud resources that can be used by the specified user through identity and policy management.

You can use CAM to bind a user or user group to a policy which allows or denies them access to specified resources to complete specified tasks. For more information on CAM policy elements, see Element Reference.

#### Authorizing sub-accounts

1. Log in to the CAM console with the root account, locate the target sub-user in the user list, and click Authorize.

Create User More *					
Username	User Type	Account ID	Creation Time	Associated Info	Operation
•	Root Account	100	2019-02-20 15:10:30	L 🛛	Jutholae More 🔻
•	Sub-user		2021-07-29 17:30:16		Authorize More 🔻

2. In the pop-up window, select the **QcloudCDBFullAccess** preset policy and click **OK** to complete the authorization.

ect Policies (667 Total)				1 selected		
upport search by policy name/description/remark	S	Q	]	Policy Name	Policy type	
Policy Name	Policy type ▼			OcloudCDBFullAccess		
AdministratorAccess This policy allows you to manage all users	Preset Policy	•		Full read-write access to TencentDB, inclu	Preset Policy	
ReadOnlyAccess This policy authorizes you with the read-o	Preset Policy		↔			
QCloudResourceFullAccess This policy allows you to manage all clou	Preset Policy					
QCloudFinanceFullAccess This policy allows you to manage all finan	Preset Policy					
OcloudAdvisorFullAccess		•				

#### **Policy syntax**

The following policy syntax is used to authorize a sub-account to set backup download rules for TencentDB for MySQL instances:

```
"resource":["resource"]
}
```

version is required. Currently, only the value "2.0" is allowed.

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

effect is required. It describes the result of a statement. The result can be "allow" or an "explicit deny". action is required. It specifies whether to allow or deny the operation. The operation can be an API (prefixed with "cdb:"). which can be an API (prefixed with name) or a feature set (a set of specific APIs prefixed with permid). resource is required. It describes the details of authorization.

#### **API** operations

In a CAM policy statement, you can specify any API operation from any service that supports CAM. For database audit, the API prefixed with name/cdb: should be used. To specify multiple operations in a single statement, separate them by comma:

"action":["name/cdb:action1", "name/cdb:action2"]

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

```
"action":["name/cdb:Describe*"]
```

#### **Resource path**

Resource paths are generally in the following format:

qcs::service\_type::account:resource

service\_type: Describes the product abbreviation, such as cdb here.

account: Describes the root account of the resource owner, such as uin/326xxx46.

resource: Describes the detailed resource information of the specific service. Each TencentDB for MySQL instance (instanceId) is a resource.

Below is a sample:

```
"resource": ["qcs::cdb::uin/326xxx46:instanceId/cdb-kfxxh3"]
```

Here, cdb-kfxxh3 is the ID of the TencentDB for MySQL instance resource, i.e., the resource in the CAM policy statement.



#### Example

The following example only shows the usage of CAM. For the complete list of APIs used to set MySQL backup download rules, see the API documentation.

#### Customizing CAM policy for setting MySQL backup download rules

1. Log in to the CAM console with the root account, select **Policies** on the left sidebar, and click **Create Custom Policy**.

P	Policy						
	Associate users or user groups with policies to grant permissions.						
	Create Custo	m Policy Delete					
		Policy Name	Service Type 🔻	Description			
		AdministratorAccess		This policy allows you to manage all users under your account and their permissions, financial information and dou			

2. In the pop-up dialog box, select Create by Policy Generator.

3. On the Select Service and Action page, select configuration items, click Add Statement, and click Next.

Effect: Select **Allow** or **Deny** for the action.

Service: Select TencentDB for MySQL.

Action: Select all APIs of setting MySQL backup download rules. For more information, see the API documentation. Resource: For more information, see Resource Description Method. You can enter \* to indicate that the backup download rules of TencentDB for MySQL instances in the specified region can be set.

1 Edit Policy > 2	Associate Users/User Groups
Visual Policy Generator	JSON
▼ Cloud Database(91 actions	)
Effect *	O Allow O Deny
Service *	Cloud Database (cdb)
Action * Collapse	Select actions         All actions (cdb:*)       Show More         Read Show More         Write (91 selected) Show More         List Show More
Resource *	All resources (*)
Condition	Source IP (i) Add other conditions.

4. On the Edit Policy page, enter the Policy Name (such as BackupDownloadRestriction ) as required and Description and click Done.

Basic Info	
Policy Name *	BackupDownloadRestriction
Description	Please enter the policy description
Associate Users/Use	er
Groups	
Authorized Users	Select Users
Authorized User Groups	Select User Groups

Associate users or user groups with policies to grant permissions.					
Create Custor	m Policy Delete				
	Policy Name	Service Type <b>T</b>	Descript		
	BackupDownloadRestriction	-	-		

# Viewing Backup Space

Last updated : 2024-07-23 10:11:56

This document describes how to view the TencenDB for MySQL instance backup space and free tier in the console.

# Overview

For two-node and three-node instances of local disk edition, the backup space occupied by TencentDB for MySQL instance backup files is allocated by region. It is equivalent to the total storage capacity used by all database backups in a region, including automatic data backups, manual data backups, and log backups. Increasing the backup retention period or manual backup frequency will use more database backup space.

For single-node instances of cloud disk edition, the backup space occupied by TencentDB for MySQL instance backup files is allocated by instance. The free backup space of one instance is equivalent to 200% of its storage space.

# Viewing the backup space and free tier of an instance of local disk edition

1. Log in to the TencentDB for MySQL console and select Database Backup on the left sidebar.

2. Select a region at the top to view its backup information on the **Overview** tab, including total backup, backup trend, and real-time backup statistics.

Total Backup: This section displays the size and quantity of all backups, data backups, and log backups as well as the free tier occupied by all the backups.

#### Note:

Green: The total backup space used does not exceed the free tier.

Orange: The total backup space used has exceeded the free tier and incurred fees. For more information, see Backup Space Billing.





Backup Trend: This section displays the trends of all backups, data backups, and log backups as well as the free space.

Real-Time Backup Statistics: This section displays the instance IDs/names (you can click an ID/name to enter the instance details page), backup space (which can be sorted by size), and data/log backups in the selected region. You can search for an instance by ID/name in the search box in the top-right corner.

3. Select **Backup List** at the top, where the backup list is divided into data backups and log backups. Click an instance name in the list to enter the instance details page. The backup list supports filtering by time period and fuzzy search by instance ID/name.

Real-ti	me Backup Statistics 🕄	)			Enter insta	nce ID to search <b>Q</b>
	Instance ID/Name	Backup Space ↓	Data Backup 💲	Log Backup 💲	Auto Backup \$	Manual Backup \$
	de 1.	4.77MB	4.77MB 7items	0B Oitems	4.77MB 7items	0B Oitems
		4.7MB	4.7MB 7items	0B Oitems	4.7MB 7items	0B 0items
Total	2 items				Lines per page 20 💌 🔫	< 1 /1 page

#### Data Backup List

Filtering by information field is supported:

Type: All, logical cold backup, physical cold backup.

Backup Mode: All, Automatic, Manual.

Backup Method: Currently, only full backup is supported.

Backups can be sorted by backup time, task start time, task end time, or backup size.

Click **Details** in the **Operation** column to enter the **Backup and Restoration** tab, where you can click **Download** to download backups. You can only **delete** manual backups

#### Log Backup List

Backups can be sorted by log data start time or end time.

Click **Details** in the **Operation** column to enter the **Backup and Restoration** tab, where you can click **Download** to download logs.

# Viewing the backup space and free tier of an instance of cloud disk edition

1. Log in to the TencentDB for MySQL console. In the instance list, click the ID of the target instance or **Manage** in the **Operation** column to enter the instance management page.



2. On the instance management page, select **Backup and Restoration** to view the used backup space and free space of the instance.

	Instance Details Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	
				Used spa	ace/Total free space	
	Clone Manual Backup Auto-Ba	ackup Settings Confi	gure Local Binlog	Used Space 27.88 GB / 100	GB (i) Snapshot Backup 27.88 GE	3 Log Backup 1.6
Data Backup List Cloned Instance List						

### FAQs

How will I be charged for backup space beyond the free tier? How can I reduce the costs of backup space?

For more information, see Backup Space Billing.
# **Configuring Local Binlog Retention Policy**

Last updated : 2025-06-10 10:06:45

This document describes how to configure the local binlog retention period for a TencentDB for MySQL instance in the console.

#### Note:

Single-node local disk instances (read-only instances) do not support configuring local binlog retention.

Cluster edition instances, dual-node instances, three-node instances, single-node cloud disk instances, and disaster recovery instances support configuring local binlog retention. The retention period policies are as follows.

The default local binlog retention period (in hours) for cluster edition instances, dual-node instances, and three-node instances is 120, with a range from 6 to 168.

The default local binlog retention period (in hours) for disaster recovery instances is 120, with a range from 120 to 168.

The default local binlog retention period (in hours) for single-node cloud disk instances is 120, with a range from 0 to 168.

If there is no disaster recovery instance for dual-node or three-node instances, the local binlog retention period (in hours) of the primary instance ranges from 6 to 168. If there is a disaster recovery instance or you plan to add one, to avoid synchronization exceptions, the local binlog retention period (in hours) of the primary instance should be no less than 120 hours, with a range from 120 to 168.

### **Binlog Description**

Binlog grows fast when a TencentDB for MySQL instance executes large transactions or lots of DML operations. Binlog is split every 256 MB and uploaded to COS. You can see the uploaded binlog files in the log list in the console.



### Overview

Before being uploaded to COS, binlog files are stored on the instance disk (i.e., locally). You can set the local binlog retention period, control the maximum percentage of disk space binlog can take up, or expand disk space in the console. We recommend that you clear data no longer used to keep the disk utilization below 80%.

MySQL's data sync is based on binlog. To ensure database restorability, stability, and high availability, binlog cannot be disabled in TencentDB for MySQL.

The generated binlog files are automatically backed up to COS through the automatic backup feature provided by TencentDB for MySQL. The binlog files whose backups are already uploaded to COS will be cleared according the local binlog retention policy. To prevent exceptions, the binlog files in use cannot be cleared even they expire. Therefore, the local binlog clearing has a delay.

#### Note:

Rule for clearing expired binlog files:

The local binlog files are checked once every 60 seconds. If a binlog file's start time or space usage does not meet the set retention rule, it will be added to the queue of to-be-deleted files. The binlog files in the queue will be sorted by time and deleted starting from the oldest file one by one until the queue is cleared.

### Directions

1. Log in to the TencentDB for MySQL console and click an instance ID in the instance list to enter the instance management page.

2. On the instance management page, select the **Backup and Restoration** and click **Configure Local Binlog**.

3. In the pop-up dialog box, enter the desired retention period and space utilization, confirm the settings, and click OK

. You can refer to the corresponding instance operations below for settings.

Operations for Dual-Node and Three-Node Instances

Single-Node Cloud Disk Instance Operations

**Disaster Recovery Instance Operations** 

Cluster Edition Instance Operations



Retention Period (hours): The maximum retention time for local binlog files after log backup is enabled. The default value is 120. You can set it to an integer in the range from 6 to 168.

#### Note:

If the local binlog retention period configured for dual-node or three-node instances is less than 120 hours, disaster recovery instances cannot be purchased for the primary instance. To mount a disaster recovery instance, the local binlog retention period of the primary instance should be no less than 120 hours.

Space Utilization Threshold (%): The default value is 30. You can set it to an integer in the range from 30 to 50. Local binlog space utilization = Local binlog size/Total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.



Configure Local Binlog		×					
Retention Period (hour)	- 120 + Range of integers: 0 - 168						
	Specify the longest retention period of local binlogs after the log backup feature is enabled.						
Space Utilization Threshold (%)	- 30 + Range: integers from 30 to 50						
	Local binlog space utilization = local binlog size / total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.						
	Restore Defaults						
	OK Cancel						

Retention Period (hours): The maximum retention time for local binlog files after log backup is enabled. The default value is 120. You can set it to an integer in the range from 0 to 168.

Space Utilization Threshold (%): The default value is 30. You can set it to an integer in the range from 30 to 50.

Local binlog space utilization = Local binlog size/Total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.

Configure Local Binlog					×
Retention Period (hour)	_	120	+	Range of integers: 120 - 168	
	Not durat	e To ens ion of this	ure the s instar	e disaster recovery synchronization works properly, the local retention nce should be at least 120 hours.	
Space Utilization Threshold (%)	—	30	+	Range: integers from 30 to 50	
	Local space clear are u	binlog sp e is recycl ed until th nexpired.	ace ut able. I e binlo	tilization = local binlog size / total purchased instance space. The binlog f the binlog space is used up, some of the earliest binlogs will be og space utilization is lower than the threshold and all remaining binlogs	
	R	estore Def	faults		
				OK Cancel	

Retention Period (hours): The maximum retention time for local binlog files after log backup is enabled. The default value is 120. You can set it to an integer in the range from 120 to 168.

Space Utilization Threshold (%): The default value is 30. You can set it to an integer in the range from 30 to 50.

Local binlog space utilization = Local binlog size/Total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.

Configure Local Binlog				×					
Retention Period (hour)	- 120	+	Range of integers: 6 - 168						
	Specify the lo	ngestı	retention period of local binlogs after the log backup feature is enabled.						
Space Utilization Threshold (%)	- 30	+	Range: integers from 30 to 50						
	Local binlog s space is recyc cleared until t are unexpired	Local binlog space utilization = local binlog size / total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.							
	Restore De	efaults							
			OK Cancel						

Retention Period (hours): The maximum retention time for local binlog files after log backup is enabled. The default value is 120. You can set it to an integer in the range from 6 to 168.

Space Utilization Threshold (%): The default value is 30. You can set it to an integer in the range from 30 to 50. Local binlog space utilization = Local binlog size/Total purchased instance space. The binlog space is recyclable. If the binlog space is used up, some of the earliest binlogs will be cleared until the binlog space utilization is lower than the threshold and all remaining binlogs are unexpired.

### FAQs

### Will database restoration be affected if the local binlog retention period is too short?

No, because the generated binlog files will be uploaded to COS through the automatic backup feature as soon as possible, and those not uploaded yet cannot be cleared. However, too short a retention period will affect the speed of rollback.

### What is the default retention policy for local binlog?

The default local binlog retention period is 120 hours, and the space utilization will not exceed 30%. You can set the local binlog retention period according to your needs.

#### Will binlog files take up the instance disk space?

Yes. Before the binlog files are uploaded to COS and cleared according to the retention policy, they will be stored on the instance disk.

#### Why can I only set the minimum local binlog retention period to 120 hours?

Check whether your instance is a dual-node or three-node instance with a disaster recovery instance mounted. If a disaster recovery instance is mounted, to avoid synchronization exceptions, the minimum local binlog retention period of the primary instance can only be set to 120 hours.

Check whether the instance for which you are trying to modify the local binlog retention period is a disaster recovery instance. If it is a disaster recovery instance, to ensure proper synchronization, you need to set the retention period within the range from 120 to 168 hours.

#### Why can't I purchase a disaster recovery instance for my primary instance?

To create a disaster recovery instance for the primary instance, the following conditions should be met. Check accordingly.

The primary instance should be a dual-node or three-node instance.

The local binlog retention period of the primary instance should be greater than or equal to 120 hours.

The primary instance should have at least 1 GB memory and 50 GB hard disk space, and the version should be MySQL 5.6 or later (for MySQL 5.6, submit a ticket to apply for this feature), with the engine being InnoDB.

# Restoring Database from Backup File Restoring Database from Physical Backup

Last updated : 2024-11-27 15:17:21

### Overview

#### Note:

To save the storage space, physical and logical backups in TencentDB for MySQL will be compressed with qpress and then packed with xbstream offered by Percona.

The open-source Percona XtraBackup can be used to back up and restore databases. This document describes how to use XtraBackup to restore a physical backup file of TencentDB for MySQL instance to a CVM-based self-built database.

#### Note:

If you use the TDE or Instant DDL feature, you cannot restore data from a physical backup in a self-built system. XtraBackup only supports Linux but not Windows.

For more information about how to restore data in Windows, see Offline Data Migration > Data Migration with Command Line Tool.

### Prerequisites

Download and install XtraBackup.

For MySQL 5.6 and 5.7, download Percona XtraBackup 2.4.6 or later at Percona official website. For more information on installation, see Percona XtraBackup 2.4 documentation.

For MySQL 8.0, download Percona XtraBackup 8.0.22-15 or later at Percona official website. For more information on installation, see Percona XtraBackup 8.0 documentation.

Supported instance architectures: two-node or three-node MySQL

Instances with TDE enabled cannot be restored from a physical backup.

#### Note:

This document takes a CVM instance on CentOS and a MySQL 5.7 instance as an example.

### Step 1. Download the backup file



You can download data backups and log backups of TencentDB for MySQL instances in the console.

#### Note:

Each IP can have up to 10 download links by default, with a download speed limit of 20–30 Mbps each.

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.

2. On the **Backup and Restoration** > **Data Backup List** tab, locate the backup file to be downloaded and click **Download** in the **Operation** column.

3. Copy the download address in the pop-up dialog box, log in to the Linux CVM in the same VPC as the TencentDB instance as instructed in Customizing Linux CVM Configurations, and run wget to download the file over the high-speed private network.

#### Note:

You can also click **Download** to download it directly. However, this may take a longer time.

wget command format: wget -c 'backup file download address' -O custom filename.xb

Below is a sample:

```
wget -c 'https://mysql-database-backup-sh-1218.cos.ap-
nanjing.myqcloud.com/12427%2Fmysql%2F0674-ffba-11e9-b592-
70bd%2Fdata%2Fautomatic-delete%2F2019-12-
03%2Fautomatic%2Fxtrabackup%2Fbk_61_156758150%2Fcdb-
293fl9ya_backup_20191203000202.xb?sign=q-sign-algorithm%3Dsha1%26q-
ak%3DAKzxfbLJ1%26q-sign-time%3D1575374119%3B1575417319%26q-key-
time%3D1575374119%3B1575417319%26q-header-list%3D%26q-url-param-list%3D%26q-
signature%3Dba959757&response-content-
disposition=attachment%3Bfilename%3D%22yuan177685_backup_20191203000202.xb%22&r
esponse-content-type=application%2Foctet-stream' -0 /data/test.xb
```

# Step 2. Download the backup decryption key (required only if the backup encryption feature is enabled)

You can download data backups and decryption keys of TencentDB for MySQL instances in the console. **Note:** 

A decryption key is generated for each database backup separately. If the backup encryption feature is enabled, you need to download and save the backup file together with the decryption key.

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.

2. On the **Backup and Restoration** > **Data Backup List** tab, locate the decryption key of the backup file to be downloaded and click **Download Key** in the **Operation** column.

Data B	Data Backup List         Log Backup List         Cloned Instance List											
All	Today Last 7 days	Last 15 days	Last 30 days	Select date Select date				Download over I	Public Network Enabled			
File Nam	ne Backup Time	Task Start Time Task End Time	Backup Size	Туре	Backup Mode	Backup Method	Backup Region	Status	Operation			
								Successful	Download Clone Download Key Delete			

### Step 3. Restore data

### 3.1 Unpack the backup file

Unpack the data backup file to the target directory using the xbstream command.

```
xbstream -x --decrypt=AES256 --encrypt-key-file=<backup key file> --parallel=2
-C /data/mysql < /data/test.xb</pre>
```

#### Note:

The target directory /data/mysql is used as an example in this document. You can replace it with the directory you actually use to store the backup file.

Replace /data/test.xb with your backup file.

The unpacking result is as shown below:

[root@		my	ysql]# l1	l/d	lata/mysql
2920					
-rw-r 1	root	root	459	4	20 10:12 backup-my.cnf.qp
drwxr-x 2	root	root	4096	4	20 10:12 <u>cdb_recycle_bin_</u>
-rw-r 1	root	root	991	4,	20 10:12 ib_buffer_pool.qp
-rw-r 1	root	root	161172	4	20 10:12 ibdata1.qp
drwxr-x 2	root	root	4096	4	20 10:12 mysql
-rw-r 1	root	root	222	4,	20 10:12 mysql-bin.000033.qp
-rw-r 1	root	root	125	4	20 10:12 mysql-bin.index.qp
-rw-r 1	root	root	2366987	4	20 10:12 mysql.ibd.qp
drwxr-x 2	root	root	4096	4	20 10:12 performance_schema
drwxr-x 2	root	root	4096	4, 1	20 10:12 sys
-rw-r 1	root	root	199604	4, [	20 10:12 undo_001.qp
-rw-r 1	root	root	198976	4	20 10:12 undo_002.qp
-rw-r 1	root	root	158	4	20 10:12 xtrabackup_binlog_info.qp
-rw-r 1	root	root	95	4	20 10:12 xtrabackup_checkpoints
-rw-r 1	root	root	872	4	20 10:12 xtrabackup_info.qp
-rw-r 1	root	root	540	4	20 10:12 xtrabackup_logfile.qp
-rw-r 1	root	root	191	4)	20 10:12 xtrabackup_slave_info.qp
-rw-r 1	root	root	130	4厂	20 10:12 xtrabackup_tablespaces.qp

If the backup file to be unpacked is a binlog backup, use the following command.

```
openssl enc -d -aes256 -k <encrypt_key> -in <download_binlog_name> -out
<output_binlog_name>;
```

<encrypt\_key>: Replace it with the actual backup key file content.

<download\_binlog\_name>: Replace it with the actual encrypted binlog file name.

<output\_binlog\_name>: Replace it with the actual decrypted binlog file name.

#### 3.2 Decompress the backup file

1. Download qpress by running the following command.

```
wget -d --user-agent="Mozilla/5.0 (Windows NT x.y; rv:10.0) Gecko/20100101
Firefox/10.0" https://docs-tencentdb-1256569818.cos.ap-
guangzhou.myqcloud.com/qpress-11-linux-x64.tar
```

#### Note:

If an error is displayed during the wget download, you can click here to download qpress locally and upload it to the Linux CVM instance. For more information, see Uploading Files from Linux or MacOS to Linux CVM via SCP.

2. Extract the qpress binary files by running the following command.

```
tar -xf qpress-11-linux-x64.tar -C /usr/local/bin
source /etc/profile
```

3. Then, decompress all . qp files in the target directory by running the following command:



xtrabackup --decompress --target-dir=/data/mysql

#### Note:

/data/mysql is the target directory where the backup file was previously stored. You can replace it with the directory you actually use.

The --remove-original option is supported only in Percona Xtrabackup 2.4.6 and later.

xtrabackup won't delete the original files during decompression by default. If you want to delete them upon the completion of decompression, add the --remove-original parameter to the above command.

[root@VM_17_57_centos data]# xtrabackupdecompresstarget-dir=/data
xtrabackup: recognized server arguments:datadir=/var/lib/mysqllog_bin=/var/lib/mysql/mysql-binserver-id=1573
xtrabackup: recognized client arguments:datadir=/var/lib/mysqllog_bin=/var/lib/mysql/mysql-binserver-id=1573decompress=1target-dir=/data
xtrabackup version 2.4.12 based on MySQL server 5.7.19 Linux (x86_64) (revision id: 170eb8c)
181204 10:39:13 [01] decompressing ./test/fcl.ibd.qp
181204 10:39:13 [01] decompressing ./test/fc2.ibd.qp
181204 10:39:13 [01] decompressing ./test/sbtest4.frm.qp
181204 10:39:13 [01] decompressing ./test/sbtest3.ibd.qp
181204 10:39:14 [01] decompressing ./test/fcl.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest5.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest7.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest2.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtestl.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest4.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest9.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest5.ibd.qp
181204 10:39:14 [01] decompressing ./test/fc2.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest10.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest7.ibd.qp
181204 10:39:14 [01] decompressing ./test/mytest.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest2.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest6.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtestl.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest10.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest3.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest8.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest6.ibd.qp
181204 10:39:14 [01] decompressing ./test/sbtest8.frm.qp
181204 10:39:14 [01] decompressing ./test/mytest.frm.qp
181204 10:39:14 [01] decompressing ./test/sbtest9.frm.qp
181204 10:39:14 [01] decompressing ./xtrabackup_logfile.qp
181204 10:39:14 [01] decompressing ./ibdatal.qp
181204 10:39:14 [01] decompressing ./undo001.qp
181204 10:39:14 [01] decompressing ./mysql/general_log.frm.qp
181204 10:39:14 [01] decompressing ./mysql/procs_priv.MYI.qp
181204 10:39:14 [01] decompressing ./mysql/ndb_binlog_index.MYI.qp
181204 10:39:14 [01] decompressing ./mysql/slow_log.frm.qp
181204 10:39:14 [01] decompressing ./mysql/general_log.CSV.qp
181204 10:39:14 [01] decompressing ./mysql/time_zone_leap_second.MYD.qp

#### 3.3 Prepare the backup file

After a backup file is decompressed, perform the apply log operation by running the following command.

xtrabackup --prepare --target-dir=/data/mysql

If the execution result contains the following output, it means that the prepare operation succeeded.

```
InnoDB: Starting shutdown...
InnoDB: Shutdown completed; log sequence number 922626089
181204 10:47:24 completed OK!
```

#### 3.4 Modify the configuration file

1. Run the following command to open the backup-my.cnf file.

```
vi /data/mysql/backup-my.cnf
```

#### Note:

The target directory /data/mysql is used as an example in this document. You can replace it with the directory you actually use.

2. Given the existing version issues, the following parameters need to be commented out from the extracted

backup-my.cnf file.

innodb\_checksum\_algorithm

innodb\_log\_checksum\_algorithm

innodb\_fast\_checksum

innodb\_page\_size

innodb\_log\_block\_size

redo\_log\_version

server\_uuid

master\_key\_id

1	# This MySQL options file was generated by innobackupex.
2	
3	# The MySQL server
4	[mysqld]
5	innodb_data_file_path=ibdata1:12M:autoextend
6	innodb_log_files_in_group=2
7	innodb_log_file_size=536870912
8	<pre>innodb_undo_directory=.</pre>
9	innodb_undo_tablespaces=0
10	server_id=0
11	<pre>#innodb_checksum_algorithm=innodb</pre>
12	<pre>#innodb_log_checksum_algorithm=innodb</pre>
13	<pre>#innodb_fast_checksum=false</pre>
14	<pre>#innodb_page_size=16384</pre>
15	<pre>#innodb_log_block_size=512</pre>
16	<pre>#redo_log_version=0</pre>

#### 3.5 Modify file attributes

Modify file attributes and check whether files are owned by the mysql user.

chown -R mysql:mysql /data/mysql

-rw-r	1	mysql	mysql	424	11	9	18:15	backup-my.cnf
-rw-r	1	mysql	mysql	12582912	11	9	18:12	ibdata1
-rw-r	1	mysql	mysql	12582912	11	9	16:30	ibtmp1
-rw-r	1	mysql	mysql	10485760	11	9	16:30	undo001
drwxr-x	2	mysql	mysql	4096	11	9	16:30	mysql
drwxr-x	2	mysql	mysql	4096	11	9	16:30	test
drwxr-x	2	mysql	mysql	4096	11	9	16:30	performance_schema
drwxr-x	2	mysql	mysql	4096	11	9	16:30	sbtest

### Step 4. Start the mysqld process and log in for verification

1. Start the mysqld process.

```
mysqld_safe --defaults-file=/data/mysql/backup-my.cnf --user=mysql --
datadir=/data/mysql &
```

2. Log in to the MySQL client for verification.

mysql -uroot

```
Welcome to the MySQL monitor.
                             Commands end with ; or \g.
Your MySQL connection id is 2
Server version: 5.6.23 Source distribution
Copyright (c) 2000, 2011, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
mysql> show databases;
 Database
 information schema
 test
 rows in set (0.00 sec)
mysql>
```

### FAQs About Backup



See FAQs and Troubleshooting.

# **Restoring Database from Logical Backup**

Last updated : 2024-07-23 10:13:16

### Overview

#### Note:

To save the storage space, physical and logical backups in TencentDB for MySQL will be compressed with qpress and then packed with xbstream offered by Percona.

TencentDB for MySQL supports logical backup as described in Backing up Databases. In the console, you can manually create logical backup files of an entire instance or specified databases/tables and download them. This document describes how to manually restore data from logical backup files.

The restoration method described in this document only applies to Linux.

For more information about how to restore data in Windows, see Offline Data Migration > Data Migration with

#### Command Line Tool.

Supported instance architectures: two-node or three-node MySQL

### Directions

### Step 1. Download the backup file

1. Log in to the TencentDB for MySQL console. In the instance list, click an instance ID or **Manage** in the **Operation** column to enter the instance management page.

2. On the **Backup and Restoration** > **Data Backup List** tab, locate the backup file to be downloaded and click **Download** in the **Operation** column.

3. We recommend that you copy the download link in the pop-up dialog box, log in to a (Linux) CVM instance in the same VPC as the TencentDB instance, and run the wget command for download over the private network at a higher speed. For more information, see Customizing Linux CVM Configurations

### Note:

You can also click **Download** to download it directly. However, this may take a longer time.

wget command format: wget -c 'backup file download address' -O custom filename.xb

#### Below is a sample:

```
wget -c 'https://mysql-database-backup-bj-118.cos.ap-
beijing.myqcloud.com/12427%2Fmysql%2F42d-11ea-b887-6c0b82b%2Fdata%2Fautomatic-
delete%2F2019-11-28%2Fautomatic%2Fxtrabackup%2Fbk_204_10385%2Fcdb-
1pe7bexs_backup_20191128044644.xb?sign=q-sign-algorithm%3Dsha1%26q-ak%3D1%26q-
sign-time%3D1574269%3B1575417469%26q-key-time%3D1575374269%3B1517469%26q-
```

### S Tencent Cloud

```
header-list%3D%26q-url-param-list%3D%26q-signature%3Dfb8fad13c4ed&response-
content-
disposition=attachment%3Bfilename%3D%2141731_backup_20191128044644.xb%22&respon
se-content-type=application%2Foctet-stream' -0 test0.xb
```

### Step 2. Unpack the backup file

Unpack the backup file with xbstream.

### Note:

xbstream can be downloaded at Percona official website. Select Percona XtraBackup 2.4.6 or later. For more information on installation, see Installing Percona XtraBackup on Red Hat Enterprise Linux and CentOS.

xbstream -x < test0.xb</pre>

### Note:

Replace test0.xb with your backup file.

The unpacking result is as shown below:

```
ubuntu@VM-15-106-ubuntu:~$ xbstream -x < test0.xb
ubuntu@VM-15-106-ubuntu:~$ 11
total 1090720
                                 4096 Dec 2 18:30 ./
drwxr-xr-x 9 ubuntu ubuntu
drwxr-xr-x 3 root
                                 4096 Oct 26
                                              2016 ../
                    root
                               702901 Nov 23 08:37 1254408587%2Fmysql%2F7f129d66-ac5f-11e9-b8bc-
-rw-rw-r-- 1 ubuntu ubuntu
cuoSr
-rw-rw-r-- 1 ubuntu ubuntu
                                  419 Nov 25 21:40 backup-my.cnf
                                  396 Nov 25 20:06 backup-my.cnf.qp
rw-rw---- 1 ubuntu ubuntu
 rw-r--r-- 1 ubuntu ubuntu
                                23576 Dec
                                           2 18:30 .bash_history
                                  220 Oct 26
                                              2016 .bash_logout
 rw-r--r-- 1 ubuntu ubuntu
     --r-- 1 ubuntu ubuntu
                                 3773 Mar
                                          16
                                              2018 .bashrc
   -\mathbf{r}
           2 ubuntu ubuntu
                                 4096
                                      Nov 28
                                              2017
irwx-
                                             18:30 cdb-jp0zua5k_backup_20191202182218.sql.qp
           1 ubuntu ubuntu
                               408724 Dec
                                           2
  w-rw----
           1 ubuntu ubuntu
                             12582912 Nov 25
                                             21:42
                                                    ibdatal
   -rw-r--
```

### Step 3. Decompress the backup file

1. Download qpress by running the following command.

```
wget -d --user-agent="Mozilla/5.0 (Windows NT x.y; rv:10.0) Gecko/20100101
Firefox/10.0" https://docs-tencentdb-1256569818.cos.ap-
guangzhou.myqcloud.com/qpress-11-linux-x64.tar
```

#### Note:

If an error is displayed during the wget download, you can click here to download qpress locally and upload it to the Linux CVM instance. For more information, see Uploading Files from Linux or MacOS to Linux CVM via SCP. 2. Extract the qpress binary files by running the following command.

tar -xf qpress-11-linux-x64.tar -C /usr/local/bin

```
source /etc/profile
```

3. Decompress the backup file with qpress.

qpress -d cdb-jp0zua5k\_backup\_20191202182218.sql.qp .

