

# TencentDB for MySQL

# Tencent Kernel TXSQL



## Copyright Notice

©2013–2026 Tencent Cloud. All rights reserved.

The complete copyright of this document, including all text, data, images, and other content, is solely and exclusively owned by Tencent Cloud Computing (Beijing) Co., Ltd. ("Tencent Cloud"); Without prior explicit written permission from Tencent Cloud, no entity shall reproduce, modify, use, plagiarize, or disseminate the entire or partial content of this document in any form. Such actions constitute an infringement of Tencent Cloud's copyright, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

## Trademark Notice



This trademark and its related service trademarks are owned by Tencent Cloud Computing (Beijing) Co., Ltd. and its affiliated companies ("Tencent Cloud"). The trademarks of third parties mentioned in this document are the property of their respective owners under the applicable laws. Without the written permission of Tencent Cloud and the relevant trademark rights owners, no entity shall use, reproduce, modify, disseminate, or copy the trademarks as mentioned above in any way. Any such actions will constitute an infringement of Tencent Cloud's and the relevant owners' trademark rights, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

## Service Notice

This document provides an overview of the as-is details of Tencent Cloud's products and services in their entirety or part. The descriptions of certain products and services may be subject to adjustments from time to time.

The commercial contract concluded by you and Tencent Cloud will provide the specific types of Tencent Cloud products and services you purchase and the service standards. Unless otherwise agreed upon by both parties, Tencent Cloud does not make any explicit or implied commitments or warranties regarding the content of this document.

## Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

# Contents

## Tencent Kernel TXSQL

Overview

Kernel Version Release Notes

TXSQL Kernel Release Notes

TXRocks Kernel Release Notes

Database Proxy Kernel Update Dynamics

Functionality Features

Killing Idle Transactions Automatically

Parallel Replication

Dynamic thread pool

NOWAIT

RETURNING

Column Compression

Flashback Query

Performance Features

Parallel Query

Overview

Supported Statement Scenarios and Restricted Scenarios

Enabling/Disabling Parallel Query

HINT Statement Control

Viewing Parallel Query

Large Transaction Replication

Execution Plan Cache for Optimizing UK/PK Queries

fdatasync()

Auto-Increment Column Persistence

Buffer Pool Initialization

FAST DDL

Invisible Index

CATS Transaction Scheduling Algorithm

Computation Pushdown

Security Features

Transparent Data Encryption

Audit

Stability Features

Second-Level Column Addition

**Async Deletion of Big Tables**

**Hotspot Update**

**SQL throttling**

**Statement Outline**

**TXRocks Engine**

**Overview**

**Instructions**

**Cost Performance**

**Best Practices**

**Checking and Fixing Kernel Issues**

**Summary of Critical Kernel-related Issues**

**Corruption in the Tablespace fseg not full list Caused by Purging or Rollbacks on Two Blob Fields**

**Rollback Failure After Upgrading to a Later Version When Instant ADD COLUMN Has Caused Page Data Usage to Exceed 50%**

**Crash Caused by Adding More Than 1024 Columns Using Instant ADD COLUMN**

**Crash Caused by Instant ADD COLUMN Followed by Online DDL on the Same Table**

**Crash Caused by Instant ADD COLUMN of a Column with the Same Name after Instant DROP**

**Crash Caused by a Redo Buffer Calculation Defect in Instant CHANGE COLUMN**

**Data Loss Caused by Table Rebuild Operations (ALTEROPTIMIZE) in Specific MySQL Versions**

**Instant DDL Cannot Properly Handle Column Name Case Inconsistency**

**Triggered Memory Corruption When DDL Statements Are Executed, Caused Crash**

**The Out-of-Bounds Access to ha\_alter\_info->key\_info\_buffer Causes Instance Crash**

# Tencent Kernel TXSQL

## Overview

Last updated: 2023-08-31 16:39:24

TXSQL is a MySQL kernel branch maintained by the TencentDB team and is fully compatible with native MySQL. It provides various features similar to those in the MySQL Enterprise Edition, such as enterprise-grade transparent data encryption (TDE), auditing, dynamic thread pool, encryption function, backup and restoration, and parallel query.

TXSQL not only extensively optimizes InnoDB storage engine, query optimization, and replication performance, but also enhances the usability and maintainability of TencentDB for MySQL. In addition to offering all MySQL features, it provides enterprise-grade advanced features such as disaster recovery, monitoring, performance optimization, read-write separation, transparent data encryption (TDE), and auditing.

The following provides more information about TXSQL:

- For details on the TencentDB for MySQL TXSQL kernel version updates, see [TXSQL Kernel Release Notes](#).
- For details on the TencentDB for MySQL TXRocks kernel version updates, see [TXRocks Kernel Release Notes](#).
- The kernel minor versions of TencentDB for MySQL can be upgraded automatically or manually. For more information, see [Upgrading Kernel Minor Version](#).
- You can use a CVM instance to log in to a TencentDB for MySQL instance and check its kernel minor version. For more information, see [Kernel Upgrade](#).

# Kernel Version Release Notes

## TXSQL Kernel Release Notes

Last updated: 2025-11-19 10:05:31

This document describes the version updates of the TXSQL kernel.

### Note

- For more information on how to upgrade the minor kernel version of a TencentDB for MySQL instance, see [Upgrading Kernel Minor Version](#).
- When you upgrade the minor version, some minor versions may be under maintenance and cannot be selected. The minor versions available in the console shall prevail.
- To facilitate comparison with the database version, the table below introduces the Community Edition, denoting the open-source version of MySQL.

### MySQL 8.0 Kernel Version Release Notes



TXSQL Kernel Version	Community Version	Description
20241005	8.0.30	<p><b>Note:</b> Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3 and will report the error "Character Set 'utf8mb3' is not supported by .Net Framework" if utf8mb3 is used. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before you upgrade the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported displaying the query plan with the Iterator tree structure in JSON format for the explain command.</li> <li>○ Supported encoding query plans and displaying the query plan ID in slow logs and the explain result.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue of missing sql_mode in subqueries of parallel queries.</li> <li>○ Fixed the issue of crash caused by the SIGABRT signal due to the std::out_of_range exception during InnoDB statistics collection under high concurrency.</li> </ul> </li> </ul>
		<p><b>Note:</b> Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3 and will report</p>

20241001	8.0.30	<p>the error "Character Set 'utf8mb3' is not supported by .Net Framework" if utf8mb3 is used. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before you upgrade the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Adapted to the scenario of disabling session-level binlog when hotspot update is enabled. sql_bin_log cannot be modified to disable session-level binlog if a session has a hotspot update transaction. It can be modified to disable session-level binlog if a session has no such transaction. If binlog is disabled for a session, subsequent hotspot update transactions in this session will not be executed.</li> </ul> </li> </ul>
		<p><b>Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3 and will report the error "Character Set 'utf8mb3' is not supported by .Net Framework" if utf8mb3 is used. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before you upgrade the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Support "show full binary logs" to show the last modification time.</li> <li>○ Support the parallel execution feature of multi-stage aggregation operation and execution.</li> <li>○ Parallel Copy DDL supports partition tables.</li> <li>○ Eliminate duplicate aggregate functions.</li> </ul> </li> </ul>

20240930	8.0.30	<ul style="list-style-type: none"> <li>○ Eliminate common subexpressions.</li> <li>○ Expand subquery cache to support EXISTS/IN subqueries.</li> <li>○ Correlation statistics of multiple columns.</li> <li>○ Index dive pruning optimization.</li> <li>○ SQL traffic throttling supports startup loading.</li> <li>● Performance optimization <ul style="list-style-type: none"> <li>○ Fix query cache bugs and optimize performance.</li> <li>○ Parallelize the process of checking indexes to speed up the check during backup.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fix the fast cleanup AHI crash issue.</li> <li>○ Fix the issue that "drop database" can cause replication interruption during the use of the recycle bin when the database contains foreign key tables.</li> <li>○ Fix the compatibility issue that the instant metadata column of the partition table is upgraded from 5.7 to 8.0.</li> <li>○ Limit the execution of XA COMMIT/ROLLBACK on RO to prevent replication interruption caused by executing related operations.</li> <li>○ Fix the issue that TRUNCATE SUBPARTITION TEMPLATE causes a crash.</li> <li>○ Fix the issue that performing the rollback operation on the partition table causes a crash.</li> <li>○ Fix the issue that subpartition by range columns causes a crash for having multiple same columns.</li> <li>○ Fix the issue that statistical information may be cleared after COPY DDL is executed.</li> <li>○ Fix the issue that the query rewrite plugin causes a crash when it is used in multi-statement mode.</li> </ul> </li> </ul>
		<div style="border: 1px solid #00a88f; padding: 10px;"> <p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3, and will throw an error if utf8mb3 is used: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, upgrade</p> </div>

20230704	8.0.30	<p>Connector/Net to version 8.0.28 or later before upgrading the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed an issue where rebuilding a table may cause data loss in some extreme cases. For details, see <a href="#">Bug#110706</a>.</li> </ul> </li> </ul>
20230703	8.0.30	<p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3, and will throw an error if utf8mb3 is used: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before upgrading the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the memory allocation method at the location of the memory leak to prevent memory leakage.</li> </ul> </li> </ul>
		<p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3, and will throw an error if utf8mb3 is used: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before</p>

20230702	8.0.30	<p>upgrading the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• Performance optimizations <ul style="list-style-type: none"> <li>○ Reduced the overhead caused by frequent notifying of parallel query exchange operators.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the abnormal issue of the instance crash caused by an instant DDL's modifying a column position.</li> </ul> </li> </ul>
20230701	8.0.30	<p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3, and will throw an error if utf8mb3 is used: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, upgrade Connector/Net to version 8.0.28 or later before upgrading the TencentDB for MySQL version. For details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> <ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Supports initiating TXSQL physical backup using the community version.</li> <li>○ Replaced all keywords in the AWR feature with "TXSQL_AWR".</li> <li>○ Fixed an issue where the background thread competing with RENAME on the dict cache could cause a failure to open the underlying table.</li> </ul> </li> </ul>
		<p><b>ⓘ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3</p>

tables in the information schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3; encountering utf8mb3 will result in an error: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, please upgrade Connector/Net to 8.0.28 or later before upgrading the TencentDB for MySQL version. For more details, see:

- [MySQL 8.0.29 Character Set Support](#)
- [Changes in MySQL Connector/NET 8.0.28](#)

- New features
  - Supported Nonblocking DDL feature.
  - Supported xa commit to record the maximum gts instance TP/AP load statistics in relay log.
  - Supported selecting Innodb temporary tables for parallel query of worker thread sharing.
  - Supported using partition tables as parallel tables for parallel queries.
  - Supported the flashback version query feature.
  - Supports persistence for flashback query.
  - Supported virtual indexes.
  - Supported the range/list secondary partition feature.
  - Supported the automatic relay log recovery feature.
  - Supported the default algorithm for DDL, with options INPLACE/INSTANT.
  - Supported Fast Query Cache.
  - Supported the conversion of partition tables from MyISAM to InnoDB.
  - Supported the correlated subquery cache feature.
- Performance Optimization
  - Optimized the BINLOG LOCK\_done lock conflict.
  - Optimized thread pool performance.
  - Optimized the issue where disabling eq\_ref cache in outer join leads to performance regression.
  - Enhanced parallel query:
    - Subquery \derived table executed in parallel independently: Optimized and executed plan for subquery \derived table in parallel.

independent of main query execution.

- Nested loop join inner table in parallel: When the NLJ outer table is small, the inner table can be selected as the parallel table for parallel execution with ROLL UP.
- Hash join (in memory) executed in parallel: Work threads build complete hash tables separately, with parallel scanning on the probe end.
- Parallel query supports global aggregation optimization.
- Parallel query supports pushdown parallelism under having condition.
- Optimized binlog submission for large transactions.
- Optimized performance fluctuation caused by binlog purge.
- Supported for hot updates, merge and optimization.
- Bug Fixes
  - Fixed the issue of assertion failure when rolling back transactions of Parallel Copy DDL.
  - Fixed the issue where EXPLAIN FORMAT=TREE does not print subqueries in the condition of HashJoin.
  - Fixed the issue where redundant format causes instance running exception after instant add.
  - Fixed the issue of global transaction id rollback after upgrading from an older version.
  - Fixed the issue where index merge intersect causes incorrect query results.
  - Fixed the issue where cross-machine statistical information collection may block the shutdown process.
  - Fixed the issue where historical histogram versions might crash in a primary-secondary environment.
  - Fixed the issue of check index holding a large number of page locks.
  - Fixed the deadlock issue in cross-machine histograms under concurrent DDL operations.
  - Fixed the issue where the lock was not released when the cross-machine histogram task included too many columns.
  - Fixed several instant DDL issues.

20230630

8.0.30

- Fixed the issue where update returning caused the client to disconnect.
- Fixed the null pointer dereference vulnerability found by the vulnerability scan.
- Fixed the issue where changes in the storage layer table structure under parallel execution may cause instance running exception.
- Fixed the performance degradation issue caused by using WHERE column IN (list) in prepare statements.
- Fixed the issue where statistical information might be empty when importing mysqldump logical backups.
- Fixed the primary key conflict issue that occurs when using Parallel Copy DDL for table changes that include auto-increment columns.
- Fixed the issue where the build branch of the hash join in parallel queries cannot be parallelized when the hash join is present.
- Fixed the issue where the non-parallel branches of a parallel query join cannot be parallelized when there is a UNION.
- Fixed two memory leak issues in parallel queries.
- Fixed the partition exit issue for empty range in parallel queries.
- Fixed the error issue with Outline IN-list.
- Fixed the issue where partition\_id overflow leads to truncate partition crash.
- Fixed the issue in parallel queries where related subqueries referencing worker table fields resulted in incorrect query results.
- Fixed the issue in parallel DDLs regarding obtaining an incorrect offset.
- Fixed the issue in parallel DDLs where adding a unique key to a column with duplicate data caused the instance to run abnormally.
- Fixed the issue of assertion in parallel hash join debug.
- Fixed the issue in parallel cost calculation where NDV was 0.
- Fixed the issue with JSON import accuracy.
- FORCE INDEX ORDER BY statement skips the

		<p>index alive bug.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue where the official subquery plan was displayed multiple times.</li> <li>○ Fixed the issue where the disk-based temporary table quantity does not increase.</li> <li>○ Fixed the deadlock issue caused by proxy change user.</li> <li>○ Fixed the issue where calling a stored procedure in a trigger due to permission verification optimization caused permission checks to be bypassed.</li> </ul>
20221221	8.0.30	<div style="border: 1px solid #00a88f; padding: 10px; margin-bottom: 10px;"> <p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3; encountering utf8mb3 will result in an error: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, please upgrade Connector/Net to 8.0.28 or later before upgrading the TencentDB for MySQL version. For more details, see:</p> <ul style="list-style-type: none"> <li>• <a href="#">MySQL 8.0.29 Character Set Support</a></li> <li>• <a href="#">Changes in MySQL Connector/NET 8.0.28</a></li> </ul> </div> <ul style="list-style-type: none"> <li>• Bug Fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where after enabling <code>log_slave_updates</code> on a secondary node, the <code>thread_id</code> of the event written to the binlog on the secondary node changed.</li> </ul> </li> </ul>
20221220	8.0.30	<div style="border: 1px solid #00a88f; padding: 10px;"> <p><b>⚠ Note:</b></p> <p>Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3; encountering utf8mb3 will result in an error: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, please</p> </div>

upgrade Connector/Net to 8.0.28 or later before upgrading the TencentDB for MySQL version. For more details, see:

- [MySQL 8.0.29 Character Set Support](#)
- [Changes in MySQL Connector/NET 8.0.28](#)

- Bug Fixes
  - Fixed the instant DDL bug.

**Note:**

Starting from MySQL 8.0.29, the query results of tables in the Information Schema will use utf8mb3 instead of utf8. Versions of Connector/Net earlier than 8.0.28 do not support utf8mb3; encountering utf8mb3 will result in an error: Character Set 'utf8mb3' is not supported by .Net Framework. If the application uses Connector/Net, please upgrade Connector/Net to 8.0.28 or later before upgrading the TencentDB for MySQL version. For more details, see:

- [MySQL 8.0.29 Character Set Support](#)
- [Changes in MySQL Connector/NET 8.0.28](#)

- New Features
  - Merges official changes from [8.0.23](#) to [8.0.30](#).
  - Supported user connection status monitoring feature, which can be viewed through show detail processlist for connection monitoring.
  - Supported the update wait N syntax.
  - Supported nvl(), to\_number(), to\_char() function feature syntax.
  - Supported cdb\_kill\_user\_extra regular expression.
- Performance Optimization
  - Optimized binlog rotate implementation method and improved binlog write speed.
  - Optimized TencentDB for MySQL startup speed.
  - Optimized binlog checksum calls, and reduced unnecessary CPU performance overhead.
  - Optimized ha\_innopart::external\_lock lock

20221215	8.0.30	<p>hotspots, and reduced lock holding time.</p> <ul style="list-style-type: none"> <li>○ Optimized xa::Transaction_cache, and reduced lock conflicts.</li> <li>○ Reduced ha_innopart::clear_blob_heaps time consumption.</li> <li>○ Optimized purge threads lock hotspots, and reduced tasks_mutex and thread conflicts.</li> <li>○ Optimized Buffer Pool initialization, supporting parallel initialization, and accelerating initialization speed.</li> <li>○ Optimized read_only and select performance under high concurrency.</li> <li>○ Optimized permission validation for prepared statement and stored procedure.</li> <li>○ Optimized access to change buffer.</li> <li>○ Avoided unnecessary calls to fil_space_get, and reduced FAQs in extreme scenarios.</li> <li>○ Optimized lock conflict for GTID during transaction commit when binlog_order_commits is disabled.</li> <li>○ Applied Lock Free Hash to optimize trx_sys mutex conflict.</li> <li>○ Optimized the overhead of taking a snapshot in the transaction system.</li> <li>○ Optimized Writeset and improved performance.</li> <li>○ Replaced index drill-down with histogram.</li> <li>○ Supports Parallel DDL.</li> </ul> <ul style="list-style-type: none"> <li>● Bug Fixes <ul style="list-style-type: none"> <li>○ Fixed issues with abnormal statistical values such as innodb_row_lock_current_waits.</li> <li>○ Fixed the issue of excessively high memory usage with Group concat with group by.</li> <li>○ Fixed the issue of statistical information being severely underestimated in long records.</li> <li>○ Fixed the issue in parsing stored procedure syntax.</li> <li>○ Fixed the issue in FAST DDL optimization of flush list to release page concurrency.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported setting the MySQL version dynamically.</li> <li>○ Supported transparent column encryption. When creating a table, you can specify the encryption attribute for the `varchar` field, and the storage</li> </ul> </li> </ul>

system will encrypt the column. This capability is expected to be commercialized in 2023.

- Fixed the exception of the third-party data subscription tool caused by subscription to the comparison SQL for internal data consistency during tool usage.

**Note:**

After the database instance is migrated, upgraded, or recovered after failure, the system will compare the data consistency to ensure the consistency of data. When comparison SQL is in `statement` mode, exceptions are easy to occur in response of some third-party subscription tools to the SQL in `statement` mode. When the instance is upgraded to its kernel, the third-party data subscription tool can't subscribe the comparison SQL for internal data consistency.

- Supported adding NO\_WAIT | WAIT [n] for DDL operations. This enables such operations to be rolled back immediately if they cannot obtain the MDL lock and must wait or if they have waited the specified time for the MDL lock.
- Supported the fast query cache feature, which is suitable for scenarios with more reads than writes. If there are more writes than reads, the data is updated very frequently, or the result set of the query is very large, we recommend that you not enable this feature.
- Supported enhanced MTS deadlock detection.
- Supported [parallel query](#). After this feature is enabled, large queries can be automatically identified. The parallel query capability leverages multiple compute cores to greatly shorten the response time of large queries.
- Performance optimizations
  - Optimized the overheads of the transaction system to take snapshots. The Copy Free Snapshot method is adopted, the transaction delay is deleted from the global active transaction back, and the

20220831	8.0.22	<p>from the global active transaction hash, and the snapshot taking method is optimized to determine the logical timestamp of the snapshot event. As tested by sysbench, the extreme performance is increased by 11% in the read-write scenario.</p> <ul style="list-style-type: none"><li>○ Optimized permission check for prepared statements. A variable is used globally to indicate the permission version number, a prepared statement records the version number after being prepared, and the system checks whether the version number has changed during execution. If there is no permission change, the system will skip the permission check; otherwise, it will check the permission and record the version number again.</li><li>○ Optimized the accuracy of time acquisition in the thread pool.</li><li>○ Optimized record offset acquisition. A record offset is cached for each index. When the conditions are met, the cached offset will be directly used, saving the computing overheads of invoking the <code>`rec_get_offsets()`</code> function.</li><li>○ Optimized parallel DDL.<ul style="list-style-type: none"><li>1.1.1 When the index field is small, the sampled memory size is reduced to lower the sampling frequency.</li><li>1.1.2 The K-way merge algorithm is used for sorting, which effectively reduces the number of rounds of merging and sorting to lower the number of IOs.</li><li>1.1.3 When records are read, the fixed-length offset is cached in order to avoid generating offsets for each record each time.</li></ul></li><li>○ Optimized the undo log information recording logic to improve the INSERT performance.</li><li>○ Improved the performance after semi-sync was enabled.</li><li>○ Optimized the audit performance to reduce the system overheads.</li></ul> <ul style="list-style-type: none"><li>● Bug fixes<ul style="list-style-type: none"><li>○ Fixed the issue where the displayed value of <code>`Thread_memory`</code> was abnormal sometimes.</li><li>○ Fixed the issue where the timestamp was inaccurate during batch statement audit.</li><li>○ Fixed issues related to column modification at the</li></ul></li></ul>
----------	--------	---

		<p>second level.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue where the `CREATE TABLE t1 AS SELECT ST_POINTFROMGEOHASH("0123", 4326);` statement caused source-replica disconnection.</li> <li>○ Fixed the issue where the replica failed to retry during concurrent requests at the table level.</li> <li>○ Fixed the `Malformed packet` error reported when `show slave hosts` was executed.</li> <li>○ Fixed recycle bin issues.</li> <li>○ Fixed the issue where the jemalloc mechanism easily triggered OOM on ARM device models.</li> <li>○ Fixed the issue where `truncate pfs account table` caused the failure to collect statistics.</li> <li>○ Fixed the exception that occurred while restoring the child table first and then restoring the parent table when the recycle bin had a foreign key constraint.</li> <li>○ Fixed sql_mode log issues.</li> <li>○ Fixed the occasional issue where a procedure became abnormal when `CREATE DEFINER` was executed.</li> <li>○ Fixed Copy Free Snapshot issues.</li> <li>○ Fixed the performance fluctuation of the thread pool.</li> <li>○ Fixed the issue where the result of `hash join+union` might be empty.</li> <li>○ Fixed memory issues.</li> </ul>
20220401	8.0.22	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where the stage variable error in Parallel DDL caused the stage null pointer to crash when creating FTS indexes.</li> <li>○ Fixed the possible crash when adding full-text indexes.</li> </ul> </li> </ul>
20220331	8.0.22	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the crash caused by dereferencing wild pointers in the thread pool.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Enabled writeset parallel replication by default.</li> </ul> </li> </ul>

20220330

8.0.22

- Supported extended resource groups to control the I/O, memory utilization, and SQL timeout policy by user.
- Supported flashback query to query data at any time point within the UNDO time range.
- Supported `RETURNING` in a `DELETE`, `INSERT`, or `REPLACE` statement to retrieve the data rows modified by the statement.
- Supported the GTID replication feature extension in row mode.
- Supported transaction lock optimization.
- Enhanced the recycle bin to support TRUNCATE TABLE and automatic cleanup of tables in the recycle bin.
- Supported parallel DDL to speed up DDL operations for which to create indexes through three-phase parallel operations.
- Supported quick index column modification.
- Supported automatic statistics collection and cross-server statistics collection.
- Performance optimizations
  - Optimized the GTID lock conflicts when transactions were committed if `binlog\_order\_commits` was disabled.
  - Accelerated MySQL startup by changing the InnoDB startup phase from single-threaded creation of Rsegs to multi-threaded creation.
- Bug fixes
  - Fixed the issue where the transaction did not end when the connection was closed after deadlock or lock wait.
  - Fixed the issue where the `innodb\_row\_lock\_current\_waits` value was abnormal.
  - Fixed the SQL type error in the audit plugin without USE DATABASE.
  - Fixed the issue where tables smaller than `innodb\_async\_table\_size` were also renamed during async drop of big tables.
  - Fixed the issue with incorrect escape characters in the audit plugin.
  - Fixed the issue of rollback after quick column modification.

