

TencentDB for MySQL

Kernel Features

Product Documentation



Copyright Notice

©2013-2025 Tencent Cloud. All rights reserved.

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

Trademark Notice



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

Service Statement

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

Contents

Kernel Features

Overview

Kernel Version Release Notes

TXSQL Engine Kernel Version Release Notes

TXRocks Engine Kernel Version Release Notes

TXSQL Engine Kernel Version Release Notes

Database Proxy Version Release Notes

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
- Second-Level Column Modification
- Async Deletion of Big Tables
- Hotspot Update
- SQL Throttling
- Statement Outline

TXRocks Engine

- Overview
- Instructions
- Cost Performance
- Practical Tutorial of TXRocks

LibraDB Engine

- LibraDB Engine Feature Highlights
- ETL Writeback Acceleration
- Ordered Pagination Feature

Kernel Features

Overview

Last updated : 2025-05-09 14:14:41

The database team at Tencent Cloud supports three distinct kernel branches for MySQL series products: the TXXSQL kernel, the TXRocks kernel, and the LibraDB kernel.

TXXSQL is a MySQL kernel branch maintained by the Tencent Cloud database team, fully compatible with the native MySQL version. It offers many features similar to those in the MySQL Enterprise Edition, including enterprise-level transparent data encryption, auditing, dynamic thread pooling, encryption functions, backup and recovery, and parallel queries.

TXXSQL not only enhances the InnoDB engine, query performance, replication capabilities, and more but also improves the usability and maintainability of TencentDB for MySQL. In addition to offering all MySQL features, it provides enterprise-level advanced features such as disaster recovery, restoration, monitoring, performance optimization, read-write separation, transparent data encryption, and auditing.

TXRocks is a transactional storage engine developed by Tencent's TXXSQL team, built upon the RocksDB engine. By leveraging the LSM tree storage structure of RocksDB, TXRocks minimizes wastage caused by partially filled InnoDB pages and fragmentation, while also supporting a compact data format. Consequently, TXRocks can achieve storage efficiency, using only half or even less space compared to InnoDB, without compromising on performance. It is particularly well-suited for applications that demand high transaction read/write performance and manage large volumes of data.

LibraDB is an analytical database engine developed independently by Tencent. It is compatible with the MySQL protocol in a columnar storage environment and supports features such as a highly scalable parallel and vectorized execution engine. Primarily designed to overcome the performance challenges associated with complex SQL queries, LibraDB also features integrated HTAP capabilities, making it suited for real-time reporting, complex analysis, and big data processing scenarios.

More information about product kernel:

For details on the TXXSQL kernel updates for TencentDB for MySQL, see [TXXSQL Kernel Release Notes](#).

For details on the TXRocks kernel updates for TencentDB for MySQL, see [TXRocks Kernel Release Notes](#).

For details on the LibraDB kernel updates for TencentDB for MySQL, see [LibraDB Engine Kernel Version Release Notes](#).

TencentDB for MySQL supports automatic or manual upgrades of kernel minor versions. For details, see [Upgrade Kernel Minor Version](#).

TencentDB for MySQL allows you to log in to the MySQL instance through CVM to view the kernel's minor version. For details, see [View Kernel Minor Version](#).

Kernel Version Release Notes

TXSQL Engine Kernel Version Release Notes

Last updated : 2025-06-03 18:25:40

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

MySQL 5.7 Kernel Version Release Notes

MySQL 5.6 Kernel Version Release Notes

TXSQL Kernel Version	Community Version	Description
20241001	8.0.30	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 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: MySQL 8.0.29 Character Set Support Changes in MySQL Connector/NET 8.0.28 Bug fixes 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.
20240930	8.0.30	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 and will report the error "Character Set 'utf8mb3' is not supported by .Net Framework" if utf8mb3 is used. If the