#### Note:

Find the backup file with .sql.qp extension by decompression time and replace cdb-

jp0zua5k\_backup\_20191202182218 with its filename.

The decompressing result is as shown below:

ubuntu@VM-15	-106-ub	untu:~\$	qpress -d	cdb	-jp(	)zua5k	_backup_20191202182218.sql.qp .			
ubuntu@VM-15-106-ubuntu:~\$ 11										
total 1091792										
drwxr-xr-x 9	ubuntu	ubuntu	4096	Dec	2	18:32	./			
drwxr-xr-x 3	root	root	4096	Oct	26	2016	/			
-rw-rw-r 1	ubuntu	ubuntu	702901	Nov	23	08:37	1254408587%2Fmysql%2F7f129d66-ac5f-11e			
kuoSr										
-rw-rw-r 1	ubuntu	ubuntu	419	Nov	25	21:40	backup-my.cnf			
-rw-rw 1	ubuntu	ubuntu	396	Nov	25	20:06	backup-my.cnf.qp			
-rw-rr 1	ubuntu	ubuntu	23657	Dec	2	18:32	.bash_history			
-rw-rr 1	ubuntu	ubuntu	220	0ct	26	2016	.bash_logout			
-rw-rr 1	ubuntu	ubuntu	3773	Mar	16	2018	.bashrc			
drwx 2	ubuntu	ubuntu	4096	Nov	28	2017	.cache/			
-rw-rw-r 1	ubuntu	ubuntu	1097176	Dec	2	18:32	cdb-jp0zua5k_backup_20191202182218.sql			
-rw-rw 1	ubuntu	ubuntu	408724	Dec	2	18:30	cdb-jp0zua5k_backup_20191202182218.sql			
-rw-rw-r 1	ubuntu	ubuntu	12582912	Nov	25	21:42	ibdatal			
-rw-rw 1	ubuntu	ubuntu	179769	Nov	25	20:06	ibdatal.gp			

### Step 4. Import the backup file into the target database

Import the .sql file into the target database by running the following command:

mysql -uroot -P3306 -h127.0.0.1 -p < cdb-jp0zua5k\_backup\_20191202182218.sql</pre>

#### Note:

This document takes importing into a local MySQL instance with port 3306 as an example. You can replace it as needed.

**Replace** cdb-jp0zua5k\_backup\_20191202182218.sql with the .sql file extracted by qpress.

# **Downloading Backups**

Last updated : 2025-06-10 10:16:10

The TencentDB for MySQL console provides the list of backup files that can be downloaded. You can use the downloaded backups to restore data from one database to another (such as a self-built one). This document describes how to download a backup file in the console.

#### Note:

The download rule varies by region. After the cross-region backup feature is enabled, you need to configure the download rule for each region selected for backup.

Data backups generated by manual backup for single-node instances of the cloud disk edition are snapshot backups, which cannot be downloaded.

### File types supported for download, decompression, and deletion

Category	Backup Type	Method	Download	Decompression Required After Download	Deletion
Data	Logical backup	Manual	$\checkmark$	$\checkmark$	5
backup	Physical	Manual	1	1	1
	backup	Automatic	1	1	×
Log backup	Physical backup	Automatic	✓	✓	×

### Downloading a data backup file

- 1. Log in to the TencentDB for MySQL console.
- 2. Select the region at the top, find the target instance, and click the **instance ID** or **Manage** in the **Operation** column
- to enter the instance management page.
- 3. On the instance management page, select the **Backup and Restoration** tab and click **Data Backup List**.



4. Click **Download** in the **Operation** column in the backup list to enter the download page. Then, select the backup file in the target region, click **Copy Download URL** for high-speed download with the wget command, or directly click **Download**.

#### Note:

We recommend that you copy the download address, log in to a Linux CVM instance in the same VPC as the TencentDB instance as instructed in Customizing Linux CVM Configurations, and run the wget command for fast download over the private network.

The download address is valid for 12 hours, after which you will need to enter the download page again to get a new one.

The URL must be enclosed with quotation marks when the wget command is used to download.

wget command format: wget -c 'backup file download address' -O custom filename.xb . 5. View the downloaded file.

### Downloading a log backup file

1. Log in to the TencentDB for MySQL console.

2. Select the region at the top, find the target instance, and click the **Instance ID** or **Manage** in the **Operation** column to enter the instance management page.

3. On the instance management page, select the **Backup and Restoration** tab and click **Log Backup List**.

4. Click **Download** in the **Operation** column in the backup list to enter the download page. Then, select the backup file in the target region, click **Copy Download URL** for high-speed download with the wget command, or directly click **Download**.

#### Note:

We recommend you copy the download address, log in to a Linux CVM instance in the same VPC as the TencentDB instance as instructed in Customizing Linux CVM Configurations, and run the wget command for fast download over the private network.

The download address is valid for 12 hours, after which you will need to enter the download page again to get a new one.

The URL must be enclosed with quotation marks when the wget command is used to download.

wget command format: wget -c '<backup file download address>' -0 <custom filename> .
5. View the downloaded file.

### Decompressing a backup file

For more information on how to decompress a physical backup file, see Restoring Database from Physical Backup. For more information on how to decompress a logical backup file, see Restoring Database from Logical Backup.

# **Rolling back Databases**

Last updated : 2025-06-10 10:18:39

### Overview

TencentDB for MySQL will not modify any of your data. Data corrupted due to personal reasons can be recovered through rollback in a self-service manner. A rollback feature is provided to roll back databases or tables in Tencent Cloud based on data backup and binlog backup. Real-time data rollback is supported.

By reconstructing periodical images and real-time transactions, the TencentDB for MySQL rollback feature can roll back a database or table to the specified point in time, and the time slices of all data are guaranteed to be identical. A new database or table will be generated in the original instance, and during the process, the original database or table can be accessed normally. Upon the completion of rollback, you can see both the new and original databases or tables.

### How Rollback Works

The rollback feature can roll back databases or tables to a specified point in time based on cold data backup and corresponding binlog backup .

1. Data is exported from the MySQL replica and imported to the cold backup system daily.

2. To roll back databases or tables, request for a temp rollback instance from the rollback system. Export the cold data from the cold backup system and import it to the temp rollback instance (types of imported data vary by rollback method).

3. Establish a source-replica relationship between the rollback instance and MySQL source instance, set the rollback time, and specify the databases or tables to be rolled back.

4. Replicate the rollback databases or tables to the MySQL source instance.

### Feature Limits

Single-node instances of the cloud disk edition cannot be rolled back.

Only source instances can be rolled back. Read-only instances and disaster recovery instances cannot be rolled back. Only specified databases or tables can be rolled back. The databases or tables after rollback will be renamed and replicated to the source instance. MySQL 5.6, 5.7, and 8.0 support rollback, while MySQL 5.5 does not support rollback. If you are using MySQL 5.5, we recommend that you upgrade it to a later version as soon as possible. For details, see Upgrading Database Engine Versions.

For database-level rollback, the following object types will not be imported: foreign key constraint, check constraint, trigger, stored procedure, user view, and event.

### Note

The rollback feature is subject to the backup cycle and retention days set for automatic backup. It enables data rollback based on data backup and binlog backup according to the configured retention days and backup cycle. For the backup cycle settings, see Backing up Databases. To ensure MySQL data security, set the automatic backup cycle to at least twice a week.

If a monthly subscribed instance has not expired but your account has overdue payments, the backup service will be downgraded with rollback disabled. If you need to roll it back, top up your account to a positive balance. If the database or table to be rolled back does not exist or has been dropped, you need to log in to the TencentDB instance and create one first before performing rollback in the console.

If the cold backup before rollback does not contain the table, the rollback will fail.

### Directions

1. Log in to the TencentDB for MySQL console, select one or multiple target instances in the instance list, and select **More** > **Roll Back**.

### Note:

If rollback is to be performed on only one instance, you can also go to the instance management page and click **Roll Back** in the top-right corner.

Up to 5 rollback tasks can be initiated at a time under the same APPID.

2. On the rollback page, select the original database or table and click **Next: set the rollback time and database table name**.

### Rollback methods:

Fast mode: Import full backup of the instance, and then roll back the selected databases and tables. This rollback mode is slower than the other modes but has no limit.

Faster mode: Full backup + database-level binlog. For cross-database operation, if associated database is not selected at the same time, the database callback will fail.

Ultrafast mode: Full backup + table-level binlog. For cross-table operation, if associated table is not selected at the same time, the table rollback will fail.



#### Note:

Only databases/tables with names containing digits, letters, underscores, or their combinations can be rolled back. Databases/tables with names containing special symbols are not supported.

In the mode where specified databases/tables can be rolled back only, a maximum of 500 databases/tables in the same instance can be rolled back at a time.

If the rollback involves composite operations on other databases or tables during the execution of binlogs, the SQL statements may fail.

If the rollback involves foreign keys and other constraints of the table during the execution of binlogs, the SQL statements may fail.

3. Set the post-rollback database or table names and rollback time, and then click Roll Back.

#### Note:

Only one rollback time can be set for each instance.

If you choose to set a batch rollback time, all databases or tables will be rolled back at the specified time.

If you choose to set a single-table rollback time, tables will be rolled back at their respective rollback time.

The database or table name after rollback can contain up to 64 letters, digits, or symbols (.-\_\$).

4. After submission, go to **Operation Log** > **Rollback Log** and view the rollback progress. Click **View Details** to view the rollback log in real time.

5. Once the rollback is completed, go to **Database Management** > **Database List** to view the new databases or tables after rollback in the original instance.

# **Deleting Backups**

Last updated : 2025-06-10 10:21:00

You can delete database backups to reduce backup space costs.

#### Note:

Manual backups can be deleted from the backup list in the console, but automatic backups cannot be manually deleted.

### Directions

1. Log in to the TencentDB for MySQL console, click an **Instance ID** in the instance list, and enter the instance management page.

2. On the **Backup and Restoration** tab, locate the backup you want to delete in the backup list, and click **Delete** in the **Operation** column.

3. In the pop-up dialog box, confirm the backup file to be deleted and click **OK**.

#### Note:

Once deleted, backup files cannot be recovered.

# **Cloning Instances**

Last updated : 2025-06-10 10:23:19

This document describes how to clone a TencentDB for MySQL instance and quickly restore data to the newly purchased clone in the console.

### Overview

TencentDB for MySQL provides the instance cloning feature. You can now restore an instance to any point in time within the log backup retention period or from a specific physical backup set by cloning. The clone is a new instance created from the backup data according to the restoration point in time you have specified. After the clone has been verified, you can migrate its data back to the original instance with DTS or you can start using the clone.

#### **Clone mode**

Clone an instance and restore the clone to any point in time within the log backup retention period you set.

Clone an instance and restore the clone from a specific physical backup set within the data backup retention period you set.

#### Clone billing

The clone adopts the pay-as-you-go billing mode. For more information on this billing mode and fee calculation, see Billing Overview.

The clone will not be billed until the clone process is completed.

### Prerequisites

Supported instance architectures: TencentDB for MySQL single-node, two-node, and three-node architectures.

The original instance must be in **Running** status.

If the clone mode is set to **By backup set**, the original instance must have created at least one physical backup. You can log in to the console and view the backup status on the **Backup List** tab.

Your account balance must be positive.

### **Functional Limitations**

Cloning a read-only instance is not supported.

MySQL 5.6, 5.7, and 8.0 support cloning, while MySQL 5.5 does not support cloning. If you are using MySQL 5.5, we recommend that you upgrade it to a later version as soon as possible. For details, see Upgrading Database Engine Versions.

### Note

The specification of the new instance must be greater than or equal to that of the original instance.

The hard disk space of the clone instance must be larger than the amount of the data to be cloned from, or else the clone task may fail.

The AZ, database version, replication mode, and default database parameters of the clone must be the same as those of the original instance.

The clone will not be displayed in the instance list in the console until the clone process is completed.

### Directions

1. Log in to the TencentDB for MySQL 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 **Backup and Restoration** > **Data Backup List** and click **Clone** in the top-left corner, or locate the target backup and click **Clone** in the **Operation** column.

3. On the displayed purchase page, specify the clone mode and other configuration items and click **Buy Now**.

By time point: You can restore an instance to any time point within the log backup retention period.

**By backup set**: You can restore data from a backup set to a new instance. The available backup sets depend on the data backup retention period.

### Note:

You can log in to the console, select **Database Backup** on the left sidebar, and view backup retention period on the **Backup List** tab.

 After successful purchase, you can check the clone details on the Backup and Restoration > Cloned Instance List tab.

5. After the clone process is completed, you can view the clone in the instance list.

### **Relevant Documentation**

For more information on restoration of a single database or table, see Rolling back Databases.

For more information on how to restore data to a self-created instance, see Restoring Database from Physical Backup and Restoring Database from Logical Backup.

### FAQs

### Will the access to the source instance be affected during the clone process?

The original backup set and binlogs uploaded to COS are used for cloning, which will not affect the access to the source instance.

# Data Migration Migrating with DTS

Last updated : 2024-07-23 10:15:54

DTS can smoothly migrate structures, full data, and incremental data of MySQL databases in local IDCs, third-party cloud vendors, and Tencent Cloud to TencentDB for MySQL without causing any server downtime. Cross-Account TencentDB Instance Migration Migration from MySQL to TencentDB for MySQL FAQs about MySQL Data Migration

# Importing SQL Files

Last updated : 2024-11-27 11:42:14

### Overview

TencentDB for MySQL supports importing SQL files via the console, allowing you to execute SQL statements in the selected database. You can also use this feature to create databases/tables and change table structures to initialize or modify the instance.

#### Note:

Only two-node and three-node TencentDB for MySQL instances support importing SQL files.

### Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an **Instance ID** or **Manage** in the **Operation** column to access the instance management page.

2. On the instance management page, select Database Management > Database List and click Import Data.

Instance Details	Instance Monitoring	Database Management	Security Group	Backup and Restoration	Operation Log	Read-Only Instance	Database Proxy	Data Encry
Database List	Parameter Settings	s Account Managemer	nt					
Import Data	Create Database							Import Recor
Database Name			Status		Database Chara	cter Set	Server Character Set	

3. In the pop-up window, click Add File to import the file. After the upload is completed, click Next.

### Note:

To avoid database unavailability caused by corruption of system tables, do not import data from system tables such as the mysql.user table.

Only incremental import is supported. Clear any obsolete data in the database before importing.

You can import an unencrypted .zip file. The unzipped files must be .sql files with a total size of no more than 5 GB.

A single file should not exceed 12 GB (if it is zipped, the file should not exceed 5 GB after being unzipped), and the file name can contain only letters, digits, and underscores.

Uploaded files are valid for 14 days and will be automatically deleted after expiration.

Select the file to be imported Select the target database Select the target									
To avoid database unavailability due to system table corruption, please do not import system table data, such as the mysql.user table. Incrementally importing data is supported only. If there is burn-in data in the database, clear data before the importing operation.									
Add File A single file does not exceed 2 GB. The file name allows English letters, numbers, and underlines.									
	File Name \$	Time \$	Size \$						
		Click "New File" to upload.							
Files are valid for 14 days and will be automatically deleted after expiration.									
Cancel Next									

- 4. On the **Select the target database** page, select the target database and click **Next**.
- 5. On the confirmation page, confirm that the imported data is correct, enter the account password, and click **Import**.

### Note:

The import cannot be rolled back. Confirm the import information before proceeding.

If you forgot the password, reset it as instructed in Resetting Password.

Select the file to	be imported > Select the target databas	e > 🕢 Confirm Importing		
The import cannot	be rolled back. Please confirm the import information.			
Import File				
File Name	Time		Size	
demo.sql	2020/04/26 11:44:50		21Byte	
Target Database Database Name	Instance			
Do not specify datab	ase Lab1-CDB01			
Database account *	Output the database account with operation permiss			
Database password	Please enter the password			
Cancel Pre	v Import			

### FAQs

# What Should I Do When Importing a SQL File Failed, with the Error Message "Import Initiation Failure (InternalError.UndefinedError)"?

The steps for setting the data import feature are: select a file to import, select the target database, and confirm to import. During the confirmation step, you need to enter the account and password of the current instance. If the input is incorrect, an import initiation failure will be prompted. Please enter the correct account and password during the confirmation step and try importing again. If you forget the password, you can reset the password and try again.

## **Migrating Data Offline**

Last updated : 2024-07-23 10:17:45

This document describes how to migrate data in the console or with the command line tool.

### Data Migration Through Console

There are two modes for migrating data through the console: physical backup and logical backup. For more information, see the following documents:

Restoring Database from Physical Backup

Restoring Database from Logical Backup

### Data Migration with Command Line Tool

1. Generate the SQL file to be imported with the MySQL command line tool "mysqldump" in the following way: Note:

The data files exported using mysqldump must be compatible with the SQL specification of your purchased

TencentDB for MySQL version. You can log in to the database and get the MySQL version information by running the select version(); command. The name of the generated SQL file can contain letters, digits, and underscores

but not "test".

Make sure that the same source and target database versions, source and target database character sets, and mysqldump tool versions are used. You can specify the character set using the parameter --default-character-set .

shell > mysqldump [options] db\_name [tbl\_name ...] > bak\_pathname

Here, options is the export option, db\_name is the database name, tbl\_name is the table name, and bak\_pathname is the export path.

For more information on how to export data with mysqldump, see MySQL official documentation.

2. Import data to the target database with the MySQL command line tool as follows:

shell > mysql -h hostname -P port -u username -p < bak\_pathname</pre>

Here, hostname is the target server for data restoration, port is the port of target server, username is the username of the database on the target server, and bak\_pathname is the full path to the backup file.

### Migrating data (Windows)

1. Use the Windows version of mysqldump to generate the SQL file to be imported. For more information, see the description in Data Migration with Command Line Tool.

2. Enter the command prompt and import the data into the target database with the MySQL command line tool.

	>mysq1	-h	-р	-u	-p < '	
Enter password:	****					

3. Log in to the target MySQL database, run the show databases; command, and you can see that the backup database has been imported into the target database.



### Migrating data (Linux)

This document uses a Linux CVM instance as an example. For more information on how to access a database from a CVM instance, see Accessing MySQL Database.

1. Log in to the CVM instance and generate the SQL file to be imported with the MySQL command line tool "mysqldump". Take the db\_blog database in TencentDB as an example:



2. Use the MySQL command line tool to restore the data to the target database.

3. Log in to the target MySQL database, run the show databases; command, and you can see that the backup database has been imported into the target database.

[root@UM_74_55_centos lib]# mysql -h localhost -u root -p db_blog < /home/db_blog.bak
Inter password. [root@UM_74_55_centos lib]# mysql -h localhost -u root -p
Enter password:
Welcome to the MariaDB monitor. Commands end with ; or \g. Your MariaDB connection id is 7
Server version: 5.5.52-MariaDB MariaDB Server
Copyright (c) 2000, 2016, Oracle, MariaDB Corporation Ab and others.
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
MariaDB [(none)]> show databases;
i Database i
++
i information_schema i
i mysql
i performance_schema i
l test
5 rows in set (0.00 sec)
Mania DB [(nema)]

### Issues with Character Set of Imported Data Files

1. If no character set is specified during data file import into TencentDB, the one set by the database will be used.

- 2. Otherwise, the specified character set will be used.
- 3. If the specified character set is different from that of TencentDB, garbled text will be displayed.



For more information, see the character set description in Use Limits.

# Monitoring and Alarms Monitoring

Last updated : 2025-05-26 21:45:56

To make it easier for you to view and stay up to date with how instances work, TencentDB for MySQL provides a wide variety of performance monitoring metrics and convenient monitoring features (custom view, time comparison, merged monitoring metrics, etc). You can log in to the TencentDB for MySQL console, and view them in **Instance Monitoring** on the instance management page.

#### Note:

You can get instance monitoring metrics by calling the GetMonitorData API or using the TencentDB for MySQL monitoring metrics in Tencent Cloud Observability Platform (TCOP).

You can create dashboards for monitoring metrics to analyze monitored data dynamically.

If one instance contains more than 300,000 tables, it can negatively impact database monitoring. Therefore, it is crucial to manage the number of tables effectively and ensure that the total number of tables in each instance stays below this threshold.

### Types of Instances for Monitoring

TencentDB for MySQL source, read-only, and disaster recovery instances as well as database proxy nodes can be monitored, and each instance is provided with a separate monitoring view for easy query.

### Monitoring Types

Four types of monitoring are available for TencentDB for MySQL: resource monitoring, engine monitoring (general), engine monitoring (extended), and deployment monitoring. You can view the metrics of different monitoring types to gain a quick and accurate understanding of how instances perform and operate. **Note:** 

TencentDB for MySQL single-node instances of cloud disk edition currently support resource monitoring and engine monitoring (general) but not engine monitoring (extended) and deployment monitoring.

**Resource monitoring** provides monitoring data of CPU, memory, disk, and network.

**Engine monitoring (general)** provides monitoring data of the number of connections, locks, hotspot tables, and slow queries, helping you troubleshoot issues and optimize the performance.

**Engine monitoring (extended)** provides a wider variety of engine-related monitoring metrics so as to assist you in identifying existing or potential database problems as much as possible.
**Deployment monitoring** provides monitoring metrics with regard to source-replica delay. It divides into source monitoring and replica monitoring:

If the instance is a source instance, the object of instance deployment monitoring is the linkage between the source instance and its hidden replica. Deployment monitoring displays the IO and SQL thread status of the hidden replica. The source-replica delay (in MB or in seconds) refers to the delay between the source instance and its hidden replica.

If the instance is a read-only instance, the object of instance deployment monitoring is the linkage between the source instance and the read-only instance. Deployment monitoring displays the IO and SQL thread status of the read-only instance. The source-replica delay (in MB or in seconds) refers to the delay between the read-only instance and the source instance.

If the instance is a disaster recovery instance:

The monitoring of disaster recovery instance deployment is mainly focused on the linkage between a disaster recovery instance and the primary instance. This monitoring system offers valuable insights into the I/O and SQL thread statuses of the disaster recovery instance, while the primary-replica delay distance and primary-replica delay time reflect the relationship between these two instances.

The monitoring of disaster recovery instance deployment is mainly focused on the linkage between a disaster recovery instance and the hidden secondary server. This monitoring system offers valuable insights into the I/O and SQL thread statuses of the hidden secondary server, while the primary-replica delay distance and primary-replica delay time reflect the relationship between the disaster recovery instance and the hidden secondary server.

### Monitoring Granularity

TencentDB for MySQL has adopted an adaptive policy for monitoring granularity since August 11, 2018, which means that you cannot select a monitoring granularity as desired for the time being. The adaptive policy is as follows:

Time Span	Monitoring Granularity	Adaptation Description	Retention Period
(0h, 4h]	5 seconds	The time span is below 4 hours, and the monitoring granularity is 5 seconds.	1 day
(4h, 2d]	1 minute	The time span is above 4 hours but below 2 days, and the monitoring granularity is 1 minute.	15 days
(2d, 10d]	5 minutes	The time span is above 2 days but below 10 days, and the monitoring granularity is 5 minutes.	31 days
(10d, 30d]	1 hour	The time span is above 10 days but below 30 days, and the monitoring granularity is 1 hour.	62 days
(30d,	1 day	The time span is above 30 days but within 180 days, the	180 days



180d]

monitoring granularity is adjusted to 1 day.

#### Note:

Currently, you can view monitoring data of TencentDB for MySQL in the past 180 days.

### **Monitoring Metrics**

Tencent Cloud Observability Platform (TCOP) offers an extensive range of monitoring metrics for TencentDB for MySQL instances, available at the instance level. For details, see the table below.

Besides mastering the running status of an instance through monitoring metrics, TencentDB for MySQL also supports monitoring events and event alarms. By setting event rules, event targets and push users, when an abnormal event is detected, the event alarm will send alarm notifications to users through various optional channels, helping users more comprehensively control the running information of the instance.

For details about event alarms, see EventBridge.

#### Note:

When you obtain monitoring data through calling APIs, the monitoring metric names used in the call parameter items should follow the Callable Metric Name column in the table below.

Monitoring Metric Name	Callable Metric Name	Unit	Metric Description
CPU Utilization	CpuUseRate	%	May exceed 100% because some idle CPU resources can be overused.
Memory Utilization	MemoryUseRate	%	May exceed 100% because some idle CPU resources can be overused.
Memory Usage	MemoryUse	MB	May exceed purchase specifications because some idle CPU resources can be overused.
Disk utilization	VolumeRate	%	Not on the allowlist: Disk utilization = Data usage space/Purchased instance space. Added to the allowlist: Disk utilization = (Data usage space + Log usage space)/Purchased instance space. <b>Note:</b> The disk utilization is calculated by the following formula: Disk utilization = (Data usage space + Log usage space)/Purchased instance space. If



			you need to display disk utilization, contact after-sales to enable this feature.
Disk usage space	RealCapacity	MB	Used space on the disk.
Disk occupied space	Capacity	MB	Space currently occupied on the disk.
IOPS	lops	Times/second	Number of I/O operations per second.
IOPS utilization	lopsUseRate	%	IOPS utilization.
Private network outbound traffic.	BytesSent	Byte/second	Number of bytes sent per second.
Private network inbound traffic.	BytesReceived	Byte/second	Number of bytes received per second.
Queries per second	Qps	Times/second	Number of SQL statements executed per second by the database (including operations such as insertion, selection, update, deletion, and replacement) primarily reflects the actual processing capability of the TencentDB for MySQL instance.
Transactions per second	Tps	Times/second	Number of transactions processed per second by the database.
Connection utilization	ConnectionUseRate	%	Current number of established connections/Maximum number of allowed connections.
Maximum Connections	MaxConnections	Count	The maximum connections.
Threads Connected	ThreadsConnected	Count	Current number of established connections.
Number of	SlowQueries	Times	Number of queries that exceed the

🔗 Tencent Cloud

Slow Queries			long_query_time threshold, measured in seconds.
Number of full-table scans	SelectScan	Times/second	Number of full-table scans per second.
Number of Queries	SelectCount	Times/second	Number of queries per second.
Number of Updates	ComUpdate	Times/second	Number of updates per second.
Number of Deletions	ComDelete	Times/second	Number of deletions per second.
Number of Insertions	ComInsert	Times/second	Number of insertions per second.
The number of Replaces	ComReplace	Times/second	Number of replacements per second.
Total Number of Requests	Queries	Times/second	All executed SQL statements, including set and show.
Query utilization	QueryRate	%	Queries per second/Recommended queries per second.
Number of temporary tables	CreatedTmpTables	Second	Number of temporary memory tables.
Number of table lock waits	TableLocksWaited	Times/second	Number of table lock waits per second.
Buffer Cache Hit Ratio	InnodbCacheHitRate	%	InnoDB cache hit rate.
Cache utilization	InnodbCacheUseRate	%	InnoDB cache utilization.
Number of disk reads	InnodbOsFileReads	Times/second	Number of InnoDB disk reads
Number of disk writes	InnodbOsFileWrites	Times/second	Number of InnoDB disk writes

🔗 Tencent Cloud

Number of Fsync operations	InnodbOsFsyncs	Times/second	Number of InnoDB Fsync operations.
Current number of open InnoDB tables	InnodbNumOpenFiles	Count	Current number of open InnoDB tables.
Buffer Cache Hit Ratio	KeyCacheHitRate	%	MyISAM cache hit rate.
Cache utilization	KeyCacheUseRate	%	MyISAM cache utilization.
Number of submissions	ComCommit	Times/second	Number of submissions per second.
Number of rollbacks	ComRollback	Times/second	Number of rollbacks per second.
Number of threads created	ThreadsCreated	Times/second	Number of threads created.
Number of running threads	ThreadsRunning	Count	Number of running threads.
Number of temporary disk tables	CreatedTmpDiskTables	Times/second	Number of temporary disk tables.
Number of temporary files	CreatedTmpFiles	Times/second	Number of temporary files.
Number of requests for next row reads	HandlerReadRndNext	Times/second	Number of requests for next row reads.
Number of internal rollbacks	HandlerRollback	Times/second	Number of transaction rollbacks per second.
Number of	HandlerCommit	Times/second	Number of internal submissions.

🔗 Tencent Cloud

internal submissions			
Number of InnoDB free pages	InnodbBufferPoolPagesFree	Count	Number of InnoDB free pages.
Total number of InnoDB pages	InnodbBufferPoolPagesTotal	Count	Total number of InnoDB pages.
InnoDB logical reads	InnodbBufferPoolReadRequests	Times/second	Number of InnoDB buffer pool reads ahead.
InnoDB physical reads	InnodbBufferPoolReads	Times/second	Number of InnoDB disk page reads.
Number of InnoDB data reads	InnodbDataRead	Byte/second	Number of bytes read per second by the InnoDB engine.
Total number of InnoDB data reads	InnodbDataReads	Times/second	Total number of InnoDB data reads.
Total number of InnoDB data writes	InnodbDataWrites	Times/second	Total number of InnoDB data writes.
InnoDB data written	InnodbDataWritten	Byte/second	Number of bytes written per second by the InnoDB engine.
InnoDB rows deleted	InnodbRowsDeleted	Times/second	Number of rows deleted per second by the InnoDB engine.
InnoDB rows inserted	InnodbRowsInserted	Times/second	Number of rows inserted per second by the InnoDB engine.
InnoDB rows updated	InnodbRowsUpdated	Times/second	Number of rows updated per second by the InnoDB engine.
InnoDB rows read	InnodbRowsRead	Times/second	Number of rows read per second by the InnoDB engine.
Average time to acquire	InnodbRowLockTimeAvg	Millisecond	Average duration of row locks in the InnoDB engine.

row locks in InnoDB			
Number of InnoDB row lock waits	InnodbRowLockWaits	Times/second	Number of row lock waits per second in the InnoDB engine.
Number of unused key cache blocks	KeyBlocksUnused	Count	Number of unused key cache blocks in the MyISAM engine.
Number of used key cache blocks	KeyBlocksUsed	Count	Number of used key cache blocks in the MyISAM engine.
Number of key cache block reads	KeyReadRequests	Times/second	Number of key cache block reads per second in the MyISAM engine.
Number of disk block reads	KeyReads	Times/second	Number of disk block reads per second in the MyISAM engine.
Number of key cache block writes	KeyWriteRequests	Times/second	Number of key cache block writes per second in the MyISAM engine.
Number of disk block writes	KeyWrites	Times/second	Number of disk block writes per second in the MyISAM engine.
Number of tables opened	OpenedTables	Count	Number of tables opened at the instance level.
Number of table locks released immediately	TableLocksImmediate	Times/second	Number of table locks released immediately.
Total number of open files	OpenFiles	Count	Total number of open files.
Log usage	LogCapacity	MB	Log usage space.
I/O thread status	SlaveloRunning	Status values: 0 - Yes 1 - No	I/O thread status of the primary instance. <b>Note:</b>

		2 - Connecting	In the status values, 0 indicates "running," 1 indicates "disconnected," and 2 indicates "connecting."
SQL thread status	SlaveSqlRunning	Status values: 0 - Yes 1 - No	SQL thread status of the primary instance. <b>Note:</b> In the status values, 0 indicates "running" and 1 indicates "disconnected."
Primary- Replica delay distance	MasterSlaveSyncDistance	MB	Primary-Replica delay distance.
Primary- Replica delay time	SecondsBehindMaster	Second	Primary-Replica delay time.