#### Improvements

- Fixed the issue where the transaction system (trx\_sys) may crash if it contains XA transactions when it is closed.
- Fixed the crash when merging derived tables.
- Fixed the issue where `binlog\_format` was modified after writeset was enabled.
- Fixed the error (error code: 1032) caused by hash scans with A→B→A→C update on the same row.
- Fixed the issue where the sort index might be invalid in prepared statement mode.
- Fixed the issue where the operator that consumed the materialized result might be merged into the returned value path of the materialized operator and result in incorrect comprehension and display of the execution plan.
- Fixed exceptions in extreme cases for async drop of big tables.
- Fixed the abnormal error message when setting a SQL filter.
- Fixed the syntax error reported during stored procedure parsing.
- Fixed the issue where historical histograms couldn't be applied.
- Fixed the role column display compatibility issue caused by `SHOW SLAVE HOSTS(show replicas)`.
- Fixed the crash of `Item\_in\_subselect::single\_value\_transformer` when the number of columns was incorrect.
- Fixed the crash caused by memory leaks during cascading update if a subtable contained virtual columns and foreign key columns.

- **New features**

- Supported quick column modification.
- Supported histogram versioning.
- Supported SQL:2003 TABLESAMPLE (single table) sampling control syntax for obtaining random samples of physical tables.
- Added non-reserved keywords: TABLESAMPLE BERNOULLI.
- Added the `HISTOGRAM()` function to build a histogram for a given input field.

20211202

8.0.22

- Supported compressed histograms.
- Supported SQL throttling.
- Supported MySQL cluster role configuration (default role: CDB\_ROLE\_UNKNOWN).
- Added a new `Role` column to the `show replicas` command's display results to display roles.
- Supported proxy.
- Performance optimizations
  - Optimized the hotspot update problem caused by `insert on duplicate key update`.
  - Accelerated the application of hash scan by aggregating multiple identical binlog events.
  - Greatly reduced the memory usage by the `PREPARE` statement in point queries in the thread pool mode when the plan cache was enabled.
- Bug fixes
  - Fixed the error of unstable performance after hotspot update optimization was enabled.
  - Fixed the issue where `select count(\*)` parallel scans caused full-table scans in extreme cases.
  - Fixed performance issues caused by execution plan changes due to reading zero statistics in various cases.
  - Fixed the bug where queries were in the `query end` status for a long time.
  - Fixed the bug where statistics were severely underestimated in long records.
  - Fixed the bug where an error was reported when the Temptable engine was used and the number of aggregate functions in the selected column exceeded 255.
  - Fixed the case sensitivity issue of column names in the `json\_table` function.
  - Fixed the bug that caused correctness issues in window functions because expressions returned early during `return true`.
  - Fixed the correctness issue caused by the pushdown by `derived condition pushdown` when it contained user variables.
  - Fixed the issue where SQL filters were prone to crash when no namespaces were added in a rule.
  - Fixed the QPS jitters when the thread pool was

		<p>enabled under high concurrency and high conflict.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue where source-replica buffer pool sync leaked file handles in extreme cases (when host file systems were corrupted).</li> <li>○ Fixed the index mapping issue.</li> <li>○ Fixed the statistics cache sync issue.</li> <li>○ Fixed the crash when information was not cleared during execution of the `UPDATE` statement or stored procedures.</li> </ul>
20210830	8.0.22	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported limiting the number of preloaded rows.</li> <li>○ Supported optimizing plan cache point query.</li> <li>○ Supported extended ANALYZE syntax (UPDATE HISTOGRAM c USING DATA 'json') and direct writes to histograms.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Replaced index seek with histogram to reduce evaluation errors and I/O overheads (this capability is not enabled by default).</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where there might be no statistics information during online DDL.</li> <li>○ Fixed the issue where generated columns on replica instances were not updated.</li> <li>○ Fixed the issue where the instance hung when binlog was compressed.</li> <li>○ Fixed the issue of missing GTID in the previous_gtid event of the newly generated binlog file.</li> <li>○ Fixed possible deadlocks when system variables were modified.</li> <li>○ Fixed the issue where the information of the SQL thread of the replica instance in SHOW PROCESSLIST was incorrectly displayed.</li> <li>○ Implemented the bug fix related to hash join provided in MySQL 8.0.23.</li> <li>○ Implemented the bug fix related to writeset provided in MySQL.</li> <li>○ Implemented the bug fix related to the query optimizer provided in MySQL 8.0.24.</li> <li>○ Fixed the concurrency bugs of optimizing flush list</li> </ul> </li> </ul>

		<ul style="list-style-type: none"> <li>and releasing pages in FAST DDL.</li> <li>○ Optimized the memory usage during data dictionary update in instances with a large number of tables.</li> <li>○ Fixed the crash caused by new primary key creation after INSTANT ADD COLUMN.</li> <li>○ Fixed the OOM caused by memory growth in full-text index query.</li> <li>○ Fixed the issue where -1 was included in the TIME field in the result set returned by SHOW PROCESSLIST.</li> <li>○ Fixed the issue where tables might fail to be opened due to histogram compatibility.</li> <li>○ Fixed the floating point accumulation error when Singleton histograms were constructed.</li> <li>○ Fixed the replication interruption caused by using many Chinese characters in the table name of a row format log.</li> </ul>
20210330	8.0.22	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported source-replica buffer pool sync: After a high-availability (HA) source-replica switch occurs, it usually takes a long time to warm up the replica, that is, to load hotspot data into its buffer pool. To accelerate the replica's warmup, TXSQL now supports the buffer pool sync between the source and the replica.</li> <li>○ Supported sort-merge join.</li> <li>○ Supported FAST DDL operations.</li> <li>○ Supported querying the value of the `character_set_client_handshake` parameter.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the mechanism of scanning and flushing the dirty pages tracked in the flush list, so as to solve the performance fluctuation issue while creating indexes and thus improve the system stability.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the deadlocks caused by the modification of the `offline_mode` and `cdb_working_mode` parameters.</li> <li>○ Fixed the persistent concurrency issue of the `max_trx_id` field in `trx_sys` table.</li> </ul> </li> </ul>

20201230	8.0.22	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported the official updates of MySQL <a href="#">8.0.19</a>, <a href="#">8.0.20</a>, <a href="#">8.0.21</a>, and <a href="#">8.0.22</a>.</li> <li>○ Supported dynamic setting of thread pooling mode or connection pooling mode by using the <code>`thread_handling`</code> parameter.</li> </ul> </li> <li>• Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the <code>`BINLOG LOCK_done`</code> conflict to improve write performance.</li> <li>○ Optimized the <code>`trx_sys mutex`</code> conflict by using lock-free hash to improve performance.</li> <li>○ Optimized redo log flushing.</li> <li>○ Optimized the buffer pool initialization time.</li> <li>○ Optimized the clearing of adaptive hash indexes (AHI) during the <code>`drop table`</code> operations on big tables.</li> <li>○ Optimized audit performance.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed performance fluctuation when cleaning InnoDB temporary tables.</li> <li>○ Fixed the read-only performance decrease when the instance has many cores.</li> <li>○ Fixed the error (error code: 1032) caused by hash scans.</li> <li>○ Fixed concurrency security issues caused by hotspot update.</li> </ul> </li> </ul>
20200630	8.0.18	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported async drop of big tables. You can clear files asynchronously and slowly to avoid business performance fluctuation caused by dropping big tables. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported automatic killing of idle tasks to reduce resource conflicts. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported transparent data encryption (TDE).</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where switch failed due to inconsistent positions between <code>`relay_log_pos`</code> and <code>`master_log_pos`</code>.</li> <li>○ Fixed the data file error caused by asynchronously storing data in the disk.</li> </ul> </li> </ul>

- Fixed the hard error when `fsync` returned `EIO` and retries were made repeatedly.
- Fixed the crash caused by phrase search under multi-byte character sets in full-text index.

## MySQL 5.7 Kernel Version Release Notes



TXSQL Kernel Version	Community Version	Description
20250803	5.7.44	<ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue of allowing insertion of incompatible GEOMETRY type data when the INSERT ... SELECT statement is used.</li> </ul> </li> </ul>
20250802	5.7.44	<div style="border: 1px solid #00aaff; padding: 10px; margin-bottom: 10px;"> <p><b>Note:</b> The OOM probability of the instance can be reduced by using this kernel minor version. It is strongly recommended to upgrade to this version or later versions.</p> </div> <ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Reduced the number of generated memory fragments to reduce the OOM probability of the instance.</li> </ul> </li> </ul>
20250730	5.7.44	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported displaying the query plan with the Iterator tree structure in JSON format for the explain command.</li> <li>○ Supported encoding query plans and displaying the query plan ID in slow logs and the explain result.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue of missing sql_mode in subqueries of parallel queries.</li> <li>○ Fixed the issue of crash caused by the SIGABRT signal due to the std::out_of_range exception during InnoDB statistics collection under high concurrency.</li> </ul> </li> </ul>
20250510	5.7.44	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Added the capability of skipping transactions with GTID when the secondary server is abnormal, reducing the probability of synchronization interruption due to unexpected operations on read-only or disaster recovery instances.</li> </ul> </li> </ul>

		<ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Optimized the error log generation logic in some scenarios to reduce useless logs.</li> </ul> </li> </ul>
20240331	5.7.44	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Merged Official <a href="#">5.7.36</a>, <a href="#">5.7.37</a>, <a href="#">5.7.38</a>, <a href="#">5.7.39</a>, <a href="#">5.7.40</a>, <a href="#">5.7.41</a>, <a href="#">5.7.42</a>, <a href="#">5.7.43</a>, <a href="#">5.7.44</a> Changes.</li> <li>○ Added update returning feature.</li> <li>○ Added a DDL progress display. During DDL operations, the <code>show detail processlist</code> command can be run to view it.</li> <li>○ Added the default table encryption feature.</li> </ul> </li> <li>• Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the SQL_TYPE audit accuracy.</li> <li>○ Optimized index drill-down pruning.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed an issue where insert returning statements caused slave crashes.</li> <li>○ Fixed an issue where the slave replay thread got stuck due to MASTER_DELAY.</li> <li>○ Fixed an issue where setting custom variables using a function caused session track errors.</li> <li>○ Fixed an issue where hash scans consumed a large amount of memory.</li> <li>○ Fixed an inconsistency when dropping/truncating a non-existent table with the recycle bin enabled versus disabled.</li> <li>○ Fixed an issue where pulling binlog resulted in garbled characters when binlog_checksum was disabled.</li> <li>○ Fixed the Parallel Copy DDL bug.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>• New Features <ul style="list-style-type: none"> <li>○ Supports persistence for flashback query.</li> <li>○ Supported drop table force, enabling drop innodb metadata.</li> <li>○ Supported Parallel Copy DDL.</li> <li>○ Supported limit in subquery.</li> <li>○ Supported the conversion of partition tables from MyISAM to InnoDB.</li> </ul> </li> <li>• Bug Fixes</li> </ul>

20230601	5.7.36	<ul style="list-style-type: none"> <li>• <b>Bug Fixes</b> <ul style="list-style-type: none"> <li>○ Fixed the issue of index anomaly in primary-secondary BP synchronization feature.</li> <li>○ Fixed the issue where killing connections during large transactions caused anomalies.</li> <li>○ Fixed the issue of obtaining user-defined variable string errors in session track.</li> <li>○ Fixed the issue of failure to create index when parallel DDL is enabled and innodb_disable_sort_file_cache is set.</li> <li>○ Fixed some errors with instant modify column.</li> </ul> </li> </ul>
20230115	5.7.36	<ul style="list-style-type: none"> <li>• <b>New Features</b> <ul style="list-style-type: none"> <li>○ Supported Nonblocking DDL feature.</li> <li>○ Supported validate password plugin.</li> <li>○ Supported for storing historical deadlock information.</li> </ul> </li> <li>• <b>Performance Optimization</b> <ul style="list-style-type: none"> <li>○ Asynchronous deletion of large tables: Temporary tables also use the innodb_async_table_size filter table, and only tables exceeding innodb_async_table_size are deleted asynchronously, improving the processing efficiency.</li> </ul> </li> <li>• <b>Bug Fixes</b> <ul style="list-style-type: none"> <li>○ Fixed the issue where creating a user with grant identified by failed, causing primary/standby interruption.</li> <li>○ Fixed the issue where GROUP_CONCAT did not correctly set USED_TABLES when the DERIVED_MERGE switch was enabled.</li> <li>○ Fixed the issue where the gtid_subset function failed to correctly handle null_value.</li> <li>○ Fixed the issue where dummy index cache failed to initialize system columns.</li> <li>○ Fixed the issue of instant add column in partition table exceeding the maximum number of columns.</li> <li>○ Fixed the issue where canal pulling binlog may cause OOM.</li> <li>○ Fixed the error in proxy when reusing connections with different users.</li> <li>○ Fixed the proxy's incorrect responses for row count. found rows. and db settings.</li> </ul> </li> </ul>

		<ul style="list-style-type: none"> <li>○ Fixed the issue where the error messages during binlog sending and receiving were incomplete.</li> <li>○ Fixed anomalies in paging and pushdown calculations.</li> <li>○ Fixed potential invalidity of m_page after creating a subtree with Parallel DDL.</li> <li>○ Fixed the crash issue with instant modify under certain character sets.</li> </ul>
20220716	5.7.36	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported auto-increment column persistence for InnoDB.</li> <li>○ Supported precise memory statistics.</li> <li>○ Supported query-level memory monitoring.</li> <li>○ Supported recycle bin.</li> <li>○ Supported parallel DDL statements.</li> <li>○ Supported flashback query.</li> <li>○ Supported async rollback for internal XA transactions.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized async drop of big tables. The original definition of big table is 50 GB, which can now be controlled by the <code>`innodb_async_table_size`</code> to make it more flexible.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where <code>`ERROR 1878 (HY000): Temporary file write failure`</code> was reported when <code>`alter table`</code> was executed to create indexes.</li> <li>○ Fixed the issue where <code>buf/buf/pool</code> couldn't be viewed in PFS memory monitoring data.</li> <li>○ Fixed the issue where the returning statement might cause exceptions in some scenarios due to permission checks.</li> <li>○ Fixed the issue where an error was reported because the parser did not correctly handle semicolons in statements.</li> <li>○ Fixed the issue where single quotation marks in audit statements were not escaped.</li> <li>○ Fixed the issue of sudden memory usage increase on the ARM platform.</li> <li>○ Fixed the issue of source-replica inconsistency caused by modifying <code>`binlog_format`</code> after <code>writeset</code></li> </ul> </li> </ul>

		<p>caused by modifying <code>binlog_format</code> after <code>binlog</code> was enabled.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue of high CPU usage caused by exiting a large number of threads at the same time.</li> <li>○ Fixed bugs related to <code>`drop table partition force`</code>.</li> <li>○ Fixed the issue where binlog dump got stuck and caused the instance restart to become slow.</li> <li>○ Fixed the issue where the source-replica sync failed because <code>`create table like temporary table`</code> did not inherit the character set in the binlog.</li> <li>○ Fixed the issue where <code>`show detail processlist`</code> displayed illegal characters.</li> <li>○ Fixed the issue where the <code>`thread_group`</code> lock was not released when the thread pool was closed in some cases.</li> <li>○ Fixed the issue where updating the parent table at the parallel table level caused the instance to run abnormally.</li> <li>○ Fixed the issue where virtual columns were calculated incorrectly on the replica.</li> <li>○ Fixed the issue where <code>`gtid_subset`</code> did not set <code>`null_value`</code> to <code>`false`</code> after executing a row.</li> </ul>
20211230	5.7.36	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported the official updates of MySQL <a href="#">5.7.19-5.7.36</a>.</li> <li>○ Supported source-replica buffer pool sync to speed up the performance recovery after HA switch (around 90 seconds faster than that in native mode).</li> <li>○ Added the backup lock feature to provide lightweight metadata locks to improve the service availability during backup.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Made functions related to <code>`utf8/utf8mb4 my_charpos`</code> inline to optimize the performance of UTF_8 functions in read_write scenarios.</li> <li>○ Upgraded jemalloc to v5.2.1.</li> <li>○ Optimized file number acquisition during binlog rotation.</li> <li>○ Optimized semi-sync replica I/O.</li> <li>○ Optimized hash scan aggregation.</li> <li>○ Accelerated the startup of crash recovery for large transactions.</li> </ul> </li> </ul>

20211102	5.7.18	<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Fixed the exception of the third-party data subscription tool caused by subscription to the comparison SQL for internal data consistency during tool usage.</li> </ul> </li> </ul> <div style="border: 1px solid #add8e6; padding: 10px; margin-top: 10px;"> <p><b>ⓘ Note:</b> After the database instance is migrated, upgraded, or recovered after failure, the system will compare the data to ensure data consistency. When comparison SQL is in `statement` mode, exceptions are prone to occur in response of some third-party subscription tools to the SQL in `statement` mode. When the instance is upgraded to its kernel, the third-party data subscription tool can't subscribe the comparison SQL for internal data consistency.</p> </div>
		<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported writeset replication.</li> </ul> </li> <li>• Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the checkpoint mechanism to increase the backup success rate.</li> <li>○ Optimized the hash scan index selection.</li> <li>○ Optimized the hotspot update performance to support `insert on duplicate key update`.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the error of unstable performance after hotspot update was enabled.</li> <li>○ Fixed the crash caused by rolling back the UPDATE operation after an instant DDL.</li> <li>○ Fixed the issue where the `CREATE TABLE AS SELECT` statement didn't inherit the compression attribute after column compression was enabled.</li> <li>○ Fixed the instance crash caused by the `show variables like 'tencent_root%'` statement after the `skip-grant-table` option was enabled.</li> <li>○ Fixed the crash of the Query Rewriter plugin in read-only mode.</li> </ul> </li> </ul>

20211031	5.7.18	<p>read-only mode.</p> <ul style="list-style-type: none"> <li>○ Fixed the error (error code: 1032) caused by hash scans in partitioned tables.</li> <li>○ Fixed the issue where the first large transaction's SBM was 0 in MTS mode.</li> <li>○ Fixed the crash of `stop slave` caused by `slave_preserve_commit_order=ON, slave_transaction_retries=0`.</li> <li>○ Fixed several XA transaction bugs.</li> <li>○ Fixed the issue where SQL splicing went wrong during `show create` after a JSON field with a default value was created.</li> <li>○ Fixed the issue where disconnected transactions could not be rolled back after transactions were blocked.</li> <li>○ Fixed the issue where there might be no statistics information for long records in InnoDB persistent mode.</li> <li>○ Ported 8.0 to fix the issue where `ANALYZE TABLE` might cause query retention.</li> <li>○ Fixed the issue where the InnoDB statistics couldn't be synced to the server layer in time after change.</li> <li>○ Fixed the issue where statistical sampling might block writes for too long and cause a crash (bug# 31889883).</li> <li>○ Fixed the possibility of reading zero rows during the InnoDB statistics update process (bug# 105224).</li> <li>○ Fixed the possible <math>O(N^2)</math> behavior in MVCC (bug# 28825617).</li> <li>○ Fixed the crash caused by closing a temp table and triggering binlog rotation when a connection was released.</li> </ul>
		<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Added the new command SHOW SLAVE DETAIL [FOR CHANNEL channel] for displaying the binlog timestamp that the current replica has replayed.</li> <li>○ Supported transaction_read_only/transaction_isolation parameters.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Accelerated the application of hash scan on replicas by aggregating multiple identical binlog</li> </ul> </li> </ul>

20210630	5.7.18	<p>events.</p> <ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where duplicate primary keys existed, columns couldn't be found, and columns were too long in temp tables caused by the `UPDATE` statement.</li> <li>○ Fixed the issue where there might be no statistics information during the DDL process.</li> <li>○ Fixed the inaccurate undo log size in connection status statistics.</li> <li>○ Fixed the instance crash caused by querying the metadata_locks table.</li> <li>○ Modified `of` as a non-reserved keyword.</li> <li>○ Fixed the issue where the dynamically modified version number was not invalidly displayed in new connections.</li> <li>○ Fixed the issue where the wild pointer was accessed during page_cache cleaning.</li> <li>○ Fixed the issue where the execution of ALTER TABLE might report the "Incorrect key file for table" error.</li> <li>○ Fixed the excessive memory usage by partitioned tables.</li> <li>○ Fixed the issue where -1 was included in the TIME field in the result set returned by SHOW PROCESSLIST.</li> <li>○ Fixed the lock wait of XA transaction replication on replica nodes.</li> <li>○ Fixed the incorrect lock of partitioned tables in equal range query.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>• New features <ul style="list-style-type: none"> <li>○ Supported `RETURNING` clause in a `DELETE`, `INSERT`, or `REPLACE` statement to return information about the rows that were deleted or modified by the statement. For `DELETE`, undo data is returned, while for `INSERT` or `UPDATE`, redo data is returned.</li> <li>○ Supported column compression: Row compression and data page compression are already supported, but if small fields in a table are read and written frequently while big fields are not, both of the compression methods waste a lot of computing</li> </ul> </li> </ul>

20210331

5.7.18

resources. In contrast, column compression can compress big fields that are infrequently accessed and reduce the space for storing whole rows of fields, so as to improve read and write access efficiency.

- Supported querying the value of the ``character_set_client_handshake`` parameter.
- Supported the manual cleaning of page cache occupied by log files by using the ``posix_fadvise()`` function based on the sliding window technique, so as to lower the memory pressure on the operating system and improve instance stability.
- Performance optimizations
  - Optimized the parallelism of CREATE INDEX: A merge sort is needed in a temp table in the process of creating indexes, which is time-consuming. The parallel temp-table merge sort algorithm is now supported to reduce the time by more than 50%.
  - Optimized the mechanism of scanning and flushing the dirty pages tracked in the flush list, so as to solve the performance fluctuation issue while creating indexes and thus improve the system stability.
- Bug fixes
  - Fixed the memory leak issue.
  - Implemented the JSON bug fixes provided in MySQL 8.0 to improve the stability of using JSON.
  - Fixed the error (error code: 1032) caused by hash scans.
  - Fixed concurrency security issues caused by hotspot update.
  - Implemented the gcol bug fixes provided by MySQL in batches.
  - Fixed the failure to compare DateTime data with String data in some cases.
  - Fixed the bug where file handles cannot be released if source-replica buffer pool sync is enabled.
  - Fixed the deadlocks caused by setting the ``offline_mode`` parameter and creating connections at the same time.
  - Fixed the crashes caused by the ``m_end_range`` parameter incorrectly set in concurrent range

		<p>queries.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue where it takes a long time to execute an `UPDATE` statement on a temp table if a JSON column appears in the `GROUP BY` clause.</li> </ul>
20201231	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported using `NOWAIT` and `SKIP LOCKED` in `SELECT FOR UPDATE/SHARE`.</li> <li>○ Supported dynamic setting of thread pooling mode or connection pooling mode by using the `thread_handling` parameter.</li> <li>○ Supported source-replica buffer pool sync.</li> <li>○ Supported monitoring of user connection status. Monitoring items include sync/async IO, memory, log size, CPU time, and lock duration.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the transaction subsystem to improve the high concurrency performance.</li> <li>○ Optimized the time to start crash recovery for large transactions.</li> <li>○ Optimized redo log flushing.</li> <li>○ Optimized the buffer pool initialization time.</li> <li>○ Optimized UTF8/UTF8MB4 string efficiency.</li> <li>○ Optimized audit performance.</li> <li>○ Revoked the restriction on the value of `gtid_purged` being empty.</li> <li>○ Optimized the backup lock. `LOCK TABLES FOR BACKUP`, `LOCK BINLOG FOR BACKUP`, and `UNLOCK BINLOG` are supported. `FLUSH TABLES WITH READ LOCK` is used to take a backup of the database, but it blocks the whole database from providing service. In contrast, the three statements above use a lightweight backup lock to ensure data consistency during physical/logical backup while allowing the database to providing service.</li> <li>○ Optimized the `drop table` operations on big tables.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the hang issue when querying `performance_schema`.</li> <li>○ Fixed the overflow issue of the `digest_add_token` function.</li> <li>○ Fixed the crash caused by ibuf access when the</li> </ul> </li> </ul>

		<ul style="list-style-type: none"> <li>- Fixed the crash caused by ibuf access when the `TRUNCATE TABLE` command was executed.</li> <li>○ Fixed the query correctness issue caused by const propagation when `LEFT JOIN` statement is used.</li> </ul>
20200930	5.7.18	<ul style="list-style-type: none"> <li>• Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the backup lock. `FLUSH TABLES WITH READ LOCK` is used to take a backup of the database, but it blocks the whole database from providing service. Therefore, a lightweight backup lock is provided in this version.</li> <li>○ Optimized the `drop table` operations on big tables. The `innodb_fast_ahi_cleanup_for_drop_table` parameter helps significantly reduce the time it takes to clean up adaptive hash indexes when dropping big tables.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the crash caused by ibuf access when TRUNCATE TABLE was executed.</li> <li>○ Fixed cold backup failures when the quick column adding feature was enabled.</li> <li>○ Fixed performance degradation caused by frequently releasing InnoDB memory table objects.</li> <li>○ Fixed the query correctness issue caused by const propagation when `LEFT JOIN` statement is used.</li> <li>○ Fixed the core issue caused by rule class name conflict between SQL throttling and query rewrite.</li> <li>○ Fixed the concurrent update issue caused by the `INSERT ON DUPLICATE KEY UPDATE` statement in multiple sessions.</li> <li>○ Fixed the `duplicate key error` caused by concurrent INSERTs when `auto_increment_increment` is used.</li> <li>○ Fixed the crashes caused by evicting InnoDB memory objects.</li> <li>○ Fixed concurrency security issues caused by hotspot update.</li> <li>○ Fixed the coredump issue when enabling the thread pool after jemalloc was upgraded to v5.2.1.</li> <li>○ Fixed the incomplete audit log issue caused by fwrite error-free handling.</li> <li>○ Fixed the issue where `mysqld_safe` failed to print logs when it was started by a root user.</li> </ul> </li> </ul>