		<p>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: MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>New features</p> <p>Support "show full binary logs" to show the last modification time.</p> <p>Support the parallel execution feature of multi-stage aggregation operation and execution.</p> <p>Parallel Copy DDL supports partition tables.</p> <p>Eliminate duplicate aggregate functions.</p> <p>Eliminate common subexpressions.</p> <p>Expand subquery cache to support EXISTS/IN subqueries.</p> <p>Correlation statistics of multiple columns.</p> <p>Index dive pruning optimization.</p> <p>SQL traffic throttling supports startup loading.</p> <p>Performance optimization</p> <p>Fix query cache bugs and optimize performance.</p> <p>Parallelize the process of checking indexes to speed up the check during backup.</p> <p>Bug fixes</p> <p>Fix the fast cleanup AHI crash issue.</p> <p>Fix the issue that "drop database" can cause replication interruption during the use of the recycle bin when the database contains foreign key tables.</p> <p>Fix the compatibility issue that the instant metadata column of the partition table is upgraded from 5.7 to 8.0.</p> <p>Limit the execution of XA COMMIT/ROLLBACK on RO to prevent replication interruption caused by executing related operations.</p> <p>Fix the issue that TRUNCATE SUBPARTITION TEMPLATE causes a crash.</p> <p>Fix the issue that performing the rollback operation on the partition table causes a crash.</p> <p>Fix the issue that subpartition by range columns causes a crash for having multiple same columns.</p> <p>Fix the issue that statistical information may be cleared after COPY DDL is executed.</p> <p>Fix the issue that the query rewrite plugin causes a crash when it is used in multi-statement mode.</p>
20230704	8.0.30	<p>Note:</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: MySQL 8.0.29 Character Set Support</p>

		<p>Changes in MySQL Connector/NET 8.0.28</p> <p>Bug fixes</p> <p>Fixed an issue where rebuilding a table may cause data loss in some extreme cases. For details, see Bug#110706.</p>
20230703	8.0.30	<p>Note:</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: MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>Bug fixes</p> <p>Fixed the memory allocation method at the location of the memory leak to prevent memory leakage.</p>
20230702	8.0.30	<p>Note:</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: MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>Performance optimizations</p> <p>Reduced the overhead caused by frequent notifying of parallel query exchange operators.</p> <p>Bug fixes</p> <p>Fixed the abnormal issue of the instance crash caused by an instant DDL's modifying a column position.</p>
20230701	8.0.30	<p>Note:</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: MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>Bug fixes</p> <p>Supports initiating TXSQL physical backup using the community version.</p> <p>Replaced all keywords in the AWR feature with "TXSQL_AWR".</p>

		Fixed an issue where the background thread competing with RENAME on the dict cache could cause a failure to open the underlying table.
20230630	8.0.30	<p>Note :</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> <p>MySQL 8.0.29 Character Set Support Changes in MySQL Connector/NET 8.0.28</p> <p>New features</p> <p>Supported Nonblocking DDL feature.</p> <p>Supported xa commit to record the maximum gts instance TP/AP load statistics in relay log.</p> <p>Supported selecting Innodb temporary tables for parallel query of worker thread sharing.</p> <p>Supported using partition tables as parallel tables for parallel queries.</p> <p>Supported the flashback version query feature.</p> <p>Supports persistence for flashback query.</p> <p>Supported virtual indexes.</p> <p>Supported the range/list secondary partition feature.</p> <p>Supported the automatic relay log recovery feature.</p> <p>Supported the default algorithm for DDL, with options INPLACE/INSTANT.</p> <p>Supported Fast Query Cache.</p> <p>Supported the conversion of partition tables from MyISAM to InnoDB.</p> <p>Supported the correlated subquery cache feature.</p> <p>Performance Optimization</p> <p>Optimized the BINLOG LOCK_done lock conflict.</p> <p>Optimized thread pool performance.</p> <p>Optimized the issue where disabling eq_ref cache in outer join leads to performance regression.</p> <p>Enhanced parallel query:</p> <p>Subquery \derived table executed in parallel independently: Optimized and executed plan for subquery \derived table in parallel, independent of main query execution.</p> <p>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.</p> <p>Hash join (in memory) executed in parallel: Work threads build complete hash tables separately, with parallel scanning on the probe end.</p> <p>Parallel query supports global aggregation optimization.</p> <p>Parallel query supports pushdown parallelism under having condition.</p>

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.

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.