### FAQs

### How Do You Configure Monitoring for Primary-Replica Delays?

You can configure monitoring for primary-replica delays as required in the following two scenarios:

Scenario 1: Configure monitoring for primary-replica delays for the primary instance.

1. Log in to the TCOP console, select Alarm Management in the left sidebar, and choose Alarm Configuration > Cloud Product Monitoring > Alarm Policy > Create Policy.

2. On the Alarm Policy page, choose CDB > MySQL > SLAVE under Policy Type.

### Note:

To configure monitoring for primary-replica delays for the primary instance, select secondary server monitoring as the policy type to monitor the delay information from the secondary server to the primary server.

3. Under Configure Alarm Rules, set the trigger conditions for monitoring the primary-replica delay distance and primary-replica delay time, complete the settings for other configuration items as needed, and click **Complete**.

4. After the settings are completed, an alarm will be triggered when the primary-replica delay distance and primaryreplica delay time meet the trigger conditions.

Scenario 2: Configure monitoring for primary-replica delays for RO and disaster recovery instances.

1. Log in to the TCOP console, select Alarm Management in the left sidebar, and chooseAlarm Configuration >

### Cloud Product Monitoring > Alarm Policy > Create Policy.

2. On the Alarm Policy page, choose **CDB** > **MySQL** > **MASTER** for the policy type.

### Note:

To configure monitoring for primary-replica delays for an RO instance, select host monitoring as the policy type to monitor the delay information from the RO instance to its primary instance.

When monitoring for primary-secondary replication delays is configured for a disaster recovery instance, selecting Monitoring CVM Instances as the policy type will monitor the delay between the disaster recovery instance and its primary instance. On the other hand, if you select Secondary Server Monitoring, it will track the delay information from the secondary instance back to the disaster recovery instance.

3. Under Configure Alarm Rules, set the trigger conditions for monitoring the primary-replica delay distance and primary-replica delay time, complete the settings for other configuration items as needed, and click **Complete**.

4. After the settings are completed, an alarm will be triggered when the primary-replica delay distance and primary-replica delay time meet the trigger conditions.

## Alarm Policies (TCOP)

Last updated : 2024-07-23 10:19:48

This document describes how to create alarm policies and associate alarm objects in the Tencent Cloud Observability Platform (TCOP) console.

### Overview

You can create alarm policies to trigger alarms and send alarm notifications when the Tencent Cloud service status changes. The created alarm policies can determine whether an alarm needs to be triggered according to the difference between the monitoring metric value and the given threshold at intervals.

You can take appropriate precautionary or remedial measures in a timely manner when the alarm is triggered by changed product status. Therefore, properly created alarm policies can help you improve the robustness and reliability of your applications. For more information on alarms, see Creating Alarm Policy in TCOP.

To send an alarm for a specific status of a product, you need to create an alarm policy at first. An alarm policy is composed of three compulsory components, that is, the name, type and alarm triggering conditions. Each alarm policy is a set of alarm triggering conditions with the logical relationship "OR", that is, as long as one of the conditions is met, an alarm will be triggered. The alarm will be sent to all users associated with the alarm policy. Upon receiving the alarm, the user can view the alarm and take appropriate actions in time.

#### Note:

Make sure that you have set the default alarm recipient; otherwise, no notifications will be sent based on the default alarm policy of TencentDB.

### Directions

### Creating an alarm policy

1. Log in to the TCOP console, select Alarm Policy under Alarm Management module on the left sidebar, and click the Policy Management tab.

2. In the alarm policy list, click **Create**.

3. On the Create Alarm Policy page, set the policy name, policy type, alarm object, and trigger condition.

**Policy Type**: It divides into source monitoring and replica monitoring, which are applicable to different types of instances.

Deploy monitoring on the source: When the monitored instance is a source instance which is not a replica of any instance, replication-related monitoring data is invalid for the source, and the IO and SQL threads are disabled.

Replication-related monitoring data is valid and the IO and SQL threads can be enabled only when the monitored instance is a disaster recovery or a read-only instance.

Deploy monitoring on the replica: The two-node or three-node source instance and disaster recovery instance come in a source/replica architecture by default. As a result, replication-related monitoring data is valid for the replica only when the monitored instance is a source or disaster recovery instance. Such monitoring data can reflect the replication delay distance and time between the source or disaster recovery instance and its hidden replica nodes. We recommend that you keep an eye out for such monitoring data of the replica. If the source or disaster recovery instance fails, its monitored hidden replica nodes can be promoted to the source instance quickly.

**Alarm Object**: The instance to be associated with the policy alarm. You can find the desired instance by selecting the region where it is located or searching for its ID.

**Trigger Condition**: An alarm trigger is a semantic condition composed of metric, comparison, threshold, statistical period, and duration. For example, if the metric is disk utilization, the comparison is >, the threshold is 80%, the statistical period is 5 minutes, and the duration is two statistical periods, then the data on disk utilization of a database will be collected once every five minutes, and an alarm will be triggered if the disk utilization exceeds 80% for two consecutive times.

**Configure Alarm Notification**: You can select a preset or custom notification template. Each alarm policy can be bound to three notification templates at most. For more information, see Creating Notification Template.

Basic Info	
Policy Name	It can contain up to 30 characters
Remarks	It can contain up to 100 characters
Monitoring Type	Cloud Product Monitoring
Policy Type	CD8 / MySQL / MASTER 🔹
Project	DEFAULT PROJECT • 1 exists. You can create 299 more static threshold policiesThe current account has 0 policies for dynamic alarm thresholds, and 20 more policies can be created.
Alarm Policy	
Alarm Object	Instance ID × Select object ×
Trigger Condition	Select template O Configure manually (♥ Apply preset trigger conditions ①)
	Metric Alarm
	When meeting any * of the following metric conditions, the metric will trigger an alarm.
	Threshold Ostatic Opynamic () Type ()
	▶ If CpuUseRate ▼ (statistical perior ▼ > ▼ 80 at 5 consecutive r ▼ then Alarm once a hour ▼ ① □

4. After confirming that everything is correct, click **Complete**.



### Associating an alarm object

After the alarm policy is created, you can associate some alarm objects with it. When an alarm object satisfies an alarm trigger condition, an alarm notification will be sent.

- 1. In the alarm policy list, click the name of an alarm policy to enter the alarm policy management page.
- 2. Click Add Object in the Alarm Object section.
- 3. In the pop-up window, select the target alarm object and click **OK** to associate it with the alarm policy.

### FAQs

#### How can I configure the source-replica delay monitoring?

You can do so by following the steps below based on the actual scenarios.

Scenario 1: Configure source-replica delay monitoring for the source instance

1.1 Log in to the TCOP console select Alarm Policy under Alarm Management module on the left sidebar, and click Create Policy on the Policy Management tab.

1.2 On the Create Alarm Policy page, select Policy Type: Cloud Database > MySQL > Replica Monitoring. Note:

When configuring source-replica delay monitoring for the source instance, you must select replica monitoring as the policy type, so that the delay information from the replica to the source is monitored.

1.3 Complete the trigger condition settings for the monitoring items "Source-Replica Delay (in MB)" and "Source-

Replica Delay (in Seconds)" as well as other configuration items as needed, and click **Complete**.

1.4 After the setting is completed, the alarm can be triggered when the monitoring items "Source-Replica Delay (in MB)" and "Source-Replica Delay (in Seconds)" meet the trigger conditions.

Scenario 2: Configure source-replica delay monitoring for the read-only and disaster recovery instances

1.1 Log in to the TCOP console select Alarm Policy under Alarm Management module on the left sidebar, and click Create Policy on the Policy Management tab.

### 1.2 On the Create Alarm Policy page, select Policy Type: Cloud Database > MySQL > Host Monitoring. Note:

When configuring source-replica delay monitoring for the read-only instance, you must select host monitoring as the policy type, so that the delay information from the read-only instance to the source instance is monitored. ?When configuring source-replica delay monitoring for the disaster recovery instance, you can select host monitoring as the policy type, so that the delay information from the disaster recovery instance to the source instance is monitoring is the policy type, so that the delay information from the disaster recovery instance to the source instance is monitored. If you select replica monitoring as the policy type, then the delay information from the disaster recovery instance to the disaster recovery instance is monitored.

1.3 Complete the trigger condition settings for the monitoring items "Source-Replica Delay (in MB)" and "Source-Replica Delay (in Seconds)" as well as other configuration items as needed, and click **Complete**.

1.4 After the setting is completed, the alarm can be triggered when the monitoring items "Source-Replica Delay (in MB)" and "Source-Replica Delay (in Seconds)" meet the trigger conditions.

## Alarm Notification (DBbrain)

Last updated : 2024-07-23 10:20:22

This document describes how to view exception alarm messages from DBbrain in the TencentDB for MySQL console. DBbrain's exception alarms notification service pushes MySQL instance exception alarm messages to you in real time, allowing you to conveniently and promptly discover database exception diagnosis problems. All pushed exception alarm messages are displayed in the historical message list, so you can quickly view and locate previously pushed exception diagnosis problems.

### Viewing an alarm

### **Option 1**

Log in to the TencentDB for MySQL console. If an exception diagnosis problem occurs on an instance when you are in the console, a window will pop up in the top-right corner of the console in real time to push the exception alarm message notification, which contains the database instance information such as instance ID, instance name, diagnosis item, and start time. This allows you to quickly learn the diagnosis of the database instance. You can click **View Exception Diagnosis Details** in the message notification to view the specific diagnosis details and optimization advice for the instance.

If you check **No alarm again today** in the message notification, when an exception diagnosis problem occurs in a database instance under your account, no exception alarm messages will be pushed to you in a pop-up window.

MySQL - Instance List													MySQL Exception Alarm	No alar	m again today
													Diagnosis item: Replica	ation IO t	hread error
Guangzhou(1)	Qingyuan(0)	Shanghai(0)	Beijing(0)	Chengdu(0)	Chongqing(0)	Hong Kong (China)(0)	Taipei (China)(0)	Singapore(0)	Bangkok(0)	Mumbai(0)	Seoul(0)		Instance		
Toronto(0)	Frankfurt(0) N	Noscow(0)											View Exception Diagnosis [	)etails 📛	
Create	omparative Monitorin	ng Restart	More	Ŧ			M	iltiple keywords are se	eparated by the ve	rtical bar "[", and m	ultiple filter tags	are se	eparated by the enter key.	Q,	¢¢
Instance ID	/Name T	Monitoring	/Status/Task 🔻	Avai	ilability Zone 🔻	Configuration <b>T</b>	Version T	Private IP		Billing Mode		Pr	roject 🔻 Operation		

### **Option 2**

Log in to the TencentDB for MySQL console, select Instance List, Task List, Parameter Template, Recycle Bin or Placement Group on the left sidebar, and click Exception Alarm in the top-right corner to expand the list of historical exception alarm messages. The number of alarms generated in the instances under your account is displayed next to the button.

MySQL - Instance List													
Guangzhou(1) Qin	yuan(0) Shanghai(0)	Beijing(0)	Chengdu(0)	Chongqing(0)	Hong Kong (China)(0)	Taipei (China)(C	) Singapore(0)	Bangkok(0)	Mumbai(0)	Seoul(0)	Tokyo(0)	Silicon Valley(0)	Virginia(0)
Toronto(0) Frankfu	t(0) Moscow(0)												
Create Compara	ive Monitoring Restart	More	Ŧ				Multiple keywords are s	eparated by the ve	rtical bar " ", and m	ultiple filter tags a	are separated by	the enter key. C	¢¢
Instance ID/Name	Monitoring	g/Status/Task 🔻	Ava	iilability Zone 🍸	Configuration T	Version	Private IF	)	Billing Mode		Project 🔻	Operation	
	ı <b>lı</b> ⊙ Runnin	g	Gua	angzhou Zone 4	High-Availability Edi 1core1000MB/25GE Network: Default-VF Default-Subnet	tion MySQL 3 PC -	5.7		Pay as you go	)	Default Projec	t Login Manage	More 🔻

In the expanded list of historical exception alarm messages, you can view all pushed exception alarm messages. You can view them by region and click a message to view the diagnosis details of the exception alarm event.

MySQL - Instance List									MySQL Exception Alarm 🔇 Guangzhou (1) Other regions (0) 💌	Alarm De
Guangzhou(	1) Qingyuan(0)	Shanghai(0)	Beijing(0)	Chengdu(0)	Chongqing(0)	Hong Kong (China)(0)	Taipei (China)(0)	Singapore(	There is 1 exception alarm event in the current region, with 1 instance involved	Last 3 hou
Toronto(0)	Frankfurt(0)	Moscow(0)							Critical Diagnosis item: Replication IO thread error Instance ID / name: In Start time:	
Create	Comparative Monito	ring Restart	More	v			М	ultiple keywords	View Exception Diagnosis Details	
Instan	ice ID/Name 🔻	Monitorin	g/Status/Task 🔻	Ava	ilability Zone 🔻	Configuration <b>T</b>	Version 1	' Priva	1	
		ı <b>lı</b> ⊙ Runnin	9	Gua	ingzhou Zone 4	High-Availability Ed 1core1000MB/25Gf Network: Default-Vf Default-Subnet	ition MySQL5.7 3 PC -			

## **Setting Event Alarms**

Last updated : 2024-06-18 11:11:49

TencentDB for MySQL supports the event alarm feature. When you set an event pattern, event target, and recipient, if an exception event is detected, the event alarm will send alarm notifications to the recipient through various optional channels. This document explains how to set up event alarms.

### Step 1: Enabling EventBridge

Tencent Cloud EventBridge implements permission management through Cloud Access Management (CAM). CAM is a permission and access management service provided by Tencent Cloud, mainly used to help customers securely manage access rights to resources under Tencent Cloud accounts. Users can create, manage, and terminate users (groups) with CAM, and control the access rights of other users to Tencent Cloud resources using identity and policy management. Before using EventBridge, you need to activate the service on the product page. For the method to activate with a main account and to authorize sub-accounts to use this service, see Activate EventBridge.

### Step 2: Setting Event Alarms

#### Note:

For alarm events generated by TencentDB for MySQL, all will be delivered to the **Tencent Cloud service event set**, with this delivery being the default and not supporting modification or editing.

After enabling the Tencent Cloud EventBridge service, it will automatically create a default cloud service event set in the **Guangzhou region**. Alarm events generated by TencentDB for MySQL will be automatically delivered to it.

### Method 1: Tencent Cloud EventBridge Configuration

- 1. Log in to Tencent Cloud Observability Platform > Event Set.
- 2. Select **Guangzhou** as the region at the top.
- 3. Click on the **default** event set under the cloud service event set.

Event bus name	Event bus configuration	Event bus description	Last update time	Operation
m	Platform Event Bus		2023-11-27 16:06:55	Publish event Edit
default	Tencent Cloud service event bus		2022-11-09 17:17:15	Publish event Edit

4. Click Manage Event Rules, then on the redirection page, click Create (taking creating an event rule as an

example, if an event rule already exists, you can directly modify it with no need to recreate it).

5. After completing the following configuration on the new event rule page, click **Next**.

### S Tencent Cloud

Parameter	Description
Rule name	Fill in the rule name, which can only include letters, digits, underscores (_), and hyphens (-), must start with a letter, end with a digit or letter, and contain 2-60 characters.
Rule description	Fill in the rule description, which can only include digits, Chinese and English characters and common punctuation marks, and must be no more than 200 characters.
Tag	Customize whether to enable tags. Once enabled, tags can be added to the event rule.
Data conversion	Event data conversion helps you easily process the event content. For example, you can extract and parse fields within the event, perform mapping reorganization, and then deliver it to the event target.
Event Example	Provides an example of event structure for reference in configuring event matching rules. You can find the target template under event example selection for reference, and select TencentDB for MySQL as the cloud service type.
Rule pattern	Supports form mode and custom events, where form mode is recommended for convenience.
Tencent Cloud service	Select TencentDB for MySQL.
Event Type	Select the alarm event type as needed.
Subject to match	Select the matching object for event alarms.

6. Complete relevant settings on the event target page, check **Enable Event rules now**, and click **Complete**. (There are various trigger methods available, which you can choose as needed. The following example selects "Message Push" as the trigger method. The channel push will be sent to the designated contacts. You need to set the recipient user separately at the recipient area. The notification period shall be agreed upon with the account user. The receiving channels can include Email, message center, SMS, and Telephone).

#### Note:

To configure multiple event targets, click Add.

elivery target	
Trigger method *	Notification message(i) v
Message template *	Monitoring alert template O General notification template
Alert content *	Chinese O English
Notification method *	publishing channel 💌
publishing channel	
Recipients *	User *
Notification period *	09:30:00 ~ 23:30:00
Delivery method • (j)	Ernail SMS Phone Message center
ld	
Chable event rules nov	v

7. After an event message push is set, if an event alarm set in the configuration occurs, the system will notify via the receiving channels set for your specified recipients.

### Method 2: TCOP Alarm Policy Configuration

#### **Create an Alarm Policy**

1. Log in to the TCOP Console, select Alarm Management > Policy Management page from the left navigation.

2. On the alarm policy list page, click **Create Policy**.

Alarm Records Policy Management Basic Configuration			
If you have any questions or suggestions, scan QR code to join our community on WeChat or WeCom.			
Create Policy Dette More •	Advanced Filter	Search by Tag, Policy Name/ID	Q

3. On the policy creation page, complete the settings for basic information, alarm rules, and alarm notifications.

**Policy Type**: Divided into host monitoring and standby monitoring, with each applicable to different instance types.

**Alarm Object**: The object instance to be associated can be found by selecting the object's region or searching for the object's instance ID.

Trigger Condition: Find event alarms, click Add Event, and add alarm events according to business needs.



Trigger Condition	Select Template O Configure manually Apply preset trigger conditions (
	Metric Alarm Event Alarm
	The database agent mount nod 🔻 頂
	MasterHealthCheckError v
	MasterHealthCheckRecovery v
	SlaveHealthCheckError v
	SlaveHealthCheckRecovery *
	Add Event

**Configure Alarm Notification**: Supports selection of system preset notification templates and user-defined notification templates. Each alarm policy can be bound to a maximum of three notification templates. For user-defined notification templates, see Create New Notification Template.

Select a system preset template.

Select notification template		
You have selected 1 notification template, and 2 more can be sele	cted.	
Search for notification template	c	2 ¢
Notification Template Name	Included Operations	
Preset Notification Template	Recipient: 1	
bl	Recipient: 1	
	Recipient: 1	
Total items: 3	20 💌 / page 🖂 🖂 1 / 1 page 🕨	×
ок	Cancel	

Create a template.

Basic Info									
Template Name	Up to 60 chara	icters							
Notification Type 🛈	✓ Alarm Trigger	- Alar	rm Recovery						
Notification Language	English			•					
Tag	Tag Key	•	Tag Value		▼ ×				
Notifications	(Fill in at least one You can add a us	item) ser only for receiv	ving messages.						
	Recipient	User	*				¢	Add User	Delete
	Notification Cycle	🗸 Mon 🗸	Tue 🔽 Wed	🔽 Thu 📘	🖊 Fri 🔽 Sat	🖌 Sun			
	Notification Period	00:00:00 ~ 2	3:59:59	<b>(</b> )					
	Receiving Channel	🔽 Email 🛛	SMS						

4. After confirmation, click Complete.

#### Associate Alarm Objects

After creating an alarm policy, you can also associate it with other alarm objects (other instances that need this alarm policy). When objects meet the alarm trigger conditions, an alarm will be sent.

1. On the Alarm Policy List Page, click Alarm Policy Name to enter the Manage Alarm Policy page.

2. In the Alarm Object column of the Alarm Policy Management page, click Add Object.

3. In the pop-up dialog box, select the alarm object you wish to associate, and click **OK** to associate the alarm object.

# Operation Logs Operation Logs

Last updated : 2023-11-20 15:45:38

### Overview

A SQL statement query that takes more time than the specified value is referred to as a "slow query", and the corresponding statement is called a "slow query statement". The process where a database administrator (DBA) analyzes slow query statements and finds out the reasons why slow queries occur is known as "slow query analysis". In the Operation Log tab, you can view the slow logs details, error logs details, rollback logs, and logs delivery of the instance and download slow logs. You can also view and download database logs on the command line interface (CLI) or through TencentDB APIs. For more information, see DescribeSlowLogs and DescribeBinlogs.

#### Note:

Currently, slow log download, rollback and log delivery features are not supported for TencentDB for MySQL singlenode instances of the cloud disk edition.

#### Notes on MySQL slow queries

long\_query\_time: Slow query threshold parameter that is accurate to the microsecond level. The default value is 1 second. When a SQL statement takes more time than the threshold to execute, it will be recorded in a slow log. When the long\_query\_time parameter is adjusted, existing slow logs will not be affected. For example, if the slow log threshold parameter is 1s, slow logs will be reported for exceeding queries; after this value is modified to 2s, the previously reported logs will still be displayed.

log\_queries\_not\_using\_indexes: Whether to log unindexed queries. The default value is OFF.

### The Directions

1. Log in to the TencentDB for MySQL console. In the instance list, click an **Instance ID** or **Manage** in the **Operation** column to enter the instance management page.

2. On the **Operation Log** tab, you can view the slow logs details, error logs details, rollback logs and logs delivery of the instance and download slow logs.

Feature	Description
Slow Log Details	Records SQL statements that took more than 1 second to execute in the database for the past month
Download Slow	Downloads slow logs



Log	
Error Log Details	Records the detailed information of each startup and shutdown as well as all the serious warnings and errors during operation
Rollback Log	Records the status and progress of rollback tasks
Log Delivery	Collects data from slow logs and error logs sourced from TencentDB for MySQL instances, and deliver them to the Cloud Log Service (CLS) for analysis

3. On the **Download Slow Log** tab, click **Download** in the **Operation** column to download slow logs.

4. We recommend that you copy the download address in the pop-up window, log in to a (Linux) CVM instance in the same VPC as the TencentDB instance, and run the wget command to download over the private network at a higher speed. For more information, see Customizing Linux CVM Configurations. Note :

Logs with a size of 0 KB cannot be downloaded.

You can also click Download to download it directly. However, this may take a longer time.

wget command format: wget -c 'log file download address' -O custom filename.log

Below is a sample:

wget -c 'http://szx.dl.cdb.tencentyun.com:303/cfdee? appid=1210&time=1591&sign=aIGM%3D' -O test.log

## Log Delivery

Last updated : 2024-06-18 16:23:02

TencentDB for MySQL offers log delivery feature, which can collect slow log and error log details sourced from TencentDB for MySQL instances and deliver it to Cloud Log Service (CLS) for analysis to quickly monitor and troubleshoot business issues. This document describes how to enable or disable the log delivery feature in the console.

### Prerequisites

Before using this feature, please ensure that you have enabled CLS.

### **Use Limits**

Only two-node/three-node local disk instances and read-only instances support the use of log delivery.

### **Slow Log Definition**

Field Value	Туре	Description
TIMESTAMP	-	Reserved field of CLS, represents the time when the log was generated.
instance_id	String	Database instance ID (e.g., cdb- xxx).
db_name	String	Database name.
rows_examined	Long	Number of scanned rows.
rows_sent	Long	Number of returned rows.
query_time	Double	Execution time (in seconds).
lock_time	Double	Lock wait time (in seconds).
user_host	String	Client information.



user_name	String	Username used by the client to connect to the database instance.
query_sql	String	Slow og SQL.

### Error Log Definition

Field Value	Туре	Description
TIMESTAMP	_	Reserved field of CLS, represents the time when the log was generated.
instance_id	String	Database instance ID (e.g., cdb- xxx).
content	String	Content of the error log.
level	String	Level of the error log. Such as warning, note and error. For logs without a level, the entry would be none, such as deadlock error logs.

### Enabling Slow Log Delivery

1. Log in to the TencentDB for MySQL console. In the instance list, click the **Instance ID** to enter the management page.

2. On the instance management page, select **Operation Log** > **Log Delivery**.



- 3. Under **Slow Log Delivery**, click the Edit icon.
- 4. In the pop-up window, select **Enable**, configure the following items, then click **OK**.

	Slow Log Deli	very ×				
	Enable	Close				
	Logset Managem	ent * O Select an existing logset O Create a logset				
	Logset *	Please select				
	Log Topic Manage	ement * O Select an existing log topic O Create a log topic				
	Log Topic *	cloud_ Please enter a name _topic				
	Retention Period	- 30 + day(s) (j				
	Index Creation *					
		OK Cancel				
Para	meter	Description				
Regi	on	Select the region for log shipping, supporting offsite shipping.				
Logset A Management S		A logset is a classification of log topics for your convenience in managing them. You can Select an existing logset or Create a logset.				
Logset		Select an existing logset: You can filter from the existing logset in the search box to categorize your slow log delivery. Create a logset: You can create a new logset to categorize your slow log delivery. The naming format for creating a logset is <b>cloud_custom_logset</b> , where the custom field supports only letters, digits, and underscores, with a length of no more than 20 characters.				
Log <sup>-</sup> Mana	Topic agement	A log topic is the fundamental unit for log data collection, storage, search, and analysis. You can Select an existing log topic or Create a log topic.				
Log <sup>-</sup>	g Topic Select an existing log topic: This option can be configured only when an existing log selected in logset operations. You may filter the log topics under the selected logset search box. Create a log topic: Create a new log topic under the selected logset. The naming for for creating a log topic is <b>cloud_custom_topic</b> , where the custom field only support letters, digits, and underscores, with a length of no more than 20 characters. <b>Note:</b> You can manage log topic. For more details, see Managing Log Topic.					
Rete	ntion Period	Select the retention period for the slow log data delivery, which defaults to 30 days. It allows you to choose between 1 to 3600 days, and the data will be automatically cleared				



	after the logs expire. If an existing log topic is selected, the retention time will default to match the existing log topic's retention period.
Index Creation	Enabled by default. Index configuration is critical for log retrieval and analysis in CLS. Only if the index is enabled, the log can be analyzed and retrieved. Learn More. If an existing log topic is selected, the index status will default to match the existing log topic's index status.

5. After you successfully enable slow log delivery, you can see the Delivery Status is enable under Slow Log Delivery. Clicking on the Log Topic Name will take you to the CLS Console for further analysis and management.

Slow Log Details	s Download Slow Log	Error Log Details	Rollback Log	Log Delivery
Slow Log Deli	very			
Delivery Status	Enable 🧨			
Logset Name	cloud_ilogset			
Log Topic Name	cloudtopic 🖸			

### Disable Slow Log Delivery

#### Note:

After disabling the slow log delivery, the generated slow log delivery data will continue to be stored according to the retention period selected when slow log delivery is enabled, and data will be automatically cleared after expiration. 1. Log in to the TencentDB for MySQL console. In the instance list, click the **Instance ID** to enter the management page.

- 2. On the instance management page, select **Operation Log** > **Log Delivery**.
- 3. Under **Slow Log Delivery**, click the Edit icon.
- 4. In the pop-up window, select **Close** and click **OK**.

### Enabling Error Log Delivery

1. Log in to the TencentDB for MySQL console. In the instance list, click the **Instance ID** to enter the management page.

- 2. On the instance management page, select **Operation Log** > **Log Delivery**.
- 3. Under Error Log Delivery, click the Edit icon.
- 4. In the pop-up window, select **Enable**, configure the following items, then click **OK**.

	Error Log Deliv	very ×				
	Enable	Close				
	Logset Manageme	nt * O Select an existing logset O Create a logset				
	Logset *	cloud_ Please enter a name _logset				
	Log Topic Manage	ment * Select an existing log topic O Create a log topic				
	Log Topic *	cloud_ Please enter a name _topic				
	Retention Period *	- 30 + day(s) (i)				
	Index Creation *					
		OK Cancel				
Para	meter	Description				
Regi	on	Select the region for log shipping, supporting offsite shipping.				
LogsetA logset is a classification of log topics for your convenience in managing themManagementSelect an existing logset or Create a logset.		A logset is a classification of log topics for your convenience in managing them. You can Select an existing logset or Create a logset.				
Logset Cre The sup		Select an existing logset: You can filter through existing logsets in the search box to categorize your error log delivery. Create a logset: You can create a new logset as the categorization for error log delivery. The naming format for creating a logset is <b>cloud_custom_logset</b> , where the custom field supports only letters, digits, and underscores, with a length of no more than 20 characters.				
Log Man	Topic agement	<ul><li>A log topic is the fundamental unit for log data collection, storage, search, and analysis.</li><li>You can Select an existing log topic or Create a log topic.</li></ul>				
Log TopicSSelect an existing log topic: This option can be configured only when an existing selected in logset operations. You may filter the log topics under the selected logsed search box.Log TopicCreate a log topic: Create a new log topic under the selected logset. The naming f creating a log topic is cloud_custom_topic, where the custom field only supports digits, and underscores, with a length of no more than 20 characters.Note: You can manage log topic. For more details, please refer to Managing Log Topic.		SSelect an existing log topic: This option can be configured only when an existing logset is selected in logset operations. You may filter the log topics under the selected logset in the search box. Create a log topic: Create a new log topic under the selected logset. The naming format for creating a log topic is <b>cloud_custom_topic</b> , where the custom field only supports letters, digits, and underscores, with a length of no more than 20 characters. <b>Note:</b> You can manage log topic. For more details, please refer to Managing Log Topic.				
Retention Period Select the retention period for the slow log data delivery, which defaults to 30 data		Select the retention period for the slow log data delivery, which defaults to 30 days. It allows you to choose between 1 to 3600 days, and the data will be automatically cleared after the logs expire. If an existing log topic is selected, the retention time will default to match the existing log topic's retention period.				

Index Creation	Enabled by default. Index configuration is critical for log retrieval and analysis in CLS. Only if the index is enabled, the log can be analyzed and retrieved. Learn More. If an existing log topic is selected, the index status will default to match the existing log topic's index status.
----------------	--

5. After you successfully enable slow log delivery, you can see the Delivery Status is enable under Error Log Delivery. Clicking on the Log Topic Name will take you to the CLS Console for further analysis and management.

Error Log Delivery			
Delivery Status	Enable 🧨		
Logset Name	cloudlogset		
Log Topic Name	cloudtopic 🗳		

### Disable Error Log Delivery

#### Note:

After disabling the error log delivery, the generated error log data will continue to be stored according to the retention period selected when the error log delivery is enabled, and data will be automatically cleared after expiration.

1. Log in to the TencentDB for MySQL console. In the instance list, click the **Instance ID** to enter the management page.

2. On the instance management page, select **Operation Log** > **Log Delivery**.

3. Under Error Log Delivery, click the Edit icon.

4. In the pop-up window, select Close and click OK.

### **Related Documents**

Logset Managing Log Topic Chart Overview Creating Processing Task Syntax and Rules

# Read-Only Analysis Engine Introduction

Last updated : 2025-05-20 10:20:22

This document introduces a new feature of TencentDB for MySQL based on the LibraDB engine: the read-only analysis engine.