		<ul style="list-style-type: none"> <li>○ Fixed the increase in the size of the DDL log file caused by `ALTER TABLE EXCHANGE PARTITION`.</li> </ul>
20200701	5.7.18	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the INNOBASE_SHARE index mapping error.</li> </ul> </li> </ul>
20200630	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported using `NOWAIT` and `SKIP LOCKED` in `SELECT FOR UPDATE/SHARE` statements.</li> <li>○ Supported large transaction optimization, which can solve such problems as source-replica delay and backup failures caused by large transactions.</li> <li>○ Optimized audit performance to support async audit.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the overflow of the `digest_add_token` function.</li> <li>○ Fixed the instance crash caused by `insert blob`.</li> <li>○ Fixed the source-replica replication interruption when a hash scan failed to find the record while updating the same row in an event.</li> <li>○ Fixed the hang issue when querying `performance_schema`.</li> </ul> </li> </ul>
20200331	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Added the official MySQL 5.7.22 JSON series functions.</li> <li>○ Supported the hotspot update feature as described in <a href="#">Real-Time Session</a> for ecommerce flash sale scenarios.</li> <li>○ Supported the SQL throttling feature as described in <a href="#">Real-Time Session</a>.</li> <li>○ Supported encryption with custom KMS keys.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the crash caused by phrase search under multi-byte character sets in full-text index.</li> <li>○ Fixed the crash of the CATS lock scheduling module in high-concurrency scenarios.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported skipping the corrupted data and continuing to parse when a binlog is corrupted. If the source instance and binlog are both damaged,</li> </ul> </li> </ul>

		<p>this feature helps restore data from the replica database for use as much as possible.</p> <ul style="list-style-type: none"> <li>○ Supported syncing data from non-GTID to GTID mode.</li> <li>○ Supported querying the "user thread memory usage" by executing the `SHOW FULL PROCESSLIST` statement.</li> <li>○ Supported quick column adding for tables as described in <a href="#">Overview</a>. This feature does not replicate the data or use disk capacity/IO, and can implement changes in real time during peak hours.</li> <li>○ Supported persistent auto-increment values.</li> </ul> <ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where replication would be interrupted if the column name in a `GRANT` statement contained reserved words.</li> <li>○ Fixed the issue where SQL execution efficiency dropped when reverse scan was performed on a partitioned table.</li> <li>○ Fixed the issue where the query result had an exception due to data inconsistency when using virtual column index and primary key.</li> <li>○ Fixed the issue where data was missing due to InnoDB primary key range queries.</li> <li>○ Fixed the issue where the system crashed when a DDL statement was executed for a table with spatial indexes.</li> <li>○ Fixed the issue where source-replica disconnection occurred when the binlog size was too large and the file length in the heartbeat information exceeded the limit.</li> <li>○ Fixed the issue where other events could not be executed as scheduled when an event was deleted.</li> <li>○ Fixed the issue where the aggregate query result was incorrect.</li> </ul> </li> </ul>
20190830	5.7.18	
20190615	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported transparent data encryption (TDE).</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where null pointer reference occurred when the LONGTEXT feature was used in subqueries.</li> </ul> </li> </ul>

20190430	5.7.18	<ul style="list-style-type: none"> <li>○ Fixed the issue where source-replica disconnection occurred due to hash scan.</li> <li>○ Fixed the issue where the replica I/O thread was interrupted due to source binlog switch.</li> <li>○ Fixed the crash caused by the use of `NAME_CONST`.</li> <li>○ Fixed the illegal mix of collation error caused by character set.</li> </ul>
20190203	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported async drop of big tables. You can clear files asynchronously and slowly to avoid business performance fluctuation caused by dropping big tables. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported CATS lock scheduling.</li> <li>○ Supported creating and dropping temp tables and CTS syntax in transactions when GTID is enabled. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported implicit primary keys. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported users without super privileges to kill sessions of other users by configuring the `cdb_kill_user_extra` parameter (default value: `root@%`).</li> <li>○ Supported enterprise-grade encryption functions. To apply for this feature, <a href="#">submit a ticket</a>.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where replication was interrupted when binlog cache file ran out of space.</li> <li>○ Fixed the hard error when `fsync` returned `EIO` and retries were made repeatedly.</li> <li>○ Fixed the issue where replication was interrupted and could not be recovered due to GTID holes.</li> </ul> </li> </ul>
		<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported automatic killing of idle transactions to reduce resource conflicts. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported automatically changing the storage engine from MEMORY to InnoDB: If the global variable `cdb_convert_memory_to_innodb` is `ON`, the engine will be changed from MEMORY to InnoDB when a table is created or modified.</li> </ul> </li> </ul>

20180918	5.7.18	<ul style="list-style-type: none"> <li>○ Supported invisible indexes.</li> <li>○ Supported memory management with jemalloc, which can replace the jlibc memory management module to reduce memory usage and improve allocation efficiency.</li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized binlog switch to reduce the `rotate` lock duration and improve system performance.</li> <li>○ Accelerated the crash recovery.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where an event became invalid due to source-replica switch.</li> <li>○ Fixed the crash caused by `REPLAY LOG RECORD`.</li> <li>○ Fixed the issue where the query result was incorrect due to loose index scans.</li> </ul> </li> </ul>
20180530	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported SQL auditing.</li> <li>○ Supported table-level concurrent replication. To apply for this feature, <a href="#">submit a ticket</a>.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized replica instance locks to improve the sync performance of replica instances.</li> <li>○ Optimized the pushdown of the `SELECT ... LIMIT` statement.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where switch failed due to inconsistent positions between `relay_log_pos` and `master_log_pos`.</li> <li>○ Fixed the crash caused by `Crash on UPDATE ON DUPLICATE KEY`.</li> <li>○ Fixed the `Invalid escape character in string.` error when a JSON column was imported.</li> </ul> </li> </ul>
20171130	5.7.18	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported the `information_schema.metadata_locks` view to query the MDL grant and wait status in the current instance. Supported the `ALTER TABLE NO_WAIT   TIMEOUT` syntax to grant DDL operations wait timeout. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported thread pool. To apply for this feature,</li> </ul> </li> </ul>

[submit a ticket.](#)

- Bug fixes
  - Fixed the error of ``innodb_buffer_pool_pages_data`` parameter overflow by calculating it based on ``bytes_data``.
  - Fixed the issue where speed limit plugin became unavailable in async mode

## MySQL 5.6 Kernel Version Release Notes



TXSQL Kernel Version	Description
20220303	<ul style="list-style-type: none"> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the abnormal release when the memory allocated by <code>`mem_strdup`</code> was used for <code>`row_mysql_truncate_t::file_name`</code> during async drop of big tables.</li> </ul> </li> </ul>
20220302	<ul style="list-style-type: none"> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the memory leak issue in <code>`sql_update.cc`</code>.</li> </ul> </li> </ul>
20220301	<ul style="list-style-type: none"> <li>• New features               <ul style="list-style-type: none"> <li>○ Supported dynamically configuring the spin cycle (0-100) with the dynamic parameter <code>`innodb_spin_wait_pause_multiplier`</code>. This parameter is used for temporary adjustment and does not support fixing the change through the console.</li> <li>○ Supported printing deadlock loop information. After this feature is enabled through the parameter <code>`innodb_print_dead_lock_loop_info`</code>, when a deadlock occurs, you can run <code>`show engine innodb status`</code> to view the deadlock loop information.</li> </ul> </li> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the issue where anonymous GTID transactions were generated in memory tables after replica restart.</li> <li>○ Fixed the issue where upgrade failed due to the missing <code>`root@localhost`</code> permission.</li> <li>○ Fixed the issue where the values of monitoring variables such as <code>`innodb_row_lock_current_waits`</code> were abnormal.</li> <li>○ Fixed the SQL type mapping error in the audit plugin.</li> </ul> </li> </ul>
20211030	<ul style="list-style-type: none"> <li>• New features               <ul style="list-style-type: none"> <li>○ Supported large transaction replication optimization.</li> </ul> </li> <li>• Performance optimizations               <ul style="list-style-type: none"> <li>○ Accelerated the application of hash scan.</li> </ul> </li> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the OOM caused by a large number of table queries.</li> <li>○ Fixed the infinite loop error caused by setting <code>`innodb_thread_concurrency`</code> to 0.</li> <li>○ Fixed the issue where there were no statistics information for long records.</li> <li>○ Fixed the OOM issue caused by a large number of table queries.</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>○ Fixed the SBRM jump error.</li> <li>○ Fixed the `LOCK_binlog_end_pos hang` error.</li> </ul>
20210630	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported large transaction replication optimization.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the incorrect copy when Index Merge was enabled.</li> <li>○ Fixed the issue where the replication would be interrupted if the execution of CREATE TABLE SELECT was interrupted when `cdb_more_gtid_feature_supported` was enabled in row mode.</li> <li>○ Fixed the bug that `max(id)` was greater than AUTO_INCREMENT in SHOW CREATE TABLE.</li> </ul> </li> </ul>
20201231	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the error (error code: 1032) caused by hash scans.</li> <li>○ Fixed the issue where the source-replica auto-increment values were inconsistent due to the `REPLACE INTO` statement in `ROW` format.</li> <li>○ Fixed the memory leak caused by not freeing up the memory requested for parsing SQL statements.</li> <li>○ Fixed the issue where the sql_mode check is skipped when running `CREATE TABLE AS SELECT`.</li> <li>○ Fixed the issue where the `sql_mode` check was skipped when inserting default values.</li> <li>○ Fixed the issue where the `sql_mode` check was skipped when running UPDATE with bound parameters.</li> </ul> </li> </ul>
20200915	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported the SQL throttling feature as described in <a href="#">Real-Time Session</a>.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the initialization acceleration of buffer pool.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the hang issue of `rename table` on both source and replica.</li> <li>○ Fixed the crash when `event_scheduler` was set to `disable` and `cdb_skip_event_scheduler` was changed from `on` to `off`.</li> <li>○ Fixed the `sync_wait_array` assertion failure when the maximum number of connections of `tencentroot` was not counted in `srv_max_n_threads`.</li> <li>○ Fixed the crash of source-replica parallel replication caused</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>○ Fixed the crash of source-replica parallel replication caused by the system table structure inconsistency between TencentDB for MySQL 5.6 and other cloud vendors' MySQL 5.6.</li> <li>○ Fixed the `INSERT ON DUPLICATE KEY UPDATE THE WRONG ROW` error.</li> <li>○ Fixed the `index_mapping` error.</li> <li>○ Fixed the MTR failure.</li> <li>○ Fixed the source-replica replication interruption when a hash scan failed to find the record while updating the same row in an event.</li> </ul>
20190930	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported querying the user thread memory usage by executing the `SHOW FULL PROCESSLIST` statement.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed GTID holes caused by the replication filter of the replica.</li> <li>○ Fixed the issue where source-replica disconnection occurred when the binlog size was too large and the file length in the heartbeat information exceeded the limit.</li> <li>○ Fixed the illegal mix of collation error caused by character set.</li> <li>○ Fixed the issue where the source-replica disconnection occurred due to hash scan.</li> <li>○ Fixed the crash caused by the use of `NAME_CONST`.</li> <li>○ Fixed the issue where the replica I/O thread was interrupted due to source binlog switch.</li> <li>○ Fixed the error of incompatible backups due to `innodb_log_checusum`.</li> </ul> </li> </ul>
20190530	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where dirty data might be read in RC mode.</li> <li>○ Fixed the issue where replica instance replay might fail due to the drop of temp table.</li> <li>○ Fixed the deadlock issue under high concurrency.</li> </ul> </li> </ul>
	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported async drop of big tables. You can clear files asynchronously and slowly to avoid business performance fluctuation caused by dropping big tables. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported users without super privileges to kill sessions of other users by configuring the `cdb_kill_user_extra` parameter (default value: `root@%`).</li> </ul> </li> </ul>

20190203	<ul style="list-style-type: none"> <li>○ Supported creating and dropping temp tables and CTS syntax in transactions when GTID is enabled. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Optimized the replication and replay of partitioned tables to improve efficiency.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the source-replica data inconsistency issue caused by insufficient temporary space.</li> <li>○ Fixed the issue of suspended hot record updates.</li> <li>○ Fixed the issue where the value of `Seconds_Behind_Master` was abnormal during concurrent replication.</li> </ul> </li> </ul>
20180915	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported automatically changing the storage engine from MEMORY to InnoDB: If the global variable `cdb_convert_memory_to_innodb` is `ON`, the engine will be changed from MEMORY to InnoDB when a table is created or modified.</li> <li>○ Supported automatic killing of idle transactions to reduce resource conflicts. To apply for this feature, <a href="#">submit a ticket</a>.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the crash caused by `REPLAY LOG RECORD`.</li> <li>○ Fixed the error of time data inconsistency between source and replica due to decimal precision issues.</li> </ul> </li> </ul>
20180130	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported thread pool. To apply for this feature, <a href="#">submit a ticket</a>.</li> <li>○ Supported dynamically modifying replication filtering rules for replica nodes.</li> </ul> </li> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Reduced performance fluctuation caused by `DROP TABLE`.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the database crash caused by authentication password strings.</li> </ul> </li> </ul>
20180122	<ul style="list-style-type: none"> <li>● New features <ul style="list-style-type: none"> <li>○ Supported SQL auditing.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the integer overflow issue.</li> <li>○ Fixed the error caused by queries using full-text index.</li> </ul> </li> </ul>

	<ul style="list-style-type: none"> <li>○ Fixed the issue where the replica crashed during replication.</li> </ul>
20170830	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where binlog speed limit became invalid in async mode.</li> <li>○ Fixed the issue where the `buffer_pool` status was abnormal.</li> <li>○ Fixed the issue where `SEQUENCE` and implicit primary key conflicted.</li> </ul> </li> </ul>
20170228	<ul style="list-style-type: none"> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the character encoding bug in `DROP TABLE`.</li> <li>○ Fixed the issue where a table contained symbols like decimal points or `replicate-wild-do-table` couldn't be used to filter databases correctly.</li> <li>○ Fixed the issue where SQL threads exited too early after the replica had a `rotate` event.</li> </ul> </li> </ul>
20161130	<ul style="list-style-type: none"> <li>● Performance optimizations <ul style="list-style-type: none"> <li>○ Split the `lock_log` lock to reduce the time used by lock logs and improve the concurrency performance.</li> <li>○ Separated the ACK thread of the source to reduce the response time.</li> <li>○ Prohibited the user thread from being killed while waiting for ACK in order to prevent phantom reads.</li> <li>○ Fixed the unnecessary `lock_sync` lock when `sync_binlog != 1`.</li> </ul> </li> </ul>

# TXRocks Kernel Release Notes

Last updated: 2023-08-31 16:42:35

This document describes the version updates of the TXRocks kernel.

## ⓘ Note

- For more information on how to upgrade the minor kernel version of a TencentDB for MySQL instance, see [Upgrading Kernel Minor Version](#).
- When you upgrade the minor version, some minor versions may be under maintenance and cannot be selected. The minor versions available in the console shall prevail.

## MySQL 8.0 Kernel Version Update Notes

Minor Version	Note
20230401	<b>New features:</b> By default, the TokuDB engine in the table creation statement is converted to the RocksDB engine.

## MySQL 5.7 Kernel Version Update Notes

Minor Version	Note
20230401	<b>New features:</b> By default, the TokuDB engine in the table creation statement is converted to the RocksDB engine.

# Database Proxy Kernel Update Dynamics

Last updated: 2023-08-31 16:42:47

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

## Note

If the MySQL kernel version requirements are not met, you can upgrade the kernel version of your database first. For detailed directions, see [Upgrading Kernel Minor Version](#).

Database Proxy Version	MySQL Kernel Version Requirements	Note
1.4.1	MySQL 5.7 20211230 or later. MySQL 8.0 20221215 or later.	<ul style="list-style-type: none"> <li>• New features               <ul style="list-style-type: none"> <li>Supports transaction-level connection pool feature.</li> </ul> </li> <li>• Bug fixes               <ul style="list-style-type: none"> <li>Resolved the issue of DDL statements getting stuck on the replica.</li> </ul> </li> </ul>
1.3.7	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the routing error of the <code>select for update</code> statement in some cases.</li> <li>○ Modified the <code>select @@read_only</code> statement and made it possible to be routed to the source database. This prevents some frameworks that use <code>read_only</code> flags from misjudging the database proxy as unwritable.</li> <li>○ Fixed the database proxy node exceptions caused by a database instance HA in some scenarios.</li> </ul> </li> </ul>
1.3.4	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the issue where the <code>show processlist</code> command returned incomplete data.</li> </ul> </li> </ul>
1.3.3	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Bug fixes               <ul style="list-style-type: none"> <li>○ Fixed the issue where an error was reported when the session connection pool reused connections to send <code>change_user</code> to the backend, and the issue where the <code>PREPARE</code> statement was not correctly handled by the database proxy after a new connection was established.</li> </ul> </li> </ul>
	MySQL 5.7	

1.3.2	5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where the <code>execute</code> statement lacked a parameter type.</li> </ul> </li> </ul>
1.3.1	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Feature updates <ul style="list-style-type: none"> <li>○ Allowed instances with a weight of 0 to sustain read requests when the weight of all valid instances under the database proxy is 0.</li> <li>○ Supported the multi-AZ deployment architecture where read-only instances can be mounted across AZs.</li> <li>○ Provided the read-only mode.</li> <li>○ Supported transaction split.</li> <li>○ Supported momentary disconnection prevention, i.e., connection persistence, where the client will not be disconnected when a database instance HA switch occurs because of a scheduled task.</li> </ul> </li> </ul>
1.2.1	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Feature updates <ul style="list-style-type: none"> <li>○ Supported the <code>db_lower_case_table_names</code> parameter, indicating not to verify the letter case by default.</li> <li>○ Supported returning error messages in the query stage when errors occur during database proxy connection establishment.</li> </ul> </li> </ul>
1.1.3	MySQL 5.7 20211030 or later. MySQL 8.0 20211202 or later.	<ul style="list-style-type: none"> <li>• Feature updates <ul style="list-style-type: none"> <li>○ Supported the use of hint routing information in the <code>COM _PREPARE</code> packet preprocessed by MySQL. After hint is used in <code>PREPARE</code> to specify the routing target, subsequent <code>execute</code> packets will be sent to the specified backend node.</li> </ul> </li> <li>• Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where the frontend connection was reset immediately after source-replica switch of the source instance on the database proxy was performed.</li> <li>○ Fixed the issue where load balancing might fail when read-only instances exceeded the latency threshold. Routing will resume normally when the read-only instance latency falls below the threshold.</li> </ul> </li> </ul>

		<p>instance latency tails below the threshold.</p> <ul style="list-style-type: none"> <li>○ Fixed the issue where MySQL 8.0 might return incorrect handshake information and cause connection failures.</li> </ul>
1.1.2	<p>MySQL 5.7 20211030 or later. MySQL 8.0 20211130 or later.</p>	<ul style="list-style-type: none"> <li>● Feature updates <ul style="list-style-type: none"> <li>○ Supported MySQL 8.0.</li> <li>○ Supported the connection pool feature at the connection level, which is useful in scenarios where non-persistent connections to the database are frequently established. The database proxy will save connections and reuse them during subsequent connection establishments.</li> <li>○ Supported the reconnection feature for read-only instances. In persistent connection scenarios, when a read-only instance is restarted or added, the database proxy will automatically establish a connection and restore routing to it.</li> <li>○ The internal memory management mechanism has been updated, resulting in lower memory consumption in the new version.</li> </ul> </li> <li>● Bug fixes <ul style="list-style-type: none"> <li>○ Fixed the issue where the client connection persisted after the backend connection was closed due to timeout.</li> <li>○ Fixed the issue where the internal cache might cause excessively rapid increase of the memory utilization.</li> <li>○ Fixed the occasional issue where packets in an incorrect format were returned.</li> </ul> </li> </ul>
1.0.1	<p>MySQL 5.7 2020123</p>	<ul style="list-style-type: none"> <li>● Feature updates <ul style="list-style-type: none"> <li>○ Supported MySQL 5.7.</li> <li>○ Supported read/write separation.</li> <li>○ Supported read weight assignment in read/write separation.</li> <li>○ Supported the configuration of source-replica replication delay threshold. Routing to a read-only instance will be stopped if its delay exceeds the thresholds and will be recovered after it drops below the threshold. If the source-replica replication is interrupted, disconnected read-only instances will be removed directly.</li> <li>○ Supported the configuration of the least number of read-only instances. When read-only instances are</li> </ul> </li> </ul>

	0 or later.	<p>read-only instances. When read-only instances are removed, if the number is set to N, at least N instance(s) will be retained for routing.</p> <ul style="list-style-type: none"><li>○ Supported the failover configuration, which is enabled by default. If it is disabled and the read weight of the source instance is 0, after all read-only instances are removed, an error will be reported for read requests. If failover is enabled and the read weight of the source instance is not 0, requests will be routed to the source instance.</li><li>○ Supported using the HINT syntax to specify routing nodes.</li></ul>
--	-------------	---

# Functionality Features

## Killing Idle Transactions Automatically

Last updated: 2025-02-19 16:06:18

### Feature Overview

This feature kills transactions that have been idle for the specified time period to release resources in time.

### Supported Versions

- Kernel version: MySQL 5.6 20180915 and later.
- Kernel version: MySQL 5.7 20180918 and and later.
- Kernel version: MySQL 8.0 20200630 and and later.

### Scenarios

If a connection starts a transaction (explicitly using `begin / start transaction` or implicitly) but no new statement has been executed for the specified threshold period, the connection will be killed.

### Notes

Use the `cdb_kill_idle_trans_timeout` parameter to enable or disable the feature. If it is `0`, the feature is disabled; otherwise, a connection idle for `cdb_kill_idle_trans_timeout` or `wait_timeout` seconds, whichever is smaller, will be killed. (`wait_timeout` is a session parameter.)

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>cdb_kill_idle_trans_timeout</code>	YES	ulong	0	[0, 3153600]	If it is <code>0</code> , the feature is disabled; otherwise, a transaction idle for <code>cdb_kill_idle_trans_timeout</code> seconds will be killed.



# Parallel Replication

Last updated: 2023-08-31 16:55:27

## Feature Overview

Prior to MySQL 5.6, the source node syncs binlogs and the replica node replays binlogs, both in the single-thread mode. MySQL 5.6 and later versions support the DATABASE/LOGICAL\_CLOCK parallel replication scheme, but the granularity is too large to achieve expected parallel replication in many cases.

Tencent Cloud's TXSQL kernel team has optimized the parallel replication scheme. Table parallel replication is now supported, improving parallelism and reducing source-replica delay.

## Supported Versions

- Kernel version: MySQL 8.0 20201230 and later.
- Kernel version: MySQL 5.7 20180530 and later.
- Kernel version: MySQL 5.6 20170830 and later.

## Scenarios

This feature is suitable for use cases where optimizing the parallelism of some loads can speed up the binlog replay at the replica node, thus reducing the source-replica delay.

## Notes

**Note:** To enable parallel replication capability, set the parameter "slave\_parallel\_workers" to a non-zero value. In the high-stability parameter template of MySQL 5.6, 5.7, and 8.0, the default value for this parameter is 0; while in the high-performance parameter template, the default value for this parameter is 8.

To perform parallel replication at the table level in MySQL 5.6 and 5.7, set the "slave\_parallel\_type" to the newly added value "TABLE" when "slave\_parallel\_workers" is a non-zero value. MySQL 8.0 does not support the TABLE mode.

Additionally, the `cdb_slave_thread_status` table is added to the `information_schema` database to display the thread status of the replica node.