		<p>Fixed the issue in parallel DDLs where adding a unique key to a column with duplicate data caused the instance to run abnormally.</p> <p>Fixed the issue of assertion in parallel hash join debug.</p> <p>Fixed the issue in parallel cost calculation where NDV was 0.</p> <p>Fixed the issue with JSON import accuracy.</p> <p>FORCE INDEX ORDER BY statement skips the index dive bug.</p> <p>Fixed the issue where the official subquery plan was displayed multiple times.</p> <p>Fixed the issue where the disk-based temporary table quantity does not increase.</p> <p>Fixed the deadlock issue caused by proxy change user.</p> <p>Fixed the issue where calling a stored procedure in a trigger due to permission verification optimization caused permission checks to be bypassed.</p>
20221221	8.0.30	<p>Note :</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> <p>MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>Bug Fixes</p> <p>Fixed the issue where after enabling log_slave_updates on a secondary node, the thread_id of the event written to the binlog on the secondary node changed.</p>
20221220	8.0.30	<p>Note :</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> <p>MySQL 8.0.29 Character Set Support</p> <p>Changes in MySQL Connector/NET 8.0.28</p> <p>Bug Fixes</p> <p>Fixed the instant DDL bug.</p>
20221215	8.0.30	<p>Note :</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</p>

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 nl(), 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 hotspots, and reduced lock holding time.

Optimized xa::Transaction_cache, and reduced lock conflicts.

Reduced ha_innopart::clear_blob_heaps time consumption.

Optimized purge threads lock hotspots, and reduced tasks_mutex and thread conflicts.

Optimized Buffer Pool initialization, supporting parallel initialization, and accelerating initialization speed.

Optimized read_only and select performance under high concurrency.

Optimized permission validation for prepared statement and stored procedure.

Optimized access to change buffer.

Avoided unnecessary calls to fil_space_get, and reduced FAQs in extreme scenarios.

Optimized lock conflict for GTID during transaction commit when binlog_order_commits is disabled.

Applied Lock Free Hash to optimize trx_sys mutex conflict.

Optimized the overhead of taking a snapshot in the transaction system.

Optimized Writerset and improved performance.

Replaced index drill-down with histogram.

Supports Parallel DDL.

Bug Fixes

Fixed issues with abnormal statistical values such as innodb_row_lock_current_waits.

Fixed the issue of excessively high memory usage with Group concat with group by.

		<p>Fixed the issue of statistical information being severely underestimated in long records.</p> <p>Fixed the issue in parsing stored procedure syntax.</p> <p>Fixed the issue in FAST DDL optimization of flush list to release page concurrency.</p>
20220831	8.0.22	<p>New features</p> <p>Supported setting the MySQL version dynamically.</p> <p>Supported transparent column encryption. When creating a table, you can specify the encryption attribute for the `varchar` field, and the storage system will encrypt the column. This capability is expected to be commercialized in 2023.</p> <p>Fixed the exception of the third-party data subscription tool caused by subscription to the comparison SQL for internal data consistency during tool usage.</p> <p>Note :</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Supported enhanced MTS deadlock detection.</p> <p>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.</p> <p>Performance optimizations</p> <p>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 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> <p>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</p>

permission change, the system will skip the permission check; otherwise, it will check the permission and record the version number again.

Optimized the accuracy of time acquisition in the thread pool.

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 ``rec_get_offsets()`` function.

Optimized parallel DDL.

1. When the index field is small, the sampled memory size is reduced to lower the sampling frequency.
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.
3. When records are read, the fixed-length offset is cached in order to avoid generating offsets for each record each time.

Optimized the undo log information recording logic to improve the INSERT performance.

Improved the performance after semi-sync was enabled.

Optimized the audit performance to reduce the system overheads.

Bug fixes

Fixed the issue where the displayed value of ``Thread_memory`` was abnormal sometimes.

Fixed the issue where the timestamp was inaccurate during batch statement audit.

Fixed issues related to column modification at the second level.

Fixed the issue where the ``CREATE TABLE t1 AS SELECT ST_POINTFROMGEOHASH("0123", 4326);`` statement caused source-replica disconnection.

Fixed the issue where the replica failed to retry during concurrent requests at the table level.