#### Note:

The read-only analytics engine is currently available for free trial, with official commercialization and billing set to commence on May 19, 2025. For product pricing details, please refer to Billing Overview.

### Background

TencentDB for MySQL has been extensively optimized to support high concurrency, strong consistency, and enterprise-grade database features. Built on the TXSQL engine, it delivers high-performance online transaction processing (OLTP). However, beyond supporting high-QPS transactional workloads, many business systems also require data mining and analysis capabilities to enable better business decision-making, drive iterative innovation, and quickly adapt to changing market conditions.

Traditional databases, in order to support high-performance online transaction processing capabilities and ensure the stability of business queries, typically adopt row-based storage structures and the volcano execution model. However, these choices render them inefficient in handling analytical queries. Naturally, some businesses opt for a "traditional database + data warehouse" solution to accommodate mixed workloads involving both transactions and analytics. Yet, this approach imposes significant maintenance costs on customers, requiring them to develop their own ETL tools for transferring data from the database to the data warehouse. Moreover, it fails to adequately address the requirements for data timeliness and consistency.

To address this, TencentDB for MySQL now offers support for the new read-only analysis engine, providing users with efficient and real-time data analysis capabilities.

### What Is the Read-Only Analysis Engine

The read-only analysis engine is a new feature of TencentDB for MySQL designed for real-time data analysis and complex query scenarios. Its pluggable engine architecture enables flexible and convenient creation and termination, while providing users with the capability to process massive volumes of data and perform efficient, real-time complex analysis.

### Supported Regions and Availability Zones

For details on the regions and availability zones supported by the Read-Only Analytics Engine, please refer to Regions and AZs.

### Feature Strengths

#### High-performance analysis engine LibraDB

The LibraDB engine delivers high-performance complex query analysis, enabling your business analysis systems to efficiently extract valuable insights from massive databases in a timely manner. It supports analysis-oriented acceleration features such as a vectorized execution engine and massively parallel processing (MPP). Whether it is large-scale multi-table JOINs, data aggregation and sorting, or complex nested SQL queries, the LibraDB engine provides an exceptional performance experience.

#### Pluggable analysis engine

The LibraDB engine is compatible with MySQL protocols and syntax, allowing users to run complex queries directly in LibraDB without modifying their business logic. Users can enable the read-only analysis engine based on actual business needs, and disable it at any time when analytical acceleration is not required, helping to control costs effectively.

#### Real-time columnar data loading capability

With the built-in data synchronization component of the LibraDB engine, existing data in TencentDB for MySQL can be quickly loaded into the read-only analysis engine. After the initial data load, all subsequent changes made to the data in the read-write instance can be synchronized in real time, ensuring consistency between row-based and columnar data. In addition, to address the inefficiency of data changes in traditional columnar storage under high-concurrency data update and deletion scenarios, the LibraDB engine offers columnar storage capabilities optimized for high-concurrency data updates, enabling real-time synchronization and achieving zero-latency performance.

#### Targeted data loading capability

In traditional read-only instances, all data from the primary database should be fully synchronized to the secondary database. However, with the read-only analysis engine, it is possible to load only specified objects into the engine, rather than requiring full synchronization. Users can choose to load only those databases and tables that need acceleration through the analysis engine, or those with analytical value, enabling flexible control over the disk space used by the read-only analysis engine.

#### Ultra-high data compression ratio

Leveraging a columnar storage structure, the engine delivers ultra-high data scan performance while also achieving an average compression ratio of 4 to 5 times, significantly reducing storage costs.

#### Comprehensive cloud-hosted capability

No Ops is required for complex ETL logic or backend database management. With a fully managed product design, you gain an out-of-the-box experience for data analysis capabilities. Additionally, the monitoring feature provides a

carefully curated set of core metrics, covering everything from TXSQL to the analysis engine, and from the linkage layer to the storage layer. This simplifies complexity and helps you quickly assess instance health using key metrics, offering effective optimization guidance for your business systems. You can also define custom alarm thresholds to proactively prevent potential exceptions.

### Applicable Scenarios

The read-only analysis engine is designed to provide users with real-time, high-performance data analysis, helping to eliminate the complex Ops challenges of building custom ETL tools. With its built-in features, users can easily create data analysis instances with a single click, using them as a foundation for business decision-making and fully unlocking the value of their data.

#### Report analysis and real-time dashboards

For report systems designed for internal enterprise analysis and management, users can view the real-time operational status of online business systems. It also applies to data analysis tasks for business operations. In such scenarios, the SQL queries are complex and variable in pattern, requiring high throughput and involving large volumes of online data. The read-only analysis engine meets the real-time and high-performance requirements of these types of workloads.

#### User profiling and behavior analysis

In advertising and game operations scenarios, in-depth analysis of user behavior and user profiling is often required, with the results used to support real-time business decisions. These scenarios typically involve large volumes of data, require timely responses, and have high query QPS. By using the read-only analysis engine, users can quickly obtain the necessary data for analysis, enabling accurate insights into user behavior that serve as a decision-making foundation for precise business targeting.

#### Real-time data warehouse

The read-only analysis engine can also be used in scenarios such as order analysis during major e-commerce promotions, waybill analysis in the logistics industry, performance analysis and metric computation in the financial sector, live streaming quality analysis, ad delivery analysis, intelligent dashboards, and probe analysis, providing ultrahigh performance for complex queries.

#### Big data reconciliation and batch computing

In certain online services, especially those involving financial transactions, periodic data aggregation and reconciliation are required. Performing batch reconciliation on traditional row-based data is often inefficient and resource-intensive, making it difficult to meet business expectations in a timely manner. By leveraging the high-concurrency computing capabilities of the read-only analysis engine, users can fulfill these business needs with significantly improved efficiency.

# Feature Limits and Compatibility Compatibility and Use Limits

Last updated : 2025-05-09 11:51:17

This document introduces the use limits and compatibility of the read-only analysis engine feature.

#### Note:

The read-only analysis engine does not support other special scenarios except for the ones mentioned in this document.

### Availability Difference

The read-only analysis engine temporarily does not support configuration changes.

The read-only analysis engine does not support high availability in a single-node scenario. If a node fails, this readonly analysis engine will become unavailable. If you need the read-only analysis engine to provide continuous service, please apply for multiple read-only analysis engines.

### Syntax Limits

In the read-only analysis engine, only read-only query statements can be executed, and no data change operations can be performed, including DDL and DML operations.

The read-only analysis engine only supports SELECT query statements. However, a few keywords and syntaxes are still not supported in SELECT statements. For details, see <u>SELECT Statement Description</u>.

In the read-only analysis engine, CTE syntax and window functions are supported. However, the use of CTE syntax and window function still faces some limits. For details, see CTE Syntax Use Instructions and Window Function Use Instructions.

The read-only analysis engine temporarily does not support full-text search syntax.

If the read-only analysis engine is enabled for TencentDB for MySQL, performing certain data change operations in the "primary instance" may cause exceptions in data loading of the read-only analysis engine. For details, see Data Loading Restrictions.

### Table Limits

The recommended maximum number of tables in the read-only analysis engine varies in different versions. **1.2404.x**: Maximum number of tables is 2,000.

2.2410.x: Maximum number of tables is 20,000.

### Value Limits

In the read-only analysis engine, for kernel versions earlier than 1.2404.16.0 (excluding 1.2404.16.0), the default value size limit for regular columns is 5 MB. If there are tables with column values exceeding 5 MB, an error will occur and the loading into the analysis engine will be paused. At this time, you can submit a ticket to upgrade the kernel version to 1.2404.16.0 or later.

If the Linux kernel version is 1.2404.16.0 or later versions, the maximum value limit supported by ordinary columns is 16 MB by default. If there are tables with column values exceeding 16 MB, it is recommended to remove these tables and load them into the read-only analysis engine.

### SQL\_MODE

Similar to MySQL, the read-only analysis engine supports setting the SQL Mode globally or at the session level by using the SET [SESSION | GLOBAL] sql\_mode='modes' statement. You can also use SELECT @@sql\_mode to query the current SQL Mode.

The read-only analysis engine supports the following common MySQL system SQL\_MODEs. Any unmentioned SQL\_MODEs are not supported. However, it should be noted that although the read-only analysis engine supports these SQL\_MODEs, some of them are not applicable in the read-only analysis engine, such as

NO	AUTO	CREATE	USER and NO	ENGINE	SUBSTITUTION.

Name	Description		
PIPES_AS_CONCAT	Treat    as the string concatenation operator $(+)$ (same as CONCAT()), rather than OR.		
ANSI_QUOTES	Treat " as an identifier. If ANSI_QUOTES is enabled, only those within single quotation marks will be regarded as String Literals, and double quotes are treated as identifiers. Therefore, double quotes cannot be used to quote strings.		
IGNORE_SPACE	If this mode is enabled, the system ignores spaces. For example, "user" and "user " are the same.		
ONLY_FULL_GROUP_BY	If a column that is not processed by an aggregate function or not included in GROUP BY appears in SELECT, HAVING, or ORDER BY, this SQL is invalid.		
NO_UNSIGNED_SUBTRACTION	In a subtraction operation, if an operand has no sign, do not mark the result as UNSIGNED (supported).		

NO_BACKSLASH_ESCAPES	If this mode is enabled, the backslash symbol ( $\)$ only represents itself.		
STRICT_TRANS_TABLES	Enable strict mode for the transaction storage engine and roll back the entire statement after inserting an illegal value.		
STRICT_ALL_TABLES	For transactional tables, roll back the entire transaction statement after writing an illegal value.		
NO_ZERO_IN_DATE	In strict mode, dates with 0 in the month or day part are not accepted. If the IGNORE option is used, we insert "0000-00-00" for similar dates. In non-strict mode, such dates can be accepted, but a warning will be generated.		
NO_ZERO_DATE	In strict mode, do not treat "0000-00-00" as a valid date. You can still insert 0 dates with the IGNORE option. In non-strict mode, this date can be accepted, but a warning will be generated.		
ALLOW_INVALID_DATES	Do not check the legitimacy of all dates. Only check whether the month value is between 1 and 12 and whether the date value is between 1 and 31. This only applies to DATE and DATETIME columns. All legitimacy of TIMESTAMP columns needs to be checked.		
ERROR_FOR_DIVISION_BY_ZERO	Enable this mode. During the INSERT or UPDATE process, when the dividend is 0, the system generates an error. If this mode is not enabled, when the dividend is 0, the system generates a warning and replaces it with NULL.		
REAL_AS_FLOAT	Treat REAL as a synonym for FLOAT rather than a synonym for DOUBLE.		
NO_DIR_IN_CREATE	When a table is created, ignore all INDEX DIRECTORY and DATA DIRECTORY instructions. This option is only useful for replication servers.		
NO_AUTO_CREATE_USER	Prevent GRANT from automatically creating users, except when specifying a password (but this has no actual effect in a read-only analysis engine).		
NO_ENGINE_SUBSTITUTION	If the required storage engine is disabled or not compiled, it can prevent automatic replacement of the storage engine (but this has no actual effect in a read-only analysis engine).		

### Character Set and Collation

A character set is a collection of symbols and codes. The default character set in a read-only analysis engine is utf8mb4.

A collation is a rule for comparing characters and the sort order of characters within a character set. For example, the result of comparing A and a is different in a binary collation.

The character sets and collation currently supported by the read-only analysis engine are as follows in the table:

Character Set	Description	Supported Collation	Maximum Size
utf8	UTF-8 Unicode	utf8_bin	3
utf8mb4	UTF-8 Unicode	utf8mb4_bin	4

#### Note:

When the objects in the read-write instance use other character sets, there is no impact on data loading into the readonly analysis engine. However, some special characters will have exceptions when they are queried in the read-only analysis engine. At the same time, inconsistent sorting results may occur due to different collations.

### Other Behavior Descriptions

The return results of executing SELECT...GROUP BY expr in a read-only analysis engine are consistent with MySQL 8.0. By default, there is no sorting, which is different from MySQL 5.7 where sorting is enabled by default. Therefore, this logic applies to read-only analysis engines built on both MySQL 5.7 and MySQL 8.0.

## **Data Loading Restrictions**

Last updated : 2025-05-09 11:51:17

Since the read-only analysis engine builds data through columnar storage, some special MySQL usage scenarios are not supported, as shown below:

#### Description of supported tables without primary keys and unique keys

**In version 1.2404.x**: When a table has neither a primary key nor a unique key, the table cannot be loaded into the read-only analysis engine. This requires that the table contain a primary key or a unique key. In the LibraDB engine, columnar data building will use the table's primary key or unique key by default. At the same time, in version 1.2404.x, any form of DDL statement to modify the table's primary key is not supported. If you modify the primary key in TencentDB for MySQL, this table in the read-only analysis engine will stop loading, and queries cannot be performed. If you need to recover it for use, you need to unload this table and then reload the data.

**In version 2.2410.x**: Tables can also be loaded into the read-only analysis engine without a primary key or unique key. Changes to the table's primary key are also supported, but the following special scenarios are not supported: The table only has fields such as time, date, timestamp, datetime, float, and double, and no other field types. When primary key DDL is performed to a table, no other field types except the above ones exist in the table.

# Tables using float or double field types as primary keys are not supported for loading into a read-only analysis engine

Tables using float or double field types (floating-point field types) as primary keys are not supported for loading data into a read-only analysis engine.

# Storage process, user-defined functions, triggers, foreign key constraints, events, and indexes will not be loaded in the read-only analysis engine.

It is not supported to build the above special objects in columnar storage.

# Tables with spatial type fields cannot be loaded into the read-only analysis engine. Json field types can be loaded into the read-only analysis engine but cannot be queried.

Tables with spatial field types cannot be loaded into the read-only analysis engine. When tables with Json fields are queried in the read-only analysis engine, the Json column values will be empty.

#### Loading temporary tables into the read-only analysis engine is not supported.

Since data modification of temporary tables is not logged, the data of temporary tables cannot be loaded into a readonly analysis engine.

#### Behavior of tables loaded into the analysis engine after being renamed

When a table is loaded into the analysis engine, after the rename statement is used to rename the table, the renamed table will also be automatically loaded into the analysis engine. At this time, if you create a table with the same name as the renamed table, this new table will also be automatically loaded into the analysis engine. For example, if table A has been loaded as columnar storage in the read-only analysis engine, and then table A is renamed to table B, in the read-only analysis engine, this table will also be renamed to table B. Meanwhile, if a new table named table A is added in the read-write instance, table A will also be automatically loaded into the read-only analysis engine.


#### **Unsupported table creation operations**

When an entire database or instance is preloaded into the analysis engine, executing the create table ... select statement in the primary instance of TencentDB for MySQL will prevent this table from being automatically loaded into the analysis engine. If you want this table to be loaded into the analysis engine, you need to remove this table from the analysis engine load list after the table is created and then load this table into the analysis engine through the console.

Loading tables with primary key fields of ultra-long numeric types into columnar storage is not supported In version 1.2404.x, loading a table with a primary key of Decimal exceeding 128 characters into columnar storage is not supported.

In version 2.2410.x, loading tables with a primary key of Decimal exceeding 256 characters into columnar storage is not supported.

#### Column-level permissions are not supported

By default, the read-only analysis engine synchronizes all users' permissions to query objects in the primary instance but does not synchronize column-level permissions. Therefore, it is impossible to control column-level permissions in the read-only analysis engine.

#### **Unsupported data types**

The read-only analysis engine has certain limitations regarding Data Type. If an object has an unsupported data type, this table will not be able to be loaded into the analysis engine.

#### **Unsupported table structure**

The read-only analysis engine does not support generated column syntax, whether it is a virtual column or a physical column. When the table definition statement contains generated column syntax, data loading will be interrupted. If such a scenario exists, please remove the table from the table load list of the read-only analysis engine.

#### Loading of views

If you need to load a view, directly load the database where the view is located into the read-only analysis engine. The view objects in this database will be automatically loaded into the analysis engine. If you do not need to load the entire database into the analysis engine but want to use the view in the read-only analysis engine, please submit a ticket for assistance.

#### Unsupported field type conversion

Some field type conversions are not supported in the read-only analysis engine. If table type conversion is performed in the primary instance of TencentDB for MySQL, it may cause the data loading task of the read-only analysis engine to terminate, and the loading statuses of all tables may become paused. For detailed information on supported type conversion, see Description of Supported Type Conversion Functions.

#### Note:

If you modify the field type of a table in TencentDB for MySQL, and this type modification is not supported in the readonly analysis engine, the loading statuses of all tables will become paused. At this time, if you need to recover it for use, you need to unload this table and then reload the data for this table.

#### Description of DDL synchronization for partition tables

**In version 1.2404.x**: Partition tables can be loaded into the read-only analysis engine by default and support querying. However, synchronization of related DDL operations on partitions of partition tables is not supported, such as rebuilding partitions, OPTIMIZE on partitions, repairing partitions, CHECK on partitions, swapping partitions, deleting partitions, and merging partitions. At the same time, querying a specific subpartition alone is not supported in the read-only analysis engine either.

#### Note:

When you drop a subpartition, truncate a partition, or execute an exchange partition for a partition table in the primary instance of TencentDB for MySQL, this table cannot be queried in the read-only analysis engine. If you need to recover it for use, you need to unload this table and then reload the data.

**In version 2.2410.x**: Partition tables can be loaded into the read-only analysis engine, and the following DDL changes can be made to partition tables on the source side.

Support the Drop Partition/Subpartition template of Range/List partition tables.

Support the Add Partition/Subpartition template of Range/List partition tables.

In the scenario without functions, the supported partition data types include Uint8, Uint16, Uint32, Uint64, Int8, Int16, Int32, and Int64.

Supported partition functions include year, month, and day. Data types supported by these functions include DATE, DATETIME, and TIMESTAMP.

# Description of Functions and Operators Description of Supported Bitwise Operations

Last updated : 2025-05-09 11:51:17

This document outlines the support for bitwise operations in the read-only analysis engine and lists the type restrictions for input parameters. Any input types or usage methods not explicitly mentioned in this document are not supported.

# Supported Bitwise Operations

## & Bitwise AND

Syntax: value1 & value2

Supported data types for value1: int, bigint, double, decimal, varchar, and char. Supported data types for value2: int, bigint, double, decimal, varchar, and char.

#### | Bitwise OR

Syntax: value1 | value2

Supported data types for value1: int, bigint, double, decimal, varchar, and char. Supported data types for value2: int, bigint, double, decimal, varchar, and char.

#### ~ Bitwise NOT

```
Syntax: ~value1
```

Supported data types for value1: int, bigint, double, decimal, varchar, and char.

#### ^ Bitwise XOR

Syntax: value1 ^ value2

Supported data types for value1: int, bigint, double, decimal, varchar, and char. Supported data types for value2: int, bigint, double, decimal, varchar, and char.

# Description of Supported Arithmetic Operations

Last updated : 2025-05-09 11:51:18

The read-only analysis engine currently supports most of the arithmetic operations logic in MySQL. However, certain limitations and incompatibilities still exist in specific scenarios. This document outlines the support for arithmetic operations in the read-only analysis engine and lists the type restrictions for input parameters. Any input types or usage methods not explicitly mentioned in this document are not supported.

# Arithmetic Operation Use Limits

**Integer division**: If the result has more than four decimal places, the read-only analysis engine will, by default, truncate the value and retain only the last four decimal places.

**Division with decimal type**: The number of decimal places in the result will be truncated based on the number of decimal places in the dividend plus 4. For example, in Decimal(9, 1) / Decimal(9, 1), the result will be truncated to 5 decimal places.

Addition, subtraction, multiplication, and division with decimal: If the total length of the result exceeds 65, an out of range error will be reported. This is because the maximum supported length for the Decimal type in the read-only analysis engine is 65.

**Mathematical operations between any numeric type and floating-point types (FLOAT, DOUBLE)**: These operations are subject to precision errors, primarily because floating-point types such as DOUBLE are not strictly precise by design. In practice, MySQL also discourages performing mathematical operations between numeric types and floating-point types.

**Arithmetic operations involving non-numeric types**: These operations involve implicit type conversion. Due to limitations in the read-only analysis engine's support for the CAST function, certain arithmetic operations may not be supported if the required CAST conversion is not supported. For details on supported CAST conversions, see Description of Supported Type Conversion Functions.

# Supported Arithmetic Operations

## + Addition Operator

#### Syntax: op1 + op2

Supported value types for op1: int, bigint, float, double, and decimal. Supported value types for op2: int, bigint, float, double, and decimal.

### - Subtraction Operator

#### Syntax: op1 - op2

Supported value types for op1: int, bigint, float, double, and decimal. Supported value types for op2: int, bigint, float, double, and decimal.

#### x Multiplication Operator

#### Syntax: op1 x op2

Supported value types for op1: int, bigint, float, double, and decimal. Supported value types for op2: int, bigint, float, double, and decimal.

#### / Division Operator

#### Syntax: op1 / op2

Supported value types for op1: int, bigint, float, double, and decimal. Supported value types for op2: int, bigint, float, double, and decimal.

## % Modulo Operator

#### Syntax: op1 % op2

Supported value types for op1: int, bigint, float, double, and decimal. Supported value types for op2: int, bigint, float, double, and decimal.

# **Comparison Operation Support Description**

Last updated : 2025-05-09 11:51:18

The read-only analysis engine currently supports most MySQL comparison operation logic, but there are still limitations and incompatibilities in some scenarios. This document introduces the support for comparison operations in the read-only analysis engine and lists the type limitations in input parameters. Input parameters and usage methods not mentioned in the document are not supported.

# Supported Comparison Operations

## >= Greater Than or Equal to Comparison

#### Syntax: op1 >= op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

#### <= Less Than or Equal to Comparison

#### Syntax: op1 <= op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

#### = Equal to Comparison

#### Syntax: op1 = op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

#### **!= Not Equal to Comparison**

#### Syntax: op1 != op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

#### < Less Than Comparison

#### Syntax: op1 < op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

#### > Greater Than Comparison



#### Syntax: op1 > op2

Supported value types for op1: int, bigint, float, double, decimal, char, varchar, date, and datetime. Supported value types for op2: int, bigint, float, double, decimal, char, varchar, date, and datetime.

# **Description of Supported Logical Operations**

Last updated : 2025-05-09 11:51:18

The read-only analysis engine currently supports most MySQL logical operations, but there are still limitations and incompatibilities in some scenarios. This document introduces the support for logical operations in the read-only analysis engine and lists the type limitations in input parameters. Input parameters and usage methods not mentioned in the document are not supported.

# Supported Logical Operations

#### and Operation

Syntax: op1 and op2

Supported value types for op1: Only the bool data type is supported. Supported value types for op2: Only the bool data type is supported.

#### or Operation

Syntax: op1 or op2

Supported value types for op1: Only the bool data type is supported. Supported value types for op2: Only the bool data type is supported.

#### not Operation

Syntax: not op1 Supported value types for op1: Only the bool data type is supported.

#### case Operation

Syntax: case when search\_condition then statement\_list else statement\_list end Supported value types for search\_condition: Only the bool data type is supported. Supported value types for statement\_list: int, bigint, double, decimal, varchar, and char.

#### if Operation

**Syntax**: if(expr1,expr2,expr3)

Supported value types for expr1: Only the bool data type is supported. Supported value types for expr2: int, bigint, double, decimal, varchar, and char. Supported value types for expr3: int, bigint, double, decimal, varchar, and char.

#### ifnull Operation

## Syntax: ifnull(expr1, expr2)

Supported value types for expr1: int, bigint, double, decimal, varchar, and char. Supported value types for expr2: int, bigint, double, decimal, varchar, and char.

# isnull Operation

Syntax: isnull(expr1)

Supported value types for expr1: int, bigint, double, decimal, varchar, and char.

# in Operation

Syntax: expr IN (value,...)

Supported value types for expr: int, bigint, double, decimal, varchar, and char. Supported value types for value: int, bigint, double, decimal, varchar, and char.

#### Note:

In scenarios with multiple value values, the value should be a constant and cannot be a variable or field name. For example, it can only be used in the format of column in ('1','2','3') rather than columnA in (columnB,varA).

Syntax: expr IN (sub\_query)

Supported value types for expr: int, bigint, double, decimal, varchar, and char.

Supported value types for sub\_query: Subquery and single field.

## like Operation

Syntax: expr like pattern

Supported value types for expr: int, bigint, double, decimal, varchar, and char.

Supported value types for pattern: int, bigint, double, decimal, varchar, and char.

#### Note:

The pattern only supports constants and does not support the substitution of variables or field names.

## **Coalescence Operation**

Syntax: coalesce(value,...)

Supported value types for value: int, bigint, double, decimal, varchar, and char.

# **Description of Supported String Functions**

Last updated : 2025-05-09 11:51:18

The read-only analysis engine currently supports most MySQL string functions, but there are still limitations and incompatibilities in some scenarios. This document introduces the supported list of string functions in the read-only analysis engine and lists the type limitations in input parameters. Input parameters and functions not mentioned in the document are not supported.

# Supported String Functions

#### substr Function

The string extraction function can be used in the following ways. **Syntax**: substr(str, start) Supported value types for str: char and varchar data. Supported value types for start: bigint, double, and decimal. **Syntax**: substr(str from start) Supported value types for str: char and varchar data. Supported value types for start: bigint, double, and decimal. **Syntax**: substr(string, start, length) Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal. Supported value types for length: bigint, double, and decimal. Supported value types for str: char and varchar data. Supported value types for str: char and varchar data. Supported value types for str: char and varchar data. Supported value types for str: char and varchar data. Supported value types for str: char and varchar data. Supported value types for str: char and varchar data. Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal. Supported value types for start: bigint, double, and decimal.

## char\_length Function

Decision function for string length, which returns the number of characters in a string.

**Syntax:** char\_length(str)

Supported value types for str: char and varchar data.

#### replace Function

String replacement function.

Syntax: replace(str, from\_str, to\_str)

Supported value types for str: char and varchar data.

Supported value types for from\_str: char and varchar data. Supported value types for to\_str: char and varchar data.

#### concat Function

String connection function.

Syntax: concat (str1, str2, ...) Supported value types for str1: char and varchar data. Supported value types for str2: char and varchar data.

## concat\_ws Function

String connection function with separators.
Syntax: concat\_ws(str1, str2, ...)
Supported value types for str1: char and varchar data.
Supported value types for str2: char and varchar data.

## left Function

Return the specified number of characters starting from the first character of the string.

Syntax: left(str, len)

Supported value types for str1: char and varchar data. Supported value types for len: int, bigint, double, and decimal data.

## right Function

Return the specified number of characters starting from the last character of the string. **Syntax**: right(str, len) Supported value types for str1: char and varchar data. Supported value types for len: int, bigint, double, and decimal data.

## ascii Function

Function that converts characters into ascii code values.

Syntax: ascii(str)

Supported value types for str1: char and varchar data.

## **length Function**

Function that determines string length and returns the number of bytes in a string.

Syntax: length(str)

Supported value types for str1: char and varchar data.

#### trim Function

Function that removes unnecessary characters before and after a string.

Syntax: trim([{BOTH | LEADING | TRAILING} [remstr] FROM] str)

Supported value types for remstr: char and varchar data.

Supported value types for str: char and varchar data.

#### **Itrim Function**

Remove unnecessary characters before a string.

```
Syntax: ltrim(str)
```

Supported value types for str: char and varchar data.

#### rtrim Function

Remove unnecessary characters after a string.

Syntax: rtrim(str)

Supported value types for str: char and varchar data.

#### position Function

Return the function corresponding to the position of the substring.

Syntax: position(substr IN str)

Supported value types for str: char and varchar data.

Supported value types for substr: char and varchar data.

#### instr Function

Search for the position of the first occurrence of a string in another string. **Syntax**: instr(str, substr) Supported value types for str: char and varchar data. Supported value types for substr: char and varchar data.

## **locate Function**

Return the position of the first occurrence of a substring in a string.

Syntax: locate(str, substr)

Supported value types for str: char and varchar data. Supported value types for substr: char and varchar data.

#### **lower Function**

Return the lowercase string of a string.

Syntax: lower(str)

Supported value types for str: char and varchar data.



#### ucase Function

Convert all lowercase letters in a given string to uppercase letters.

**Syntax**: ucase(str)

Supported value types for str: char and varchar data.

#### upper Function

Convert all lowercase letters in a given string to uppercase letters.

Syntax: upper(str)

Supported value types for str: char and varchar data.

#### substring\_index Function

Support reverse string extraction functions.

**Syntax**: substring\_index(str, delim, count)

Supported value types for str: char and varchar data.

Supported value types for delim: char and varchar data types. It can only be a constant, and variable values such as field values, are not accepted.

Supported value types for count: int, bigint, double, and decimal data types. It supports negative numbers. It can only be a constant and variable values such as field values are not accepted.

#### **Ipad Function**

Pad the beginning of a string to achieve the specified length.

**Syntax**: lpad(str, len, padstr)

Supported value types for str: char and varchar data.

Supported value types for len: int, bigint, double, and decimal data. It can only be a constant and variable values like field values are not accepted.

Supported value types for padstr: char and varchar data.

#### rpad Function

Pad the end of a string to achieve the specified length.

Syntax: rpad(str, len, padstr)

Supported value types for str: char and varchar data.

Supported value types for len: int, bigint, double, and decimal data. It can only be a constant.

Supported value types for padstr: char and varchar data.

#### strcmp Function

Compare two strings and return an integer based on the comparison result.

**Syntax**: strcmp(expr1, expr2)



Supported value types for expr1: char and varchar data.

Supported value types for expr2: char and varchar data.

## quote Function

Function that references a string containing special characters in a query.

```
Syntax: quote(str)
```

Supported value types for str: char and varchar data.

# Description of Supported Mathematical Operation Functions

Last updated : 2025-05-09 11:51:18

The read-only analysis engine currently supports most MySQL mathematical operation functions, but there are still limits and incompatibilities in some scenarios. This document introduces the supported list of mathematical operation functions in the read-only analysis engine and lists the type limitations in input parameters. Input parameters and functions not mentioned in the document are not supported.

# Supported Mathematical Operation Functions

#### round Function

Function that rounds a number.Syntax: round(x)Supported value types for x: Decimal data type.

### abs Function

Function that returns the absolute value.Syntax: abs (x)Supported value types for x: int, bigint, double, and decimal data types.

#### floor Function

Function that gives the nearest integer down.Syntax: floor(x)Supported value types for x: int, bigint, double, and decimal data types.

#### ceil Function

Function that gives the nearest integer up.
Syntax: ceil(x)
Supported value types for x: int, bigint, double, and decimal data types.

## ceiling Function

Function that rounds a number up.

Syntax: ceiling(x)

Supported value types for x: int, bigint, double, and decimal data types.



## sqrt Function

Square root function.

Syntax: sqrt(x)

Supported value types for x: int, bigint, double, and decimal data types.

### log Function

Logarithmic function. **Syntax:** log(x) Supported value types for x: int, bigint, double, and decimal data types.

## **log2** Function

Logarithmic function to base 2.

Syntax: log2(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### log10 Function

Logarithmic function to base 10.

Syntax: log10(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### In Function

Function that returns the natural logarithm of a number.

Syntax: ln(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### exp Function

Function that returns the e raised to the power of x.

Syntax: exp(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### pow Function

Function that calculates x raised to the power of y.

Syntax: pow(x, y)

Supported value types for x: int, bigint, double, and decimal data types. Supported value types for y: Support int, bigint, double, and decimal data types.