### MySQL 5.6 Version Related Parameter Descriptions

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
slave_parallel_type	YES	char*	SCHEMA	SCHEMA/TABLE	<p>The level of parallel replication on the replica node:</p> <ul style="list-style-type: none"> <li>• SCHEMA: Replication events of different schemas can be executed in parallel.</li> <li>• TABLE: Replication events of different tables can be executed in parallel.</li> </ul>

## MySQL 5.7 Version Parameter Descriptions

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
slave_parallel_type	YES	char*	LOGICAL_CLOCK	DATABASE/TABLE/LOGICAL_CLOCK	<p>The level of parallel replication on the replica node:</p> <ul style="list-style-type: none"> <li>• <b>DATABASE:</b> Replication events of different databases can be executed in parallel.</li> <li>• <b>TABLE:</b> Replication events of different tables can be executed in parallel.</li> <li>• <b>LOGICAL_CLOCK:</b> Replication events of the same logical clock on the host can be executed in parallel.</li> </ul>

### MySQL 8.0 Version Related Parameter Descriptions

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
slave_parallel_type	YES	char*	LOGICAL_CLOCK	DATABASE/LOGICAL_CLOCK	<p>The level of parallel replication on the replica node:</p> <ul style="list-style-type: none"> <li>• DATABASE: Replication events of different databases can be executed in parallel.</li> <li>• LOGICAL_CLOCK: Replication events of the same logical clock on the host can be executed in parallel.</li> </ul>

# Dynamic thread pool

Last updated: 2023-08-31 17:00:54

## Feature Overview

The thread pool (Thread\_pool) uses a certain number of worker threads to process connection requests, which is typically used in scenarios with online transaction processing (OLTP) workloads. However, when many requests are slow queries, worker threads will be blocked by high-latency operations and fail to quickly respond to new requests. As a result, the system throughput of the thread pool mode is lower than that of the traditional one-thread-per-connection (Per\_thread) mode.

The Per\_thread and Thread\_pool modes have their advantages and disadvantages, so the system needs to flexibly switch between them based on business types. Unfortunately, the mode switch must be completed by restarting the server (during peak hours in most cases), adversely affecting the business.

To allow users to flexibly switch between Per\_thread and Thread\_pool, TencentDB for MySQL has introduced the optimization of thread pool dynamic switch, that is, to enable or disable the thread pool without restarting the database service.

## Supported Versions

- Kernel version: MySQL 8.0 20201230 and later.
- Kernel version: MySQL 5.7 20201230 and later.

## Scenarios

This feature is suitable for the business which is sensitive to performance and needs to flexibly change the database working mode based on the business type.

## Impact on Performance

- Switching from the thread pool mode to the one-thread-per-connection mode won't block queries or affect database performance.
- Switching from the one-thread-per-connection mode to the thread pool mode under extremely high QPS and persistent high pressure may block requests because the thread pool is disabled before the switch.
  - **Solution 1:** you can increase the `thread_pool_oversubscribe` parameter and decrease the `thread_pool_stall_limit` parameter to quickly enable the thread pool. After the blocked SQL queries are processed, you can restore the parameters to their original values as needed.

- Solution 2: if SQL queries start to be blocked, you can suspend or reduce service traffic for a few seconds, wait for thread pool enablement to complete, and then resume the continuous high-pressure service traffic.

## Notes

You can use the `thread_handling_switch_mode` parameter to control whether to dynamically change the thread working mode. Parameter values are described as follows:

Value	Description
disabled	The mode cannot be changed dynamically.
stable	The mode can only be changed for new connections.
fast	(Default value) The mode can be changed for new connections and new requests.
sharp	Active connections will be killed in order to force the user to reconnect so that the mode can be changed quickly.

The `show threadpool status` command displays the following new status:

- `connections_moved_from_per_thread`: the number of connections switched from `Per_thread` to `Thread_pool`.
- `connections_moved_to_per_thread`: the number of connections switched from `Thread_pool` to `Per_thread`.
- `events_consumed`: the total number of events consumed by the worker thread group in each thread pool. After the thread working mode is switched from `Thread_pool` to `Per_thread`, the total number of events won't increase any more.
- `average_wait_usecs_in_queue`: the average time each event waits in the queue.

The following status has been added to the `show full processlist` command:

- `Moved_to_per_thread`: the number of times that the connection is switched to `Per_thread`.
- `Moved_to_thread_pool`: the number of times that the connection is switched to `Thread_pool`.

## Parameter Status Description

Thread pool parameters are described as follows:

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>thread_pool_idle_timeout</code>	Yes	uint	60	[1, UINT_MAX]	If there are no network event to handle, the worker thread will wait for <code>thread_pool_idle_timeout</code> seconds before being killed.
<code>thread_pool_over_subscribe</code>	Yes	uint	3	[1,1000]	The maximum number of worker threads allowed in a worker thread group
<code>thread_pool_size</code>	Yes	uint	The number of CPUs on the current machine	[1,1000]	Number of security groups.
<code>thread_pool_stall_limit</code>	Yes	uint	500	[10, UINT_MAX]	The timer thread checks all thread groups every <code>thread_pool_stall_limit</code> milliseconds. If a thread group has no listener, neither high nor low priority queue is empty, and no new I/O network event occurs, the thread group is considered stalled. The timer thread will wake up a sleeping thread or create a new thread to relieve the thread group's pressure.
<code>thread_pool_max</code>	Yes	uint	100000	[1,100000]	The total number of worker threads in the

<code>_threads</code>				<code>UJ</code>	thread pool
<code>thread_pool_high_prio_mode</code>	Yes, session	enum	transactions	transactions\statement\none	<p>The high priority queue supports three modes:</p> <ul style="list-style-type: none"> <li><b>transactions:</b> Only when there is one statement that already starts a transaction and the ticket (<code>thread_pool_high_prio_tickets</code>) held by the connection of this statement is not zero, the connection can enter the high priority queue. After a connection has entered the priority queue <code>thread_pool_high_prio_tickets</code> times, it will be placed into the ordinary queue.</li> <li><b>statement:</b> All connections enter the high priority queue.</li> <li><b>none:</b> All connections enter the low priority queue.</li> </ul>
<code>thread_pool_high_prio_tickets</code>	Yes, session	uint	UINT_MAX	[0, UINT_MAX]	The value of the ticket assigned to each connection in <code>transactions</code> mode
<code>threadpool_work_around_epoll_bug</code>	Yes	bool	false	true/false	Whether to bypass the epoll bug of Linux 2.x (which was fixed in Linux 3)

The `show threadpool status` command displays the following status:

Status	Note
groupid	Thread group ID
connection_count	The number of user connections in the thread group
thread_count	The number of worker threads in the thread group
havelistener	Whether the thread group has a listener
active_thread_count	The number of active worker threads in the thread group
waiting_thread_count	The number of worker threads calling wait_begin in the thread group
waiting_threads_size	The number of sleeping worker threads waiting to be woken up in the thread group when there is no network event to handle (such worker threads will wait for <code>thread_pool_idle_timeout</code> seconds before being automatically killed).
queue_size	The length of the ordinary queue of the thread group
high_prio_queue_size	The length of the high priority queue of the thread group
get_high_prio_queue_num	The total number of times that events in the thread group are removed from the high priority queue
get_normal_queue_num	The total number of times that events in the thread group are removed from the ordinary queue
create_thread_num	The total number of worker threads created in the thread group
wake_thread_num	The total number of worker threads in the thread group awakened from the waiting_threads queue
oversubscribed_num	The number of times that worker threads are ready to go to sleep because the thread group is oversubscribed
mysql_cond_timedwait_num	The total number of times that worker threads in the thread group enter the waiting_threads queue
check_stall_nolistener	The total number of times that no listener is detected in the thread group in the stall check performed by the timer thread
check_stall_stall	The total number of times that the thread group is considered stalled in the stall check performed by the timer thread

<code>max_req_latency_us</code>	The maximum time in milliseconds for a user connection to wait in the queue in the thread group
<code>conns_timeout_killed</code>	The total number of times that user connections in the thread group are killed because there has been no new message on the client for the threshold period
<code>connections_moved_in</code>	The total number of connections migrated from other thread groups to this thread group
<code>connections_moved_out</code>	The total number of connections migrated from this thread group to other thread groups
<code>connections_moved_from_per_thread</code>	The total number of connections switched from the one-thread-per-connection mode to this thread group
<code>connections_moved_to_per_thread</code>	The total number of connections switched from this thread group to the one-thread-per-connection mode
<code>events_consumed</code>	The total number of events processed by the thread group
<code>average_wait_usecs_in_queue</code>	The average waiting time of all events in the queue in the thread group

# NOWAIT

Last updated: 2023-08-31 17:01:10

## Feature Overview

- DDL statements support `NO_WAIT` and `WAIT` options. If a DDL statement with `WAIT` enabled fails to obtain an MDL lock, it will wait for `WAIT` seconds before it directly returns the query result. If a DDL statement with `NO_WAIT` enabled, it will directly return the query result without waiting for the MDL lock.
- `SELECT FOR UPDATE` statements support `NOWAIT` and `SKIP LOCKED` options. If target rows are locked by another transaction, a `SELECT FOR UPDATE` statement is supposed to wait for the transaction to release the lock. But in some use cases like flash sales, you do not want to wait for a lock. You can use `SKIP LOCKED` to skip locked rows (as a result, the locked rows won't be returned in the query result set) or `NOWAIT` to return an error without waiting for the lock.

Note that `NO_WAIT` and `NOWAIT` are different keywords.

## Supported Versions

- `NO_WAIT` and `WAIT` in DDL statements are supported in kernel version MySQL 5.7 20171130 and above.
- `NOWAIT` and `SKIP LOCKED` in `SELECT FOR UPDATE` statements are supported in kernel version MySQL 5.7 20200630 and above (not just limited to MySQL 8.0 that natively supports the feature).

## Scenarios

Currently, DevAPI/XPlugin does not support using `SKIP LOCKED` or `NOWAIT` in `SELECT FOR UPDATE/SHARE` statements. Note that `NO_WAIT` in DDL statements and `NOWAIT` in `SELECT FOR UPDATE` statements are different keywords for historical reasons.

## Notes

### SELECT FOR UPDATE NOWAIT/SKIP LOCKED

```
#####session 1#####
MySQL [test]> create table t1(seat_id int, state int, primary
key(seat_id)) engine=innodb;
Query OK, 0 rows affected (0.03 sec)
```

```
MySQL [test]> INSERT INTO t1 VALUES(1,0), (2,0), (3,0), (4,0);
Query OK, 4 rows affected (0.01 sec)
Records: 4  Duplicates: 0  Warnings: 0

MySQL [test]> begin;
Query OK, 0 rows affected (0.01 sec)

MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR SHARE;
+-----+-----+
| seat_id | state |
+-----+-----+
|      1 |     0 |
|      2 |     0 |
+-----+-----+
2 rows in set (0.00 sec)

#####session 2#####
MySQL [test]> SET SESSION innodb_lock_wait_timeout=1;
Query OK, 0 rows affected (0.00 sec)

MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR UPDATE;
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting
transaction
MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR UPDATE
NOWAIT;
ERROR 5010 (HY000): Do not wait for lock.
MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR UPDATE SKIP
LOCKED;
+-----+-----+
| seat_id | state |
+-----+-----+
|      3 |     0 |
|      4 |     0 |
+-----+-----+
2 rows in set (0.00 sec)

MySQL [test]> SELECT * FROM t1 WHERE seat_id > 0 LIMIT 2 FOR UPDATE
NOWAIT;
ERROR 5010 (HY000): Do not wait for lock.
MySQL [test]> SELECT * FROM t1 WHERE seat_id > 0 LIMIT 2 FOR UPDATE SKIP
LOCKED;
```

```
+-----+-----+
| seat_id | state |
+-----+-----+
|      3 |     0 |
|      4 |     0 |
+-----+-----+
2 rows in set (0.00 sec)

MySQL [test]> commit;
Query OK, 0 rows affected (0.00 sec)
```

## SELECT FOR SHARE NOWAIT/SKIP LOCKED

```
#####session 1#####
MySQL [test]> begin;
Query OK, 0 rows affected (0.01 sec)

MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR UPDATE;
+-----+-----+
| seat_id | state |
+-----+-----+
|      1 |     0 |
|      2 |     0 |
+-----+-----+
2 rows in set (0.00 sec)

#####session 2#####
MySQL [test]> SET SESSION innodb_lock_wait_timeout=1;
Query OK, 0 rows affected (0.00 sec)

MySQL [test]> begin;
Query OK, 0 rows affected (0.00 sec)

MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 LOCK IN SHARE
MODE;
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting
transaction
MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR SHARE;
```

```
ERROR 1205 (HY000): Lock wait timeout exceeded; try restarting
transaction
MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR SHARE NOWAIT;
ERROR 5010 (HY000): Do not wait for lock.
MySQL [test]> SELECT * FROM t1 WHERE state = 0 LIMIT 2 FOR SHARE SKIP
LOCKED;
+-----+-----+
| seat_id | state |
+-----+-----+
|      3 |     0 |
|      4 |     0 |
+-----+-----+
2 rows in set (0.00 sec)

MySQL [test]> commit;
Query OK, 0 rows affected (0.00 sec)
```

## NO\_WAIT and WAIT in DDL statements

```
ALTER TABLE table [NO_WAIT | WAIT [n]] command;
DROP TABLE table [NO_WAIT | WAIT [n]];
TRUNCATE TABLE table [NO_WAIT | WAIT [n]];
OPTIMIZE TABLE table [NO_WAIT | WAIT [n]];
RENAME TABLE table_src [NO_WAIT | WAIT [n]] TO table_dst;
CREATE INDEX index ON table.columns [NO_WAIT | WAIT [n]];
CREATE FULLTEXT INDEX index ON table.columns [NO_WAIT | WAIT [n]];
CREATE SPATIAL INDEX index ON table.columns [NO_WAIT | WAIT [n]];
DROP INDEX index ON table [NO_WAIT | WAIT [n]];
```

# RETURNING

Last updated: 2023-08-31 17:01:26

## Feature Overview

In some scenarios, you need to retrieve the rows manipulated by DML statements. There are generally two ways to do so:

- Add a SELECT statement after the DML statement if the transaction is enabled.
- Use a trigger or other complex operations.

However, running a SELECT statement increases query costs, and creating a trigger makes SQL implementation more complex and inflexible.

Therefore, TXSQL supports the RETURNING keyword to optimize such scenarios. The above requirements can be flexibly and efficiently met by appending RETURNING to a DML statement.

## Supported Versions

- Kernel version: MySQL 5.7 20210330 and above.
- Kernel version MySQL 8.0 20220330 or later.

## Scenarios

MySQL 5.7 20210330 and above support INSERT ... RETURNING, REPLACE ... RETURNING, and DELETE ... RETURNING. The RETURNING keyword returns all rows that have been manipulated by an INSERT/REPLACE/DELETE statement. RETURNING can also be used in prepared statements and stored procedures.

In MySQL 8.0 20220330 or later kernel versions, the following syntaxes are supported: DELETE ... RETURNING, INSERT ... RETURNING, REPLACE ... RETURNING, and UPDATE ... RETURNING. These can return the rows manipulated by the respective statement.

Notes:

1. For DELETE ... RETURNING, the returned data rows are pre-images, while for INSERT/REPLACE ... RETURNING, they are post-images.
2. For INSERT/REPLACE ... RETURNING, columns in the outer table are currently invisible to the subquery in the RETURNING clause.
3. INSERT/REPLACE ... RETURNING only returns the value of `last_insert_id()` before the statement is executed successfully. To obtain the true value of `last_insert_id()`, you should use RETURNING to return the auto-increment column ID of the table.

## Notes

## INSERT... RETURNING

```
MySQL [test]> CREATE TABLE t1 (id1 INT);
Query OK, 0 rows affected (0.04 sec)

MySQL [test]> CREATE TABLE t2 (id2 INT);
Query OK, 0 rows affected (0.03 sec)

MySQL [test]> INSERT INTO t2 (id2) values (1);
Query OK, 1 row affected (0.00 sec)

MySQL [test]> INSERT INTO t1 (id1) values (1) returning *, id1 * 2, id1
+ 1, id1 * id1 as alias, (select * from t2);
+-----+-----+-----+-----+-----+
| id1 | id1 * 2 | id1 + 1 | alias | (select * from t2) |
+-----+-----+-----+-----+-----+
| 1 | 2 | 2 | 1 | 1 |
+-----+-----+-----+-----+-----+
1 row in set (0.01 sec)

MySQL [test]> INSERT INTO t1 (id1) SELECT id2 from t2 returning id1;
+-----+
| id1 |
+-----+
| 1 |
+-----+
1 row in set (0.01 sec)
```

## REPLACE ... RETURNING

```
MySQL [test]> CREATE TABLE t1(id1 INT PRIMARY KEY, val1 VARCHAR(1));
Query OK, 0 rows affected (0.04 sec)

MySQL [test]> CREATE TABLE t2(id2 INT PRIMARY KEY, val2 VARCHAR(1));
Query OK, 0 rows affected (0.03 sec)

MySQL [test]> INSERT INTO t2 VALUES (1, 'a'), (2, 'b'), (3, 'c');
Query OK, 3 rows affected (0.00 sec)
Records: 3 Duplicates: 0 Warnings: 0
```

```
MySQL [test]> REPLACE INTO t1 (id1, val1) VALUES (1, 'a');
Query OK, 1 row affected (0.00 sec)

MySQL [test]> REPLACE INTO t1 (id1, val1) VALUES (1, 'b') RETURNING *;
+-----+-----+
| id1 | val1 |
+-----+-----+
| 1 | b |
+-----+-----+
1 row in set (0.01 sec)
```

## DELETE ... RETURNING

```
MySQL [test]> CREATE TABLE t1 (a int, b varchar(32));
Query OK, 0 rows affected (0.04 sec)
```

```
MySQL [test]> INSERT INTO t1 VALUES
-> (7, 'ggggggg'), (1, 'a'), (3, 'ccc'),
-> (4, 'dddd'), (1, 'A'), (2, 'BB'), (4, 'DDDD'),
-> (5, 'EEEEEE'), (7, 'GGGGGGG'), (2, 'bb');
Query OK, 10 rows affected (0.03 sec)
Records: 10 Duplicates: 0 Warnings: 0
```

```
MySQL [test]> DELETE FROM t1 WHERE a=2 RETURNING *;
+-----+-----+
| a | b |
+-----+-----+
| 2 | BB |
| 2 | bb |
+-----+-----+
2 rows in set (0.01 sec)
```

```
MySQL [test]> DELETE FROM t1 RETURNING *;
+-----+-----+
| a | b |
+-----+-----+
| 7 | ggggggg |
| 1 | a |
+-----+-----+
```

```
| 3 | ccc |
| 4 | dddd |
| 1 | A |
| 4 | DDDD |
| 5 | EEEEE |
| 7 | GGGGGG |
+-----+-----+
8 rows in set (0.01 sec)
```

## Stored procedure

```
MySQL [test]> CREATE TABLE t (id INT);
Query OK, 0 rows affected (0.03 sec)

MySQL [test]> delimiter $$
MySQL [test]> CREATE PROCEDURE test(in param INT)
  -> BEGIN
  ->     INSERT INTO t (id) values (param) returning *;
  -> END$$
Query OK, 0 rows affected (0.00 sec)
MySQL [test]> delimiter ;

MySQL [test]> CALL test(100);
+-----+
| id |
+-----+
| 100 |
+-----+
1 row in set (0.01 sec)

Query OK, 0 rows affected (0.01 sec)
```

# Column Compression

Last updated: 2025-07-22 19:05:23

## Feature Overview

Row compression and data page compression are already supported, but if small fields in a table are read and written frequently while big fields are not, both of the compression methods waste a lot of computing resources.

In contrast, column compression can compress big fields that are infrequently accessed while ignoring the frequently accessed small fields, which not only reduces the space for storing whole rows of fields but also improves the read and write access efficiency.

For example, in the employee table `create table employee(id int, age int, gender boolean, other varchar(1000) primary key (id))`, if access is frequent for small fields such as `id`, `age`, and `gender` but infrequent for the large field `other`, you can compress the `other` column. Generally, only read/write of the `other` column rather than other columns will trigger the compression and decompression of this column, which further reduces the size of the stored row data. In this way, frequently accessed small fields can be accessed more quickly, while infrequently accessed large fields can be compress to use less storage space.

## Supported Versions

Kernel version: MySQL5.7 20210330 and later.

### Note

Column compression for disabled by default. To enable it, [submit a ticket](#).

## Scenarios

If a table has many frequently accessed small fields and infrequently accessed large fields, you can compress the large field columns.

## Notes

### Supported data types

- `BLOB` (including `TINYBLOB`, `MEDIUMBLOB`, and `LOBLOB`)
- `TEXT` (including `TINYTEXT`, `MEDIUMTEXT`, and `LONGTEXT`)
- `VARCHAR`
- `VARBINARY`

**Note**

The maximum length supported for `LONGBLOB` and `LONGTEXT` is  $2^{32}-2$ , which is one byte less than the official [String Type Storage Requirements](#) support of  $2^{32}-1$ .

## Supported DDL syntax types

Different from the [table creation syntax](#) of native MySQL, the definition of `COLUMN_FORMAT` in `column_definition` is changed in TencentDB for MySQL. In addition, column compression is supported only for tables with the InnoDB storage engine.

```
column_definition:
    data_type [NOT NULL | NULL] [DEFAULT default_value]
        [AUTO_INCREMENT] [UNIQUE [KEY]] [[PRIMARY] KEY]
        [COMMENT 'string']
        [COLLATE collation_name]
        [COLUMN_FORMAT {FIXED|DYNAMIC|DEFAULT}|COMPRESSED=[zlib]] #
COMPRESSED keyword for compressed columns
        [STORAGE {DISK|MEMORY}]
        [reference_definition]
```

Below is a simple example:

```
CREATE TABLE t1(
    id INT PRIMARY KEY,
    b BLOB COMPRESSED
);
```

Here, as the compression algorithm is not specified, the `zlib` algorithm will be selected by default. You can also specify the compression algorithm keyword, but only `zlib` is supported currently.

```
CREATE TABLE t1(
    id INT PRIMARY KEY,
    b BLOB COMPRESSED=zlib
);
```

The following DDL syntaxes are supported:

### CREATE TABLE:

DDL	Whether the Compression Attribute is Inherited
<code>CREATE TABLE t2 LIKE t1;</code>	Y
<code>CREATE TABLE t2 SELECT * FROM t1;</code>	Y
<code>CREATE TABLE t2(a BLOB) SELECT * FROM t1;</code>	N

### Alter table aspects:

DDL	Description
<code>ALTER TABLE t1 MODIFY COLUMN a BLOB;</code>	Alters a compressed column into a non-compressed one
<code>ALTER TABLE t1 MODIFY COLUMN a BLOB COMPRESSED;</code>	Alters a non-compressed column into a compressed one

### New variable description

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>innodb_column_compression_zlib_wrap</code>	Yes	bool	TRUE	true/false	If it is set to <code>TRUE</code> , the data zlib header and tail will be generated, and Adler-32 check will be performed.
<code>innodb_column_compression_zlib_strategy</code>	Yes	Integer	0	[0,4]	Column compression policy. Valid values: 0: <code>Z_DEFAULT_STRATEGY</code> ; 1: <code>Z_FILTERED</code> ; 2: <code>Z_HUFFMAN_ONLY</code> ; 3: <code>Z_RLE</code> ; 4: <code>Z_FIXED</code> . Generally, <code>Z_DEFAULT_STRATEGY</code> is the best choice for text data, while <code>Z_RLE</code> for image data.
<code>innodb_column_compression_zlib_level</code>	Yes	Integer	6	[0,9]	Column compression level. Value range: 0-9. 0 indicates not to compress. The higher the value, the smaller the data size after compression, and the longer the compression duration.
<code>innodb_column_compression_threshold</code>	Yes	Integer	256	[0, 0xffffffff]	Column compression threshold in bytes. Value range: 1-0xffffffff. Only data whose length is at or above this threshold will be compressed; otherwise, the original data will stay unchanged with only a compression header added.
					The compression ratio for column compression ranges from 1 to 100, with the unit being a percentage. Data will only be

innodb_column_compression_pct	Yes	Integer	100	[1, 100]	compressed if the <b>compressed data size / uncompressed data size</b> is lower than this value; otherwise, the original data remains unchanged, and only a compression header is added.
-------------------------------	-----	---------	-----	----------	--

**Note**

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

## New status description

Name	Local Disk Types	Note
<code>Innodb_column_compressed</code>	Integer	Number of column compressions, including compressions for non-compressed data and compressed data.
<code>Innodb_column_decompressed</code>	Integer	Number of column decompressions, including decompressions for non-compressed data and compressed data.

## New error description

Name	Description	Note
<code>Compressed column '%-.192s'</code> <code>can't be used in key specification</code>	Name of the column specified for compression	The compression attribute cannot be specified for a column with an index.
<code>Unknown compression method: %s"</code>	Name of the compression algorithm specified in the DDL statement	An invalid compression algorithm other than <code>zlib</code> is specified in the <code>CREATE TABLE</code> or <code>ALTER TABLE</code> statement.
<code>Compressed column '%-.192s'</code> <code>can't be used in column format specification</code>	Name of the column specified for compression	If the <code>COLUMN_FORMAT</code> attribute has been specified for a column, other attributes cannot be specified, and <code>COLUMN_FORMAT</code> can be used only in NDB.
<code>Alter table ... discard/import tablespace not support column compression</code>	\	The <code>ALTER TABLE ... DISCARD/IMPORT TABLESPACE</code> statement cannot be executed for tables with column compression enabled.

## Performance

The performance varies by DDL and DML statements:

For DDL statements, sysbench is used for testing:

- Column compression compromises much performance of DDL statements with the COPY algorithm, and the performance after compression is 7-8 times lower than before.
- The impact of column compression on INPLACE DDL statements is subject to the data volume after compression. If the overall data size is reduced after compression, the DDL performance will be improved; otherwise, it will be compromised.