Fixed the ``Malformed packet`` error reported when ``show slave hosts`` was executed.

Fixed recycle bin issues.

Fixed the issue where the jemalloc mechanism easily triggered OOM on ARM device models.

Fixed the issue where ``truncate pfs account table`` caused the failure to collect statistics.

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.

Fixed `sql_mode` log issues.

Fixed the occasional issue where a procedure became abnormal when ``CREATE DEFINER`` was executed.

Fixed Copy Free Snapshot issues.

Fixed the performance fluctuation of the thread pool.

Fixed the issue where the result of ``hash join+union`` might be empty.

Fixed memory issues.

20220401	8.0.22	<p>Bug fixes</p> <p>Fixed the issue where the stage variable error in Parallel DDL caused the stage null pointer to crash when creating FTS indexes.</p> <p>Fixed the possible crash when adding full-text indexes.</p>
20220331	8.0.22	<p>Bug fixes</p> <p>Fixed the crash caused by dereferencing wild pointers in the thread pool.</p>
20220330	8.0.22	<p>New features</p> <p>Enabled writeset parallel replication by default.</p> <p>Supported extended resource groups to control the I/O, memory utilization, and SQL timeout policy by user.</p> <p>Supported flashback query to query data at any time point within the UNDO time range.</p> <p>Supported `RETURNING` in a `DELETE`, `INSERT`, or `REPLACE` statement to retrieve the data rows modified by the statement.</p> <p>Supported the GTID replication feature extension in row mode.</p> <p>Supported transaction lock optimization.</p> <p>Enhanced the recycle bin to support TRUNCATE TABLE and automatic cleanup of tables in the recycle bin.</p> <p>Supported parallel DDL to speed up DDL operations for which to create indexes through three-phase parallel operations.</p> <p>Supported quick index column modification.</p> <p>Supported automatic statistics collection and cross-server statistics collection.</p> <p>Performance optimizations</p> <p>Optimized the GTID lock conflicts when transactions were committed if `binlog_order_commits` was disabled.</p> <p>Accelerated MySQL startup by changing the InnoDB startup phase from single-threaded creation of Rsegs to multi-threaded creation.</p> <p>Bug fixes</p> <p>Fixed the issue where the transaction did not end when the connection was closed after deadlock or lock wait.</p> <p>Fixed the issue where the `innodb_row_lock_current_waits` value was abnormal.</p> <p>Fixed the SQL type error in the audit plugin without USE DATABASE.</p> <p>Fixed the issue where tables smaller than `innodb_async_table_size` were also renamed during async drop of big tables.</p> <p>Fixed the issue with incorrect escape characters in the audit plugin.</p> <p>Fixed the issue of rollback after quick column modification.</p> <p>Fixed the issue where the transaction system (trx_sys) may crash if it contains XA transactions when it is closed.</p> <p>Fixed the crash when merging derived tables.</p> <p>Fixed the issue where `binlog_format` was modified after writeset was enabled.</p>

		<p>Fixed the error (error code: 1032) caused by hash scans with A->B->A->C update on the same row.</p> <p>Fixed the issue where the sort index might be invalid in prepared statement mode.</p> <p>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.</p> <p>Fixed exceptions in extreme cases for async drop of big tables.</p> <p>Fixed the abnormal error message when setting a SQL filter.</p> <p>Fixed the syntax error reported during stored procedure parsing.</p> <p>Fixed the issue where historical histograms couldn't be applied.</p> <p>Fixed the role column display compatibility issue caused by `SHOW SLAVE HOSTS(show replicas)`.</p> <p>Fixed the crash of `Item_in_subselect::single_value_transformer` when the number of columns was incorrect.</p> <p>Fixed the crash caused by memory leaks during cascading update if a subtable contained virtual columns and foreign key columns.</p>
20211202	8.0.22	<p>New features</p> <p>Supported quick column modification.</p> <p>Supported histogram versioning.</p> <p>Supported SQL:2003 TABLESAMPLE (single table) sampling control syntax for obtaining random samples of physical tables.</p> <p>Added non-reserved keywords: TABLESAMPLE BERNOULLI.</p> <p>Added the `HISTOGRAM()` function to build a histogram for a given input field.</p> <p>Supported compressed histograms.</p> <p>Supported SQL throttling.</p> <p>Supported MySQL cluster role configuration (default role: CDB_ROLE_UNKNOWN).</p> <p>Added a new `Role` column to the `show replicas` command's display results to display roles.</p> <p>Supported proxy.</p> <p>Performance optimizations</p> <p>Optimized the hotspot update problem caused by `insert on duplicate key update`.</p> <p>Accelerated the application of hash scan by aggregating multiple identical binlog events.</p> <p>Greatly reduced the memory usage by the `PREPARE` statement in point queries in the thread pool mode when the plan cache was enabled.</p> <p>Bug fixes</p> <p>Fixed the error of unstable performance after hotspot update optimization was enabled.</p> <p>Fixed the issue where `select count(*)` parallel scans caused full-table scans in extreme cases.</p>