## sign Function

Used to determine the sign of a real number.



#### Syntax: sign(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### radians Function

Convert degrees to radians.

**Syntax**: radians(x)

Supported value types for x: int, bigint, double, and decimal data types.

#### degrees Function

Used to convert radians to degrees. **Syntax**: degrees(x) Supported value types for x: int, bigint, double, and decimal data types.

#### conv Function

Convert a number from one numeral base system to another and return the result as a string value.

**Syntax**: conv(n, from\_base, to\_base)

Supported value types for n: int, bigint, double, and decimal data types.

Supported value types for from\_base: int, bigint, double, and decimal data types.

Supported value types for to\_base: int data type.

#### crc32 Function

Cyclic redundancy check function.

Syntax: crc32(expr)

Supported value types for expr: int, bigint, double, and decimal data types.

#### greatest Function

Return the maximum value among N values.

Syntax: greatest (value1, value2, ...)

Supported value types for value1: int, bigint, double, and decimal data types. Supported value types for value2: int, bigint, double, and decimal data types.

#### least Function

Return the minimum value among N values.

Syntax: least(value1, value2, ...)

Supported value types for value1: int, bigint, double, and decimal data types.

Supported value types for value2: int, bigint, double, and decimal data types.

# **Description of Supported Date Functions**

Last updated : 2025-05-09 11:51:19

The read-only analysis engine currently supports most MySQL date functions, but there are still limitations and incompatibilities in some scenarios. This document introduces the supported list of date functions in the read-only analysis engine and lists the type limitations in input parameters. If input parameters and functions are not mentioned in the document, they are not supported.

# Supported Date Functions

## datediff Function

A function that calculates the number of specified time intervals between two dates.

Syntax: datediff(date1, date2)

Supported value types for date1: Date and datetime data.

Supported value types for date2: Date and datetime data.

#### year Function

A function that calculates the year. **Syntax**: year(date1) Supported value types for date1: Date, datetime, char, and varchar data.

## month Function

A function that calculates the month. **Syntax:** month(date1) Supported value types for date1: Date, datetime, char, and varchar data.

## day Function

A function that calculates the current day number.

Syntax: day(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

## hour Function

A function that calculates the hour.

Syntax: hour(date1)

Supported value types for date1: Date and datetime data.



### minute Function

A function that calculates the minute. **Syntax:** minute(date1) Supported value types for date1: Date and datetime data.

#### second Function

A function that calculates the second. **Syntax:** second(date1) Supported value types for date1: Date and datetime data.

#### microsecond Function

A function that returns the microsecond part of the time/date and time.

Syntax: microsecond(date1)

Supported value types for date1: Date and datetime data.

#### quarter Function

A function that returns the quarter of the year for the given date value (a number between 1 and 4).

**Syntax**: quarter(date1)

Supported value types for date1: Date, datetime, varchar, and char data.

#### date Function

Date function.
Syntax: date(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

## sysdate Function

A function that returns system time.

Syntax: sysdate()

#### date\_add Function

A function that adds dates. **Syntax**: date\_add(date1, interval expr type) Supported value types for date1: Date and datetime data. Supported value types for expr: Int, long, double, and decimal. Supported value types for type: DAY, which is a keyword.

#### date\_sub Function

A function that subtracts dates.

Syntax: date\_sub(date1, interval expr type) Supported value types for date1: Date and datetime data. Supported value types for expr: Int, long, double, and decimal. Supported value types for type: DAY, which is a keyword.

### dayofyear Function

A function that returns the day of the year for a given date. **Syntax:** dayofyear(date1) Supported value types for date1: Date, datetime, char, and varchar data.

#### dayofmonth Function

A function that returns the day of the month for a given date.

Syntax: dayofmonth(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

#### dayofweek Function

A function that returns the day of the week for a given date.

Syntax: dayofweek(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

#### dayname Function

A function that returns the name of the day of the week corresponding to the specified date.

Syntax: dayname(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

#### monthname Function

A function that returns the name of the month corresponding to the specified date.

**Syntax:** monthname(date1)

Supported value types for date1: Date, datetime, char, and varchar data.

#### to\_seconds Function

A function that returns the number of seconds counting from year 0 to the specified date/date and time.

**Syntax**: to\_seconds (date1)

Supported value types for date1: Date, datetime, char, and varchar data.

## last\_day Function

A function that returns the last day of the month.

Syntax: last\_day(date1)



Supported value types for date1: Date, datetime, char, and varchar data.

#### **SLEEP Function**

A function that pauses for a few seconds during execution.

Syntax: sleep(N)

Supported value types for N: int data.

#### **NOW Function**

A function that obtains the current time. **Syntax**: NOW() Obtain the current time without input parameters.

## CURRENT\_TIMESTAMP Function

A function that obtains the current time, equivalent to the NOW function.

Syntax: CURRENT\_TIMESTAMP()

Obtain the current time without input parameters.

#### **CURDATE Function**

A function that obtains the current date, equivalent to the CURRENT\_DATE function.

Syntax: CURDATE()

Obtain the current date without input parameters.

#### **CURRENT\_DATE** Function

A function that obtains the current date, equivalent to the CURDATE function.

Syntax: CURRENT\_DATE()

Obtain the current time without input parameters.

#### **CURRENT\_TIME** Function

A function that obtains the current time.

Syntax: CURRENT\_TIME()

Obtain the current time without input parameters.

#### **CURTIME Function**

A function that obtains the current time, equivalent to the CURRENT\_TIME function.

Syntax: CURTIME()

Obtain the current time without input parameters.

#### FROM\_UNIXTIME Function

A function that converts timestamp to time format.

Syntax: FROM\_UNIXTIME(unix\_timestamp[,format])

unix\_timestamp supports int type numbers.

The formats supported by format are shown in the table of date\_format function below. If format is empty, the default format is %Y-%m-%d %H:%i:%s.

## UNIX\_TIMESTAMP Function

A function that converts date and time to timestamp.

**Syntax:** UNIX\_TIMESTAMP([date])

date supports the date type of the DATE standard.

#### Note:

Only 1.2404.17.0 and later versions support the FROM\_UNIXTIME function and the UNIX\_TIMESTAMP function.

#### date\_format Function

A function that formats a date.

Syntax: date\_format(date1, format)

Supported value types for date1: Date, datetime, char, and varchar data, and do not support date1 containing null scenarios. format specifies the output format of date/time. The available formats are shown in the table below.

Format	Description
%a	Abbreviated names of the days of the week.
%b	Abbreviated names of the months.
%с	Month and numeric value.
%D	Days of the month with English prefixes.
%d	Day of the month; numeric value (00 - 31).
%e	Day of the month; numeric value (0 - 31).
%f	Microsecond.
%Н	Hour (00 - 23).
%h	Hour (01 - 12).
%I	Hour (01 - 12).
%i	Minute; numeric value (00 - 59).
%j	Day of the year (001 - 366).

🔗 Tencent Cloud

%k	Hour (0 - 23).
%	Hour (1 - 12).
%M	Names of the months.
%m	Month; numeric value (00 - 12).
%р	AM or PM.
%r	Time; 12-hour clock (hh:mm:ss AM or PM).
%S	Second (00 - 59).
%S	Second (00 - 59).
%T	Time; 24-hour clock (hh:mm:ss).
%U	Week (00 - 53). Sunday is the first day of the week.
%u	Week (00 - 53). Monday is the first day of the week.
%V	Week (01 - 53). Sunday is the first day of the week, used with %X.
%v	Week (01 - 53). Monday is the first day of the week, used with %x.
%W	Names of the days of the week.
%w	Day of the week $(0 = Sunday; 6 = Saturday)$ .
%X	Year, where Sunday is the first day of the week. The value contains 4 digits, used with %V.
%x	Year, where Monday is the first day of the week. The value contains 4 digits, used with %v.
%Y	Year (4 digits).
%y	Year (2 digits).

# **Description of Supported Aggregate Functions**

Last updated : 2025-05-09 11:51:19

The read-only analysis engine currently supports most MySQL aggregate functions, but there are still limitations and incompatibilities in some scenarios. This document introduces the supported list of aggregate functions in the read-only analysis engine and lists the type limitations in input parameters. If input parameters and functions are not mentioned in the document, they are not supported.

# Supported Aggregate Functions

#### count Function

Syntax: count (val)

Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### sum Function

**Syntax:** sum(val) Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### avg Function

Syntax: avg(val) Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### max Function

Syntax: max(val) Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### min Function

Syntax: min(val) Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### group\_concat Function

Syntax: group\_concat(val)

Supported value types for val: int, bigint, double, decimal, varchar, and char.

#### any\_value(arg) Function

Syntax: any\_value(val)



Supported value types for val: int, bigint, double, decimal, varchar, and char.

# **JSON** Function

Last updated : 2025-05-09 11:51:19

In 2.2410.x and later versions, the read-only analysis engine adds support for JSON (JavaScript Object Notation) data types that conform to RFC 7159 standard. Users can save semi-structured JSON data and can access and modify data in JSON documents.

JSON documents stored in JSON columns can be converted into a format that is easy to read and access. When the server needs to read a JSON value stored in binary format, it does not need to parse the value from text. Besides JSON data types, there is a set of SQL functions that can be used to create, search, and perform other operations on JSON values. The supported functions are as follows:

# JSON\_OBJECT Function

Create a JSON object. Syntax: JSON\_OBJECT (key1, value1, key2, value2...) Create a JSON object based on the input parameter value.

# JSON\_REMOVE Function

Delete the corresponding path from the specified JSON object. **Syntax:** JSON\_REMOVE(json\_doc, path[, path] ...) json\_doc is a JSON field, and path is a path. Multiple paths are supported.

# JSON\_EXTRACT Function

Extract elements at the specified path from a JSON field. If the specified path does not exist, NULL is returned. Multiple paths can be specified, and multiple matched values will be returned as an array. **Syntax:** JSON\_EXTRACT(json\_doc, path[, path] ...) is synonymous with column-> path . json\_doc is a JSON field, and path is a path.

# JSON\_UNQUOTE Function

Escape JSON and output it as a string.

**Syntax**: JSON\_UNQUOTE(column -> path) is synonymous with column->> path.



json\_doc is a JSON field, and path is a path.

# JSON\_CONTAINS Function

Specify whether the path contains the specified data. If it does, 1 is returned; if not, 0 is returned. If any parameter is NULL or the path does not exist, NULL is returned.

**Syntax:** JSON\_CONTAINS(json\_doc, candidate[, path])

json\_doc is a JSON field, candidate is a judgment value, and path is a path.

# JSON\_CONTAINS\_PATH Function

Check whether a specified path exists and whether one or all conditions are met. If it exists, 1 is returned; if not, 0 is returned. If any parameter is NULL, NULL is returned.

**Syntax**: JSON\_CONTAINS\_PATH(json\_doc, one\_or\_all, path[, path] ...)

json\_doc is a JSON field. In one\_or\_all, one means return once one is queried, and all means return after all are queried. path is a path. Multiple paths are supported.

# Type Conversion

Support conversion from JSON to INT, UINT, FLOAT, STRING, DATE, DATETIME, and DURATION. Support conversion from INT, UINT, FLOAT, STRING, DATE, DATETIME, and DURATION to JSON.

# Description of Supported Type Conversion Functions

Last updated : 2025-05-09 11:51:19

In the read-only analysis engine, performing SQL operations may result in implicit conversions caused by mismatches in field types, or explicit type conversions may be specified using the CAST function. The support for type conversions varies across different versions of the read-only analysis engine. This document will provide a detailed description of the supported type conversion.

#### Note:

When values of different types are compared or calculated in an SQL expression, implicit conversion will occur. All implicit conversion logic that may occur during usage is subject to the conversion support conditions shown in the table below. If a value does not support implicit conversion in a comparison or calculation, an error message will be generated indicating that the conversion is not supported.

**Syntax:** CAST (expression AS data\_type)

In the read-only analysis engine of **version 1.2404.x**, the detailed information about supported conversion types is shown in the table below.

Source/Destination	Binary	Char	Date	DateTime	Decimal	Double	Float	Real	Sign
Bigint	1	1	×	×	1	1	×	1	1
Binary	1	×	×	×	1	1	×	1	1
Date	1	×	1	1	×	×	×	×	×
Datetime	1	×	1	1	×	×	×	×	×
Decimal	1	1	×	×	1	1	×	1	1
Double	1	1	×	×	1	1	×	1	×
Float	1	1	×	×	1	×	1	1	1
Real	1	1	×	×	1	1	×	1	×
Signed	1	1	×	×	1	1	×	1	1
Time	×	×	×	×	×	×	×	×	×
Unsigned	1	1	×	×	1	1	×	1	1
Year	1	1	×	×	1	1	×	1	1

🕗 Tencent Cloud

Varchar	$\checkmark$	$\checkmark$	×	×	1	1	×	$\checkmark$	1

In the read-only analysis engine of **version 2.2410.x**, the detailed information about supported conversion types is shown in the table below.

Source/Destination	Int /Bigint	Binary	Char	Date	DateTime	Decimal	Double	Float	Re
Int/Bigint	1	1	1	1	1	1	1	1	1
Binary	1	1	×	1	1	1	1	1	1
Date	1	1	1	1	1	1	1	1	1
Datetime	1	1	1	1	1	1	1	1	1
Decimal	1	1	1	1	1	1	1	1	1
Double	<i>✓</i>	1	1	1	1	$\checkmark$	1	1	1
Float	<i>✓</i>	1	1	1	1	$\checkmark$	1	1	1
Real	<i>✓</i>	1	1	1	1	$\checkmark$	1	1	1
Signed	1	1	1	1	1	1	1	1	1
Time	1	1	1	1	1	1	1	1	1
Unsigned	1	1	1	1	1	1	1	1	1
Year	1	1	1	1	1	1	1	1	1
Varchar	1	1	1	1	1	1	1	1	1

# Data Type

Last updated : 2025-05-09 11:51:19

This document provides a detailed list of the data types supported by the read-only analysis engine.

# Numeric Data Types

Type Classification	Туре	Support Status and Description
	TINYINT	Supported
	SMALLINT	Supported
	MEDIUMINT	Supported
	INT	Supported
	BIGINT	Supported
	DECIMAL and NUMERIC	Supported
Numeric Data Types	FLOAT	Supported
	DOUBLE	Supported
	BIT	Supported Note: In the read-only analysis engine, the BIT data type is supported by converting it to an unsigned INT type. As a result, BIT data is only supported in numeric comparison scenarios within the read-only analysis engine.
	DATE	Supported
	DATETIMESTAMP	Supported
Date and Time Data Types	TIMESTAMP	Supported
	TIME	Supported
	YEAR	Supported
Character Data	CHAR	Supported
Types	VARCHAR	Supported



	BINARY	Supported	
	BLOB	Supported	
	TEXT	Supported	
	ENUM	Supported Note: In the read-only analysis engine, the ENUM and SET data types are supported by	
	SET	converting them to the STRING type. As a result, related expressions or specialized functions such as & and FIND_IN_SET are not supported.	
	GEOMETRY	Not supported.	
Cratic Data Tyrac	POINT	Not supported.	
Spallal Data Types	LINESTRING	Not supported.	
	POLYGON	Not supported.	
JSON Data Types	JSON	Supported Note: The JSON data type is supported starting from read-only analysis engine kernel version 2.2410.1.0. It is not supported in version 1.2404.x.	

# Read-Only Analysis Engine Management Enabling Read-Only Analysis Engine

Last updated : 2025-05-09 11:51:19

This document introduces how to create a read-only analysis engine instance for TencentDB for MySQL.

# Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. The instance is running.

# **Operation Steps**

#### Notes:

Currently, you can create a read-only analysis engine only under an existing cluster in the provisioned resource instance mode, but cannot create it under a cluster being created.

The creation of a read-only analysis engine is only supported for MySQL 5.7 and MySQL 8.0 kernel versions.

The creation of a read-only analysis engine is not supported for instances using the RocksDB kernel.

The creation of a read-only analysis engine is only supported for instances with two-node or three-node architectures. 1. Log in to the MySQL console.

2. In the instance list, locate the instance to which you want to add a read-only analysis engine. Click the **Instance ID** or **Manage** in the operation column to enter the instance management page.

3. Select the **Read-only Instance** page and click **Create** to enter the read-only instance creation page.

#### Note:

You can also go to the instance detail page and click **Add a read-only analysis engine** under the instance architecture diagram to enter the **Read-Only Instance** page, and then click **Create**.

#### 4. Complete the following configuration and click **Buy Now**.

Parameter	Description
Instance Engine	Select the engine for the current read-only instance. Currently, both InnoDB and LibraDB engines are supported. In this case, select the LibraDB engine.
Specify Analytic RO Group	Supports automatic assignment by the system, creating an RO group, or selecting an existing RO group.

	Automatic assignment by the system: If multiple instances are purchased at once, the system will assign a separate RO group to each instance. The weight distribution method is set to automatic assignment by the system by default. Create an RO group: Creates an RO group. If multiple instances are purchased at once, all will be assigned to this RO group. The weight distribution method is set to automatic assignment by the system by default. Existing RO group: Specifies an existing RO group. If multiple instances are purchased at once, all will be assigned to this RO group. The weight distribution method follows the RO group configuration: If the RO group is set to automatic assignment by the system, RO groups will be added automatically based on the purchased specifications; if it is set to custom assignment, the default weight is zero. Since instances in the same RO group share the same private IP address, they will also share the same security group settings in VPC networks. If you select an existing RO group, you will not be able to customize the security group during the purchasing process. <b>Note:</b> Analytical RO groups do not support capabilities such as delayed replica elimination upon exceeding latency thresholds. All other functionalities are consistent with those of standard RO groups.
Analytic RO Group Name	When creating an RO group, you need to specify a name for the new RO group. The name can include Chinese characters, English letters, digits, hyphens (-), underscores (_), and periods (.), and should be less than 60 characters in length.
Billing Mode	If the billing mode of the primary instance is monthly subscription, the read-only analysis engine supports either monthly subscription or pay-as-you-go billing mode. However, if the primary instance is billed under the pay-as-you-go mode, the read-only analysis engine only supports pay-as-you-go billing.
Region	Select the region where the read-only analysis engine instance resides. Currently, only the same region as the primary instance of TencentDB for MySQL is supported.
AZ	Select the availability zone where the read-only analysis engine instance resides. For the specific available zones, see the options shown on the actual page.
Instance Specification	Select the instance specification for the read-only analysis engine. Different specifications offer varying levels of performance. Since data analysis and complex SQL execution can consume significant computing resources, the instance specifications for the read-only analysis engine are relatively high. In high-concurrency data write scenarios, additional resources may also be consumed to ensure data consistency between the primary and secondary nodes. Choose the specification based on your actual business needs.
Hard Disk	Select the storage capacity of the hard disk. The minimum disk size is 100 GB, and the maximum limit depends on the selected instance specification. See the actual purchase page for available options.
Data Load Mode	Specifies how the read-only analysis engine loads data from the corresponding primary instance of TencentDB for MySQL. The default option is no objects.

	All objects: Loads all data objects from the primary instance into the read-only analysis engine (unsupported objects will be automatically skipped). Partial objects: Only the selected objects will be loaded into the read-only analysis engine. After selecting this option, you should specify the objects to be loaded in the section below. No objects: No objects will be loaded. This option creates an empty read-only analysis engine instance, and you will need to configure data loading manually afterward.
Network	By default, the VPC is the same as that of the primary instance of TencentDB for MySQL, and the subnet is selected from the available subnets within the specified availability zone. It also supports deployment in a different VPC from the primary instance.
Security Group	You can select a security group different from that of the corresponding TencentDB for MySQL instance for the read-only analysis engine, in order to control its access policy.
Tag	Adds tags to the read-only analysis engine instance.
Instance Name	Customizes the instance name. You can choose <b>naming after creation</b> or <b>immediate naming</b> .
Quantity	Select the quantity to purchase based on your actual business needs. You can purchase up to 6 read-only analysis engines.

5. After the purchase is completed, you will be redirected to the instance list. Once the instance status shows running, it is ready for use.

# **Basic Operations**

Last updated : 2025-06-10 10:26:54

The operations for read-only analysis engine instances are consistent with those for other read-only instances. You can manage the read-only analysis engine instances independently. This document introduces the supported operations for read-only analysis engine instances in TencentDB for MySQL.

# Prerequisites

You have prepared a read-only analysis engine instance. See Enabling Read-Only Analysis Engine. The instance is running.

# **Related Operation Guide**

## Viewing Instance Information

1. Log in to MySQL console. In the instance list, click **Instance ID** or **manage** in the operation column to enter the instance detail page.

2. In the instance architecture diagram, you can view the read-only analysis engine instances associated with the current primary instance.

3. Click the **Instance ID** of the target read-only analysis engine instance to enter its details page and view detailed instance information.

#### **Renaming an Instance**

The method is the same as for primary and read-only instances.

#### **Modifying Security Groups**

The method is the same as for primary and read-only instances, except that you need to select the instance type as the read-only analysis engine instance. See Managing Cloud Database Security Groups.

#### Modifying the IP Address and Port

The method is the same as for primary and read-only instances. See Modify Private Network Address.

## **Changing Network**

The method is the same as for primary and read-only instances. See Switch Network.

## **Enabling/Disabling Public Network Address**

The method is the same as for primary and read-only instances. See Enabling/Public Network Connection Address. **Notes:** 

To enable public network access, you need to authorize the binding of a public network address to a private network address. Exposing the instance to the public network introduces certain security risks, so external access should be explicitly authorized before it can be enabled.

Public network access is intended only for development or auxiliary database management. For business access, use private network access.

#### Modifying the Maintenance Time

The method is the same as for primary and read-only instances. See Modify Instance Maintenance Time.

#### **Deleting and Restoring Instances**

When an instance is no longer needed, you can delete it. For details, see Terminate an Instance. After a read-only analysis engine instance is deleted, it will be inaccessible, but data will continue to be loaded from the read/write node into the engine.

For monthly subscription instances, the instance will be moved to the recycle bin and retained for 7 days after deletion. During this period, the instance will be inaccessible. If you want to restore the instance, you can renew it in the Recycle Bin.

For a pay-as-you-go instance, the instance will be moved to the recycle bin and retained for 24 hours after deletion. During this period, the instance will be inaccessible. If you want to restore the instance, you can renew it in the Recycle Bin.
## Parameter Management

Last updated : 2025-05-09 11:51:20

This document introduces the parameters that can be modified for read-only analysis engine instances.

### Prerequisites

You have prepared a read-only analysis engine instance. See Enabling Read-Only Analysis Engine. The instance is running.

### Parameter Management

Managing parameters for read-only analysis engine instances is supported, including modifying parameters, applying parameter templates, and viewing parameter change history. The operation is generally consistent with that of read-only instances based on the TXSQL engine. For details, see <u>Setting Instance Parameters</u>.

### Parameter List

The following table lists the parameters that can be modified for read-only analysis engine instances, along with parameter modification recommendations:

### **Common Parameters**

Parameter Name	Default Value	Whether to Restart	Parameter Modification Recommendations
sql_mode	Consistent with the read-write instance	No	This parameter controls SQL execution behavior in the read- only analysis engine. Its behavior is consistent with that of MySQL. However, some SQL_MODE values are either incompatible or configurable but not effective. For details, see the SQL_MODE section in Compatibility and Use Limits.

mpp-time-out	3600	Yes	This parameter controls the SQL execution timeout in the read-only analysis engine. If a query exceeds this time limit, a timeout error will be returned. You can adjust this parameter based on actual business needs.
max_connections	1024	Yes	This parameter controls the maximum number of connections to the read-only analysis engine. When the number of connections exceeds this limit, new connection attempts will result in an error. It is recommended to configure the maximum number of connections based on actual business needs. Since the read-only analysis engine is primarily used for handling complex SQL queries, it is not recommended to create too many connections. Keeping the number of connections below 200 is recommended.
max_threads	16	No	This parameter controls the maximum number of threads that can be used for SQL execution in the read-only analysis engine. By default, it is set to match the maximum number of CPUs allocated to the instance. It is recommended to set this value to half of the instance's CPU count.
libratree_enable_check_disk_space_available	ON	Yes	This parameter controls whether the disk space reservation feature is enabled. By default, users can only use up to 90% of the disk capacity allocated to the instance. Once



			disk usage reaches 90%, data synchronization and loading in the read-only analysis engine will be suspended.
libratree_reserve_space_pcnt	10%	Yes	This parameter controls whether the disk space reservation ratio is enabled. The default value is 10%, meaning that 10% of the disk space is reserved and cannot be used for data storage.

### **Scenario Parameters**

Parameter Name	Default Value	Whether to Restart	Parameter Modificatior Recommen
replication_sink_parallelism	16	No	Modifying th parameter increases th concurrency during incre data loading which can ir loading spe However, it increases th on the read- node. It is recommend adjust this parameter c when the re write node h sufficient av resources.
replication_dumper_split_chunk_parallel	1	No	This paramic controls the concurrency for splitting chunks. A h concurrency



			can speed u full data load process. Ho increasing t value will pla additional lo the read-wr node. It is recommend adjust this parameter of when the re write node k sufficient av resources.
replication_dumper_table_parallel	4	No	This paramic controls the number of ta that can be concurrently during full d loading. Wh read-write instance ha sufficient resources, increasing t number of concurrently loaded table speed up th data loading process. It i recommend adjust this parameter c when the re write node t sufficient av resources.
replication_dumper_insert_parallel	4	No	This parame controls the concurrency for writing to



			read-only ar engine durir data loading When the re only analysi engine has sufficient resources, a higher write concurrency speed up th data loading process It is recommend adjust this parameter c when the re only analysi engine has sufficient av resources.
replication_dumper_select_parallel	16	No	This parame controls the number of concurrent 1 that read tal data during data loading When the re write instan- sufficient resources, increasing t concurrency can speed t full data load process. It i recommend adjust this parameter c when the re write node h sufficient av resources.

replication_dumper_chunk_size	10000	No	This paramic controls the each chunk from the rea instance du data loading Larger chur consume m resources o read-write instance bu speed up th data loading process. It i recommend adjust this parameter c when the re only analysi engine has sufficient av resources.
libra_runtime_filter_type	MIN_MAX,BLOOM_FILTER,IN_FILTER	No	Controls the of Runtime that can be by the Runt Filter feature more details Runtime Filt User Manua
libra_enable_runtime_filter	ON	No	Controls wh to enable or disable the Runtime Fili feature. For details, see Runtime Fili User Manua
libra_enable_late_materialization	ON	No	Controls wh to enable or disable the materializat feature. For



			details, see Materializat
null_for_aggregation_by_empty_set	ON	No	This parame controls where return null v when a GR( BY operation encounters empty result this parame enabled, nu values will k included in the GROUP BY operation. If disabled, nu values will k skipped.

# Data Loading Loading Data to Read-Only Analysis Engine

Last updated : 2025-05-09 11:51:21

This document introduces how to load data into the read-only analysis engine for data analysis and querying. **Note:** 

Currently, only databases, tables, and views are supported for loading into the read-only analysis engine.

### Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. The instance is running.

### Directions

The read-only analysis engine currently supports multiple methods for data loading. The detailed procedures are as follows.

### Method 1: Loading Data When Creating the Read-Only Analysis Engine

When creating a read-only analysis engine instance, you can choose the data loading method. Currently, three options are supported: all objects, partial objects, and no object. Each option corresponds to a different data loading scheme. 1. The data loading method is all objects.

In all objects mode, all databases and tables in the TencentDB for MySQL instance are automatically synchronized to the read-only analysis engine. If certain objects do not meet the loading requirements, they will be automatically ignored and will not be loaded into the engine.

2. The data loading method is partial objects.

If partial objects is selected as the data loading method, you need to specify the objects to be loaded into the analysis engine when creating the read-only analysis engine instance. Any objects that do not meet the loading requirements will be ignored. If a view references a table that has not been loaded into the engine, an error will be reported during the view loading process indicating that the table does not exist.

3. The data loading method is no object.

If no object is selected as the data loading method, no objects will be loaded when you create the read-only analysis engine instance. You will need to manually specify the objects to be loaded during instance operation. See Method 2: Manual Loading.

### Method 2: Manual Loading

If you selected no object when creating the read-only analysis engine instance, or if you need to modify the objects loaded into the engine, you can use the following method.

#### Loading Objects via the Console

1. Log in to MySQL console. In the instance list, click **Instance ID** or **manage** in the operation column to enter the instance detail page.

2. Under the instance architecture diagram, locate the read-only analysis engine instance and click **Instance ID** to enter the details page of the read-only analysis engine.

3. Click **Modify Objects**. In the Modify Objects dialog box, **Select** or **cancel** the objects you want to load into the read-only analysis engine.

# Viewing Data Loading Status

Last updated : 2025-05-09 11:51:21

### Overview

There are two distinct phases when you load data into the read-only analysis engine: the full data loading phase, followed by the real-time data loading phase. When data is being loaded into the read-only analysis engine, queries can only be performed after the data objects have been fully built. You can check the loading status of data tables using the following methods.

### Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. The instance is running. Loading Data to Read-Only Analysis Engine.

### Directions

### Method 1: Viewing Data Loading Status in the Console

1. Log in to MySQL console. In the instance list, click **Instance ID** or **manage** in the operation column to enter the instance detail page.

2. Under the instance architecture diagram, locate the read-only analysis engine instance and click **Instance ID** to enter the details page of the read-only analysis engine.

3. On the details page of the read-only analysis engine instance, you can view object information, which includes the data loading status.

Parameter	Description
Database Name	Indicates the name of the database to which the object belongs. Supports filtering by database name. Note: If a large number of databases are loaded into the read-only analysis engine, filtering by database may result in longer console loading times.
Object Name	Indicates the name of the table loaded into the read-only analysis engine. Views can also be loaded into the engine, but view information is not displayed here. To view loaded views,

	use MySQL commands. When a large number of objects are loaded, you can search by keyword in the object name through the console.
Status	Status values include data loading, real-time increment, data sync error, and data sync completed. For details, see it in the console.
Data loading status	If an object is currently in the data loading phase, the data loading status will display the loading progress of the corresponding table.
Incremental Phase Latency	If an object is in the real-time incremental phase, its synchronization delay information will be displayed.
Other Info	If a data loading exception occurs, relevant information about the exception will be displayed under other information for the object.

### Method 2: Viewing Data Loading Status via SQL Statements

```
show replication tables;
Or
select * from information_schema.libra_table_status;
```

In the read-only analysis engine, executing this command allows you to view the status of all data objects currently being loaded. You can also apply a WHERE clause to filter the results.

The field display information is shown in the figure below:

# **Controlling Data Loading Speed**

Last updated : 2025-05-09 11:51:21

### Overview

During full data loading, a certain amount of resource consumption will be imposed on the primary instance of TencentDB for MySQL, especially when you perform one-time loading on instances with large data volumes. For instances with limited resources, this operation may affect the performance and stability of the primary instance to some extent. You can use the following methods to control the impact on the primary instance of TencentDB for MySQL during the data loading process and reduce resource consumption on the source side. However, **these operations may reduce the efficiency of converting row-based data to columnar storage**. Therefore, make adjustments based on your actual business needs.

### Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. The instance is running. Loading Data to Read-Only Analysis Engine.

### Controlling Data Loading Speed

### Adjust the number of concurrent data reads for a single table in the scenario of full data loading.

Modify the replication\_dumper\_select\_parallel parameter through the feature of adjusting parameters. This parameter indicates the concurrency for simultaneous data reads from the same object. A larger value allows for a broader range of data reads from the table. In scenarios with sufficient resources, this can improve the loading efficiency. If resources are insufficient, it is recommended to set this value smaller. See Setting Instance Parameters for operation details.