- Column compression almost has no impact on INSTANT DDL statements.

For DML statements, in an 8-column table with the most common compression ratio of 1:1.8 (where the length of the inserted data varies randomly from 1 to 6,000, the inserted characters are random within 09 and a-b, a column contains a large volume of varchar data, and the data types of other columns are either char(60) or int), the performance of insertion, deletion, and query of non-compressed columns in this table is improved by below 10%, but the performance of update of non-compressed and compressed columns is reduced by below 10% and 15% respectively. This is because that TencentDB for MySQL first reads the value of a row and then writes the updated value, which triggers one decompression/compression process, while insertion and query only trigger one compression or decompression.

## Supports and Limits

1. During logic export, CREATE TABLE statements will carry the `COMPRESSED` keyword. Therefore, TencentDB for MySQL supports such statements during import. Below are notes on official MySQL versions:
  - If the official MySQL version is below 5.7.18, data can be imported directly.
  - If the official MySQL version is 5.7.18 or above, the `COMPRESSED` keyword must be removed after logic export.
2. When DTS exports data from other cloud service providers or users, incompatibility may occur during binlog sync. In this case, you can skip DDL statements with the `COMPRESSED` keyword.

# Flashback Query

Last updated: 2023-08-31 17:07:11

## Feature Overview

Maloperations may occur in the process of database Ops and severely affect the business. Rollback and cloning are common recovery methods for maloperations, but they are error-prone and time-consuming in case of minor data changes and urgent troubleshooting, and are uncontrollable in recovery time when dealing with major data changes.

The TXSQL team has developed and implemented the flashback query feature for the InnoDB engine. It allows you to query the historical data before a maloperation with a simple SQL statement and query the data at a specified time point through specific SQL syntax. This greatly saves the data query and recovery time and enables fast data recovery for better business continuity.

## Supported Versions

- Kernel version: MySQL 5.7 20220715 and later.
- Kernel version: MySQL 8.0 20220331 and later.

For more information on how to view or upgrade the minor kernel version, see [Upgrading Kernel Minor Version](#).

## Scenarios

The flashback query feature is used to quickly query the historical data after a maloperation during database Ops.

Notes:

- Flashback query is supported only for InnoDB physical tables but not views, other engines, or functions without actual columns such as `last_insert_id()`.
- Only second-level flashback query is supported, and the accuracy cannot be fully guaranteed. If there are multiple changes within one second, any of them may be returned.
- Flashback query is supported only for primary keys (or GEN\_CLUST\_INDEX).
- Flashback query cannot be used in prepared statements or stored procedures.
- Flashback query does not support DDL. If you perform DDL on a table (such as TRUNCATE TABLE, which should be recovered through the recycle bin), the results obtained by flashback query may not be as expected.
- In the same statement, if multiple flashback query times are specified for the same table, the earliest time will be selected.

- Due to the time difference between the source and replica instances, if you specify the same time for flashback query, the results obtained for the instances may be different.
- Enabling the flashback query feature will delay undo log cleanup and increase the memory usage. We recommend you not set `Innodb_backquery_window` to a large value (preferably between 900 and 1,800), especially for instances with frequent business access requests.
- If the database instance restarts or crashes, the historical information before the restart or crash cannot be queried. The specified time should be within the supported range (which can be viewed through the status variables `Innodb_backquery_up_time` and `Innodb_backquery_low_time` by running `show status like '%backquery%'`).

## Notes

Flashback query provides a new `AS OF` syntax. You can set the `Innodb_backquery_enable` parameter to `ON` to enable the flashback query feature and then query data at the specified time through the following syntax:

```
SELECT ... FROM <table name>
AS OF TIMESTAMP <time>;
```

### Example of querying data at a specified time

```
MySQL [test]> create table t1(id int,c1 int) engine=innodb;
Query OK, 0 rows affected (0.06 sec)

MySQL [test]> insert into t1 values(1,1), (2,2), (3,3), (4,4);
Query OK, 4 rows affected (0.01 sec)
Records: 4  Duplicates: 0  Warnings: 0

MySQL [test]> select now();
+-----+
| now() |
+-----+
| 2022-02-17 16:01:01 |
+-----+
1 row in set (0.00 sec)

MySQL [test]> delete from t1 where id=4;
Query OK, 1 row affected (0.00 sec)

MySQL [test]> select * from t1;
```

```
+-----+-----+
| id   | c1   |
+-----+-----+
|  1   |  1   |
|  2   |  2   |
|  3   |  3   |
+-----+-----+
3 rows in set (0.00 sec)
```

```
MySQL [test]> select * from t1 as of timestamp '2022-02-17 16:01:01';
```

```
+-----+-----+
| id   | c1   |
+-----+-----+
|  1   |  1   |
|  2   |  2   |
|  3   |  3   |
|  4   |  4   |
+-----+-----+
4 rows in set (0.00 sec)
```

## Create an example using historical data

```
create table t3 select * from t1 as of timestamp '2022-02-17 16:01:01';
```

## Example of Inserting Historical Data into a Table

```
insert into t4 select * from t1 as of timestamp '2022-02-17 16:01:01';
```

## Description

The following table lists the configurable parameters of the flashback query feature.

Parameter	Valid Values	Local Disk Types	Default value	Value Range	Restart Required	Note
Innodb_backquery_enable	Global	Boolean	OFF	ON/OFF	Not required	The switch of the flashback query feature
Innodb_backquery_window	Global	Integer	900	1 – 86400	Not required	The time range for flashback query in seconds. The larger the value of this parameter, the longer the historical data query time supported for flashback query, and the more storage space used by the undo tablespace.
Innodb_backquery_history_limit	Global	Integer	800000	1 – 922337203685476000	Not required	The length of the undo linked list for flashback query. If this value is exceeded, <code>Innodb_backquery_window</code> will be ignored and a purge will be triggered until the historical linked list length is lower than this value.

# Performance Features

## Parallel Query

### Overview

Last updated: 2023-08-31 17:09:24

TencentDB for MySQL supports parallel query. After this feature is enabled, large queries can be automatically identified. The parallel query capability leverages multiple compute cores to greatly shorten the response time of large queries.

### Concept

Parallel query uses more computing resources to complete the query workload. The traditional query method is relatively friendly to small amounts of data (hundreds of gigabytes), but as the business grows, the data volume has reached the TB level in many cases, which exceeds the processing capacity of traditional databases. Parallel query is designed to solve this problem. During parallel query, the data is distributed to different threads at the storage layer, multiple threads on a single node process the data in parallel, the result pipelines are aggregated to the main thread, and the main thread performs a simple merge and returns the result. This greatly improves the query efficiency.

### Feature Background

TencentDB for MySQL goes beyond traditional MySQL databases in terms of computing, storage, disaster recovery, and elastic expansion; however, it still faces the following challenges:

- As the internet develops, databases become more capable of storing data, and forms are carrying more and more data. When it comes to big table queries, SQL statements tend to be slow due to existing technical bottlenecks, which adversely affects the business process.
- The current market environment sees an increasing number of report statistics and other analytical queries. Although not large in number, they involve a high data volume and are quite sensitive to query time. Gradually, data analysis capability, especially heterogeneous data processing, has become a must-have.

The above challenges are caused by the traditional technical implementation mode in the MySQL ecosystem. In particular, open-source releases support only the single-thread query mode, where only one thread (called user thread) is responsible for the parsing, optimization, and execution of a SQL statement. This mode cannot make full use of the hardware resources of modern multi-core CPUs and large memory devices, leading to a resource waste.

Therefore, it is important to streamline analysis and enhance performance by using multi-core services in the query of a large amount of data, which is also the key to query acceleration, cost reduction, and efficiency improvement.

## Strengths

- **Performance enhancement at no extra costs:** You can upgrade the kernel capabilities at no extra costs, so that you can get the most out of the instance CPU computation for quicker statement response and higher computing performance.
- **Support for common statements:** You can use most common SQL statements in virtually any business scenarios. This helps you accelerate your business smoothly.
- **Flexible parameter settings:** Multiple parameters are provided to help you control the start and stop conditions of parallel queries, making the queries more intelligent and adaptable to your business scenarios without the need for modification.

# Supported Statement Scenarios and Restricted Scenarios

Last updated: 2023-08-31 17:09:39

This document describes the supported statement scenarios and restricted scenarios of the parallel query.

## Supported statement scenarios

TencentDB for MySQL has implemented the parallel query feature for SQL statements with the following characteristics, with more to come.

- Single-table scan: Full-table scan, index scan, index range scan, and index REF query in ascending or descending order are supported.
- Multi-table join: The nested-loop join (NLJ) algorithm as well as semi join, anti join, and outer join are supported.
- Subquery: Parallel query is supported for derived tables.
- Data type: Different data types can be queried, such as integer, string, floating point, time, and overflow (with a runtime size limit).
- There are no restrictions on common operators and functions.
- COUNT, SUM, AVG, MIN, and MAX aggregate functions are supported.
- UNION and UNION ALL queries are supported.
- Traditional (default), JSON, and tree EXPLAIN formats are supported.

## Restricted scenarios

The parallel query feature of TencentDB for MySQL is not supported in the following scenarios.

Limit	Description
Statement compatibility restriction	Parallel query is not supported for non-query statements, including INSERT ... SELECT and REPLACE ... SELECT.
	Parallel query is not supported for statements in a stored program.
	Parallel query is not supported for prepared statements.
	Parallel query is not supported for statements in serial isolation-level transactions.
	Parallel query is not supported for locking reads, such as SELECT FOR UPDATE and SELECT ... FOR SHARE.
	Parallel query is not supported for CTEs.
Table/Index compatibility restriction	Parallel query is not supported for system, temp, and non-InnoDB tables.
	Parallel query is not supported for space index.
	Parallel query is not supported for full-text index.
	Parallel query is not supported for partitioned tables.
	Parallel query is not supported for tables in <code>index_merge</code> scan mode.
Expression/Field compatibility restriction	Parallel query is not supported for tables containing generated columns or BLOB, TEXT, JSON, BIT, and GEOMETRY fields.
	Parallel query is not supported for aggregate functions of the BIT_AND, BIT_OR, or BIT_XOR type.
	Parallel query is not supported for DISTINCT aggregations, such as SUM(DISTINCT) and COUNT(DISTINCT).
	Parallel query is not supported for GIS functions such as SP_WITHIN_FUNC and ST_DISTANCE.
	Parallel query is not supported for custom functions.
	Parallel query is not supported for JSON functions such as JSON_LENGTH, JSON_TYPE, and JSON_ARRAYAGG.
	Parallel query is not supported for XML functions such as XML_STR.
	Parallel query is not supported for user-lock functions such as IS_FREE_LOCK, IS_USED_LOCK, RELEASE_LOCK,

RELEASE\_ALL\_LOCKS, and GET\_LOCK.

Parallel query is not supported for SLEEP, RANDOM, GROUP\_CONCAT, SET\_USER\_VAR, and WEIGHT\_STRING functions.

Parallel query is not supported for certain statistical functions such as STD, STDDEV, STDDEV\_POP, VARIANCE, VAR\_POP, and VAR\_SAMP.

Parallel query is not supported for subqueries.

Parallel query is not supported for window functions.

Parallel query is not supported for ROLLUP.

Besides the above examples in [Supported statement scenarios](#), you can also check the parallel query execution plan and thread working status to see whether a statement can be queried parallelly. For more information, see [Viewing Parallel Query](#).

# Enabling/Disabling Parallel Query

Last updated: 2023-08-31 17:09:54

This document describes how to enable or disable the parallel query feature of TencentDB for MySQL via the console or command line.

## Preparations

Database version: MySQL 8.0 on kernel version 20220831 or later.

## Description

### Note

The parallel query feature can be enabled for both source and read-only instances, as long as their number of CPU cores is greater than or equal to 4.

You can enable the parallel query feature for the current instance by setting the `txsql_max_parallel_worker_threads` and `txsql_parallel_degree` parameters to a value other than `0` via the console or command line. Parameters and suggested settings are as follows:

### Parameter information

Category	Variable type	Scope	Default value	Value Range	Note
<code>txsql_max_parallel_worker_threads</code>	Integer	Global	$\{\text{MIN}(\text{DB InitCpu}, 0)\}$	0– $\{\text{MAX}(\text{DB InitCpu} - 2, 2)\}$	The total number of threads of the instance node that can be used for parallel query. If it is set to 0, no parallel thread is available, indicating to disable the parallel query feature.
<code>txsql_parallel_degree</code>	Integer	Global/session	4	0 – 64	The maximum number of threads (default parallelism) that can be used during the parallel query of a single statement. 0 indicates to disable the parallel query feature.

### Recommended Settings

- **Parallelism limit:** `txsql_parallel_degree` indicates the maximum number of threads for the parallel query of a single statement, i.e., the default parallelism. We recommend that you limit this value to half of the CPU core quantity of the instance. To ensure the stability, the parallel query feature is disabled for small clusters with fewer than four CPU cores, and you cannot adjust parallel query parameters via the console or command line.
- During the parallel query of a SQL statement, the parallelism set by `txsql_parallel_degree` will be used by default, which can be adjusted through the HINT statement. For more information, see [HINT Statement Control](#).
- `txsql_max_parallel_worker_threads` indicates the number of threads of the instance that can be used for parallel query, and `txsql_max_parallel_worker_threads / txsql_parallel_degree` indicates the maximum number of SQL statements allowed in a parallel query.
- `txsql_max_parallel_worker_threads` and `txsql_parallel_degree` control the status of the parallel query feature. When either of them is 0, the feature is disabled.

TencentDB for MySQL offers various parameters for you to set the execution conditions of parallel query for business adaptation and stability. After conditions are set, the database will check whether each SQL statement can be executed against such conditions like execution

cost, number of table rows, and memory usage for the parallel statement execution.  
Parameters are described as follows:

Category	Variable type	Scope	Default value	Value Range	Note
innodb_txsql_parallel_partitions_per_worker	Integer	Global/Session	13	0-256	The average number of partitions to be scanned per thread in parallel scans of partitioned data.
txsql_optimizer_context_max_memory_size	Integer	Global/Session	{MIN(DBInitMemory*52429,8388608)}	0-{DBInitMemory*52429}	The maximum memory size that a single statement can apply for in the parallel query plan environment.
txsql_parallel_cost_threshold	Integer	Global/Session	50000	0-9223372036854476000	The threshold of parallel execution cost. Only statements with an execution cost higher than this threshold will be executed parallelly.
txsql_parallel_exchange_buffer_size	Integer	Global/Session	1048576	65536-268435456	The data exchange buffer size.
txsql_parallel_table_record_threshold	Integer	Global/Session	5000	0-9223372036854476000	The threshold of parallel table row quantity. Only tables with a row quantity higher than this threshold will be selected as parallel tables.

### Note

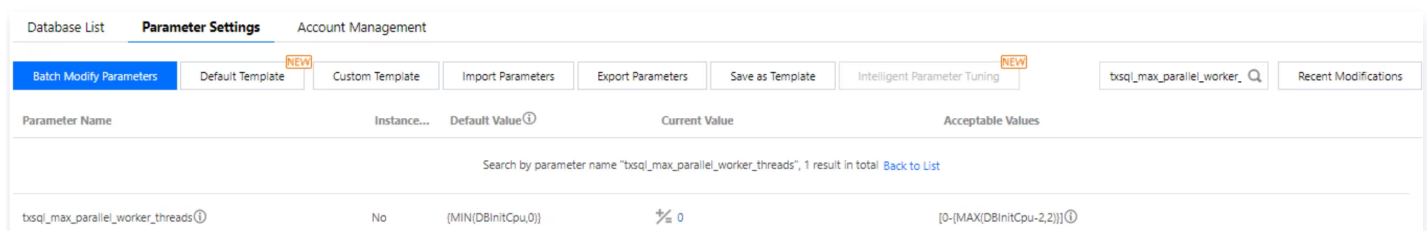
- Parallel query parameters take effect immediately after being set, with no instance restart required.
- If the scope of a parameter is session, it takes effect only for the statement.

## Enabling or disabling parallel query in the console

You can enable or disable the feature by setting parameters on the **Parameter Settings** page in the TencentDB for MySQL console.

- Set `txsql_max_parallel_worker_threads` and `txsql_parallel_degree` to a value other than 0 to enable parallel query.
- Set `txsql_max_parallel_worker_threads` or `txsql_parallel_degree` to 0 to disable parallel query.

You can also set execution conditions on the **Parameter Settings** page. For detailed directions, see [Setting Instance Parameters](#).



Parameter Name	Instance...	Default Value	Current Value	Acceptable Values
txsql_max_parallel_worker_threads	No	(MIN(DBInitCpu,0))	0	[0-(MAX(DBInitCpu-2,2))]

## Specifying the parallel execution mode of a SQL statement through the HINT statement

TencentDB for MySQL supports specifying parallel execution methods for individual SQL statements. When configuring a single SQL statement, the hint statement method is used. For more information, please refer to [Hint Statement Control](#).

## Documentation

- [View parallel query](#)
- [Hint statement control](#)

# HINT Statement Control

Last updated: 2023-08-31 17:10:10

TencentDB for MySQL allows you to enable or disable the parallel query feature by adjusting parameters. Specifically, you can enable or disable the feature for all SQL statements, set execution conditions, or specify the execution mode of a specific SQL statement through the HINT statement in the console.

## Note

The HINT statement can specify whether to execute a SQL statement and apply session parameters to the statement. In addition, it also supports querying the specified parallel table.

## HINT statement usage

SDK	Command line	Note
Enable parallel query	<pre>SELECT /*+PARALLEL(x)*/ ... FROM ...;</pre>	<code>x</code> indicates the parallelism for the SQL statement, which should be greater than <code>0</code> .
Disable parallel query	<pre>SELECT /*+PARALLEL(x)*/ ... FROM ...;</pre>	If <code>x</code> is set to <code>0</code> , it indicates to disable parallel query.
Specify the parallel table	<p>You can specify the table to be included in or excluded from the parallel query execution plan in either of the following ways:</p> <ul style="list-style-type: none"> <li>Specify the table to be included in the plan through <pre>SELECT /*+PARALLEL(t)*/ ... FROM ...;</pre></li> <li>Specify the table to be excluded from the plan through <pre>NO_PARALLEL SELECT /*+NO_PARALLEL(t)*/ ... FROM ...;</pre></li> </ul>	<code>t</code> is the table name.
Specify both the parallel table and parallelism	<pre>SELECT /*+PARALLEL(t x)*/ * ... FROM ...;</pre>	<code>x</code> indicates the parallelism for the SQL statement, which should be greater than <code>0</code> . <code>t</code> is the table name.
Set the session parameter through the HINT statement, which takes effect only for the specified SQL statement	<pre>SELECT /*+SET_VAR(var=n)*/ * ... FROM ...;</pre>	<code>var</code> is the parallel query parameter in the <code>session</code> scope.

## HINT statement use cases

**Use case 1:**

```
select /*+PARALLEL( )*/ * FROM t1, t2;
```

Set the parallelism to the value of `txsql_parallel_degree` (default) for the parallel query. If a statement does not meet the parallel query execution condition, serial query will be used.

**Use case 2:**

```
select /*+PARALLEL(4)*/ * FROM t1, t2;
```

Set the parallelism of the statement to `4` regardless of the default value, i.e., `txsql_parallel`

`_degree = 4` . If the statement does not meet the parallel query execution condition, serial query will be used.

**Use case 3:** `select /*+PARALLEL (t1) */ * FROM t1, t2;`

Include the `t1` table in the parallel query and use the default parallelism. If `t1` is smaller than the value of `txsql_parallel_table_record_threshold` , serial query will be used.

**Use case 4:** `select /*+PARALLEL (t1 8) */ * FROM t1, t2;`

Include the `t1` table in the parallel query and set the parallelism to `8` . If `t1` is smaller than the value of `txsql_parallel_table_record_threshold` , serial query will be used.

**Use case 5:** `select /*+NO_PARALLEL (t1) */ * FROM t1, t2;`

Exclude the `t1` table from the parallel query. If `t1` is greater than the value of `txsql_parallel_table_record_threshold` , serial query will be used.

**Use case 6:** `select /*+SET_VAR(txsql_parallel_degree=8)*/ * FROM t1, t2;`

Set the parallelism of the statement to `8` regardless of the default value, i.e., `txsql_parallel_degree = 8` .

**Use case 7:** `select /*+SET_VAR(txsql_parallel_cost_threshold=1000)*/ * FROM t1, t2`

Set `txsql_parallel_cost_threshold=1000` for the statement. If its execution penalty is greater than `1000` , parallel query can be used.

**Scenario 8:** `select /*+SET_VAR(txsql_optimizer_context_max_mem_size=500000)*/ * FROM t1, t2`

This sets the `txsql_optimizer_context_max_mem_size=500000` for a single statement, adjusting the maximum memory limit for the parallel query plan environment to 500,000.

## Documentation

- [Enable or Disable Parallel Query](#)
- [View parallel query](#)

# Viewing Parallel Query

Last updated: 2023-08-31 17:10:29

TencentDB for MySQL allows you to view the parallel query execution plan and threads in the plan, so that you can clearly know how parallel query takes effect in a database and quickly troubleshoot issues.

This document describes two common methods for viewing parallel queries.

## Option 1: Using the EXPLAIN statement

**Sample SQL statement:**

```
SELECT l_returnflag, l_linestatus, sum(l_quantity) as sum_qty
FROM lineitem
WHERE l_shipdate <= '1998-09-02'
GROUP BY l_returnflag, l_linestatus
ORDER BY l_returnflag, l_linestatus;
```

This sample is a simplified version of TPC-H Q1, a typical report operation.

**Execution plan print statement (EXPLAIN):**

```
EXPLAIN SELECT l_returnflag, l_linestatus, sum(l_quantity) as sum_qty
FROM lineitem
WHERE l_shipdate <= '1998-09-02'
GROUP BY l_returnflag, l_linestatus
ORDER BY l_returnflag, l_linestatus;
```

**Query result:**

```
MySQL [tpch100g]> explain SELECT l_returnflag, l_linestatus,
sum(l_quantity) as sum_qty FROM lineitem WHERE l_shipdate <= '1998-09-
02' GROUP BY l_returnflag, l_linestatus ORDER BY l_returnflag,
l_linestatus;
+----+-----+-----+-----+-----+-----+-----+
----+-----+-----+-----+-----+-----+-----+
-----+
| id | select_type | table          | partitions | type | possible_keys |
key | key_len | ref | rows          | filtered | Extra
|
```

```

+----+-----+-----+-----+-----+-----+
----+-----+-----+-----+-----+-----+-----+
-----+
| 1 | SIMPLE | lineitem | NULL | ALL | i_l_shipdate |
NULL | NULL | NULL | 593184480 | 50.00 | Parallel scan (4
workers); Using where; Using temporary |
| 1 | SIMPLE | <sender1> | NULL | ALL | NULL |
NULL | NULL | NULL | 0 | 0.00 | Send to (<receiver1>)
|
| 1 | SIMPLE | <receiver1> | NULL | ALL | NULL |
NULL | NULL | NULL | 0 | 0.00 | Receive from (<sender1>);
Using temporary; Using filesort |
+----+-----+-----+-----+-----+-----+
----+-----+-----+-----+-----+-----+
-----+
3 rows in set, 1 warning (0.00 sec)

```

### Tree-shaped execution plan print statement (EXPLAIN format=tree):

```

EXPLAIN format=tree query SELECT l_returnflag, l_linestatus,
sum(l_quantity) as sum_qty
FROM lineitem
WHERE l_shipdate <= '1998-09-02'
GROUP BY l_returnflag, l_linestatus
ORDER BY l_returnflag, l_linestatus;

```

### Query result:

```

MySQL [tpch100g]> explain format=tree SELECT l_returnflag, l_linestatus,
sum(l_quantity) as sum_qty FROM lineitem WHERE l_shipdate <= '1998-09-
02' GROUP BY l_returnflag, l_linestatus ORDER BY l_returnflag,
l_linestatus\G
***** 1. row *****
EXPLAIN: -> Sort: lineitem.L_RETURNFLAG, lineitem.L_LINESTATUS
-> Table scan on <temporary>
-> Final Aggregate using temporary table
-> PX Receiver (slice: 0; workers: 1)
-> PX Sender (slice: 1; workers: 4)
-> Table scan on <temporary>

```

```

-> Aggregate using temporary table
      -> Filter: (lineitem.L_SHIPDATE <=
DATE'1998-09-02') (cost=65449341.10 rows=296592240)
      -> Parallel table scan on lineitem
(cost=65449341.10 rows=593184480)

1 row in set (0.00 sec)

```

As can be seen from the above result:

- The parallel query plan assigns the statement to four worker threads for computing.
- Aggregate operations are split into two segments that are executed by the user and parallel threads respectively.
- The parallel scan operator is used for the `lineitem` table.
- EXPLAIN format=tree query works better than the traditional EXPLAIN.

## Option 2: Viewing in the thread list

The result of the `show processlist` command displays which threads are running. You can view not only the total number of current connections but also the connection status to identify abnormal query statements.