		<p>Fixed performance issues caused by execution plan changes due to reading zero statistics in various cases.</p> <p>Fixed the bug where queries were in the `query end` status for a long time.</p> <p>Fixed the bug where statistics were severely underestimated in long records.</p> <p>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.</p> <p>Fixed the case sensitivity issue of column names in the `json_table` function.</p> <p>Fixed the bug that caused correctness issues in window functions because expressions returned early during `return true`.</p> <p>Fixed the correctness issue caused by the pushdown by `derived condition pushdown` when it contained user variables.</p> <p>Fixed the issue where SQL filters were prone to crash when no namespaces were added in a rule.</p> <p>Fixed the QPS jitters when the thread pool was enabled under high concurrency and high conflict.</p> <p>Fixed the issue where source-replica buffer pool sync leaked file handles in extreme cases (when host file systems were corrupted).</p> <p>Fixed the index mapping issue.</p> <p>Fixed the statistics cache sync issue.</p> <p>Fixed the crash when information was not cleared during execution of the `UPDATE` statement or stored procedures.</p>
20210830	8.0.22	<p>New features</p> <p>Supported limiting the number of preloaded rows.</p> <p>Supported optimizing plan cache point query.</p> <p>Supported extended ANALYZE syntax (UPDATE HISTOGRAM c USING DATA 'json') and direct writes to histograms.</p> <p>Performance optimizations</p> <p>Replaced index seek with histogram to reduce evaluation errors and I/O overheads (this capability is not enabled by default).</p> <p>Bug fixes</p> <p>Fixed the issue where there might be no statistics information during online DDL.</p> <p>Fixed the issue where generated columns on replica instances were not updated.</p> <p>Fixed the issue where the instance hung when binlog was compressed.</p> <p>Fixed the issue of missing GTID in the previous_gtid event of the newly generated binlog file.</p> <p>Fixed possible deadlocks when system variables were modified.</p> <p>Fixed the issue where the information of the SQL thread of the replica instance in SHOW PROCESSLIST was incorrectly displayed.</p> <p>Implemented the bug fix related to hash join provided in MySQL 8.0.23.</p> <p>Implemented the bug fix related to writeset provided in MySQL.</p>

		<p>Implemented the bug fix related to the query optimizer provided in MySQL 8.0.24.</p> <p>Fixed the concurrency bugs of optimizing flush list and releasing pages in FAST DDL.</p> <p>Optimized the memory usage during data dictionary update in instances with a large number of tables.</p> <p>Fixed the crash caused by new primary key creation after INSTANT ADD COLUMN.</p> <p>Fixed the OOM caused by memory growth in full-text index query.</p> <p>Fixed the issue where -1 was included in the TIME field in the result set returned by SHOW PROCESSLIST.</p> <p>Fixed the issue where tables might fail to be opened due to histogram compatibility.</p> <p>Fixed the floating point accumulation error when Singleton histograms were constructed.</p> <p>Fixed the replication interruption caused by using many Chinese characters in the table name of a row format log.</p>
20210330	8.0.22	<p>New features</p> <p>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.</p> <p>Supported sort-merge join.</p> <p>Supported FAST DDL operations.</p> <p>Supported querying the value of the <code>`character_set_client_handshake`</code> parameter.</p> <p>Performance optimizations</p> <p>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.</p> <p>Bug fixes</p> <p>Fixed the deadlocks caused by the modification of the <code>`offline_mode`</code> and <code>`cdb_working_mode`</code> parameters.</p> <p>Fixed the persistent concurrency issue of the <code>`max_trx_id`</code> field in <code>`trx_sys`</code> table.</p>
20201230	8.0.22	<p>New features</p> <p>Supported the official updates of MySQL 8.0.19, 8.0.20, 8.0.21, and 8.0.22.</p> <p>Supported dynamic setting of thread pooling mode or connection pooling mode by using the <code>`thread_handling`</code> parameter.</p> <p>Performance optimizations</p> <p>Optimized the <code>`BINLOG LOCK_done`</code> conflict to improve write performance.</p> <p>Optimized the <code>`trx_sys mutex`</code> conflict by using lock-free hash to improve performance.</p>