### Adjust the number of rows read from a single table in the scenario of full data loading.

Modify the replication\_dumper\_chunk\_size parameter through the feature of adjusting parameters. This parameter indicates the number of data rows read in a single operation for the same object. A larger value allows for more data reads each time. In scenarios with sufficient resources, this can improve the loading efficiency. If resources are insufficient, it is recommended to set this value smaller. See Setting Instance Parameters for operation details.

### Adjust the total number of concurrent data reads in the scenario of full data loading.

Modify the replication\_dumper\_table\_parallel parameter through the feature of adjusting parameters. This parameter indicates the number of tables from which data can be read simultaneously. A larger value allows more tables to be



read simultaneously. In scenarios with sufficient resources, this can improve the loading efficiency. See Setting Instance Parameters for operation details.

# Canceling Loading Certain Objects into the Read-Only Analysis Engine

Last updated : 2025-05-21 18:06:50

### Overview

During the data loading process into the read-only analysis engine, there may be scenarios where it is necessary to cancel the loading of certain objects from the engine.

### Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. The instance is running. Loading Data to Read-Only Analysis Engine.

### Directions

### Method 1

The following section describes the steps to cancel data loading from a read-only analysis engine instance. 1. Log in to the MySQL console. In the instance list, click the **Instance ID** of the read-only analysis engine or click **manage** in the operation column to enter the instance detail page of the read-only analysis engine.

### 2. Click **Modify objects** next to **Data Object**.

3. In the pop-up window, under the selected object on the right, deselect the objects that have been loaded into the read-only analysis engine, and click **OK**.

### Method 2

The following section describes the steps to cancel data loading by navigating from the primary instance details page where the read-only analysis engine is mounted.

1. Log in to MySQL console. In the instance list, click **Instance ID** or **manage** in the operation column to enter the instance detail page.

2. Under the instance architecture diagram, locate the read-only analysis engine instance and click **Instance ID** to enter the details page of the read-only analysis engine.

3. Click **Modify objects**. In the pop-up window, under the selected object on the right, deselect the objects that have been loaded into the read-only analysis engine, and click **OK**.

# Executing Queries Collecting Statistics

Last updated : 2025-05-09 11:51:21

For the read-only analysis engine, collecting basic statistics on table objects is essential for achieving optimal SQL execution results. Only after statistics collection is completed can the system provide more accurate and efficient execution plan recommendations for user queries.

### Statistics Overview

Statistics refer to the features and distribution of data within database objects. They enable the database to estimate execution costs more accurately and optimize execution plan selection. Types of statistics include the number of rows in a table, the number of distinct values (NDV) in a column, and column histograms. Collecting statistics involves scanning database objects and storing the gathered information in the system's data dictionary. In addition, the system maintains a local cache of statistics to improve the optimizer's access efficiency. When you collect and use statistics, the following two points should be noted:

The accuracy of statistics is crucial to the optimizer's decision-making process. Therefore, regularly updating and maintaining statistics is key to ensuring consistent database performance.

Collecting statistics may have some impact on database performance, as it involves scanning and analyzing database objects. Therefore, it is important to balance performance and the accuracy of statistics during the collection process. In summary, statistics plays a vital role in database management systems by providing key information about database objects. They assist the optimizer in making more informed decisions, thereby enhancing query performance and improving the overall efficiency of the database system.

### Note:

Starting from version 2.2410.1.0, random sampling is supported. This feature allows random sampling based on the sampling rate specified by the TableScan operator. It also supports block-level sampling and enables statistics collection according to the configured sampling ratio, significantly reducing the overhead caused by full table scans during the statistics collection process.

### **Statistics Collection**

The read-only analysis engine currently supports two different methods for collecting statistics:

### Automatic collection

Starting from version 2.2410.1.0, the read-only analysis engine supports dynamic sampling and automatic statistics collection during the process of loading data into columnar storage. Users do not need to manually collect statistics, as

the system automatically generates the relevant statistics for databases and tables.

#### Manual collection

If you need to manually obtain the latest statistics for a specific table, you can log in to the read-only analysis engine instance using the MySQL client and execute the following statement to collect statistics for the specified table object:

ANALYZE TABLE ;

# **SELECT Statement Description**

Last updated : 2025-05-09 11:51:22

This document describes the query statements supported in the read-only analysis engine and related instructions.

### **Object Permission Description**

In the read-only analysis engine, only data of objects that have been fully loaded into the engine can be queried. To query data, the SELECT permission of the corresponding data table is required. If you do not have the permission, apply for object permissions through account management of TencentDB for MySQL. For details, see Modifying Account Permissions.

### SELECT Statement

In the read-only analysis engine, SELECT statements are mainly used for data queries.

```
SELECT
[ALL | DISTINCT | DISTINCTROW ]
select_expr [, select_expr] ...
[FROM table_references
[WHERE where_condition]
[GROUP BY {col_name | expr | position}]
[HAVING where_condition]
[WINDOW window_name AS (window_spec)
[, window_name AS (window_spec)] ...]
[ORDER BY {col_name | expr | position}
[ASC | DESC]]
[LIMIT {[offset,] row_count | row_count OFFSET offset}]
```

The preceding SELECT statement block is supported in the read-only analysis engine. The following section describes the **statement units** within the complete SELECT statement segment of MySQL that are not supported in the read-only analysis engine:

HIGH\_PRIORITY indicates the priority for executing SELECT statements, which does not take effect in the read-only analysis engine. However, if this keyword is included in an SQL statement, it does not affect the execution result. STRAIGHT\_JOIN indicates that the optimizer is forced to join tables based on the order listed in the clause. However, in a read-only analysis engine, the execution plan is selected only based on the columnar execution optimizer of the read-only analysis engine. Therefore, this keyword does not take effect in a read-only analysis engine.



The SQL\_SMALL\_RESULT, SQL\_BIT\_RESULT, SQL\_BUFFER\_RESULT, SQL\_NO\_CACHE, and

SQL\_CALC\_FOUND\_ROWS query options also do not take effect in a read-only analysis engine but do not affect the execution of query statements.

The PARTITION partition\_list clause is not supported in the read-only analysis engine, and this option will cause the query execution to fail.

The WITH ROLLUP syntax is not supported. If this keyword is included in an SQL statement, a query execution error occurs.

FOR UPDATE | Lock IN SHARE MODE and other clause syntaxes are not supported. If such clauses are included in an SQL statement, they do not take effect, but the SQL statement can be executed properly.

into option is not supported. If this keyword is included in an SQL statement, an error occurs when the SQL statement is executed.

# **CTE Syntax Use Instructions**

Last updated : 2025-05-09 11:51:23

In version 5.7, TencentDB for MySQL does not support the Common Table Expression (CTE) syntax. Only the row storage engine in version 8.0 and later supports the CTE syntax. However, the read-only analysis engine supports the CTE syntax in both version 5.7 and 8.0.

### Support Status

#### 1.2404.x

In version 1.2404.x, the read-only analysis engine only supports non-recursive CTEs. When the the CTE syntax is used, Hint /\*+ MERGE() \*/ needs to be added. Otherwise, an incorrect execution plan may be generated. In later versions, the read-only analysis engine will gradually support recursive CTEs and optimize the CTE execution performance.

#### 2.2410.x

In version 2.2410.x, the read-only analysis engine still only supports non-recursive CTEs. However, the CTE syntax can be executed properly without needing to add Hint /\*+ MERGE() \*/ . In addition, in version 2.2410.x, streaming execution is supported, which significantly optimizes the CTE execution performance.

### Introduction to CTEs

CTEs are parts of the SQL standard and are often referred to as "WITH clauses". CTEs were first introduced in the SQL:1999 standard and designed to provide a concise and powerful method to define temporary result sets. These result sets can be referenced multiple times within a single SQL statement, greatly improving the readability and maintainability of queries.

WITH clause usage example:

```
-- Start defining CTE.
WITH CustomerCTE AS (
    SELECT customer_id, first_name, last_name, email_address
    FROM customer
)
-- End defining CTE.
SELECT *
FROM CustomerCTE; -- Reference the CTE.
```

### **CTE Advantages**

CTEs have many advantages in complex SQL queries.

**Simplifying queries**: CTEs can organize and simplify complex SQL statements and improve maintainability. For example, in scenarios in which a subquery needs to be referenced multiple times, CTEs can avoid duplicate code, making the query clearer.

**Improving code readability**: Meaningful names are used to represent intermediate results, making SQL statements easier to understand.

**Avoiding duplicate queries**: CTEs allow definition of temporary result sets, which can be referenced multiple times in a single SQL statement, avoiding repeated execution of the same operation.

**Recursive queries**: CTEs support recursive queries and can handle the query requirements of hierarchical data (such as the employee organizational structure). They are very useful for processing tree-structured data.

### Syntax Structure

The CTE syntax structure is as follows.

```
with_clause:
WITH [RECURSIVE]
cte_name [(col_name [, col_name] ...)] AS (subquery)
[, cte_name [(col_name [, col_name] ...)] AS (subquery)] ...
```

Parameter Item	Description
WITH keyword	Indicates the start of the CTE definition.
[RECURSIVE]	An optional keyword. If RECURSIVE is included, queries can reference themselves in CTEs for creating recursive queries.
cte_name	The name specified for a CTE, which can be referenced in subsequent queries.
[(col_name [, col_name] )]	An optional column name list, which is used to specify column names for the result sets of a CTE. If it is omitted, the column names in the subquery are used.
AS (subquery)	A subquery in a CTE, which defines the CTE content.
Comma and additional CTEs	In a WITH clause, multiple CTEs can be defined, which are separated by commas. Each additional CTE follows the same structure: cte_name [(col_name)] AS (subquery).

### **Non-Recursive CTEs**

A non-recursive CTE only references other tables or previously defined CTEs and will not reference itself. It is suitable for decomposing multi-step queries and gradually constructing the final query results through intermediate-level calculations.

### **Recursive CTEs**

A recursive CTE may reference itself. Recursive CTEs are commonly used for data queries involving tree or graph structures, such as calculating factorials, generating sequences, or traversing hierarchical relationships. A recursive CTE consists of three parts: Seed Part Subquery, Union Type, and Recursive Part Subquery. The Seed Part Subquery does not reference itself, and the Recursive Part Subquery will reference itself.

```
WITH RECURSIVE cte(n, fact) AS (
    SELECT 0, 1 -- Seed Part Subquery
    UNION ALL -- Union Type
    SELECT n + 1, (n + 1) * fact FROM cte WHERE n < 5 -- Recursive Part
Subquery
)
SELECT n, fact FROM cte;</pre>
```

### Example

#### **Calculating factorials**

```
WITH RECURSIVE cte(n, fact) AS (
    SELECT 0, 1
    UNION ALL
    SELECT n + 1, (n + 1) * fact FROM cte WHERE n < 5
)
SELECT n, fact FROM cte;</pre>
```

In this example, the recursive part UNION ALL SELECT n + 1, (n + 1) \* fact FROM cte WHERE n < 5 calls itself repeatedly until n reaches 5. When the recursive part outputs an empty row, the recursion ends.

#### Traversing a tree structure

For example, we have a table named employees, which contains the hierarchical relationship of employees. id is the unique identifier of an employee, name is the employee's name, and manager\_id is the ID of the superior employee of this employee.

```
CREATE TABLE employees (
```

```
id INT PRIMARY KEY,
   name VARCHAR(100),
   manager_id INT
);
INSERT INTO employees (id, name, manager_id) VALUES
(1, 'CEO', NULL),
(2, 'Manager 1', 1),
(3, 'Manager 2', 1),
(4, 'Employee 1', 2),
(5, 'Employee 2', 2),
(6, 'Employee 3', 3);
Recursive CTEs are used to traverse the employee hierarchy and obtain all
subordinates from top to bottom.
WITH RECURSIVE employee_hierarchy AS (
   -- Basic situation: Start from the CEO.
   SELECT
       id,
       name,
       manager_id,
        1 AS level
   FROM employees
   WHERE manager_id IS NULL
   UNION ALL
    -- Recursive situation: Find the subordinates of each employee.
    SELECT
       e.id,
        e.name,
        e.manager_id,
        eh.level + 1
   FROM employees e
   INNER JOIN employee_hierarchy eh ON eh.id = e.manager_id
)
SELECT id, name, manager_id, level
FROM employee_hierarchy
ORDER BY level, manager_id;
-- Result
                     manager_id level
  id
           name
| int32 | varchar
                        int32
                                int32
     1 | CEO
                                        1
     2 Manager 1
                                        2
                                1
```

3	Manager 2	1	2
4	Employee 1	2	3
5	Employee 2	2	3
6	Employee 3	3	3
L	1		

### **Basic CTE**

```
WITH CustomerCTE AS (
    SELECT customer_id, first_name, last_name, email_address
    FROM customer
)
SELECT /*+ MERGE() */ *
FROM CustomerCTE;
```

### **Multiple CTEs**

```
WITH
CTE1 AS (
    SELECT customer_id, first_name, last_name, email_address
    FROM customer
),
CTE2 AS (
    SELECT ss_item_sk, ss_customer_sk, ss_sold_date_sk, ss_sales_price
    FROM store_sales
)
SELECT /*+ MERGE() */ CTE1.first_name, CTE1.last_name, CTE2.ss_sales_price
FROM CTE1
JOIN CTE2 ON CTE1.customer_id = CTE2.ss_customer_sk;
```

Two CTEs: CTE1 and CTE2 are defined.

In the final query, the result sets of the two CTEs are joined.

Execution Result:

+		+-		+-	+
	first_name		last_name		<pre>ss_sales_price  </pre>
T		Τ-		τ-	
	John		Doe		45.99
	Jane		Smith		32.50
	Michael		Johnson		78.25
	Emily		Brown		19.99
	David		Wilson		55.00
	John		Doe		67.75
	Jane		Smith		22.99
	Michael		Johnson		41.50



	:	Emily			Brown			8	39.9	9	
	]	David			Wilsor	ר ר			33.2	5	
-	+			+-		+	+				+
	10	rows	in	set	(0.12	sec)					

### **Nested CTEs**

```
WITH SalesSummary AS (
    SELECT ss_customer_sk, SUM(ss_net_paid) AS total_spent
   FROM store_sales
    GROUP BY ss_customer_sk
),
TopCustomers AS (
    SELECT ss_customer_sk, total_spent
   FROM SalesSummary
   WHERE total_consumption > 1000 -- Assume that a threshold is set, for
example, for customers whose consumption exceeds 1000.
),
CustomerDetails AS (
    SELECT c.customer_id, c.first_name, c.last_name, tc.total_spent
    FROM customer c
    JOIN TopCustomers tc ON c.customer_id = tc.ss_customer_sk
)
SELECT /*+ MERGE() */ *
FROM CustomerDetails;
```

SalesSummary calculates the total consumption of each customer.

TopCustomers filters out customers who have consumed more than 1000 from the SalesSummary result set. CustomerDetails connects the customer information in the customer table with the result set in TopCustomers. The final SELECT query extracts all data from CustomerDetails. Execution Result:

```
| customer_id | first_name | last_name | total_spent
              ____+
    _____
      1001 | John
                  | Doe
                           1523.75 |
1002 | Jane
                  | Smith
                                2105.50 |
1003 | Michael
                  | Johnson |
                                1789.99 |
1004 | Emily
                  Brown
                           1650.25 |
1005 | David
                  | Wilson
                           1875.00 |
                                2250.75 |
1006 | Sarah
                  | Davis
                           1007 | Robert
                  | Taylor
                           1955.50 |
      1008 | Jennifer | Anderson |
                                1725.25 |
1009 | William
                  | Thomas
                           2015.00 |
      1010 | Lisa
                  | Jackson |
                                1890.75 |
```



+----+

10 rows in set (0.15 sec)

# Window Function Use Instructions

Last updated : 2025-05-09 11:51:23

Window functions, also known as Online Analytical Processing (OLAP) functions, can process and analyze data in real-time and are very important for an analytical database.

The usage method of window functions in the read-only analysis engine is basically the same as that in MySQL 8.0. However, the read-only analysis engine enabled in MySQL 5.7 can still use window functions and the window syntax. The specific window syntax is as follows:

SQL example:

select studentid,departmentid,classid,math,
row\_number() over(partition by departmentid,classid order by math) as row\_num
from student\_scores;

This example sorts the data within the groups of departmentid and classid by math.

### WindowFrame

Frame	ROWS Scenario	RANGE Scenario
CURRENT ROW	Current row.	All rows identical to the current row.
UNBOUNDED PRECEDING	Go to the first row.	Go to the first row.
UNBOUNDED FOLLOWING	Go to the last row.	Go to the last row.
<n> PRECEDING</n>	First N rows.	Go to the row whose value is greater than or equal to the value obtained by subtracting <n> from the OrderBy column value.</n>
<n> FOLLOWING</n>	Last N rows.	Go to the row whose value is less than or equal to the value obtained by adding <n> to the OrderBy column value.</n>

ROWS | RANGE < Frame>.



ROWS | RANGE BETWEEN <Frame> AND <Frame>.

### WindowFunction

The following table describes the window functions supported in the window syntax.

Function Name	Feature Description	Function Parameter	Supported Type	
ROW_NUMBER()	Marks row numbers for data in each partition.	-	-	
RANK()	Performs non-intensive data sorting in each partition.	-	-	
DENSE_RANK()	Performs intensive data sorting in each partition.	-	-	
		[Required] <expr>: Calculated column.</expr>	All types (three parameters except the Time type).	
LEAD( <expr>, <offset>, <default>)</default></offset></expr>	Calculates the values of <offset> rows after the current row. If no rows</offset>	[Optional] <offset>: Number of rows to offset after the current row. The default value is 1.</offset>	Value type	
	<pre>default&gt; is returned.</pre>	[Optional] <default>: Default return value when the calculated row does not meet the condition. The default value is NULL.</default>	Consistent with the <expr> type.</expr>	
LAG( <expr>, <offset>, <default>)</default></offset></expr>	Calculates the values of <offset> rows before the current row. If no rows</offset>	[Required] <expr>: Calculated column.</expr>	All types (three parameters except the Time type).	
	<default> is returned.</default>	[Optional] <offset>: Number of rows to offset before the current row. The default value is 1.</offset>	Value type	
		[Optional] <default>: Default return value when the calculated row</default>	Consistent with the <expr> type.</expr>	



		does not meet the condition. The default value is NULL.	
FIRST_VALUE( <expr>)</expr>	Calculates the first value in a window in a partition.	[Required] <expr>: Calculated column.</expr>	All types.
LAST_VALUE( <expr>)</expr>	Calculates the last value in a window in a partition.	[Required] <expr>: Calculated column.</expr>	All types.
MIN( <expr>)</expr>	Calculates the <expr> value of the row corresponding to the minimum value in the OrderBy column in a window in a partition.</expr>	[Required] <expr>: Calculated column.</expr>	All types.
MAX( <expr>)</expr>	Calculates the <expr> value of the row corresponding to the maximum value in the OrderBy column in a window in a partition.</expr>	[Required] <expr>: Calculated column.</expr>	All types.
COUNT( <expr>)</expr>	Counts the total number of rows of data in a window in a partition.	[Required] <expr>: Calculated column.</expr>	All types.
SUM( <expr>)</expr>	Calculates the total data volume in a window in a partition.	[Required] <expr>: Calculated column.</expr>	Value type
AVG( <expr>)</expr>	Calculates the average value of data in a window in a partition.	[Required] <expr>: Calculated column.</expr>	Value type

Numerical types: int, bigint, float, double, and decimal.

Character types: char and varchar.

Time types: date, time, datetime, and timestamp.

### **Detailed Examples**



Example table creation statement:

```
drop table if exists test.tb_window;
create table test.tb_window (c1 int not null primary key, c2 int, c3 int);
create table test.tb_window (c1 Int32, c2 Nullable(Int32), c3 Nullable(Int32))
engine = LibraTree order by (c1);
insert into test.tb_window values (1, 1, 1), (2, 1, 1), (3, 1, 2), (4, 1, 4),
(5, 1, 6), (6, 1, 6);
```

#### **ROWS Keyword**

#### Note:

This keyword measures the window size by row and calculates the data in the window.

```
-- Sample statement
mysql> select c2, c3, COUNT(c1) over (partition by c2 order by c3 ROWS BETWEEN
CURRENT ROW AND 2 FOLLOWING) cn from test.tb_window;
+----+
| c2 | c3 | cn |
+----+
| 1 | 1 | 3 | -- Window row index range: current -> next 2 rows [0-2]
| 1 | 1 | 3 | -- Window row index range: current -> next 2 rows [1-3]
  1 | 2 | 3 | -- Window row index range: current -> next 2 rows [2-4]
1 | 4 | 3 | -- Window row index range: current -> next 2 rows [3-5]
| 1 | 6 | 2 | -- Window row index range: current -> next 1 row [4-5] (Only
one row behind)
  1 | 6 | 1 | -- Window row index range: current (no data behind)
+----+
6 rows in set (0.06 sec)
```

#### **RANGE Keyword**

#### Note:

This keyword measures the window size by value and calculates the data in the window. In the following SQL example, the current row to the row corresponding to the current row value in the C3 column plus 2 is a window.

```
-- Sample statement
mysql> select c2, c3, COUNT(c1) over (partition by c2 order by c3 RANGE BETWEEN
CURRENT ROW AND 2 FOLLOWING) cn from test.tb_window;
+----+
| c2 | c3 | cn |
+----+
| 1 | 1 | 3 | -- Window row index range: current -> row index corresponding
to 3 [0-2]
| 1 | 1 | 3 |
```

```
| 1 | 2 | 2 | -- Window row index range: current -> row index corresponding
to 4 [2-3]
| 1 | 4 | 3 | -- Window row index range: current -> row index corresponding
to 6 [3-5]
| 1 | 6 | 2 | -- Window row index range: current -> row index corresponding
to 8 [4-5]
| 1 | 6 | 2 |
+----+
6 rows in set (0.06 sec)
```

#### **ROW\_NUMBER**

#### Note:

This function numbers the data within the partition and is not restricted by <WindowFrame>.

```
-- Sample statement
mysql> select c2, c3, ROW_NUMBER() over (partition by c2 order by c3) rn from
test.tb_window;
+----+
| c2 | c3 | rn |
+----+
| 1 | 1 |
          2 |
| 1 | 2 |
          3 |
| 1 | 6 |
          5 |
| 1 | 6 |
          6 |
+----+
6 rows in set (0.04 sec)
```

#### **RANK & DENSE\_RANK**

RANK function: Performs non-intensive data sorting in a partition. This function is not restricted by <WindowFrame>. DENSE\_RANK function: Performs intensive data sorting in a partition. This function is not restricted by <WindowFrame>.

	1		1		1		1	
	1		2		3		2	
	1		4		4		3	
	1		6		5		4	
	1		6		5		4	
+-		-+		-+		+		- +
6	rows	in	set	(0.0	)5 s	ec)		

### LEAD & LAG

#### 1. One-Parameter Scenarios

LEAD(<expr>) function: Calculates the data of the row after the current row in the partition. If no next row is available, NULL is filled by default. This function is not restricted by <WindowFrame>.

LAG(<expr>) function: Calculates the data of the row before the current row in the partition. If no previous row is available, NULL is filled by default. This function is not restricted by <WindowFrame>.

```
-- Sample statement
mysql> select
  c2, c3,
  LEAD(c3) over (partition by c2 order by c3) ld,
  LAG(c3) over (partition by c2 order by c3) lg
from test.tb_window;
+----+
| c2 | c3 | ld | lg |
 ____+
   1 | 1 | 1 | NULL |
        1 |
             2 | 1 |
   1 |
1 |
             4 |
2 |
                  1 |
   1 |
        4 |
             6 |
                   2 |
4 |
        6 |
             6 |
1 |
   1 |
       6 | NULL |
                  6 |
6 rows in set (0.11 sec)
```

#### 1. Two-Parameter Scenarios

LEAD(<expr>, <offset>) function: Calculates the data of <offset> rows after the current row in the partition. If no next <offset> rows are available, NULL is filled by default. This function is not restricted by <WindowFrame>. LAG(<expr>, <offset>) function: Calculates the data of <offset> rows before the current row in the partition. If no previous <offset> rows are available, NULL is filled by default. This function is not restricted by <WindowFrame>.

```
-- Sample statement
mysql> select
    c2, c3,
```

	LE	AD (	c3,	2)	over	(]	parti	tid	on by	y cź	2 or	der	by	z c3)	) ld
	LA	.G (c	3, 2	)	over	(p	artit	ioı	n by	c2	ord	er	by	c3)	lg
fı	com t	est	.tb_	wi	ndow;										
+-		-+-		-+		-+-		-+							
	c2		с3		ld		lg								
+-		-+-		-+-		-+-		-+							
	1		1		2		NULL								
	1		1		4		NULL								
	1		2		6		1								
	1		4		6		1								
	1		6		NULL		2								
	1		6		NULL		4								
+-		-+-		-+-		-+-		-+							
6	rows	in	set	(	0.07 s	se	C)								

#### 1. Three-Parameter Scenarios

LEAD(<expr>, <offset>, <default>>) function: Calculates the data of <offset> rows after the current row in the partition. If no next <offset> rows are available, <default> is filled by default. This function is not restricted by <WindowFrame>. LAG(<expr>, <offset>, <default>>) function: Calculates the data of <offset> rows before the current row in the partition. If no previous <offset> rows are available, <default> is filled by default. This function is not restricted by <WindowFrame>.

```
-- Sample statement
mysql> select
  c2, c3,
  LEAD(c3, 2, 1000) over (partition by c2 order by c3) ld,
  LAG(c3, 2, 1000) over (partition by c2 order by c3) lg
from test.tb_window;
+----+
| c2 | c3 | ld | lg |
+----+
        1 | 2 | 1000 |
   1 |
4 | 1000 |
        1 |
   1 |
2 |
             6 | 1 |
   1 |
1 |
        4 | 6 |
                  1 |
1 |
        6 | 1000 |
                   2 |
1
   1 |
        6 | 1000 |
                  4
+----+
6 rows in set (0.10 sec)
```

### FIRST\_VALUE & LAST\_VALUE

FIRST\_VALUE(<expr>) function: Calculates the first value in a window in a partition. (If OrderBy c3 is adopted and the first value in the c3 column is duplicated with another value, first\_value(c4) may be unstable.)

LAST\_VALUE(<expr>) function: Calculates the last value in a window in a partition. (If OrderBy c3 is adopted and the last value in the c3 column is duplicated with another value, last\_value(c4) may be unstable.)

```
-- Sample statement
mysql> select
  c2, c3,
   FIRST_VALUE(c3) over (partition by c2 order by c3 ROWS BETWEEN UNBOUNDED
PRECEDING AND UNBOUNDED FOLLOWING) fv,
   LAST_VALUE(c3) over (partition by c2 order by c3 ROWS BETWEEN UNBOUNDED
PRECEDING AND UNBOUNDED FOLLOWING) 1v
from test.tb_window;
+----+
| c2 | c3 | fv | lv |
+----+
   1 | 1 | 1 |
                   6 |
   1 |
         1 |
              1 |
6 |
   1 |
        2 |
              1 |
6 |
              1 |
   1 |
         4 |
                    6 |
1 |
        6 |
              1 |
                    6 |
        6 |
              1 |
                    6 |
1 |
+----+
6 rows in set (0.07 sec)
```

### **MIN & MAX**

MIN(<expr>) function: Calculates the minimum value in a window in a partition.

MAX(<expr>) function: Calculates the maximum value in a window in a partition.

```
-- Sample statement
mysql> select
   c2, c3,
   MIN(c3) over (partition by c2 order by c3 ROWS BETWEEN UNBOUNDED PRECEDING
AND UNBOUNDED FOLLOWING) mi,
   MAX(c3) over (partition by c2 order by c3 ROWS BETWEEN UNBOUNDED PRECEDING
AND UNBOUNDED FOLLOWING) ma
from test.tb_window;
+----+
| c2 | c3 | mi | ma |
+----+
         1 |
              1 |
   1 |
                    6 |
1 |
         1 |
              1 |
                    6 |
2 |
              1 |
1 |
                    6 |
        4 |
              1 |
   1 |
6 |
         6 |
              1 |
1 |
                    6 |
1 |
         6 |
              1 |
                     6 |
          -+-
```

6 rows in set (0.07 sec)

### COUNT

#### Note:

It calculates the total number of rows of data in a window in a partition.

```
-- Sample statement
mysql> select c2, c3, COUNT(c3) over (partition by c2 order by c3 ROWS BETWEEN
UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING) c from test.tb_window;
+----+
| c2 | c3 | c |
 ----+
   1 | 1 | 6 |
1 |
        1 | 6 |
1 |
        2 | 6 |
1 |
        4 | 6 |
   1 |
        6 6 6
1 | 6 | 6 |
+----+
6 rows in set (0.04 sec)
```

#### SUM

#### Note:

It calculates the total data volume in a window in a partition.

```
-- Sample statement
mysql> select c2, c3, SUM(c3) over (partition by c2 order by c3 ROWS BETWEEN
UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING) s from test.tb_window;
```

	с2	(	23		S	I
+-		-+		-+-		+
	1		1		20	) (
	1		1		20	) (
	1		2		20	) (
	1		4		20	) (
	1		6		20	) (
	1		6		20	) (
+-		-+		-+-		+
6	rows	in	set	((	).06	sec)

+---+

### AVG

Note:

It calculates the average value of data in a window in a partition.

```
-- Sample statement
mysql> select c2, c3, AVG(c3) over (partition by c2 order by c3 ROWS BETWEEN
UNBOUNDED PRECEDING AND UNBOUNDED FOLLOWING) a from test.tb_window;
+----+
| c2 | c3 | a
                 +----+
   1 | 1 | 3.3333 |
1 |
        1 | 3.3333 |
1 |
        2 | 3.3333 |
1 | 4 | 3.3333 |
6 | 3.3333 |
1 |
   1 | 6 | 3.3333 |
+----+
6 rows in set (0.06 sec)
```
# Executing SQL and Viewing the Execution Plan

Last updated : 2025-05-09 11:51:23

The read-only analysis engine is compatible with the MySQL protocol and syntax. Users can access the read-only analysis engine through the MySQL client or the MySQL connector in the development language to execute SQL statements.

The read-only analysis engine is a read-only instance, so the SQL statements that can be executed are mainly SELECT statements, and the DML and DDL statements cannot be executed. At the same time, queries are strictly subject to object permissions. If the access account does not have the SELECT permission on an object, it cannot view this object either.

When you are executing SQL, you can directly use the MySQL syntax for syntax editing and then execute it in the read-only analysis engine. However, sometimes the execution results may not meet expectations. In this case, you can view the execution plan through the EXPLAIN feature.

## **EXPLAIN** Feature and Usage

Use EXPLAIN to view the execution plan selected when executing query statements in the read-only analysis engine. This is the optimal query plan that is finally selected after multiple stages of optimization by the internal optimizer. When you view the execution plan, SQL will not be executed actually; instead, only the execution plan is output. The EXPLAIN example is as follows:

```
EXPLAIN <SELECT_STATMENT>
explain select * from t1 left join t2 on t1.id = t2.id;
```

The returned result is as follows:



## Interpretation of EXPLAIN Results

\_\_\_\_\_

The returned result of EXPLAIN contains the following fields:

ID is the operator number displayed in the plan, which increments downward from 0, facilitating users to view details. OPERATOR is the operator name, representing the operation performed at each step in the SQL execution.