Based on the `show processlist` command, TencentDB for MySQL offers the proprietary `show parallel processlist` statement, which displays only the threads related to parallel query and filters out irrelevant threads.

**Sample SQL statement:**

```

SELECT l_returnflag, l_linestatus, sum(l_quantity) as sum_qty
FROM lineitem
WHERE l_shipdate <= '1998-09-02'
GROUP BY l_returnflag, l_linestatus
ORDER BY l_returnflag, l_linestatus;

```

This sample is a simplified version of TPC-H Q1, a typical report operation.

**show processlist query result:**

```

mysql> show processlist;
+-----+-----+-----+-----+-----+-----
+-----+-----+-----+-----+-----+-----
-----+

```

```

| Id      | User      | Host      | db      | Command | Time | State
| Info
|
+-----+-----+-----+-----+-----+-----+-----
+-----+-----+-----+-----+-----+-----+-----
-----+
|      7 | tencentroot | 127.0.0.1:49238 | NULL    | Sleep   | 0    |
| NULL
|
|     11 | tencentroot | 127.0.0.1:49262 | NULL    | Sleep   | 0    |
| NULL
|
|     13 | tencentroot | 127.0.0.1:49288 | NULL    | Sleep   | 1    |
| NULL
|
| 237062 | tencentroot | localhost      | tpch100g | Query   | 24   |
Scheduling | SELECT l_returnflag, l_linestatus, sum(l_quantity) as
sum_qty FROM lineitem WHERE l_shipdate <= '199 |
| 237107 | tencentroot | localhost      | NULL    | Query   | 0    |
init      | show processlist
|
+-----+-----+-----+-----+-----+-----+-----
+-----+-----+-----+-----+-----+-----+-----
-----+
6 rows in set (0.00 sec)

```

### Show parallel processlist query result:

```

mysql> show parallel processlist;
+-----+-----+-----+-----+-----+-----+-----
-----+
-----+
| Id      | User      | Host      | db      | Command | Time | State
| Info
|
+-----+-----+-----+-----+-----+-----+-----
-----+
-----+

```

```

| 237062 | tencentroot | localhost | tpch100g | Query | 18 |
Scheduling | SELECT l_returnflag, l_linestatus, sum(l_quantity) as
sum_qty FROM lineitem WHERE l_shipdate <= '199 |
| 237110 | | | | Task | 18 | Task
runing | connection 237062, worker 0, task 1
|
| 237111 | | | | Task | 18 | Task
runing | connection 237062, worker 1, task 1
|
| 237112 | | | | Task | 18 | Task
runing | connection 237062, worker 2, task 1
|
| 237113 | | | | Task | 18 | Task
runing | connection 237062, worker 3, task 1
|
+-----+-----+-----+-----+-----+-----+
-----+-----+-----+-----+-----+-----+
-----+
5 rows in set (0.00 sec)

```

As can be seen from the above result:

- The parallel query plan assigns queries to four worker threads. There is only one data item in the user thread (ID: 237062). The SQL statement is pushed down to four worker threads. As indicated in `info`, all these four threads are executing `task 1`.
- Each thread can be identified and located precisely.
- Compared to `show processlist`, `show parallel processlist` can precisely find all running threads of parallel query and will not be affected by other threads.

## Documentation

- [Enable or Disable Parallel Query](#)
- [Hint statement control](#)

# Large Transaction Replication

Last updated: 2023-08-31 17:11:47

## Feature Overview

If multi-row large transaction is updated by a single statement in row format, an event will be generated for each row. As a result, a large number of binlogs are created, and APPLY operations in the replica database become slower during replication, causing replication delays.

After analyzing and optimizing the large transaction replication scenarios, the Tencent Cloud kernel team developed the large transaction replication optimization feature. With this feature, large transactions are automatically identified and binlogs are converted from row format into statement format, thus reducing the quantity of binlogs and increasing the replication performance.

## Supported Versions

- Kernel version: MySQL 5.6 20210630 and above.
- Kernel version: MySQL 5.7 20200630 and above.
- Kernel version: MySQL 8.0 20200830 and above.

## Scenarios

- This feature accelerates large transaction replay for tables without a primary key in row format. It can be enabled if you are sure that the delay is caused by slow replay due to the lack of primary key.
- This feature aims to solve slow replication when there are large transactions in row format.

## Performance Statistics

The replication time is reduced by 85% for UPDATE operations and about 30% for INSERT operations.

## Notes

The large transaction replication optimization feature judges whether a transaction is large based on the historical execution statistics of the SQL statement. If a transaction is identified as large and optimizable, its isolation level will be automatically upgraded to repeatable read (RR), and the binlogs will be stored in statement format to reduce the time of executing the large transaction in the replica database. Here:

- `cdb_optimize_large_trans_binlog` is the switch of this feature.

- `cdb_sql_statistics` is the switch of SQL statement execution statistics collection.
- `cdb_optimize_large_trans_binlog_last_affected_rows_threshold` and `cdb_optimize_large_trans_binlog_aver_affected_rows_threshold` are the thresholds for judging a large transaction.
- `cdb_sql_statistics_info_threshold` is the number of legacy data entries retained in the memory.

To better monitor the transaction execution, the `CDB_SQL_STATISTICS` table is added in the `information_schema` database for you to query the statistics of the current transaction.

## New parameters

### Note

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

## Newly added `information_schema.CDB_SQL_STATISTICS` table

Name	Local Disk Types	Note
DIGEST_MD5	MYSQL_TYPE_STRING	MD5 value calculated from the digest of the SQL statement.
DIGEST_TEXT	MYSQL_TYPE_STRING	SQL statement digest text format.
SQL_COMMAND	MYSQL_TYPE_STRING	SQL command type.
FIRST_UPDATE_TIMESTAMP	MYSQL_TYPE_DATETIME	The time when the statistics information of the SQL statement is generated for the first time.
LAST_UPDATE_TIMESTAMP	MYSQL_TYPE_DATETIME	The time when the statistics information of the SQL statement is last updated.
LAST_ACCESS_TIMESTAMP	MYSQL_TYPE_DATETIME	The time when the statistics information of the SQL statement is last accessed.
EXECUTE_COUNT	MYSQL_TYPE_LONG	The number of executions of this SQL statement.
TOTAL_AFFECTED_ROWS	MYSQL_TYPE_LONG	Total number of affected rows.
AVER_AFFECTED_ROWS	MYSQL_TYPE_LONG	Average number of affected rows.
LAST_AFFECTED_ROWS	MYSQL_TYPE_LONG	The number of rows affected last time.
STMT_BINLOG_FORMAT_IF_POSSIBLE	MYSQL_TYPE_STRING	Whether binlogs for this SQL statement can be stored in statement format. Valid values: TRUE, FALSE.

# Execution Plan Cache for Optimizing UK/PK Queries

Last updated: 2023-08-31 17:12:41

## Feature Overview

In MySQL, SQL statement execution is divided into four stages: parsing, preparation, optimization, and execution. The execution plan cache feature is only available for prepared statements. After the feature is enabled, the first three stages will be skipped when executing a prepared statement, greatly boosting query performance.

In MySQL 8.0 20210830, the execution plan cache takes effect only for queries using unique keys (UKs) or primary keys (PKs). We will cover more types of queries in later versions.

## Supported Versions

Kernel version: MySQL 8.0 20210830 and later.

## Scenarios

This feature is mainly used to improve the query performance when executing short prepared statements with UKs or PKs on TencentDB instances. However, the extent to which performance may improved depends on your business.

## Impact on Performance

- For UK and PK SQL statements, the delay is reduced by 20%–30% and the throughput is improved by 20%–30% after the execution plan cache is enabled (according to the sysbench test using the point\_select.lua script).
- Memory usage will increase after the execution plan cache is enabled.

## Notes

You can use the `cdb_plan_cache` parameter to enable or disable the execution plan cache and the `cdb_plan_cache_stats` parameter to query information about cache hits. Note that only accounts with the `tencentroot` permission can use the two parameters.

Parameter	Status	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>cdb_plan_cache</code>	yes	bool	false	true/false	Whether to enable the feature. Only accounts with the feature permission can use the parameter.

**Note**

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

Use the `show cdb_plan_cache` command to view the plan cache hit status. The fields are explained as follows:

Field	Note
<code>sql</code>	A SQL statement with the question mark (?) which represents that the execution plan of this statement has been cached.
<code>mode</code>	SQL cache mode. Currently, only the prepare mode is supported.
<code>hit</code>	Number of hits for this session.

After `cdb_plan_cache_stats` is enabled, cache hit information will be recorded, affecting database performance.

## SQL execution status

You can run `show profile` to show the status at each stage of SQL statement execution. But when the execution plan cache is hit, the status of `optimizing`, `statistics`, and `preparing` will be omitted.

# fdatasync()

Last updated: 2023-08-31 17:13:48

## Feature Overview

The `fsync()` system call flushes redo logs to disk, including metadata and data. But metadata contains unimportant information such as the last modified time. You can enable the `fdatasync()` system call to skip metadata when flushing redo logs in order to reduce costs.

## Supported Versions

- Kernel version: MySQL 5.7 20201230 and above.
- Kernel version: MySQL 8.0 20201230 and above.

## Scenarios

This feature is suitable for use cases with heavy write pressure.

## Performance Statistics

TPS is improved by about 10%, according to the sysbench test in a high-concurrency continuous write scenario using the `oltp_write_only.lua` script.

## Notes

Use the `innodb_flush_redo_using_fdatasync` parameter to enable or disable `fdatasync()`. Valid values: `true` (enable), `false` (disable). Default value: `false`. If `fdatasync()` is enabled, metadata of redo logs won't be flushed to disk in real time.

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>innodb_flush_redo_using_fdatasync</code>	yes	bool	false	true/false	Whether to call <code>fdatasync()</code> to flush redo logs

### Note

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

# Auto-Increment Column Persistence

Last updated: 2023-08-31 17:14:02

## Feature Overview

The auto-increment column persistence feature can persist an auto-increment column into a page to avoid duplicate auto-increment values.

## Supported Versions

Kernel version: MySQL 5.7 20190830 and above.

## Scenarios

This feature is suitable for scenarios where you don't want duplicate auto-increment values, such as legacy data archive.

## Notes

This feature is enabled in the kernel by default.

# Buffer Pool Initialization

Last updated: 2023-08-31 17:14:16

## Feature Overview

This feature speeds up the initialization of the buffer pool, reducing the startup time of the database instance.

## Supported Versions

- Kernel version: MySQL 5.6 20200915 and above.
- Kernel version: MySQL 5.7 20200630 and above.

## Scenarios

This feature is used to speed up the startup of the database instance.

## Performance Test Data

Performance test data collected from eight instances:

buffer_pool_size	Buffer Pool Initialization Time (Before Optimization)	Buffer Pool Initialization Time (After Optimization)	Increase (%)
50GB	2.55s	0.13s	1962%
200GB	10.28s	0.52s	1977%
500GB	25.72s	1.32s	1948%

## Notes

This feature is enabled in the kernel by default.

# FAST DDL

Last updated: 2023-08-31 17:15:26

## Feature Overview

This feature speeds up the creation of secondary index. After the feature is enabled, secondary indexes can be concurrently sorted in a temp table using multiple threads. The feature also optimizes the operation of locking the flush list when loading bulk data, effectively reducing the time consumed by CREATE INDEX and the impact on concurrent DML operations.

## Supported Versions

- Kernel Version MySQL 8.0 20210330 or later.
- Kernel Version MySQL 5.7 20210331 or later.

## Scenarios

You need to perform DDL operations frequently on your database and may encounter the following DDL-related problems:

- Why does database performance fluctuate when I add indexes, which even affects business writes and reads?
- Why does it sometimes take more than 10 minutes to execute a DDL operation on a table less than one GB in size?
- Why does database performance fluctuate when I exit a connection where a temp table is used?

To solve the above common problems, the TXSQL kernel team has optimized the operation of locking the flush list when loading bulk data, based on in-depth analysis and testing in multiple scenarios. The optimization effectively reduces the time consumed by CREATE INDEX, the impact on concurrent DML operations, and the impact caused by DDL operations.

## Performance Statistics

Use sysbench to test database performance when importing two billion rows of data (about 453 GB) before and after FAST DDL is enabled.

```
mysql> set global innodb_fast_ddl=ON;
Query OK, 0 rows affected (0.00 sec)
```

When the feature is disabled, the operation takes 4,395 seconds; when the feature is enabled, the operation takes 2,455 seconds.

## Notes

Use the `innodb_fast_ddl` parameter to enable or disable this feature.

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/Value Range	Note
<code>innodb_fast_ddl</code>	Yes	bool	OFF	{ON,OFF}	Enable or disable FAST DDL

# Invisible Index

Last updated: 2023-08-31 17:17:25

## Feature Overview

Many users require the invisibility of an index to assess if it can be deleted. By making an index as invisible, you can test the impact of its deletion on query performance before deleting it. If the index is being used by any program or database user, an error will occur or be reported. This feature is now available to MySQL 5.7 and later versions, not just limited to MySQL 8.0.

## Supported Versions

Kernel version: MySQL 5.7 20180918 and above.

## Scenarios

Before deleting an index, you can make it invisible to see if it is still in use. If not, it can be securely deleted.

## Notes

Run the following statements to create an invisible index or make an index invisible:

```
CREATE TABLE t1 (  
  i INT,  
  j INT,  
  k INT,  
  INDEX i_idx (i) INVISIBLE  
) ENGINE = InnoDB;  
CREATE INDEX j_idx ON t1 (j) INVISIBLE;  
ALTER TABLE t1 ADD INDEX k_idx (k) INVISIBLE;
```

Run the following statements to make an index visible:

```
ALTER TABLE t1 ALTER INDEX i_idx INVISIBLE;  
ALTER TABLE t1 ALTER INDEX i_idx VISIBLE
```

# CATS Transaction Scheduling Algorithm

Last updated: 2023-08-31 17:18:31

## Feature Overview

TXSQL supports the Contention-Aware Transaction Scheduling (CATS) algorithm. This new algorithm automatically detects lock contention between transactions and schedules them based on their scheduling weights.

MySQL supports another transaction scheduling algorithm, aka First In First Out (FIFO), which was introduced earlier than CATS. When multiple transactions are waiting for the same lock, CATS prioritizes them by assigning a scheduling weight which is computed based on the number of transactions that a transaction blocks. The transaction with a higher scheduling weight will be executed sooner. Thus, transaction throughput is improved.

## Supported Versions

- Kernel version: MySQL 5.7 20190230 and above.
- Kernel version: MySQL 8.0 20200630 and above.

## Scenarios

This feature is suitable for use cases under high concurrency and heavy lock contention.

## Performance Statistics

TPS is improved by more than 50% under high concurrency and heavy lock contention.

- Test method: use the `oltp_read_write.lua` script of sysbench (pareto random type enabled) to test TPS on eight tables (10 MB data) at the REPEATABLE READ transaction isolation level
- Test environment: TencentDB instance with 32 cores and 128 GB memory

Threads	FCFS(FIFO)	CATS	Higher performance
128	11999	12005	0%
256	6609	10137	53%
512	3453	9365	171%
1024	2196	7015	219%

## Notes

In MySQL 5.7, you can use the global parameter `innodb_trx_schedule_algorithm` to specify the transaction scheduling algorithm. The default value is `auto`.

Valid values:

- `auto`: Automatically adjust the transaction scheduling algorithm based on current system status. If the number of threads waiting for a lock exceeds 32, adopt CATS; otherwise, adopt First Come First Serve (FCFS), an algorithm similar to FIFO.
- `fcfs`: adopt the FCFS algorithm.
- `cats`: adopt the Contention-Aware Transaction Scheduling algorithm.

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>innodb_trx_schedule_algorithm</code>	yes	string	auto	[auto,fcfs,cats]	Specify the transaction scheduling algorithm

### Note

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

In MySQL 8.0, `auto` is the only valid value.

# Computation Pushdown

Last updated: 2023-08-31 17:22:12

## Feature Overview

This feature pushes LIMIT/OFFSET and SUM operations down to the storage engine InnoDB when querying single tables, effectively reducing query latency.

- When LIMIT/OFFSET is executed using secondary indexes, this feature can avoid using the clustered index values as pointers to find the full table rows, effectively cutting the cost of scanning table data.
- This feature pushes SUM operations down to InnoDB. In other words, instead of sending rows to the MySQL server, InnoDB calculates data itself and returns the final result to the MySQL server.

## Supported Versions

- LIMIT/OFFSET optimization applies to kernel version MySQL 5.7 20180530 and later
- SUM optimization applies to kernel version MySQL 5.7 20180918 and later.

## Scenarios

- This feature is mainly used to optimize single-table queries with LIMIT/OFFSET or SUM clauses, such as `Select *from tbl Limit 10`, `Select* from tbl Limit 10,2`, and `Select sum(c1) from tbl`.
- This feature cannot optimize the following queries:
  - Queries with DISTINCT, GROUP BY, or HAVING clauses
  - Nested subqueries
  - Queries with FULLTEXT indexes
  - Queries with ORDER BY clauses, where the optimizer fails to use indexes to implement ORDER BY
  - Queries with multi-range read (MRR)
  - Queries with SQL\_CALC\_FOUND\_ROWS.

## Performance Statistics

Import one million rows of data and test query performance in sysbench:

- The execution time of `select * from sbtest1 limit 1000000,1;` decreases from 6.3 to 2.8 seconds.
- The execution time of `select sum(k) from sbtest1;` decreases from 5.4 to 1.5 seconds.

## Notes

During the execution of a SQL statement, the optimizer automatically modifies the query execution plan to implement computation pushdown according to the following parameters.

Parameters are as follows:

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>cdb_enable_offset_pushdown</code>	Yes	bool	ON	{ON,OFF}	Enable or disable LIMIT/OFFSET pushdown. It is enabled by default.
<code>cdb_enable_sumagg_pushdown</code>	Yes	bool	OFF	{ON,OFF}	Enable or disable SUM pushdown. It is disabled by default.

### Note

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

# Security Features

## Transparent Data Encryption

Last updated: 2023-08-31 17:22:28

### Feature Overview

TXSQL inherits the transparent data encryption mechanism of MySQL and provides another implementation of the keyring plugin: keyring KMS, which integrates keyring with Tencent Cloud's enterprise-grade [Key Management Service \(KMS\)](#) service.

KMS is a data and key security protection service of Tencent Cloud, where all involved processes use high-security communication protocols to guarantee high service security. In addition, it provides distributed cluster management and hot backup capabilities to ensure high service reliability and availability.

KMS uses a two-layer key system, which involves two types of keys: customer master key (CMK) and data encryption key (DEK). A CMK is used to encrypt small packet data (up to 4 KB in size), such as DEK, password, certificate, and configuration file. A DEK is used to encrypt massive amounts of business data in symmetric encryption method during storage or communication and is encrypted and protected in asymmetric encryption method with a CMK. In this way, data can be encrypted both in the memory and files.

### Supported Versions

- Kernel version: MySQL 5.7 20171130 and above.
- Kernel version: MySQL 8.0 20200630 and above.

### Scenarios

Transparent data encryption means that data encryption/decryption operations are imperceptible to users. It supports real-time I/O encryption/decryption of data files; that is, data will be encrypted before being written to the disk and decrypted when being read from the disk into the memory. This helps meet the compliance requirements for static data encryption.

### Notes

For more information, see [Enabling Transparent Data Encryption](#).

# Audit

Last updated: 2023-08-31 17:22:40

## Feature Overview

Tencent Cloud offers database auditing for TencentDB MySQL instances. With this feature, database access and SQL statement execution information, including the start time of statement execution, the number of scanned rows, lock wait time, CPU time, client IP, username, and SQL statement, will be audited, assisting enterprises in risk management and data protection.

## Scenarios

This feature is suitable for the use cases where risky database behaviors (such as SQL injection and abnormal operation) need to be recorded and alarmed.

## Impact on Performance

There are two audit modes: sync and async. Sync audit synchronously records all audit logs with an average impact of less than 6% on instance performance. But async audit has almost no impact (less than 3%, to be precise), which is industry-leading.

## Notes

For more information on how to enable TencentDB for MySQL audit, see [Enabling TencentDB for MySQL Audit](#).

# Stability Features

## Second-Level Column Addition

Last updated: 2023-08-31 17:22:56

### Feature Overview

The quick column adding feature allows you to quickly add columns to a big table by only modifying the data dictionary, which eliminates the need of data replication during column adding and greatly reduces the column adding time for big tables and the impact on the system.

### Supported Versions

- Kernel version: MySQL 5.7 20190830 and above.
- Kernel version: MySQL 8.0 20200630 and above.

### Scenarios

This feature is suitable for adding columns to a table with a high volume of data.

### Performance Statistics

In tests with a table of 5 GB data, the time for adding a column is reduced from 40 seconds to below 1 second.

### Notes

- **INSTANT ADD COLUMN** syntax

Add the `algorithm=instant` clause to `ALTER TABLE` to add a column as follows:

```
ALTER TABLE t1 ADD COLUMN c INT, ADD COLUMN d INT DEFAULT 1000,  
ALGORITHM=INSTANT;
```

- The `innodb_alter_table_default_algorithm` parameter is added, which can be set to `inplace` or `instant`.

This parameter is `inplace` by default and can be configured to adjust the default algorithm of `ALTER TABLE` as follows:

```
SET @@global.innodb_alter_table_default_algorithm=instant;
```

If no algorithm is specified, the default algorithm configured by this parameter will be used for `ALTER TABLE` operations.

## Restrictions on `INSTANT ADD COLUMN`

- A statement can contain only column addition operations.
- A new column will be added to the end, and column order cannot be changed.
- `INSTANT ADD COLUMN` is not supported in tables with the row format being `COMPRESSED`.
- `INSTANT ADD COLUMN` is not supported in tables with a full-text index.
- `INSTANT ADD COLUMN` is not supported for temp tables.

# Async Deletion of Big Tables

Last updated: 2025-09-25 09:46:10

## Overview

This feature is used to drop tables with large data files to avoid I/O fluctuation. When the DROP TABLE operation is performed, the system will first rename the original database file (.ibd) to make a new temporary file and then returns the message of successful operation promptly. The temporary file is stored in the directory specified by the `innodb_async_drop_tmp_dir` parameter and is truncated in batches by the system in the backend. The size of the file to be truncated each time is specified by the `innodb_async_truncate_size` parameter (supported in MySQL 5.7 and 8.0 but not 5.6). Users do not need to perform this operation. Instead, it is automatically completed by the kernel. Principle: A hard link is created in another directory for the data file of a table when the table is dropped. As a result, when DROP TABLE is executed, only the hard link to the file is deleted. After that, the backend thread will scan files that need to be deleted in the hard-linked directory and automatically truncate the data file of the dropped table.

## Supported Versions

- Kernel version: MySQL 5.6 20220303 and later.
- Kernel version: MySQL 5.7 20230601 and later.
- Kernel version: MySQL 8.0 20200630 and later.

## Use Cases

This feature is used to drop tables with large data files.

## Instructions

- For MySQL 5.6, you can set the `innodb_async_truncate_work_enabled` parameter to `ON` to enable the async mode of `DROP TABLE`. The default value is `OFF`.
- For MySQL 5.7 and 8.0, you can set the `innodb_table_drop_mode` parameter to `ASYNC_DROP` to enable the async mode of `DROP TABLE`. The default value is `ASYNC_DROP`.
- The size of the file to be truncated each time is specified by the `innodb_async_truncate_size` parameter. This is not supported for MySQL 5.6.
- You can make the async drop of big tables more efficient by enabling the `innodb_fast_ddl` parameter as instructed in [FAST DDL](#).

Instructions for MySQL 5.6

## Steps to Enable the Feature

1. Before you use the feature of asynchronous big table deletion, set the parameter `innodb_adaptive_hash_index` to OFF.
2. Set the parameter `innodb_async_truncate_work_enabled` to ON to enable the feature of asynchronous big table deletion. For parameter settings, see [Setting Instance Parameters](#).

## Relevant Parameters

Parameter Name	Dynami c	Typ e	Def aul t	Valu e Rang e	Description
<code>innodb_adaptive_hash_index</code>	yes	str in g	ON	ON/ OFF	Whether to enable InnoDB adaptive hash index. <ul style="list-style-type: none"> <li>• ON: Enable.</li> <li>• OFF: Disable.</li> </ul>
<code>innodb_async_truncate_work_enabled</code>	yes	str in g	OF F	ON/ OFF	Whether to enable asynchronous big table deletion. <ul style="list-style-type: none"> <li>• ON: Enable.</li> <li>• OFF: Disable.</li> </ul>

### Instructions for MySQL 5.7 and 8.0

## Steps to Enable the Feature

1. Before you use the feature of asynchronous big table deletion, set the parameter `innodb_fast_ahi_cleanup_for_drop_table` to ON.
2. Set the parameter `innodb_adaptive_hash_index` to OFF.
3. Set the parameter `innodb_table_drop_mode` to ASYNC\_DROP to enable the feature of asynchronous big table deletion. For parameter settings, see [Setting Instance Parameters](#).
4. (Optional) Set the parameter `innodb_fast_ddl` to ON, which can make asynchronous big table deletion more efficient. The FAST DDL feature and relevant parameters are

involved. For details, see [FAST DDL](#).