NAME shows the information such as tables and partitions accessed in the query.

EST.ROWS shows the number of rows that each operator is expected to process, as estimated by the read-only analysis engine based on statistics. If there are no statistics, it will be calculated and displayed based on a default value.

### **Operator Introduction**

An operator is a specific step that is performed to return query results. The following table introduces the names and functions of various operators currently supported by the read-only analysis engine:

Operator Name	Operator Introduction
SORT	The SORT operator is used to sort the input data.
TOPN	If the ORDER BY statement block is followed by a LIMIT statement, the optimizer will further optimize the execution plan, generate a TOP-N SORT operator, and use heapsort

	to select the TOP-N data.
LIMIT	The LIMIT operator is used to limit the number of rows of data output, which has the same function as the LIMIT operator in MySQL.
FILTER	The FILTER operator is used to perform filtering operations on data according to specified predicate conditions, usually appearing in WHERE/HAVING/ON clauses.
HASH JOIN	Use HASH JOIN to perform JOIN operations on large datasets. The optimizer uses tables in two datasets. The smaller dataset is used to build a HASH table in memory according to the JOIN conditions, and then the database scans the larger dataset and probes the HASH table to find the rows that meet the join conditions.
COLUMN READ	A late materialization operator. The read-only analysis engine supports pushing some filtering conditions down to the TableScan operator. Specifically, it scans the column data related to the filtering conditions first, filters the data to obtain the rows that meet the conditions, and then scans the remaining column data of these rows to continue with subsequent calculations, reducing the computing volume of I/O scans and data processing.
TABLE FULL SCAN	Performs a full-table scan on the target table.
UNION	The UNION operator is used to perform a union operation on the result sets of two queries.
WINDOW FUNCTION	The WINDOW FUNCTION operator is used to implement the analysis functions (also known as window functions) in SQL, calculating the results of relevant rows within the window. Each group of window functions can return multiple rows, and each row within the group is the result of logic calculation based on the window.
HASH GROUP BY	The GROUP BY operator is mainly used for group aggregation calculations in SQL. Ordinary aggregation functions (SUM/MAX/MIN/AVG/COUNT/STDDEV) are completed by allocating the GROUP BY operator.
PROJECTION	A projection operator, corresponding to the SELECT list in SQL statements, has the capability of mapping each input data to new output data.
EXCHANGE RECEIVER	A data receiving operator, used to receive data when various compute nodes exchange data during the MPP query.
EXCHANGE SENDER	A data sending operator, used to send data when various compute nodes exchange data during the MPP query.

### **Detail Info**

In addition to the information shown in the above table of returned results, you can see that a Detail item is displayed below. The Detail item shows some additional information carried by each operator, in the format of n - detail info, where n represents the ID of each operator and what follows is the specific information of each operator. The following table lists the specific information that each operator displays.

Operator Name	Detail Info
SORT	ORDER BY KEY: Represents the sort key used by the sort operator.
TOPN	ORDER BY KEY: Represents the sort key used by the sort operator. OFFSET: The offset specified by the Limit operation. COUNT: The number of preserved rows specified by the Limit operation.
LIMIT	OFFSET: The offset specified by the Limit operation. COUNT: The number of preserved rows specified by the Limit operation.
FILTER	CONDITIONS: Represents the predicate condition adopted by the FILTER operator when filtering data.
HASH JOIN	JOIN TYPE: The current JOIN type, such as inner join, left outer join, and semi join. NON EQUAL: Specifies whether the current JOIN is a Cartesian product. If this field does not exist, it indicates it is not a Cartesian product. EQUAL: Displays the equivalent condition adopted when two tables are joined. OTHER COND: Displays the non-equivalent condition when two tables are joined. In addition, it can be seen that in the EXPLAIN results above, the operator names with IDs 1 and 2 are followed by the marks (B) and (P) respectively, which mark the Build and Probe sides during the HASH JOIN operation, where B represents Build and P represents Probe.
COLUMN READ	COLUMN READ: The column name read by the late materialization.
TABLE FULL SCAN	STORAGE: The underlying storage type being read. Currently, two types are supported: LIBRASTORE and TDSQL. BLOCK OFFSET: The query block number of the current table in the entire SQL statement, used to assist with Hint.
UNION	No.
WINDOW FUNCTION	WINDOW FUCN DESCS: The window function name. PARTITION BY: The partition key. ORDER BY: The sorting key used for sorting. FRAME: The window definition of a window function.
HASH GROUP BY	GROUP BY: The group by key specified when executing aggregate functions. AGG FUNCS: The aggregate functions specified in SQL.
PROJECTION	EXPRS: The list of expressions executed by the PROJECTION operator, mainly including



	casts being performed and various scalarfuncs.
EXCHANGE RECEIVER	No.
EXCHANGE SENDER	<ul> <li>ExchangeType: Methods for data exchange, including:</li> <li>PASS: Sends data from multiple nodes to one node.</li> <li>BCJ: Broadcasts data from one node to multiple nodes. For example, broadcast the data from the build table to nodes during JOIN to perform the JOIN operation.</li> <li>HASH: Distributes data to nodes after it is hashed with the HASH function. For example, scatter and redistribute data from left and right tables during the JOIN operation.</li> <li>HASH (BY PARTITION): When two tables are involved in a JOIN operation, if one table has a join key that is also its partition key, the other table is shuffled according to the distribution method of the table where the join key is the partition key.</li> </ul>

In addition to the basic information of operators in the above table, the two sides of the HASH JOIN operator may have the following two kinds of special detail information when the runtime filter operator is enabled and there is a HASH JOIN operation in the plan.

Probe Runtime Filters: When this information exists on an operator, it means that the current operator applies a runtime filter to filter redundant data.

Build Runtime Filters: This information only appears on the JOIN operator, indicating that a runtime filter has been generated on the build side of the current JOIN operator to filter the redundant data on the probe side in advance. For details on the use and optimization of the runtime filter, and specific effects of the displayed information, see Runtime Filter User Manual.

After EXPLAIN is executed, different operators return different information. You can use Optimizer Hints to control the behavior of the optimizer and thus control the selection of physical operators.

For example, /\*+ HASH\_JOIN\_PROBE(t1) \*/ indicates that the optimizer forces the use of t1 table as the probe table in the hash join.

# Exclusive Features ETL Writeback Acceleration

Last updated : 2025-05-09 11:51:23

TencentDB for MySQL can use the read-only analysis engine to accelerate data queries and write the results back to the read-write instance. By using this feature, you can use columnar storage in the read-only analysis engine to accelerate the SELECT query in an INSERT...SELECT... statement, transmit the data back to the read-write instance through the internal network, and write the data into the target table. For a detailed feature description, see ETL Writeback Acceleration.

## **Application Scenario**

#### Note:

The read-only analysis engine adopts the asynchronous replication mode. Therefore, a certain time delay may exist between the query results and those in the read-write instance when there is a latency in the read-only analysis engine. This feature applies only to scenarios that are not sensitive to data latency.

Currently, only INSERT...SELECT... statements can be accelerated. UPDATE...SELECT...,

DELETE...SELECT... , and other statements cannot be accelerated.

This feature is recommended in scenarios where the query conditions are complex, the SQL statement execution time is long, and the data volume of the query result set is small. In such scenarios, this feature can significantly improve performance because the read-only analysis engine will accelerate SELECT queries.

This feature does not bring performance benefits in all scenarios. Instead, the performance may deteriorate in some scenarios. For example:

When the SELECT query in an INSERT...SELECT... statement is relatively simple, reading data from the read-only analysis engine and writing the data back to the read-write instance will incur additional network overhead. The strengths are not obvious compared with the method of directly reading data from the read-write instance. When the result set of the SELECT query in an INSERT...SELECT... statement contains a large volume of data, the main performance bottleneck lies in the process of transmitting the result set over the network and writing it to the read-write instance. In this scenario, this feature cannot be used to improve performance.

### Prerequisites

Both the read-only analysis engine instance and the read-write instance are running normally. The user should have the INSERT permission for the corresponding objects.



Only the execution of INSERT...SELECT... can be accelerated.

## Feature Parameter Description

Parameter Name	Description	Default Value	Parameter Value
libra_etl_to_rw	Controls whether to enable the writeback feature.	off	on: Enable the writeback feature. It means that writeback will be performed. off: Disable the writeback feature. When the value is off, an error of insufficient permission will be reported during INSERTSELECT execution.
libra_concurrent_etl	Controls whether to enable concurrent writeback to the read-write instance.	off	on: Enable concurrent writeback. Concurrent writeback can improve the efficiency of result writeback, but transaction consistency among multiple threads cannot be guaranteed. off: Disable concurrent writeback. When the value is off, data writeback to the read-write instance will be completed in one transaction.
libra_etl_concurrency	Controls the number of concurrent threads for data writeback to the read-write instance. This parameter takes effect only when libra_concurrent_etl is set to on.	1	Value range: [1–number of read-write instance CPU cores]. The value should be an integer, and the maximum value is the number of CPU cores of the read- write instance. <b>Note:</b> The writeback speed is high if multiple concurrent threads are used. However, a large number of concurrent threads can damage the performance of the read-write instance because it cannot withstand the write load.

## Example

1. Use a client tool to log in to the read-only analysis engine. For example, you can use the MySQL client to log in to the read-only analysis engine with the database account and password.

2. Execute the following SQL statements to use the feature of ETL writeback acceleration (writing data back to the

read-write instance from the read-only analysis engine).

Set parameters in the session:

```
set libra_etl_to_rw=on;
set libra_concurrent_etl=on;
set libra_etl_concurrency=4;
INSERT INTO database1.table1 select a.t1,b.t2,a.t2,b.t3 from a,b where
a.t1=b.t1 and a.t5='x';
set libra_concurrent_etl=off;
set libra_etl_to_rw=off;
```

Add a hint with parameters to the SQL statement:

```
INSERT /*+ SET_VAR(libra_etl_to_rw=ON) SET_VAR(libra_concurrent_etl=ON)
SET_VAR(libra_etl_concurrency=4)*/ INTO database1.table1 select
a.t1,b.t2,a.t2,b.t3 from a,b where a.t1=b.t1 and a.t5='x';
```

#### Note:

If you add a hint with parameters in the MySQL client, you need to add the -c parameter when logging in to the readonly analysis engine in the MySQL client. Otherwise, the hint will not take effect.

## Performance Comparison

#### **Test Environment**

Read-write instance: 16-core 64 GB. Read-only analysis engine: 16-core 64 GB.

#### Test Data

Use the TPC-H dataset with 10 GB of data.

#### **Complex SQL Statement Test and Result**

```
set libra_etl_to_rw=on;
set libra_concurrent_etl=off;
set libra_etl_concurrency=1;
create table t1 (
    supp_nation VARCHAR(50),
    cust_nation VARCHAR(50),
    l_year INT,
    revenue DECIMAL
);
```



```
INSERT INTO t1 SELECT
supp_nation, cust_nation, l_year, sum(volume) as revenue
from (select n1.n_name as supp_nation, n2.n_name as cust_nation,
             extract(year from l_shipdate) as l_year,
             l_extendedprice * (1 - l_discount) as volume
      from supplier, lineitem, orders, customer, nation n1, nation n2
      where s_suppkey = l_suppkey
        and o_orderkey = l_orderkey
       and c custkey = o custkey
        and s_nationkey = n1.n_nationkey
        and c_nationkey = n2.n_nationkey
        and ((n1.n_name = 'JAPAN' and n2.n_name = 'INDIA')
              or (n1.n_name = 'INDIA' and n2.n_name = 'JAPAN'))
        and l_shipdate between '1995-01-01' and '1996-12-31') as shipping
group by supp_nation, cust_nation, l_year
order by supp_nation, cust_nation, l_year;
```

When the query result contains 4 rows of data, the test results are as follows (unit: seconds).

Duration of		
SELECT	Writeback Feature Disabled	Writeback Feature Enabled
Execution on	Duration of INSERTSELECT	Duration of INSERTSELECT
Read-Only	Execution on Read-Write Instance	Execution on Read-Only Analysis Engine
Analysis Engine		
0.57	232.77	0.61

```
set libra_etl_to_rw=on;
set libra_concurrent_etl=off;
set libra_etl_concurrency=1;
CREATE TABLE t2
 (
   p_brand VARCHAR(10),
   p_type VARCHAR(25),
   p_size INTEGER,
    supplier_cnt INTEGER
);
INSERT INTO t2 SELECT
   p_brand,
   p_type,
   p_size,
    count(distinct ps_suppkey) as supplier_cnt
from
```

```
partsupp,
   part
where
   p_partkey = ps_partkey
   and p_brand <> 'Brand#45'
   and p_type not like 'MEDIUM POLISHED%'
    and p_size in (49, 14, 23, 45, 19, 3, 36, 9)
    and ps_suppkey not in (
        select
            s_suppkey
        from
            supplier
        where
            s_comment like '%Customer%Complaints%'
    )
group by
   p_brand,
   p_type,
   p_size
order by
   supplier_cnt desc,
   p_brand,
   p_type,
   p_size;
```

When the query result contains 27,840 rows of data, the test results are as follows (unit: seconds).

Duration of SELECT Execution on Read-Only Analysis Engine	Writeback Feature Disabled Duration of INSERTSELECT Execution on Read-Write Instance	Writeback Feature Enabled Duration of INSERTSELECT Execution on Read-Only Analysis Engine
0.15	8.77	0.61

#### Simple SQL Statement Test and Result

```
set libra_etl_to_rw=on;
set libra_concurrent_etl=on;
set libra_etl_concurrency=8;
CREATE TABLE IF NOT EXISTS lineitem ( L_ORDERKEY INTEGER NOT NULL,
L_PARTKEY INTEGER NOT NULL,
L_SUPPKEY INTEGER NOT NULL,
L_LINENUMBER INTEGER NOT NULL,
L_QUANTITY DECIMAL(15,2) NOT NULL,
L_EXTENDEDPRICE DECIMAL(15,2) NOT NULL,
```

```
L_DISCOUNT DECIMAL(15,2) NOT NULL,
L_TAX DECIMAL(15,2) NOT NULL,
L_RETURNFLAG CHAR(1) NOT NULL,
L_LINESTATUS CHAR(1) NOT NULL,
L_SHIPDATE DATE NOT NULL,
L_COMMITDATE DATE NOT NULL,
L_RECEIPTDATE DATE NOT NULL,
L_SHIPINSTRUCT CHAR(25) NOT NULL,
L_SHIPMODE CHAR(10) NOT NULL, primary key(L_ORDERKEY, L_LINENUMBER));
INSERT INTO lineitem_t SELECT
  *
from
    lineitem;
```

When the query result contains 59,986,051 rows of data, the test results are as follows (unit: seconds).

Number of Concurrent Threads	Writeback Feature Disabled Duration of INSERTSELECT Execution on Read-Write Instance	Writeback Feature Enabled Duration of INSERTSELECT Execution on Read-Only Analysis Engine
Concurrent writeback disabled		1441.02
8 concurrent threads	622.7	259.20
16 concurrent threads		181.79
32 concurrent threads		185.64

# **Ordered Pagination Feature**

Last updated : 2025-05-09 11:51:23

The ordered pagination feature of the read-only analysis engine can solve the problem of unstable pagination query results caused by duplicate field values of the order by clause during parallel data queries. This document describes how to use the ordered pagination feature. For a detailed feature introduction, see Paging Order Preservation Capability.

## **Ordered Pagination Scenarios**

Only the limit operator is used, and no order by operator is used. Duplicate order by keys exist, and only some fields in the output result are sorted. The subquery involves sorting, but sorting is not performed in the outer query.

## **Usage Instructions**

### Enable the ordered pagination feature globally.

Use the feature of setting instance parameters to modify the parameter libra\_preserve\_order\_for\_pagination of the read-only analysis engine. The default value of this parameter is OFF. You can set the value to ON to enable the ordered pagination feature for the entire instance.

#### Note:

The ordered pagination feature performs implicit sorting on the output results of queries by default. Therefore, it may cause performance degradation in scenarios with large result sets. It is recommended to specify to use the ordered pagination feature in SQL statements that require this feature.

### Enable the ordered pagination feature in a session.

Specify session-level parameters after you log in to the read-only analysis engine to enable/disable the ordered pagination feature in the current session.

```
mysql> set libra_preserve_order_for_pagination=on;
mysql> set libra_preserve_order_for_pagination=off;
```

### Specify to use the ordered pagination feature during SQL execution.

When a SQL statement is executed, add a hint to specify this SQL statement to use the ordered pagination feature. **Note:** 

If you use the MySQL client to access the read-only analysis engine, add the -c parameter. Otherwise, the settings of using the ordered pagination feature during SQL execution cannot take effect.



mysql> select /\*+ SET\_VAR(libra\_preserve\_order\_for\_pagination=ON) \*/ t1.c1, t1.c2

# Performance Optimization Optimizing SQL Execution With HINT

Last updated : 2025-05-09 11:51:24

## What Is a Hint

Typically, the optimizer selects the optimal execution plan for a user's SQL query. However, in certain scenarios such as estimation errors in statistics or fitting bias in the cost model - the execution plan generated by the optimizer may not be the most optimal. At this time, users can guide the optimizer to generate a better execution plan through the hint mechanism.

Hints are case-insensitive and follow the SELECT keyword as comments in the form of /\*+...\*/. Multiple hints are separated by spaces or commas. Here is an example of using a hint.

```
SELECT /*+ [hint_text] [hin_text]... */ * FROM ....
```

## Effective Scope of Hints

Hints take effect on a per query block basis. In a query statement, each query block has a QB\_NAME (Query Block Name). The read-only analysis engine generates a QB\_NAME for each query block from left to right in the format of @sel\_1, @sel\_2, and so on. Take the following SQL statement as an example.

SELECT \* FROM (SELECT \* FROM t) t1, (SELECT \* FROM t) t2;

This SQL statement contains three query blocks. The name of the query block where the outermost SELECT query is located is sel\_1, and the names of the two SELECT sub-queries are sel\_2 and sel\_3, with the numbers increasing in sequence. QB\_NAME can be used in the hint to control the scope and target of the hint. If QB\_NAME is not explicitly specified in the hint, the scope of the hint is the query block where the current hint is located. The following is an example.

```
SELECT /*+ HASH_JOIN_PROBE(@sel_2 t1) */ * FROM (SELECT t1.a, t1.b FROM t t1, t
t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
SELECT /*+ HASH_JOIN_PROBE(t1@sel_2) */ * FROM (SELECT t1.a, t1.b FROM t t1, t
t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

The above two SQL statements use two different methods to specify QB\_NAME in the hint. The first SQL statement specifies QB\_NAME in the first parameter of the hint and separates it from other parameters with a space. The second SQL specifies the effective scope of the hint by adding QB\_NAME after the parameter.

## Overview of Supported Hints

The parameter name, semantics and syntax related to the hint of the read-only analysis engine are shown in the following table.

Name	Syntax	Semantics
SHUFFLE_JOIN	SHUFFLE_JOIN([QB_NAME] tbl1_name,tbl2_name)	Specifies tha JOIN operati uses the shut method to distribute dat
BROADCAST_JOIN	BROADCAST_JOIN([QB_NAME] tbl1_name,tbl2_name)	Specifies tha JOIN operati uses the broadcast me to distribute c
HASH_JOIN_BUILD	HASH_JOIN_BUILD([QB_NAME] tbl1_name,tbl2_name)	Specifies the Build table in HASH JOIN operation.
HASH_JOIN_PROBE	HASH_JOIN_PROBE([QB_NAME] tbl1_name,tbl2_name)	Specifies the Probe table in HASH JOIN operation.
LEADING	LEADING([QB_NAME] tbl1_name,tbl2_name)	Specifies the order for the operation.
SET_VAR	SET_VAR(setting_name = value)	Sets system parameters a SQL level.
NO_PX_JOIN_FILTER_ID/PX_JOIN_FILTER_ID	NO_PX_JOIN_FILTER_ID(rf_id1,rf_id2)/ PX_JOIN_FILTER_ID(rf_id1,rf_id2)	Controls the enabling/disa of the runtime filter.

## Detailed Explanation of the Read-Only Analysis Engine Hint Syntax

### SHUFFLE\_JOIN(t1\_name, t2\_name ...)

#### **Usage instructions**

SHUFFLE\_JOIN(t1\_name, t2\_name...) is used to instruct the read-only analysis engine optimizer to use the shuffle join algorithm, during JOIN operations, to scatter and redistribute the data from the left and right tables before performing JOIN operations to return results.

#### **Reference examples**

```
EXPLAIN SELECT /*+ SHUFFLE_JOIN(t1) */ * FROM t1, t2 WHERE t1.id = t2.id;
```

In this SQL statement, the t1 and t2 tables perform a JOIN operation, and the join data distribution method is specified as the shuffle join through the SHUFFLE\_JOIN hint. The final plan is shown in the figure below. It can be seen that the EXCHANGE TYPE in Rows 2 and 5 of Details has become HASH, indicating that the hash shuffle is used.

Besides a single table, you can also specify the intermediate result in a JOIN operation for data redistribution.

```
EXPLAIN SELECT /*+ SHUFFLE_JOIN((t1@sel_2,t2@sel_2)) */ * FROM (SELECT t1.a,
t1.b FROM t t1, t t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

As shown above, you can specify that the shuffle join is used to join the intermediate result of the JOIN operation between tables t1 and t3 with table t2 by enclosing tables t1 and t3 in parentheses and specifying the QB\_NAME of each table.

#### Note:

This hint is effective only when a distributed plan is generated and is not effective for a standalone plan.

### BROADCAST\_JOIN(t1\_name [, tl\_name ...])

#### **Usage instructions**

BROADCAST\_JOIN(t1\_name,t2\_name...) is used to instruct the read-only analysis engine optimizer to use the broadcast join algorithm, during JOIN operations, to broadcast the specified table data to all nodes, perform JOIN operations, and return results.

#### **Reference examples**

```
EXPLAIN SELECT /*+ BROADCAST_JOIN(t1) */ * FROM t1, t2 WHERE t1.id = t2.id;
```

In this SQL statement, the t1 and t2 tables perform a JOIN operation, and the JOIN method is specified as the broadcast join through the BROADCAST\_JOIN hint. The final plan is shown in the figure below. It can be seen that the EXCHANGE TYPE in Rows 2 and 5 of Details has become BCJ, indicating that the broadcast is used.

Besides a single table, you can also specify the intermediate result in a JOIN operation for data broadcast.

```
EXPLAIN SELECT /*+ BROADCAST_JOIN((t1@sel_2,t3@sel_1)) */ * FROM (SELECT t1.a,
t1.b FROM t t1, t t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

As shown above, you can specify that the broadcast join is used to join the intermediate result of the JOIN operation between tables t1 and t3 with table t2 by enclosing tables t1 and t3 in parentheses and specifying the QB\_NAME of each table. The result is shown in the figure below.

#### Note:

This hint is effective only when a distributed plan is generated and is not effective for a standalone plan. The read-only analysis engine optimizer will select the Build side of the hash join as the broadcast table for the broadcast. If adjustment is needed, it can be adjusted together with HASH\_JOIN\_BUILD.

### HASH\_JOIN\_BUILD(t1\_name,t2\_name...)

#### Usage instructions

HASH\_JOIN\_BUILD(t1\_name,t2\_name...) is used to instruct the read-only analysis engine optimizer to use the HASH JOIN algorithm for specified tables, and to designate the specified tables as the Build side for the HASH JOIN algorithm. Specifically, it means using the specified tables to build a hash table.

#### **Reference examples**

```
EXPLAIN SELECT /*+ HASH_JOIN_BUILD(t2)*/ * FROM t t1, t t2 WHERE t1.a = t2.a;
```

This SQL statement specifies table t2 as the Build table in the HASH JOIN operation. The final plan is shown in the figure below.

Besides a single table, you can also specify the intermediate result in a JOIN operation as the BUILD side. The following is an example.

```
EXPLAIN SELECT /*+ HASH_JOIN_BUILD((t1@sel_2,t3@sel_1)) */ * FROM (SELECT t1.a,
t1.b FROM t t1, t t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

As shown above, enclose tables t1 and t2 in parentheses and specify a QB\_NAME for each table.

### HASH\_JOIN\_PROBE(t1\_name,t2\_name...)

#### **Usage instructions**

HASH\_JOIN\_PROBE(t1\_name,t2\_name...) is used to instruct the optimizer to use the HASH JOIN algorithm for specified tables, and to designate the specified tables as the Probe side for the HASH JOIN algorithm. Specifically, it means using the specified tables as the Probe table in the HASH JOIN operation.

#### **Reference examples**

EXPLAIN SELECT /\*+ HASH\_JOIN\_PROBE(t2)\*/ \* FROM t t1, t t2 WHERE t1.a = t2.a;



This SQL statement specifies table t2 as the Probe table in the HASH JOIN operation. The final plan is shown in the figure below.

Besides specifying a single table, you can also specify the intermediate result in a JOIN as the Probe side. The following is an example.

```
EXPLAIN SELECT /*+ HASH_JOIN_PROBE((t1@sel_2,t3@sel_1)) */ * FROM (SELECT t1.a,
t1.b FROM t t1, t t2 WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

As shown above, enclose tables t1 and t2 in parentheses and specify a QB\_NAME for each table. The final plan is shown in the figure below.

### LEADING(t1\_name,t2\_name ...)

#### Usage instructions

Leading(t1\_name,t2\_name...) is used to control the join order generated by the optimizer during the join reorder phase. The optimizer will determine the join order according to the order in the Leading hint.

#### **Reference examples**

```
EXPLAIN SELECT /*+ LEADING(t1,t3,t2,t4)*/ * FROM t1,t2,t3,t4 WHERE t1.a = t2.a and t2.a = t3.a and t3.a = t4.a;
```

With the LEADING hint, this SQL statement explicitly specifies that table t1 first performs a JOIN operation with table t3, then with table t2, and finally with table t4. The final plan is shown in the figure below.

The above method can only generate a left deep tree by specifying the JOIN sequence, however, the read-only analysis engine also provides an advanced syntax for the LEADING hit, which can generate a bushy tree with the help of parentheses. An example is as follows.

```
EXPLAIN SELECT /*+ LEADING((t1,t3),(t2,t4))*/ * FROM t1,t2,t3,t4 WHERE t1.a = t2.a and t2.a = t3.a and t3.a = t4.a;
```

The LEADING hint of this SQL statement first instructs the optimizer to perform a JOIN between table t1 and table t3 by (t1,t3), then a JOIN between table t2 and table t4 by (t2,t4), and finally a JOIN between the results of the above two JOIN operations by ((t1,t3),(t2,t4)). The final execution plan is shown in the figure below.

#### Note:

The hint will be invalid when there are multiple LEADING hints.

The hint will be invalid when the optimizer cannot perform the table join according to the LEADING hint.

### SET\_VAR(NAME="VALUE")



#### Usage instructions

SET\_VAR(XXXX="YY") is used to temporarily modify system variables during the SQL execution. After the SQL execution is completed, the specified system variables will automatically revert to their original values. The usage is as follows.

#### **Reference examples**

SELECT /\*+ SET\_VAR(max\_threads=64) \*/ \* FROM t1

This SQL statement uses the SET\_VAR hint to temporarily set the maximum thread count to 64 for the duration of the SQL execution.

#### Note:

Confirm that target parameters support for modification with the hint before use, as not all parameters support the SET\_VAR hint. For parameters that can be modified with the hint, see System Variables.

### NO\_PX\_JOIN\_FILTER\_ID(ID)/PX\_JOIN\_FILTER\_ID(ID)

#### **Usage instructions**

no\_px\_join\_filter\_id(ID)/px\_join\_filter\_id(ID) is used to instruct the optimizer to disable or enable RuntimeFilter.

#### **Reference examples**

For specific usage methods, see Runtime Filter User Manual.

### **Common Hint Issues**

#### Removal of Hint by MySQL Client Causes Ineffectiveness

The MySQL command-line client earlier than version 5.7 removed Optimizer hints by default. If you need to use the hint syntax in these earlier versions of the client, you need to add the --comments option when starting the client. For example: mysql -h 127.0.0.1 -P 4000 -u root -c.

# Cross-Database Query Without a Database Name Specified Causes the Hint Not to Take Effect

For tables that require cross-database access in the query, you need to explicitly specify the database name in the hint, otherwise the hint may not take effect. For example, for the following SQL:

SELECT /\*+ SHUFFLE\_JOIN(t1) \*/ \* FROM test1.t1, test2.t2 WHERE t1.id = t2.id;

The current t1 table is in the current database, causing the hint to be invalid. The Warning message is as follows.



```
| Level | Code | Message
|
+----+
-------+
| Warning | 1815 | There are no matching table names for (t1) in optimizer hint
/*+ SHUFFLE_JOIN(t1) */ or /*+ SHUFFLE_JOIN(t1) */. Maybe you can use the table
alias name |
+-----+
1 row in set (0.00 sec)
```

### QB\_NAME Not Specified/Incorrectly Specified Causes the Hint Not to Take Effect

For queries with multiple QB\_NAMEs, if the hint is not written in the query block of the target table, it is necessary to explicitly specify the QB\_NAME in the hint. If it is not specified, the Query hint may not take effect. For example:

```
SELECT /*+ HASH_JOIN_PROBE(t2) */ * FROM (SELECT t1.a, t1.b FROM t t1, t t2
WHERE t1.a = t2.a) t1, t t3 WHERE t1.b = t3.b;
```

The QB\_NAME is not explicitly specified for the t1 table in the hint, causing the hint not to take effect. The Warning message is as follows.

```
mysql> show warnings;
+-----+
| Level | Code | Message
|
+----+
| Warning | 1815 | There are no matching table names for (t2) in optimizer hint
/*+ HASH_JOIN_PROBE(t2) */. Maybe you can use the table alias name |
+-----+
1 row in set (0.00 sec)
```

At this time, you can query the SQL plan by keywords to confirm the QB\_NAME of the query block for each table.

### **Incorrect Position of the Hint Causes Ineffectiveness**

If the hint is not correctly placed after the specified keyword according to the syntax of Optimizer hints, it will not take effect. For example:

```
MySQL> SELECT * /*+ SET_VAR(max_threads = 64)) */ FROM t;
```



ERROR 1064 (42000): You have an error in your SQL syntax; check the manual that corresponds to your MySQL server version for the right syntax to use line 1 column 42 near "/\*+ SET\_VAR(max\_threads = 64)) \*/ FROM t"

# Runtime Filter User Manual

Last updated : 2025-05-09 11:51:24

HASH JOIN is a commonly used join algorithm in databases that accelerates the join process by using hash tables. It typically consists of two phases: Build and Probe. When the Probe side has a large volume of data but produces a small output, enabling the Runtime Filter can help pre-filter some of the data, thereby improving performance.

The Runtime Filter in the read-only analysis engine consists of two components: RF Build and RF Filter. RF Build is applied on the Build side of a HASH JOIN to construct the Runtime Filter, while RF Filter is applied on the TableScan of the Probe side of the corresponding HASH JOIN to filter data early and enhance performance.

## **Runtime Filter Types**

### Local Runtime Filter

A Local Runtime Filter is typically used in join scenarios where the data is not shuffled. In this case, the Runtime Filter built on the current node is sufficient for the Probe side, eliminating the need for network transmission. The filter data can be passed directly to the Probe side for immediate use.

As shown in the figure above, when a JOIN is performed and the Build table is not shuffled, the Runtime Filter Build operator within the same execution plan sends the constructed filter data directly to the corresponding Filter Probe component in the plan.

### **Global Runtime Filter**

When JOIN data is shuffled across different nodes for construction, the Runtime Filter built by the current node alone is not sufficient to meet the filtering requirements. In this case, the node should receive Runtime Filters from other nodes. After all Runtime Filters from the participating nodes are merged, the filter can then be used.

When a JOIN is performed and the data from the Build table is shuffled, the Runtime Filter built by the Runtime Filter Build operator in the current execution plan is incomplete. In this case, the Runtime Filter should receive and merge filter data not only from the current plan's operator but also from other operators within the same execution plan. Only after the merging is complete can the Runtime Filter be used.

## Filter Types

When you select a filter algorithm, one or more of the following filtering methods are typically chosen based on the data distribution.

### **Bloom Filter**

Bloom Filter is a classic filtering algorithm that determines data existence using multiple hash functions. In the Runtime Filter, the size of the Bloom Filter is typically determined by the data's NDV. Although Bloom Filters may produce false positives, meaning some data that should be filtered is not, such data will still be eliminated during the Probe phase of the JOIN.

#### **MIN\_MAX Filter**

The MIN\_MAX Filter collects the maximum and minimum values from the Build side data. During filtering, it checks whether the incoming data falls within this range. If the data is outside the range, it will be filtered out. This type of filter is particularly effective when the Build side data is distributed across a well-defined value range.

#### **IN Filter**

The IN Filter is designed for scenarios with a low NDV. In this case, all values of the column are directly sent to the Probe side for matching.

## Runtime Filter in the Read-Only Analysis Engine

### **Enabling or Disabling the Runtime Filter**

By default, the Runtime Filter in the read-only analysis engine is enabled. You can use the following settings to enable or disable it.

```
mysql> set libra_enable_runtime_filter=ON;
mysql> set libra_enable_runtime_filter=OFF;
```

After it is enabled, the optimizer will assess JOIN operations and automatically apply the Runtime Filter when the conditions are met.

If you want to force the Runtime Filter to be enabled for all JOIN operations, you can configure the following parameter in addition to the parameter mentioned above.

```
mysql>SET libra_enable_cost_based_runtime_filter=OFF;
```

### **Runtime Filter Plan**

As shown below, this is a Local Runtime Filter plan. Three types of Runtime Filters are assigned to the JOIN operation. In this scenario, there is no data redistribution between the Build side and the Probe side of the HASH JOIN.

The plan shown below represents a Global Runtime Filter. In this case, data is redistributed between the Build side and the Probe side. The Runtime Filter can be applied before the data is transmitted over the network, reducing network overhead and the cost of subsequent JOIN operations, thereby improving overall performance.

### Adjusting Runtime Filter Parameters

The following parameters can be adjusted for the Runtime Filter.

libra\_enable\_runtime\_filter indicates whether the Runtime Filter is enabled.

Attribute	Description
Parameter Type	BOOL.
Default Value	ON.
Value range	ON: Enables the Runtime Filter. OFF: Disables the Runtime Filter.
Scope	Global & Session.
SET_VAR Hint supported	Yes.

libra\_runtime\_filter\_type specifies the types of Runtime Filters that can be assigned.

Attribute	Description
Parameter Type	VARCHAR.
Default Value	MIN_MAX, BLOOM_FILTER, and IN_FILTER.
Value range	<ul> <li>BLOOM_FILTER: Builds a Bloom Filter on the JOIN key from the Build side to filter data on the Probe side.</li> <li>MIN_MAX: Builds the minimum and maximum values of the JOIN key from the Build side to filter data on the Probe side.</li> <li>IN: Builds a value list of the JOIN key from the Build side to filter data on the Probe side.</li> <li>Empty string: Indicates that the Runtime Filter feature is disabled.</li> </ul>
Scope	Global & Session.



SET_VAR Hint	Yes.		
supported			

libra\_enable\_cost\_based\_runtime\_filter indicates whether cost-based Runtime Filter assignment is enabled. If it is disabled, all Runtime Filters will be generated by default.

Attribute	Description
Parameter Type	BOOL.
Default Value	ON.
Value range	ON: Enables cost-based Runtime Filter assignment. OFF: Disables cost-based Runtime Filter assignment.
Scope	Global & Session.
SET_VAR Hint supported	Yes.

libra\_max\_in\_runtime\_filter\_ndv specifies the maximum NDV allowed on the Build side when an IN type Runtime Filter is generated in the cost-based Runtime Filter assignment.

Attribute	Description
Parameter Type	INT.
Default Value	1,024.00
Value range	0 - MaxValue.
Scope	Global & Session.
SET_VAR Hint supported	Yes.

# Lazy Materialization

Last updated : 2025-05-09 11:51:24

## What Is Lazy Materialization

If the primary key or index key cannot be hit when a SQL query is executed in a database, a full table scan (TableScan) is required. This process of scanning all data is usually costly. To optimize this process, Lazy Materialization can be adopted. This technology works by delaying data materialization until the query execution phase when necessary computation and storage are performed, thereby improving the query performance and system response speed.

In the process of executing SQL queries in a read-only analysis engine, predicate columns are read first, and the Filter operator is used to perform calculations to obtain the filtered results. Subsequently, the read-only analysis engine materializes other columns that need to be read based on these filtered results. This method effectively reduces the amount of data read from non-predicate columns during large-scale data filtering, thereby increasing the scanning speed.

## Advantages of Lazy Materialization

Data is compressed in the read-only analysis engine, so the materialization process requires decompression of the data. Lazy Materialization can reduce the range of the data to be decompressed, thereby reducing the CPU overhead caused by decompression.

When a SQL query involves multiple columns, early materialization may cause the database to read and combine data from all columns, even though some of the columns may not be used in the query of the final result. In contrast, Lazy Materialization delays the data combination process, allowing the database to read and process only the columns that are really needed. This approach effectively reduces unnecessary I/O operations, thereby improving the query efficiency.

Lazy Materialization makes operations like filtering and aggregation more efficient, as these operations only need to process column data rather than the entire row. The characteristics of columnar storage (such as data compression and batch processing) can be better utilized, resulting in faster query execution.

## Parameters of Lazy Materialization

Lazy Materialization can be enabled/disabled with the parameter libra\_enable\_late\_materialization, and "ON" indicates that Lazy Materialization is enabled.



Attribute	Description
Parameter Type	BOOL.
Default Value	ON.
Value range	ON: Enable Lazy Materialization. OFF: Disable Lazy Materialization.
Scope	Global & Session.
SET_VAR Hint supported	Yes.

```
#Disable Lazy Materialization at session level.
set libra_enable_late_materialization=off;
#Enable Lazy Materialization at session level.
set libra_enable_late_materialization=on;
```

## Example of Lazy Materialization

As shown in the figure below, when Lazy Materialization is enabled, you can clearly see the COLUMN READ operator in the execution plan. This operator is an example of Lazy Materialization being enabled.

## How to Set Lazy Materialization in HINT

SET\_VAR Hint allows you to enable or disable Lazy Materialization in a single SQL statement. An example is as follows.

```
select /*+ set_var(libra_enable_late_materialization=1)*/ * from t where c1=1
and c2=1;
```

# Merge Join

Last updated : 2025-05-09 11:51:24

## What Is Merge Join

Merge Join is generally called Sort Merge Join and is a commonly used join method in multi-table join queries. Specifically, it sorts the associated columns of the associated tables, respectively, then extracts data from each sorted table, and matches the data in a sorted table with the data in another sorted table.

Merge Join requires more sorting operations, and therefore, it consumes many resources. Generally speaking, in scenarios where Merge Join can be used, Hash Join can achieve better performance, that is, the effect of Hash Join is better than that of Merge Join. However, if the join keys are already sorted, there is no need to sort them when Merge Join is performed. In this case, the performance of Merge Join is better than that of Hash Join, and the performance advantage of Merge Join is effectively reflected at this time.

In earlier versions, the LibraDB engine only supports Hash Join. In version 2.2410.1.0 and later, the column storage engine also supports Merge Join based on the primary key.

## **Applicable Scenarios**

In scenarios where the join key is the primary key of both tables, Merge Join has obvious effects for equivalent association (A.a = B.a).

## Usage

Merge Join can only be used by specifying a hint, as shown in the following command.

select /\*+ merge\_join(lineitem, customer) \*/ \* from orders o, lineitem l where o.o\_or

# Columnar Storage Secondary Index

Last updated : 2025-05-09 11:51:24

## Index Introduction

An index is an important capability for database query acceleration. To fully meet the query needs of different users and improve the overall database performance, the read-only analysis engine has supported the secondary index capability based on columnar storage since version 2.2410.1.0 (including version 2.2410.1.0).

Generally speaking, creating an index can significantly reduce the data query volume for low selectivity predicate queries on high-cardinality columns, significantly optimizing the query speed.

Currently, the read-only analysis engine supports three types of indexes: Zonemap index, Bloom filter index, and Bitmap index.

### Note:

Currently, the capability to create indexes by yourself is not available. If you want to experience the index feature, submit a ticket.

## Zonemap Index

Zonemap index is a built-in index that does not need the special attention from users. It automatically maintains statistics for each column and records information such as the maximum value, the minimum value, and whether there is NULL for each data block.

For scenarios such as equality queries, range queries, and IS NULL, you can determine whether the data file and data block contain data that meets the conditions with the information such as the maximum value and minimum value. If not, skip reading the corresponding file or data block. In this way, unnecessary I/O operations can be reduced, effectively accelerating the query process.

## **Bloom Filter Index**

A Bloom filter index is a skip index based on Bloom filter, which uses Bloom filter to skip data blocks that do not meet the specified conditions of equality queries, so as to reduce I/O operations and achieve query acceleration.

A Bloom filter is a fast search algorithm for multi-hash function maps proposed by Bloom in 1970. It is usually applicable for scenarios where it is necessary to quickly determine whether an element is part of a collection, but there is no strict requirement for 100% accuracy. Bloom Filter has the following features:

A high space efficiency probabilistic data structure used to check whether an element is in a collection.

For a call to detect the existence of an element, the Bloom filter will tell the caller either of the results: it may exist or it definitely does not exist.

### **Applicable Scenarios**

The Bloom filter index can accelerate equality queries (including = and IN), and it works well for high-cardinality fields.

### Limitations

The Bloom filter index has no effect on queries other than = and IN, such as !=, NOT INT, >, and <.

The Bloom filter index only supports the INT type with a maximum length of 256, String type, Decimal type with a maximum length of 256, and Time, Date, DateTime field types.

Index creation for expressions is not supported, nor is multi-column join index.

The first column of a single primary key column or a multi-field join primary key does not support the creation of a Bloom filter index.

### Indexing

When you are executing SQL, if a Bloom filter index is created on the fields in the equality predicate or IN predicate within the where clause, the index will be automatically applied for query acceleration during querying.

## **Bitmap Index**

Bitmap index is an index represented by a bitmap which is created for every key value of the column. Compared with other indexes, the advantage of Bitmap Index is that it occupies very little storage and is very fast to create and use, while the disadvantage is that the lock granularity of modification operation is large and not suitable for scenarios with frequent updates.

### **Applicable Scenarios**

It is suitable to be created for columns with high repetition, 100 to 100,000 recommended, such as Occupation and Prefecture-Level City columns. If the repetition is too high, there is no obvious advantage compared to other types of indexes, and if too low, the spatial efficiency and performance will be greatly reduced.

Specific types of queries, such as logical operations like count, or, and, only require bit operations. For example, query with a combination of multiple conditions: select count (\*) from table where city = 'Nanjing' and job = 'doctor' and Type = 'iphone' and gender ='male'. For such scenarios, if a Bitmap index is created on each query condition column, the database can perform efficient bit operations to precisely locate the required data and reduce disk I/O operations. The smaller the filtered result set, the more pronounced the advantage of Bitmap index becomes.

It is suitable for analysis scenarios such as ad-hoc query and multi-dimensional analysis. If there is a table with 100 columns and users use 20 columns as query conditions (using any columns among these 20 columns) to create 20

Bitmap indexes on these columns, then all queries can be applied to the indexes.

### **Scenarios Not Applicable**

Columns with low repetition, for example, Identity Number and Mobile Number columns.

Columns with high repetition, for example, Gender column. For such columns, you can create a Bitmap index, but it is recommended to use it to filter in conjunction with other conditions rather than using it as the sole query condition. Columns that often need to be updated or modified.

### Limitations

The Bitmap index supports expressions such as =, !=, >, <, >=, <=, in, is null, is not null, but multiple predicates can only be connected by "and".

The Bitmap index only supports INT type with a maximum length of 256, String type, Decimal type with a maximum length of 256, and Time, Date, DateTime field types.

Index creation for expressions is not supported, nor is multi-column join index.

The first column of a single primary key column or a multi-field join primary key does not support the creation of a Bitmap index.

# System Variables

Last updated : 2025-05-09 11:51:24

This document details the list of variables that can be specified with set\_var in SQL HINT.

## **Runtime Filter Related**

System Variables	Description	Parameter Type	Default Value	Value Ranç
libra_enable_runtime_filter	Specifies whether to enable the Runtime filter.	BOOL	ON	ON: Enable Runtime filt OFF: Disat Runtime filt
libra_runtime_filter_type	Sets the Runtime filter types that can be assigned.	VARCHAR	MIN_MAX, BLOOM_FILTER, and IN_FILTER	BLOOM_F Builds a Ble Filter on the key from th side to filter on the Prok MIN_MAX: the minimu maximum v of the JOIN from the Bu side to filter on the Prok IN: Builds a list of the J key from th side to filter on the Prok Empty strin Indicates th Runtime Fi feature is disabled.
libra_enable_cost_based_runtime_filter	Enables/Disables	BOOL	ON	ON: Enable



	the cost-based Runtime filter.			cost-based Runtime filt assignmen OFF: Disak cost-based Runtime Fi assignmen
libra_max_in_runtime_filter_ndv	The maximum NDV (Number of Distinct Values) on the BUILD side when a Runtime filter of type IN is generated in the cost-based Runtime filter.	INT	1024	0 - MaxVal
runtime_filter_wait_time_ms	Indicates the maximum time to wait for RF Ready for the first time on the RF USE side, where 0 means waiting indefinitely until timeout or interrupting the query execution.	INT	100	[0, MaxValı

## Late Materialization Related

System Variables	Description	Parameter Type	Default Value	Value Range	Scope
libra_enable_late_materialization	Specifies whether to enable the late materialization.	BOOL	OFF	ON: Enables the late materialization. OFF: Disables the late	Global & Session

		materialization.	

# Parallel & Concurrency Related

System Variables	Description	Parameter Type	Default Value	Value Range	Scope	Support for SET_VAR Hint
max_threads	Queries the parallelism during the execution.	INT	The number of CPU cores in a node	0 - MaxValue (recommended not to exceed twice the number of CPU cores).	Global & Session	Yes

## **Execution Engine Related**

System Variables	Description	Parameter Type	Default Value	Value Ra
max_bytes_before_external_agg_uniq_exact	When the count distinct is using the calculation function uniqExactDisk, if the hashset exceeds this limit, it starts to switch to external memory. The default value is 0, which means no external memory calculation is performed.	INT	0	[0, MaxV;
one_bucket_max_temp_file_size	The maximum size of each file allowed to be	INT	0	[0, MaxVa



	stored into the bucket when the count distinct is using the calculation function uniqExactDisk. When this size is exceeded, it is necessary to switch to a new file to write.			
count_distinct_implementation	The name of the function the count distinct uses for calculation. By default, uniqExactDisk is used.	VARCHAR	uniqExactDisk	uniqExac An accura CountDis calculatic method w the storaç feature. uniq: An approxim CountDis calculatic method. uniqExac accurate CountDis calculatic method fo pure men
join_build_concurrency	Specify the parallelism used for the join build. The default value is 0, which means using the default value or specifying the MaxThreads.	INT	0	[0, MaxV;
enable_local_tunnel	Supports communication optimization of	BOOL	ON	ON: Enat the local channel l



	the local sender and receiver.			communi optimizat OFF: Dis the local channel le communi optimizat
mpp_max_packet_size	The maximum size of data packets sent by the sender. 0 means no limit.	INT	0	[0, MaxV
tunnel_buffer_size	The number of BUFFERs, the buffered sending queues from the sender side. The default value is 1.	INT	1	[0, MaxVi
pipeline_executor_use_thread_manager	The pipeline model uses a dynamic thread pool model.	BOOL	OFF	ON: Enak the pipeli model to dynamic thread pc OFF: Dis the pipeli model fro using a dynamic thread pc
enable_order_by_push_down	Enables/Disables the order by push-down feature.	BOOL	ON	ON: Enat the order push-dov feature. OFF: Dis the order push-dov feature.
max_block_size	Sets the rows of data blocks passed in the	INT	65409	[1, MaxV


	pipeline execution engine.			
shuffle_after_join	Sets whether to scatter data into multiple pipelines after joins.	BOOL	ON	ON: Enak the scatte feature. OFF: Dis the scatte feature.

# System Tables SLOW\_LOG

Last updated : 2025-05-09 11:51:24

### Feature

SLOW\_LOG is a system table that records slow SQL queries. In the read-only analysis engine, any SQL statement with an execution time exceeding 300ms will be logged in this table. By default, slow query records are retained for 7 days.

### Query Example

select \* from libra\_system.slow\_log;

Field Name	Туре	Nullable	Description
start_time	text	NO	The time when the SQL execution started.
tso	bigint(20) unsigned	NO	The unique identifier of the SQL statement.
sql	text	NO	The SQL text information.
duration	bigint(20) unsigned	NO	The execution time of the SQL statement.
details	text	NO	Detailed information about the SQL execution.
conn_id	bigint(20) unsigned	NO	The session ID of the client connection.

### **Field Description**



user	text	NO	The username used for execution.
db	text	NO	The database involved.
digest	text	NO	Reserved column name (currently unused).
SUCC	tinyint(3) unsigned	NO	Whether the SQL statement was executed successfully.
host	text	NO	Information about the node that executed the SQL statement.

# Performance White Paper TPC-H Testing Method

Last updated : 2025-05-09 11:51:25

The read-only analysis engine of TencentDB for MySQL is designed primarily for complex SQL queries and data analysis scenarios. In the industry, TPC-H benchmark test sets are commonly used to assess performance in such use cases. Therefore, this document uses the TPC-H benchmark to test the performance of the read-only analysis engine for user reference.

### Prerequisites

You have prepared a TencentDB for MySQL instance. See Creating a MySQL Instance. For the architecture, select either a two-node or three-node configuration.

You have created a read-only analysis engine based on the MySQL instance above, and configured it with an appropriate instance specification. See Enabling Read-Only Analysis Engine.

You have prepared a database account for data queries. See Creating an Account.

#### Note:

The test results are proportional to the compute specifications of the instance. In data analysis scenarios, achieving better performance requires more resources. You can select an appropriate instance specification based on your actual needs.

### **Building a Test Table**

The TPC-H test set includes eight data tables and one view. Use the following statements to create the tables in TencentDB for MySQL.

```
create database tpch;
use tpch;
drop table if exists customer;
create table customer (
   c_custkey bigint(20) not null,
   c_name varchar(25) not null,
   c_address varchar(40) not null,
   c_nationkey bigint(20) not null,
   c_phone char(15) not null,
   c_acctbal decimal(15,2) not null,
```

```
c_mktsegment char(10) not null,
 c_comment varchar(117) not null,
 primary key (c_custkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists lineitem;
create table lineitem (
  l orderkey bigint(20) not null,
 l_partkey bigint(20) not null,
 l_suppkey bigint(20) not null,
 l_linenumber bigint(20) not null,
 l_quantity decimal(15,2) not null,
 l_extendedprice decimal(15,2) not null,
 l_discount decimal(15,2) not null,
 l_tax decimal(15,2) not null,
 l_returnflag char(1) not null,
 l_linestatus char(1) not null,
 l_shipdate date not null,
 l_commitdate date not null,
 l_receiptdate date not null,
 l_shipinstruct char(25) not null,
 l_shipmode char(10) not null,
 l_comment varchar(44) not null,
 primary key (l_orderkey,l_linenumber)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists nation;
create table nation (
 n_nationkey bigint(20) not null,
 n_name char(25) not null,
 n_regionkey bigint(20) not null,
 n_comment varchar(152) default null,
 primary key (n_nationkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists orders;
create table orders (
 o_orderkey bigint(20) not null,
 o_custkey bigint(20) not null,
 o_orderstatus char(1) not null,
 o_totalprice decimal(15,2) not null,
 o_orderdate date not null,
 o_orderpriority char(15) not null,
 o_clerk char(15) not null,
 o_shippriority bigint(20) not null,
 o_comment varchar(79) not null,
 primary key (o_orderkey)
```

```
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists part;
create table part (
 p_partkey bigint(20) not null,
 p_name varchar(55) not null,
 p_mfgr char(25) not null,
 p brand char(10) not null,
 p_type varchar(25) not null,
 p_size bigint(20) not null,
 p_container char(10) not null,
 p_retailprice decimal(15,2) not null,
 p_comment varchar(23) not null,
 primary key (p_partkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists partsupp;
create table partsupp (
 ps_partkey bigint(20) not null,
 ps_suppkey bigint(20) not null,
 ps_availqty bigint(20) not null,
 ps_supplycost decimal(15,2) not null,
 ps_comment varchar(199) not null,
 primary key (ps_partkey,ps_suppkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists region;
create table region (
 r_regionkey bigint(20) not null,
 r_name char(25) not null,
 r_comment varchar(152) default null,
 primary key (r_regionkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
drop table if exists supplier;
create table supplier (
 s_suppkey bigint(20) not null,
 s name char(25) not null,
 s_address varchar(40) not null,
 s_nationkey bigint(20) not null,
 s_phone char(15) not null,
 s_acctbal decimal(15,2) not null,
 s comment varchar(101) not null,
 primary key (s_suppkey)
) engine=innodb default charset=utf8mb4 collate=utf8mb4_bin;
```

## View



### Generating Test Data

You can register and download the official TPC-H testing tools from the TPC Official Website. After downloading, upload the files to your server and proceed with compiling the data generation tool.

```
unzip TPC-H_Tools_v3.0.0.zip
cd TPC-H_Tools_v3.0.0/dbgen
make
```

#### Note:

If the server does not have compilation tools such as gcc or make installed, use the server's installation image to install the required components.

After the compilation is completed, the dbgen tool will be generated. Run the following command to proceed.

```
scale=100
chunk=10
for i in seq 1 $chunk
do
./dbgen -s $scale -C $chunk -S $i -f
done
```

#### Note:

In the command above, scale specifies the size of the generated data. For example, setting it to 100 generates 100 GB of data. chunk defines how many parts the data files will be split into. These two parameters can be adjusted based on actual needs. For large data set testing scenarios, it is recommended to set a higher chunk value, which results in smaller sub-files and allows for concurrent data imports.

### Importing Data

The following section explains how to import TPC-H test data into the TencentDB for MySQL instance. The table below lists the number of records for each table in the 100 GB TPC-H data set as a reference. For other data sizes,

the number of records increases or decreases proportionally. For example, a 10 GB data set contains one-tenth the number of records compared to the 100 GB data set.

Table Name	Number of Records
customer	1500000
lineitem	600037902
nation	25
orders	15000000
part	2000000
partsupp	8000000
region	5
supplier	100000

The table files generated by the TPC-H dbgen tool are stored in the same directory as the tool and use the .tbl file extension. If the data was split during generation, the files will include a numeric suffix, following a format such as partsupp.tbl.1.

```
ls *.tbl
customer.tbl lineitem.tbl nation.tbl orders.tbl partsupp.tbl part.tbl
region.tbl supplier.tbl
```

For TencentDB for MySQL instances, you can use the LOAD DATA command to import the data. First, navigate to the directory where the generated test table data is stored, then run the following command to import the data into the read-write instance:

#### Note:

The following commands navigate to the directory where the TPC-H data is stored and import the files into the database one by one. You can adjust the script as needed based on your actual environment.

```
cd /data/tpchdata
HOST=172.16.0.22
PORT=3306
USER=root
Password=xxxxx
DATABASE=tpch
ls *.tbl* | while read filename
do
      tablename=echo $filename | awk -F'.tbl' '{print $1}'
```

#### 🔗 Tencent Cloud

```
mysql -u${USER} -h${HOST} -P${PORT} -p${PASSWORD} ${DATABASE} -e "LOAD DATA
LOCAL INFILE '${filename}' INTO TABLE ${tablename} FIELDS TERMINATED BY '|'
LINES TERMINATED BY '\\n';"
done
```

After importing the data into the read-write instance, you can use SQL commands to verify that the data has been accurately and completely imported. You can also go to the console and directly configure full data loading for the tpch database.

You can also log in to the read-only analysis engine via the command line to execute SQL queries, or check the data loading status through the console.

```
mysql -uroot -p'' -h10.1.1.3 -P3306 -c
show replication tables;
```

When the value of the REPLICATION\_STEP field for a table is Change Propagation, it indicates that the data has been fully loaded.

Once the data has been fully loaded into the read-only analysis engine, you can proceed to collect table statistics.

```
ANALYZE TABLE customer;
ANALYZE TABLE lineitem;
ANALYZE TABLE nation;
ANALYZE TABLE orders;
ANALYZE TABLE part;
ANALYZE TABLE partsupp;
ANALYZE TABLE region;
ANALYZE TABLE supplier;
```

After collecting the statistics, you can log in to the read-only analysis engine to run the TPC-H test SQL queries. There are 22 TPC-H queries in total. The complete SQL statements are provided below. You can copy the SQL and paste it into the read-only analysis engine for execution.

#### Note:

Note: When using the MySQL client to log in to the read-only analysis engine and execute SQL statements, you need to include the -c parameter. This ensures that hints within the SQL statements are properly passed to the database, allowing for better performance. For example, mysql -uroot -p" -h10.1.1.3 -P3306 -c.

```
# Q1
select l_returnflag, l_linestatus, sum(l_quantity) as sum_qty, sum(l_extendedprice)
# Q2
select /*+ SET_VAR(libra_enable_runtime_filter=1) SET_VAR(libra_enable_cost_based_
# Q3
select l_orderkey, sum(l_extendedprice * (1 - l_discount)) as revenue, o_orderdate
# Q4
select o_orderpriority, count(*) as order_count from orders where o_orderdate >= '1
# Q5
select /*+ SET_VAR(libra_enable_runtime_filter=1) SET_VAR(libra_enable_cost_based_
```

S Tencent Cloud

# Q6

select sum(l\_extendedprice \* l\_discount) as revenue from lineitem where l\_shipdate # Q7 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based # Q8 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_r # Q9 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_r # Q10 select /\*+ HASH JOIN PROBE((nation, customer)), leading(nation, customer, (orders, 1 # Q11 select ps\_partkey, sum(ps\_supplycost \* ps\_availqty) as value from partsupp, supplie # 012 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_ # 013 select c\_count, count(\*) as custdist from (select c\_custkey, count(o\_orderkey) as # Q14 100.00 \* sum(case when p\_type like 'PROMO%'then l\_extendedprice \* (1 - l\_di select # Q15 select s\_suppkey, s\_name, s\_address, s\_phone, total\_revenue from supplier, revenue # Q16 select p\_brand, p\_type, p\_size, count(distinct ps\_suppkey) as supplier\_cnt from pa # Q17 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_r # Q18 select /\*+ PX\_JOIN\_FILTER\_ID(0,1,2) leading(customer, (orders, lineitem@sel\_2)) # 019 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_ # Q20 select s\_name, s\_address from supplier, nation where s\_suppkey in (select ps\_suppk # Q21 select /\*+ SET\_VAR(libra\_enable\_runtime\_filter=1) SET\_VAR(libra\_enable\_cost\_based\_ # 022 select cntrycode, count(\*) as numcust, sum(c\_acctbal) as totacctbal from (select s

# **TPC-H Performance Results**

Last updated : 2025-05-09 11:51:25

This document introduces the TPC-H performance test results for the read-only analysis engine of TencentDB for MySQL.

#### Note:

The performance of read-only analysis engine instances varies depending on the instance specification. The test results in this document are for reference only. For details on the testing method, see TPC-H Test Methodology.

#### TPC-H

Instance Specification	Disk Size	Test Data Set
General type 32 C 128 GB	200GB	100GB

#### Performance Results:

Query SQL	Execution Time (s)
Q1	3.66
Q2	0.39
Q3	2.37
Q4	1.53
Q5	2.91
Q6	0.74
Q7	1.99
Q8	1.6
Q9	7.8
Q10	2.81
Q11	0.42
Q12	1.79
Q13	3.32

Q14	1.12
Q15	1.2
Q16	1.18
Q17	1.42
Q18	3.62
Q19	1.74
Q20	1.64
Q21	5.04
Q22	0.51
Total	48.8

# Tag Overview

Last updated : 2024-07-23 10:22:09

#### Overview

A **tag** is a key-value pair provided by Tencent Cloud to identify a resource in the cloud. For more information, see Tag Overview.

You can manage TencentDB for MySQL resources in a categorized manner by using various tag types, such as business, purpose, and person-in-charge, making it easier to find the right resources. Tags have no semantic meaning for Tencent Cloud and are parsed and matches strictly based on strings. During us, you only need to pay attention to applicable use limits.

Below is a specific use case to show how a tag is used.

#### Case Background

A company owns 10 TencentDB for MySQL instances in Tencent Cloud. Distributed in three departments (ecommerce, gaming, and entertainment), these instances are used to serve internal business lines such as marketing, game A, game B, and post-production. The OPS owners of the three departments are John, Jane, and Harry, respectively.

#### Setting a Tag

To facilitate management, the company categorizes its TencentDB for MySQL resources with tags and defines the following tag key-value pairs.

Tag Key	Tag Value
Department	Ecommerce, gaming, and entertainment
Business	Marketing, game A, game B, and post-production
OPS owner	John, Jane, and Harry

These tag keys/values are bound to TencentDB for MySQL instances in the following way:

instance-id	Department	Business	OPS Owner
cdb-abcdef1	Ecommerce	Marketing	Harry
cdb-abcdef2	Ecommerce	Marketing	Harry
cdb-abcdef3	Gaming	Game A	John

cdb-abcdef3	Gaming	Game B	John
cdb-abcdef4	Gaming	Game B	John
cdb-abcdef5	Gaming	Game B	Jane
cdb-abcdef6	Gaming	Game B	Jane
cdb-abcdef7	Gaming	Game B	Jane
cdb-abcdef8	Entertainment	Post-production	Harry
cdb-abcdef9	Entertainment	Post-production	Harry
cdb-abcdef10	Entertainment	Post-production	Harry

#### Using a Tag

For more information on how to create and delete a tag, see Operation Guide.

For more information on how to edit a tag in TencentDB for MySQL, see Editing a Tag.

# **Editing Tags**

Last updated : 2024-07-23 10:22:21

You can edit resource tags by the following steps.

### Editing the tag of a single instance

1. Log in to the TencentDB for MySQL console. In the instance list, locate the desired instance and click **More** > **Edit Tag** in the **Operation** column.

2. In the pop-up dialog box, add, modify, or delete a tag, and click OK.

Edit Tags				×
The tag is used to m tag does not meet y	anage resource our requiremen	s by category fro ts, please go to	om different dimen: Manage Tags 🗹	sions. If the existing
1 resource selected				
ivy	Ŧ	t1	•	×
+ Add				
		OK Car	ncel	

### Editing the tags of multiple instances

1. Log in to the TencentDB for MySQL console. In the instance list, select multiple instances and click **More** > **Edit Tag** at the top.

2. In the pop-up dialog box, add, modify, or delete a tag, and click **OK**.