## Relevant Parameters

Parameter Name	Dynamic	Type	Default	Value Range	Description
<code>innodb_fast_ahi_cleanup_for_drop_table</code>	yes	string	ON	ON/OFF	Whether to enable quick cleanup optimization for adaptive hash indexes. <ul style="list-style-type: none"> <li>ON: Enable. Once it is enabled, it can prevent big table deletion from being blocked due to slow hash index cleanup.</li> <li>OFF: Disable.</li> </ul>
<code>innodb_adaptive_hash_index</code>	yes	string	OFF	ON/OFF	Whether to enable InnoDB adaptive hash index. <ul style="list-style-type: none"> <li>ON: Enable.</li> <li>OFF: Disable.</li> </ul>
<code>innodb_table_drop_mode</code>	yes	string	ASYNC_DROP	SYNC_DROP/ASYNC_DROP	Whether to enable asynchronous big table deletion. <ul style="list-style-type: none"> <li>ASYNC_DROP: Asynchronous mode. This value indicates that asynchronous big table deletion is enabled.</li> <li>SYNC_DROP: Synchronous mode. This value indicates that asynchronous big table deletion is disabled.</li> </ul>
<code>innodb_async_truncate_size</code>	yes	int	128	128 – 168	Size of the file to be truncated each time in the backend during asynchronous big table deletion, in MB.



# Hotspot Update

Last updated: 2023-09-01 14:44:52

## Feature Overview

For businesses with frequent updates or flash sales, the hotspot update feature greatly optimizes the performance of the UPDATE operation on frequently updated rows. If automatic hotspot update detection is enabled, the system will automatically detect whether there is a single row of hotspot update, and if so, it will queue the large number of concurrent UPDATE operations and execute them in sequence, so as to reduce the risk of concurrency performance being compromised by many row locks.

## Supported Versions

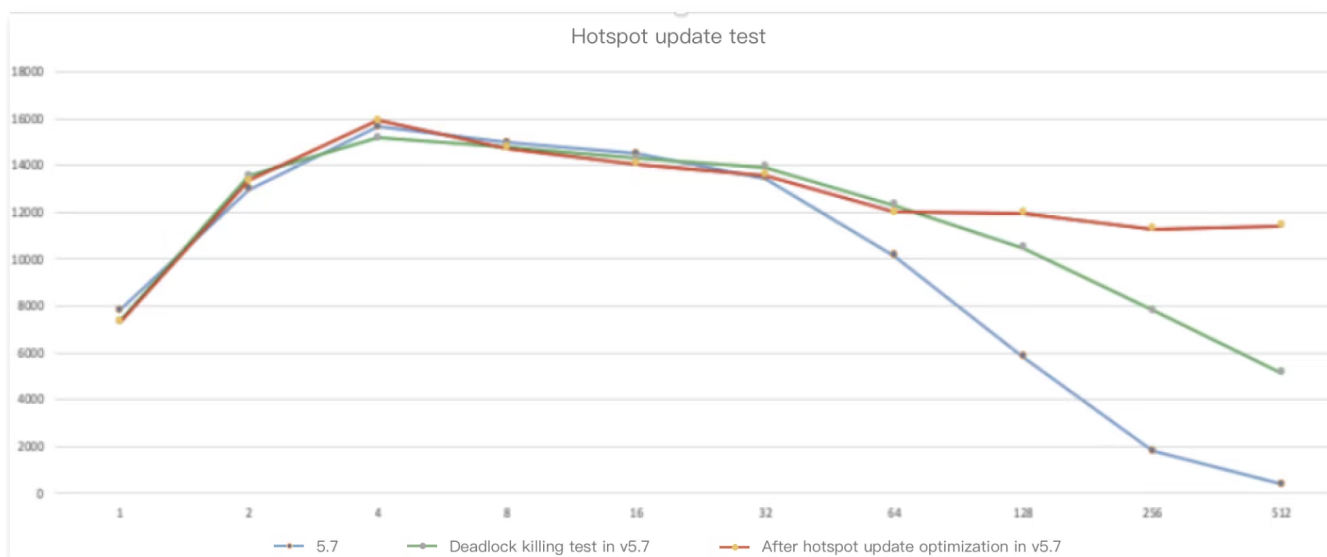
- Kernel version: MySQL 5.7 20200630 and above.
- Kernel version: MySQL 8.0 20200830 and above.

## Scenarios

This feature is suitable for scenarios where the pressure of updating a single row or multiple rows with the primary key specified is very high, such as flash sales.

## Performance Statistics

For high-concurrency UPDATE operations on a single row with the primary key specified, the performance is improved by over 10 times.



## Notes

## Hotspot Update Protection

# SQL throttling

Last updated: 2023-09-01 14:45:09

## Feature Overview

The SQL throttling feature enables you to set a keyword to limit the number of concurrent executions of the specified SQL statement.

## Supported Versions

- Kernel version: MySQL 5.7 20200330 and above.
- Kernel version: MySQL 5.6 20200915 and above.

## Scenarios

This feature is suitable for SQL statements with high concurrency and resource usage that compromise system performance.

## Use Cases

[SQL Throttling](#)

# Statement Outline

Last updated: 2023-09-01 14:48:20

## Feature Overview

SQL optimization is a crucial step in improving database performance. To avoid the impact when the optimizer fails to select an appropriate execution plan, TXSQL provides the outline feature for you to bind execution plans. MySQL allows you to use hints to manually bind execution plans. The hint information contains the optimization rule for SQL statements, algorithm to be used, and index for data scan, and an outline relies on hints to specify execution plans. Tencent Cloud provides the `mysql.outline` system table for you to add plan binding rules and the `cdb_opt_outline_enabled` switch for you to enable/disable the outline feature.

## Supported Versions

Kernel version: MySQL 8.0 20201230 and above.

## Scenarios

This feature is suitable for scenarios where an execution plan in the production environment has poor performance (for example, the index in the execution plan is incorrect), but you don't want to modify SQL statements and release a new version to fix this problem.

## Impact on Performance

- If `cdb_opt_outline_enabled` is enabled, the execution efficiency of SQL statements missing the outline will not be affected.
- The execution efficiency of SQL statements hitting the outline will be lower than that of general execution, but the performance after outline binding is generally improved by several times.
- To use `cdb_opt_outline_enabled`, you should consult the OPS or kernel engineers to avoid faulty binding and consequential performance compromise.

## Notes

The outline syntax uses a new syntax form:

- Configure outline information: `outline "sql" set outline_info "outline";`
- Clear outline information: `outline reset ""; outline reset all;`
- Refresh outline information: `outline flush;`

Below are the outline use methods with the following schemas as examples:

```
create table t1(a int, b int, c int, primary key(a));
create table t2(a int, b int, c int, unique key idx2(a));
create table t3(a int, b int, c int, unique key idx3(a));
```

Parameter	Effective Immediately	Local Disk Types	Default Value	Valid Values/ Value Range	Note
<code>cdb_opt_outline_enabled</code>	yes	bool	false	true/false	Whether to enable the outline feature.

### Note

Currently, you cannot directly modify the values of the above parameters. If needed, [submit a ticket](#) for assistance.

## Binding outline

The direct binding of OUTLINE involves replacing one SQL statement with another, without changing the semantics of the SQL, but adding some HINT information to guide the optimizer on how to execute it. The syntax is: `outline "sql" set outline_info "outline";` Note that the string following `outline_info` should start with "OUTLINE:", and the SQL with added HINT follows "OUTLINE:". For example, adding an index on column a of table t2 for the SQL statement `select * from t1, t2 where t1.a = t2.a.`

```
outline "select* from t1, t2 where t1.a = t2.a" set outline_info
"OUTLINE:select * from t1, t2 use index(idx2) where t1.a = t2.a";
```

## Binding optimizer hint

To provide more flexibility, TXSQL allows incremental addition of optimizer hints to SQL statements. The same functionality can also be achieved through direct binding of outlines. The syntax is: `outline "sql" set outline_info "outline";` Note that the string following `outline_info` should start with "OPT:", and the optimizer hint information to be added follows "OPT:". For example, specifying MATERIALIZATION/DUPSWEEDOUT SEMIJOIN for the SQL statement `select * from t1 where t1.a in (select b from t2) .`

```
outline "select* from t1 where t1.a in (select b from t2)" set
outline_info "OPT:2#qb_name(qb2)";
outline "select * from t1 where t1.a in (select b from t2)" set
outline_info "OPT:1#SEMIJOIN(@qb2 MATERIALIZATION, DUPSWEEEDOUT)";
```

You can add only one optimizer hint to the original SQL statement at a time and must comply with the following rules:

- The `OPT` keyword must follow `"`.
- `'` must be placed before the new statement to be bound.
- You must add two fields (query block number#optimizer hint string), which must be separated with `"#"` ( e.g., ``"OPT:1#max_execution_time(1000)"`` ).

## Binding index hint

To make the feature more flexible, TXSQL allows you to add index hints incrementally to SQL statements. You can also implement the same feature by directly binding an outline.

The syntax is in the format of `outline "sql" set outline_info "outline";`. Note that the string after `outline_info` must start with `"INDEX:"`, which is followed by the index hint information to be added.

For example, you can add the index `idx1` of `USE INDEX` in `FOR JOIN` type to the table `t1` in the database `test` in query block 3 for the SQL statement ``select *from t1 where t1.a in (select t1.a from t1 where t1.b in (select t1.a from t1 left join t2 on t1.a = t2.a))`` as follows:

```
outline "select* from t1 where t1.a in (select t1.a from t1 where t1.b
in (select t1.a from t1 left join t2 on t1.a = t2.a))" set outline_info
"INDEX:3#test#t1#idx1#1#0";
```

You can add only one index hint to the original SQL statement at a time and must comply with the following rules:

- The `INDEX` keyword must follow `"`.
- `'` must be placed before the new statement to be bound.
- You must add five fields (query block number#db\_name#table\_name#index\_name#index\_type#clause).
- Here, `index_type` has three valid values (0: `INDEX_HINT_IGNORE`; 1: `INDEX_HINT_USE`; 2: `INDEX_HINT_FORCE`), and `clause` also has three valid values (1: `FOR JOIN`; 2: `FOR ORDER BY`; 3: `FOR GROUP BY`), which must be separated by `"#"` ( e.g., ``"INDEX:2#test#t2#idx2#1#0"`, indicating to bind the index `idx1` in `USE INDEX FOR JOIN` type to the table `t2` in the second query block).

## Deleting the outline information of SQL statement

TXSQL allows you to delete the outline binding information from an SQL statement.

The syntax is in the format of `outline reset "sql";`. For example, to delete the outline information from `select *from t1, t2 where t1.a = t2.a`, run the following statement: `outline reset "select* from t1, t2 where t1.a = t2.a";`.

## Clearing all outline information

TXSQL allows you to clear all outline binding information in the kernel. The syntax is `outline reset all`, and the execution statement is `outline reset all;`.

There may be some specific problems in the production environment where you must bind an index. In this case, you can directly configure an outline for binding.

You should analyze the possible performance compromise after configuring an outline and bind an outline only if the compromised performance is acceptable. You can consult kernel engineers if necessary.

## Parameter Status Description

TXSQL provides multiple methods to view the outline binding of your SQL statements. You can use the `mysql.outline` table to view the information of configured outlines. You can also use the `show cdb_outline_info` and `select * from information_schema.cdb_outline_info` APIs to view the outline information in the memory. Whether the entered SQL statement is bound to the specified outline is subject to whether the outline information is in the memory. Therefore, you can use the two APIs for debugging.

The `mysql.outline` system table is added to store the records of configured outline information, which has the following fields:

Field	Note
Id	Outline number
Digest	Hash value of the original SQL statement
Digest_text	Fingerprint information text of the original SQL statement
Outline_text	Fingerprint information text of the SQL statement after the outline is bound

You can use `show cdb_outline_info` or `select * from information_schema.cdb_outline_info` to view the outline records in the memory, and execution of the corresponding SQL statement will hit the bound plan in the outline. The parameters are as follows:

Field	Note
origin	Original SQL statement fingerprint
outline	SQL statement fingerprint after the outline is bound

# TXRocks Engine Overview

Last updated: 2023-09-01 14:48:57

## TXRocks Overview

TXRocks is a transactional storage engine developed by Tencent's TXSQL team based on RocksDB, a very popular high-performance persistent key-value (KV) store.

## Why TXRocks?

By leveraging the LSM tree storage structure of RocksDB, the TXRocks transactional storage engine not only reduces wastes caused by InnoDB's half-full pages and fragments, but also uses the compact storage format. Therefore, it has a performance comparable to that of InnoDB but requires only a half or even smaller storage space. It is more suitable for businesses with a large data volume and high requirements for the transactional read/write performance.

## RocksDB's LSM Tree Architecture

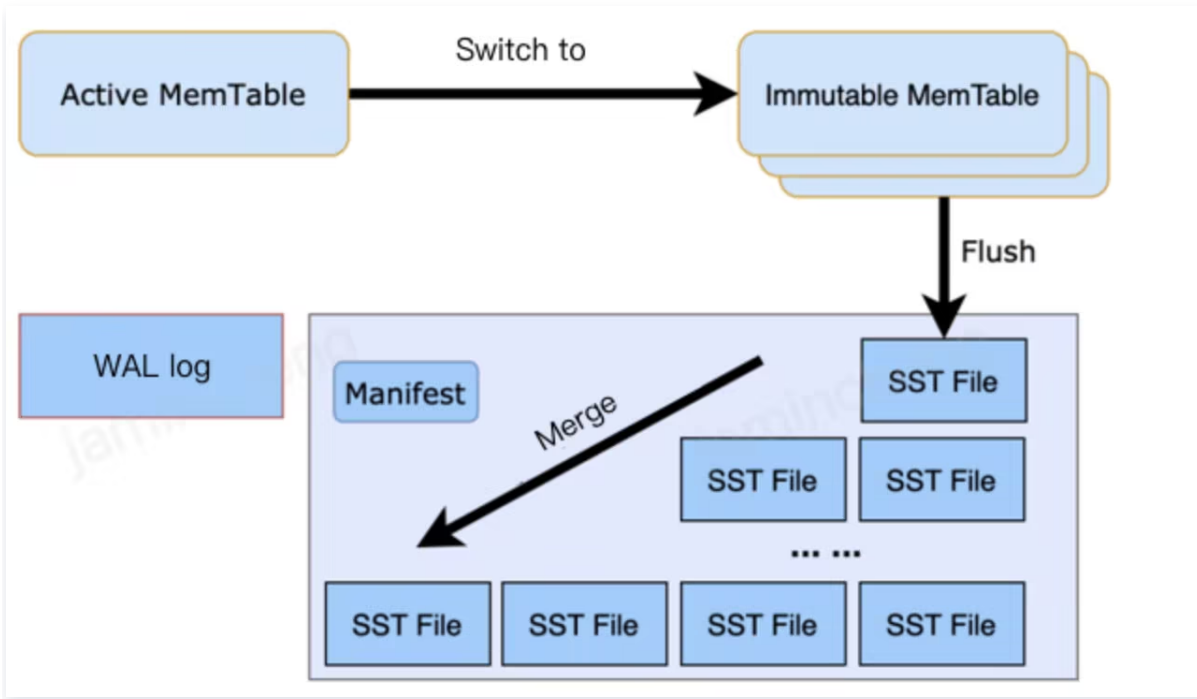
RocksDB uses the LSM tree storage structure, where data is organized as a set of MemTables in the memory and multi-level SST files on the disk.

For a write request, the new version of the record is first written to an active MemTable, and then WAL is written for data durability. After the MemTable and WAL file are written for the request, a response can be returned.

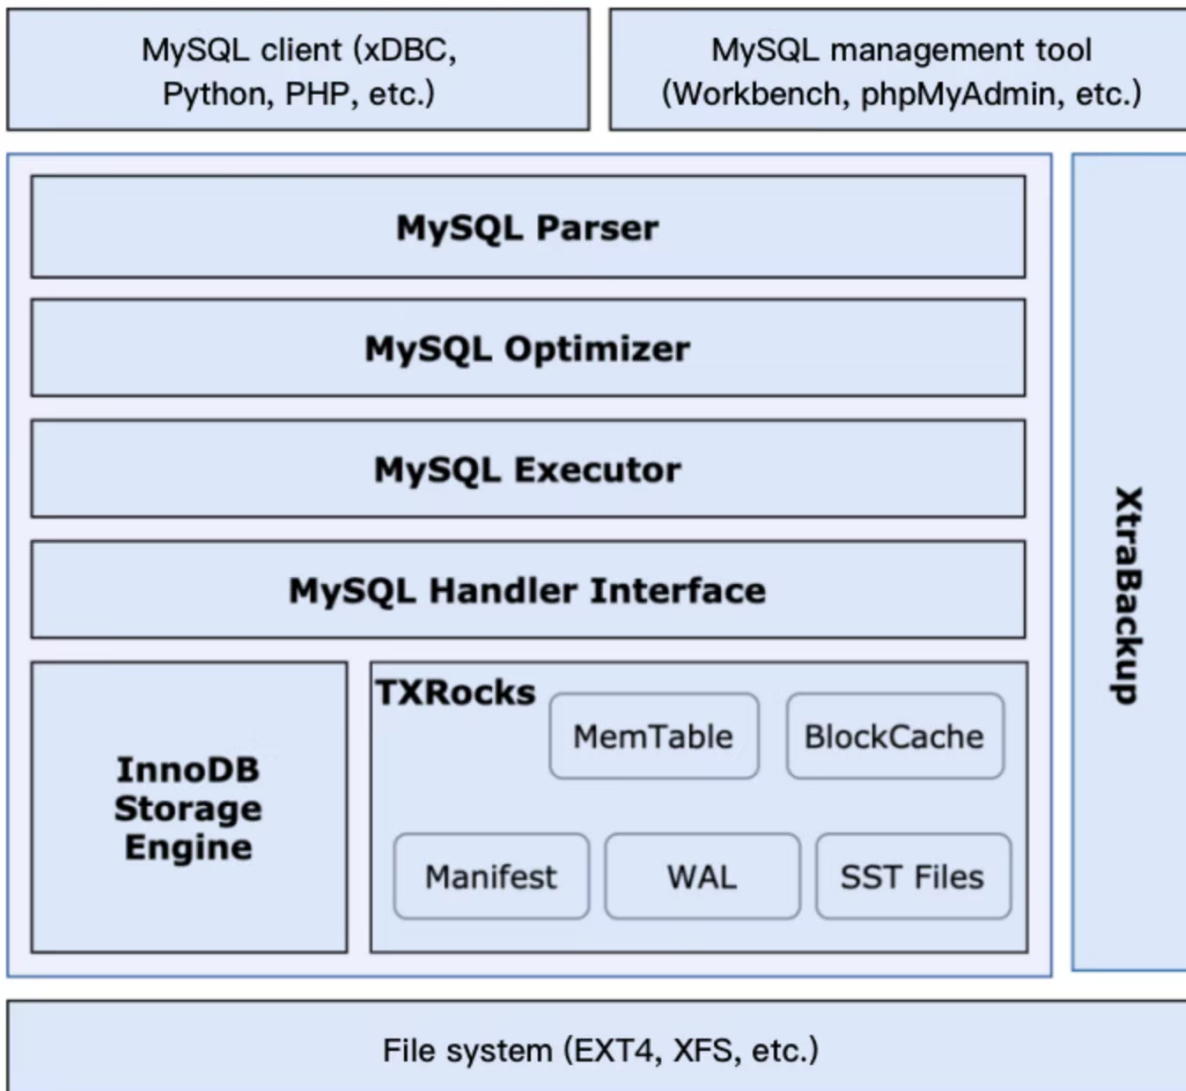
When the volume of data written into an active MemTable reaches a certain threshold, the MemTable will be switched to a frozen immutable MemTable. The backend thread flushes the immutable MemTable to the disk to generate an SST file. SST files are sorted at L0 to L6 levels in the flush order. At each level, the records of different SST files are sequential and don't overlap. However, to release the memory space occupied by immutable MemTables promptly, such records are allowed to overlap at L0.

When a record is read, it will be searched for in the active MemTable, immutable MemTable, SST files at L0, and SST files at L1-6 in sequence. Once the record is found in any component, the latest version of the record is found and will be returned immediately.

When performing a range scan, iterators are generated for each level of data, including the MemTable. These iterators merge to find the next record. As seen from the read process, if there are too many layers in the LSM Tree, the read performance, especially the range scan performance, will significantly decrease. Therefore, to maintain a better LSM Tree shape, the background continuously executes compaction operations, merging lower-level data into higher-level data, and reducing the number of layers.



## TXRocks Architecture



## TXRocks Strengths

### Saved storage space

Compared with the B+tree structure used by InnoDB, the LSM tree can save a considerable amount of storage space.

InnoDB's B+tree split often results in half-full pages, idle pages, and space waste; therefore, InnoDB has a lower effective page utilization.

The size of TXRocks SST files is generally set to dozens or hundreds of MB or a greater value. Therefore, TXRocks has much fewer wastes caused by 4K alignment. Although an SST file is divided into blocks, those blocks don't need to be aligned.

In addition, TXRocks SST files use prefix compression, so that only one record will be generated for data records with the same prefix. SST files at different levels can adopt different compression algorithms, further reducing the storage space overheads. For transaction overheads, InnoDB records must contain `trx id` and `roll_ptr` fields. By contrast, other transaction overheads don't need to be stored for SST files at the lowest level of

TXRocks (containing most data); for example, the version number in a record can be erased after a long enough period of time.

## Lower write amplification

InnoDB uses the in-place change mode, where the entire page may be flushed to the disk even when only one row of data is changed, causing a high write amplification and more random writes.

TXRocks uses the append-only change mode, which has a lower write amplification; therefore, it is more friendly to devices such as SSD with a limited number of write cycles.

## Scenarios

TXRocks is very suitable for businesses that are sensitive to the storage costs, have much more writes than reads and a large data volume, and require a high transaction read/write performance.

## How to Use TXRocks

For more information, see [Instructions](#).

## Optimization and Subsequent Development

The TXSQL team has been optimizing TXRocks based on business needs; for example, they have improved the SUM query performance by over 30 times. They are also actively exploring the integration with new hardware to use AEP as the L2 cache for higher performance and cost-effectiveness.

As the storage engine of TencentDB for MySQL, TXRocks will be continuously optimized and improved with regard to problems encountered during use. The TXSQL team will also make technical explorations based on new hardware and release TXRocks in more key services as an important supplement to InnoDB.

# Instructions

Last updated: 2023-09-01 14:52:32

TXRocks is a transactional storage engine developed by Tencent's TXSQL team based on RocksDB. It saves more storage space and has a lower write amplification.

## Introduction

By leveraging the LSM tree storage structure of RocksDB, the TXRocks transactional storage engine not only reduces wastes caused by InnoDB's half-full pages and fragments, but also uses the compact storage format. Therefore, it has a performance comparable to that of InnoDB but requires only a half or even smaller storage space. It is more suitable for businesses with a large data volume and high requirements for the transactional read/write performance.

## Preparations

The database version must be MySQL 5.7 or 8.0 on a two-node architecture.

## Purchasing TencentDB for MySQL Instance (with RocksDB Engine)

You can select RocksDB as the engine when purchasing an instance on the TencentDB for MySQL [purchase page](#). For more information on other parameters, see [Creating MySQL Instance](#).

Database Version

MySQL5.6 MySQL5.7 **MySQL8.0** [Learn More](#)

TDSQL-C, the latest TencentDB product is recommended. Being 100% compatible with MySQL, it takes just a few seconds to add read-only instances. It allows you to upgrade/downgrade without data migration, and back up/roll back data using snapshots. Massive storage and auto-scaling are also supported. Pay-as-you-go billing is adopted.

Engine ⓘ

InnoDB **RocksDB** [Learn More](#)

Key-Value storage engine, with efficient writing and high compression

### Note

RocksDB is a key-value storage engine, with efficient writing and high compression. Currently, only TencentDB for MySQL 5.7 and 8.0 instances can use the RocksDB engine.

## Creating RocksDB Table

If RocksDB is selected as the default engine during instance creation, it will be the default engine used for table creation. You can run the following command to view the default engine:

```
show variables like '%default_storage_engine%';
```

```
mysql> show variables like '%default_storage_engine%';
+-----+-----+
| Variable_name      | Value  |
+-----+-----+
| default_storage_engine | RocksDB |
+-----+-----+
1 row in set (0.00 sec)
```

If the default engine is RocksDB, you cannot specify a storage engine in table creation statements:

```
mysql> create table tencent_db (id int primary key,c1 varchar(30),c2 varchar(50));
Query OK, 0 rows affected (0.00 sec)

mysql> show create table tencent_db;
+-----+-----+
| Table      | Create Table
+-----+-----+
| tencent_db | CREATE TABLE `tencent_db` (
  `id` int(11) NOT NULL,
  `c1` varchar(30) DEFAULT NULL,
  `c2` varchar(50) DEFAULT NULL,
  PRIMARY KEY (`id`)
) ENGINE=RocksDB DEFAULT CHARSET=utf8 |
+-----+-----+
1 row in set (0.00 sec)
```

After a table is created, its data will be stored in RocksDB and can be used in the same way as in InnoDB.

## Engine Feature Limits

TXRocks has certain limits on engine features as detailed below:

Function type	Feature	TXRocks Limit
DDL	Online DDL	Not supported. For example, <code>ALTER TABLE ... ALOGRITHM=INSTANT</code> is not supported. Only the <code>COPY</code> algorithm is supported for partition management operations.
SQL	Foreign key	Not supported.
	Partitioning Table	Not supported.
	Generated column	Not supported.
	Explicit DEFAULT expression	Not supported. For example, <code>CREATE TABLE t1(c1 FLOAT DEFAULT (RAND())) ENGINE=ROCKSDB</code> will fail, with the error <code>'Specited storage engine' is not supported for default value expressions reported.</code>
Index	Encrypted table	Not supported.
	Spatial index	The spatial index and spatial data types such as <code>GEOMETRY</code> and <code>POINT</code> are not supported.
	Full-text index	Not supported.
Replication	Multi-valued index	Not supported.
	Group replication	Not supported.
	Binlog format	Only the <code>ROW</code> format is supported, while <code>STMT</code> and <code>MIXED</code> formats are not.
	Clone plugin	Not supported.
	Transportable tablespace	Not supported.
	LOCK NOWAIT and SKIP LOCKED	Not supported.

Transaction and lock	Gap lock	Not supported.
	Savepoint	Savepoint
	Partial LOB field update	Not supported.
	XA transaction	Not recommended.

## Description

### Note

When creating a TencentDB for MySQL instance, after select RocksDB as the default storage engine, you can also customize the parameter template to suit your needs by following the parameter descriptions below.

## MySQL 5.7 parameter list

Parameter name	Whether to restart the instance for the parameter to take effect	Default value	Value Range/ Valid Values	Description
rocksdb_use_direct_io_for_flush_and_compaction	Supported	ON	ON/OFF	Whether to use DIO during compaction.
rocksdb_flush_log_at_trx_commit	Not required	1	0/1/2	<p>Controls when to write logs to the disk. It is similar to <code>innodb_flush_log_at_trx_commit</code> and indicates whether transactions need to be synced when being committed.</p> <ul style="list-style-type: none"> <li>• 0: Transactions are not synced when being committed.</li> <li>• 1: Transactions are synced when being committed.</li> <li>• 2: Transactions are synced once every second.</li> </ul>
rocksdb_lock_wait_timeout	Not required	1	1-1073741824	Lock wait timeout period in seconds.
rocksdb_deadlock_detect	Not required	ON	ON/OFF	Whether to enable deadlock detection. After it is enabled, all deadlock information will be recorded in mysqld error logs.
rocksdb_max_wal_size				If the total size of WAL files exceeds <code>rocksdb_max_wal_size</code> ,

<code>rocksdb_manual_wal_flush</code>	Supported	ON	ON/OFF	RocksDB will forcibly flush the column family to the disk to ensure that the oldest WAL file can be deleted.
---------------------------------------	-----------	----	--------	--

## MySQL 8.0 parameter list

Parameter name	Whether to restart the instance for the parameter to take effect	Default value	Value Range/ Valid Values	Description
rocksdb_flush_log_at_trx_commit	Not required	1	0/1/2	<p>Controls when to write logs to the disk. It is similar to <code>innodb_flush_log_at_trx_commit</code> and indicates whether transactions need to be synced when being committed.</p> <ul style="list-style-type: none"> <li>• 0: Transactions are not synced when being committed.</li> <li>• 1: Transactions are synced when being committed.</li> <li>• 2: Transactions are synced once every second.</li> </ul>
rocksdb_lock_wait_timeout	Not required	1	1–1073741824	Lock wait timeout period in seconds.
rocksdb_merge_buf_size	Not required	524288 (=512K)	100–18446744073709551615	Size of the merge-sort buffer used during secondary index creation.
rocksdb_merge_combine_read_size	Not required	8388608 (=8M)	524288(=512K)–18446744073709551615	Size of the memory used by k-way merge during secondary index creation.
rocksdb_deadlock_detect	Not required	ON	ON/OFF	Whether to enable deadlock detection.
				If the total size of WAL files exceeds <code>rocksdb_max_wal_size</code>

rocksdb_manual_wal_flush	Supported	ON	ON/OFF	<pre>x_total_wal_size,</pre> <p>RocksDB will forcibly flush the column family to the disk to ensure that the oldest WAL file can be deleted.</p>
--------------------------	-----------	----	--------	--

## RocksDB Monitoring Metrics

RocksDB monitoring metrics are as listed below:

Metrics	Description
rocksdb_bytes_read	Data read from disk
rocksdb_bytes_written	Data written to disk
rocksdb_block_cache_bytes_read	Blocks read
rocksdb_block_cache_bytes_write	Blocks written
rocksdb_wal_log_capacity	Data written to WAL log

# Cost Performance


Last updated: 2023-09-01 14:53:08

TXRocks has a performance comparable to that of InnoDB; however, its LSM tree structure can reduce wastes caused by InnoDB half-full pages and fragments, saving more storage space and delivering an ultra high cost performance.

## Background Information

TXRocks is used in TencentDB products as an important supplement to InnoDB. With a similar performance, it is further optimized and improved to save more storage space. Below is the comparison between the two engines in terms of space usage and performance.

## TXRocks Uses Less Space Than InnoDB

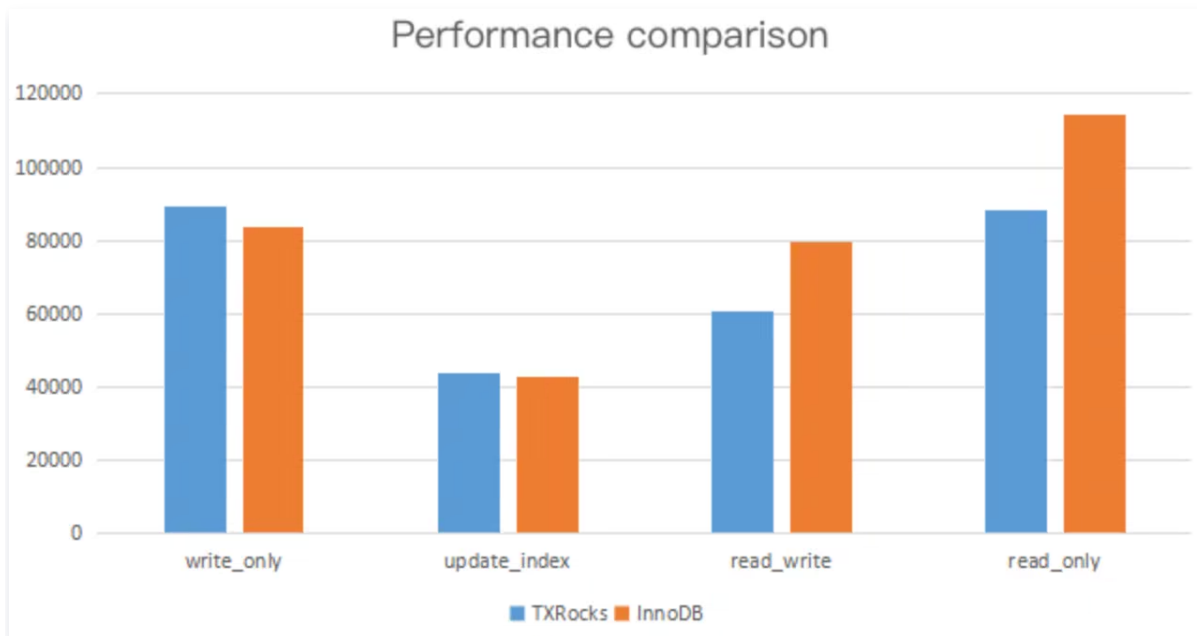


**Test scenario:** Both storage engines use the default configuration and the default table structure of sysbench. Each table contains 800,000 records, and the total number of tables gradually increases from 4 to 512.

The space usage of TXRocks and InnoDB under the specified test conditions is as shown above. The disk usage is displayed on the Y axis.

As shown in the test data, the greater the data volume, the slower the increase of disk usage by TXRocks, and the higher the storage space utilization of TXRocks (it uses only 42.71% of the space used by InnoDB in the best case). For data records with highly repetitive prefixes, TXRocks has a higher compression rate and storage cost performance.

## TXRocks and InnoDB Have a Similar Performance



**Test scenario:** An 8-core 32 GB MEM instance and six tables containing five million rows of data each are used for testing. Each test case is performed after cold instance restart and runs for 1,200 seconds.

The performance comparison between TXRocks and InnoDB under the specified test conditions is as shown above. You can see that TXRocks and InnoDB have a similar performance.

**Key parameters in the sysbench command:**

```
sysbench --table-size=5000000 --tables=6 --threads=32 --time=1200
```

## Summary

TXRocks is a TencentDB for MySQL storage engine that has a performance comparable to that of InnoDB but uses less space. It not only guarantees the business performance, but also reduces the storage costs. For more information, see [Overview](#)

# Best Practices

Last updated: 2023-09-01 14:53:35

This document describes how to accelerate the import of massive amounts of data to a database with the TXRocks engine.

## Example

- **Scenario:** The import of massive amounts of data to a database with the TXRocks engine needs to be accelerated.
- **Impact:** When importing massive amounts of data, a `Rows inserted during bulk load must not overlap existing rows` error may occur.

## Option 1

1. Delete secondary indexes and retain only the primary key index.
2. Adjust memory parameters based on the specification and data volume.

### Note

Appropriately increase the values of `rocksdb_merge_buf_size` and `rocksdb_merge_combine_read_size` parameters based on the specification and data volume.

- `rocksdb_merge_buf_size` indicates the data volume of each way in k-way merge during index creation. `rocksdb_merge_combine_read_size` indicates the total memory used in k-way merge.
- `rocksdb_block_cache_size` indicates the size of `rocksdb_block_cache`. We recommend you decrease its value temporarily during k-way merge.

3. Use bulk load to import the data.

```
SET session rocksdb_bulk_load_allow_unsorted=1;
SET session rocksdb_bulk_load=1;
...
Importing Data
...
SET session rocksdb_bulk_load=0;
SET session rocksdb_bulk_load_allow_unsorted=0;
```

### Note

If the imported data is sorted, you don't need to configure `rocksdb_bulk_load_allow_unsorted`.

#### 4. Recreate secondary indexes one by one after all data is imported.

##### Note

- Secondary index creation involves k-way merge. `rocksdb_merge_buf_size` indicates the data volume of each way, and `rocksdb_merge_combine_read_size` indicates the total memory used in k-way merge.
- For example, we recommend you set `rocksdb_merge_buf_size` to 64 MB or higher and set `rocksdb_merge_combine_read_size` to 1 GB or higher to avoid OOM. After all data is imported, you must modify the parameters to their original values.
- As a lot of memory is used during the creation of each secondary index, we recommend you not create many of them at the same time.

## Option 2

You can disable `unique_check` during data import to improve the import performance.

```
SET unique_checks=OFF;
...
Importing Data
...
SET unique_checks=ON;
```

##### Note

After the operation is completed, you must set `unique_checks` back to `ON`; otherwise, the uniqueness of INSERT operations in subsequent normal transaction writes will not be checked.

# Checking and Fixing Kernel Issues

## Summary of Critical Kernel-related Issues

Last updated: 2026-03-10 19:12:51

The self-developed kernel TXSQL used by TencentDB for MySQL is fully compatible with native MySQL and has continuously integrated and developed multiple enhancement features for diverse business scenarios such as gaming, education, finance, and industrial manufacturing. The TencentDB team regularly iterates and updates the kernel version to continuously improve performance, enhance stability, and fix known issues or official bugs. To keep your database instance in optimal operating condition, it is recommended that you stay informed and upgrade to the corresponding or latest kernel version in a timely manner. The following section lists some potential issues that may arise in kernel versions and their corresponding solutions for your reference and troubleshooting.

### Summary of Kernel-related Issues Causing Crashes

Issue	Affected Kernel Version	Fixed Kernel Version
<a href="#">Corruption in the Tablespace fseg not full list Caused by Purging or Rollbacks on Two Blob Fields</a>	MySQL 8.0 20220831 MySQL 8.0 20220401 MySQL 8.0 20220331 MySQL 8.0 20220330 MySQL 8.0 20211202 MySQL 8.0 20210830 MySQL 8.0 20210330 MySQL 8.0 20201230	MySQL 8.0 20221002
<a href="#">Rollback Failure After Upgrading to a Later Version When Instant ADD COLUMN Has Caused Page Data Usage to Exceed 50%</a>	MySQL 8.0 version 20221221 and earlier versions	MySQL 8.0 20241005
<a href="#">Crash Caused by Adding More Than 1024 Columns Using Instant ADD COLUMN</a>	MySQL 8.0 version 20221221 and earlier versions	MySQL 8.0 20241005
<a href="#">Crash Caused by Instant ADD COLUMN Followed by Online DDL on the Same Table</a>	MySQL 8.0 version 20221221 and earlier versions	MySQL 8.0 20241005
<a href="#">Crash Caused by Instant ADD COLUMN of a Column with the Same Name after Instant DROP</a>	MySQL 8.0 version 20221221 and earlier versions	MySQL 8.0 20241005
<a href="#">Crash Caused by a Redo Buffer Calculation Defect in Instant CHANGE COLUMN</a>	MySQL 8.0 version 20230704 and earlier versions	MySQL 8.0 20241005
<a href="#">Data Loss Caused by Table Rebuild Operations (ALTER/OPTIMIZE) in Specific MySQL Versions</a>	Versions from MySQL 8.0 20221215 (inclusive) to MySQL 8.0 20230703 (inclusive)	MySQL 8.0 20230704
<a href="#">Instant DDL Cannot Properly Handle Column Name Case Inconsistency</a>	MySQL 8.0 version 20241005 and earlier versions	MySQL 8.0 20241009
<a href="#">Memory Corruption Triggered When DDL Statements Are Executed Causes crash</a>	MySQL 5.7 versions before 20211102	MySQL 5.7 20250803
<a href="#">Instant DDL Caused by Out of</a>	MySQL 5.6 versions 20210830 and later	

Instance crash caused by Out-of-Bounds Access to `ha_alter_info->key_info_buffer`.

20210630 and later  
MySQL 5.7 version  
20211102 and earlier  
versions

MySQL 5.7  
20240331

# Corruption in the Tablespace fseg not full list Caused by Purging or Rollbacks on Two Blob Fields

Last updated: 2025-12-24 19:59:41

## Issue

Purging or rollbacks on two blob fields cause corruption in the tablespace fseg not full list.

## Issue Triggering Scenarios

Performing Data Manipulation Language (DML) operations on a table containing two columns of blob data may corrupt the tablespace fseg not full list. This can subsequently cause crashes during DML operations on the tablespace. Data in primary and standby databases may become corrupted simultaneously, leading to a crash of both instances.

## Affected Kernel Versions

MySQL 8.0 20220831, MySQL 8.0 20220401, MySQL 8.0 20220331, MySQL 8.0 20220330, MySQL 8.0 20211202, MySQL 8.0 20210830, MySQL 8.0 20210330, and MySQL 8.0 20201230.

## Fixed Kernel Version

MySQL 8.0 20221002.

## Check Methods

You can refer to the following commands to check for the issue.

- Query detailed column information of tables.

```
SELECT
  t1.TABLE_SCHEMA,
  t1.TABLE_NAME,
  t1.COLUMN_NAME,
  t1.DATA_TYPE
FROM
  INFORMATION_SCHEMA.COLUMNS t1
INNER JOIN (
  -- Subquery: First, find the tables that meet the condition
  (number of specific-type columns >= 2).
```

```

SELECT
    TABLE_SCHEMA,
    TABLE_NAME
FROM
    INFORMATION_SCHEMA.COLUMNS
WHERE
    DATA_TYPE IN ('blob', 'mediumblob', 'longblob', 'json',
'text', 'MEDIUMTEXT', 'LONGTEXT', 'geometry')
    AND TABLE_SCHEMA IN ('xxx', 'xxxx') -- Replace with actual
database names.
GROUP BY
    TABLE_SCHEMA,
    TABLE_NAME
HAVING
    COUNT(*) >= 2
) t2 ON t1.TABLE_SCHEMA = t2.TABLE_SCHEMA AND t1.TABLE_NAME =
t2.TABLE_NAME
WHERE
    -- Outer layer filters types again to ensure only risky columns
are listed.
    t1.DATA_TYPE IN ('blob', 'mediumblob', 'longblob', 'json', 'text',
'MEDIUMTEXT', 'LONGTEXT', 'geometry');

```

- **Query risk table information.**

```

SELECT
    TABLE_SCHEMA,
    TABLE_NAME,
    COUNT(*) as risky_column_count,
    GROUP_CONCAT(COLUMN_NAME) as risky_columns -- Optional:
Concatenate column names into a single row for display.
FROM
    INFORMATION_SCHEMA.COLUMNS
WHERE
    DATA_TYPE IN ('blob', 'mediumblob', 'longblob', 'json', 'text',
'MEDIUMTEXT', 'LONGTEXT', 'geometry')
    AND TABLE_SCHEMA IN ('xxx', 'xxxx')
GROUP BY
    TABLE_SCHEMA,

```

```
TABLE_NAME  
HAVING  
COUNT (*) >= 2;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20221002. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused by table locking, it is recommended to perform the repair during off-peak hours using tools such as pt-osc.

# Rollback Failure After Upgrading to a Later Version When Instant ADD COLUMN Has Caused Page Data Usage to Exceed 50%

Last updated: 2025-12-24 20:00:15

## Issue

After an upgrade to a later kernel version, the rollback fails when instant ADD COLUMN has caused page data usage to exceed 50%.

## Issue Triggering Scenarios

MySQL 8.0 20221221 and earlier versions allow instant ADD COLUMN operations even when the row size exceeds half the page size. After a kernel version upgrade, a rollback failure may occur due to row size exceeding the limit during the rollback.

## Affected Kernel Versions

MySQL 8.0 20221221 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241005.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241005. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused

by table locking, it is recommended to perform the repair during off-peak hours using tools such as `pt-osc`.

# Crash Caused by Adding More Than 1024 Columns Using Instant ADD COLUMN

Last updated: 2025-12-24 20:00:44

## Issue

A crash occurs when more than 1024 columns are added using instant ADD COLUMN.

## Issue Triggering Scenarios

MySQL 8.0 20221221 and earlier versions allow instant ADD COLUMN operations even when the row size exceeds half the page size. After a kernel version upgrade, a rollback failure may occur due to row size exceeding the limit during the rollback.

## Affected Kernel Versions

MySQL 8.0 20221221 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241005.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241005. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused by table locking, it is recommended to perform the repair during off-peak hours using tools such as pt-osc.

# Crash Caused by Instant ADD COLUMN Followed by Online DDL on the Same Table

Last updated: 2025-12-24 20:01:07

## Issue

Performing an instant ADD COLUMN operation followed by an online DDL operation on the same table causes a crash.

## Issue Triggering Scenarios

If an instant DDL operation is performed on a table, immediately followed by an online DDL operation such as OPTIMIZE TABLE or ALTER TABLE .. ADD PRIMARY KEY, a program crash may occur. This is because the code fails to handle the scenario correctly (failing to add columns or the REDUNDANT format).

## Affected Kernel Versions

MySQL 8.0 20221221 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241005.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241005. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused

by table locking, it is recommended to perform the repair during off-peak hours using tools such as `pt-osc`.

# Crash Caused by Instant ADD COLUMN of a Column with the Same Name after Instant DROP

Last updated: 2025-12-24 20:01:38

## Issue

Performing an instant ADD COLUMN operation on a column with the same name followed by an instant DROP operation causes a crash.

## Issue Triggering Scenarios

After an instant DROP operation is performed, the column is not actually deleted but is renamed and becomes a hidden column. If an instant ADD COLUMN operation is then performed to add a column with the same name as the hidden column, it can lead to metadata inconsistency, resulting in a crash.

## Affected Kernel Versions

MySQL 8.0 20221221 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241005.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241005. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused

by table locking, it is recommended to perform the repair during off-peak hours using tools such as `pt-osc`.

# Crash Caused by a Redo Buffer Calculation Defect in Instant CHANGE COLUMN

Last updated: 2025-12-24 20:02:04

## Issue

A redo buffer calculation defect in the instant CHANGE COLUMN operation causes a crash.

## Issue Triggering Scenarios

To calculate the size of log indexes, changes in column order should be considered, and the size of index information should also be recorded in the redo log. Originally, these columns were not accounted for. Performing an UPDATE operation immediately after executing an ALTER TABLE statement to modify column order causes a crash due to miscalculated redo log space.

## Affected Kernel Versions

MySQL 8.0 20230704 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241005.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241005. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused

by table locking, it is recommended to perform the repair during off-peak hours using tools such as `pt-osc`.

# Data Loss Caused by Table Rebuild Operations (ALTER/OPTIMIZE) in Specific MySQL Versions

Last updated: 2025-12-24 20:02:34

## Issue

Table rebuild operations (ALTER/OPTIMIZE) in specific MySQL versions cause data loss.

## Issue Triggering Scenarios

In MySQL 8.0.30, performing table rebuild operations (ALTER TABLE TABLE\_NAME ENGINE=INNODB and OPTIMIZE TABLE) may lead to data loss. For more information, see the MySQL official documentation: [MySQL Bugs: #115608](#) and [MySQL Bugs: #113812](#).

## Affected Kernel Versions

In TencentDB for MySQL, the kernel versions affected by this issue range from MySQL 8.0 20221215 (inclusive) to MySQL 8.0 20230703 (inclusive). You can query and check whether your database instance falls within this range.

## Fixed Kernel Version

MySQL 8.0 20230704.

## Fixing Methods

Upgrade the kernel version to MySQL 8.0 20230704. For the operation method, see [Upgrading the Kernel Minor Version](#).

# Instant DDL Cannot Properly Handle Column Name Case Inconsistency

Last updated: 2025-12-24 20:02:57

## Issue

Instant DDL cannot properly handle the inconsistent letter case of column names in DDL statements and CREATE TABLE statements.

## Issue Triggering Scenarios

This issue occurs when an instant DDL operation involves an INPLACE rename of a column name and its character set, and the case of the renamed column does not match the case defined in the original table schema.

## Affected Kernel Versions

MySQL 8.0 20241005 and earlier versions.

## Fixed Kernel Version

MySQL 8.0 20241009.

## Check Methods

You can refer to the following commands to check for tables with potential risks.

```
select replace(name, '/', '.') as 'table_schema.table_name' from
information_schema.innodb_tables where TOTAL_ROW_VERSIONS!=0 or
INSTANT_COLS!=0;
```

## Fixing Methods

1. Upgrade the kernel version to MySQL 8.0 20241009. For the operation method, see [Upgrading the Kernel Minor Version](#).
2. After a kernel version upgrade, rebuild and repair the corrupted tables. You can use the method `alter table xx engine = innodb;` for repair. To avoid business impact caused by table locking, it is recommended to perform the repair during off-peak hours using tools such as pt-osc.

# Triggered Memory Corruption When DDL Statements Are Executed, Caused Crash

Last updated: 2026-03-10 19:07:35

## Issue

Resolved an issue where executing DDL operations involving partitioned tables within a single SQL statement could trigger memory corruption, leading to crashes.

## Issue Triggering Scenarios

When executing DDL statements, if there are multiple DDL operations in a single SQL statement, the blob\_heap pointer from a previously executed DDL is assigned to m\_prebuilt->blob\_heap, but this pointer is released, and subsequent DDL operations then use this blob\_heap, resulting in memory corruption and causing a Crash.

## Affected Kernel Versions

MySQL 5.7 versions before 20211102.

## Fixed Kernel Version

MySQL 5.7 20250803.

## Fixing Method

To upgrade the kernel version to MySQL 5.7 20250803, see [Upgrade the Kernel Minor Version](#) for instructions.

# The Out-of-Bounds Access to `ha_alter_info->key_info_buffer` Causes Instance Crash

Last updated: 2026-03-10 19:09:42

## Issue

Concurrent execution of DDL statements and repeated INSERT operations result in out-of-bounds access to `ha_alter_info->key_info_buffer`, causing an instance Crash.

## Issue Triggering Scenarios

Rebuilding tables without primary keys may trigger out-of-bounds access to `key_info_buffer` when a duplicate row is inserted, causing an instance Crash.

## Affected Kernel Versions

- MySQL 5.6 versions 20210630 and later.
- MySQL 5.7 versions 20211102 and earlier.

## Fixed Kernel Version

MySQL 5.7 20240331.

## Fixing Method

1. To upgrade the kernel version to MySQL 5.7 20240331, see [Upgrade the Kernel Minor Version](#) and [Upgrade the Database Version for MySQL 5.6](#) for instructions.
2. Alternatively, you can add a primary key to a table without one. For the operation, see [The primary and secondary instances have inconsistent query data](#).