		<p>Optimized redo log flushing.</p> <p>Optimized the buffer pool initialization time.</p> <p>Optimized the clearing of adaptive hash indexes (AHI) during the `drop table` operations on big tables.</p> <p>Optimized audit performance.</p> <p>Bug fixes</p> <p>Fixed performance fluctuation when cleaning InnoDB temporary tables.</p> <p>Fixed the read-only performance decrease when the instance has many cores.</p> <p>Fixed the error (error code: 1032) caused by hash scans.</p> <p>Fixed concurrency security issues caused by hotspot update.</p>
20200630	8.0.18	<p>New features</p> <p>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, submit a ticket.</p> <p>Supported automatic killing of idle tasks to reduce resource conflicts. To apply for this feature, submit a ticket.</p> <p>Supported transparent data encryption (TDE).</p> <p>Bug fixes</p> <p>Fixed the issue where switch failed due to inconsistent positions between `relay_log_pos` and `master_log_pos`.</p> <p>Fixed the data file error caused by asynchronously storing data in the disk.</p> <p>Fixed the hard error when `fsync` returned `EIO` and retries were made repeatedly.</p> <p>Fixed the crash caused by phrase search under multi-byte character sets in full-text index.</p>

TXSQL Kernel Version	Community Version	Description
20240331	5.7.44	<p>New features</p> <p>Merged Official 5.7.36,5.7.37,5.7.38,5.7.39,5.7.40,5.7.41,5.7.42,5.7.43,5.7.44 Changes.</p> <p>Added update returning feature.</p> <p>Added a DDL progress display. During DDL operations, the <code>show detail processlist</code> command can be run to view it.</p> <p>Added the default table encryption feature.</p> <p>Performance optimizations</p> <p>Optimized the SQL_TYPE audit accuracy.</p> <p>Optimized index drill-down pruning.</p> <p>Bug fixes</p> <p>Fixed an issue where insert returning statements caused slave crashes.</p> <p>Fixed an issue where the slave replay thread got stuck due to MASTER_DELAY.</p>

		<p>Fixed an issue where setting custom variables using a function caused session track errors.</p> <p>Fixed an issue where hash scans consumed a large amount of memory.</p> <p>Fixed an inconsistency when dropping/truncating a non-existent table with the recycle bin enabled versus disabled.</p> <p>Fixed an issue where pulling binlog resulted in garbled characters when binlog_checksum was disabled.</p> <p>Fixed the Parallel Copy DDL bug.</p>
20230601	5.7.36	<p>New Features</p> <p>Supports persistence for flashback query.</p> <p>Supported drop table force, enabling drop innodb metadata.</p> <p>Supported Parallel Copy DDL.</p> <p>Supported limit in subquery.</p> <p>Supported the conversion of partition tables from MyISAM to InnoDB.</p> <p>Bug Fixes</p> <p>Fixed the issue of index anomaly in primary-secondary BP synchronization feature.</p> <p>Fixed the issue where killing connections during large transactions caused anomalies.</p> <p>Fixed the issue of obtaining user-defined variable string errors in session track.</p> <p>Fixed the issue of failure to create index when parallel DDL is enabled and innodb_disable_sort_file_cache is set.</p> <p>Fixed some errors with instant modify column.</p>
20230115	5.7.36	<p>New Features</p> <p>Supported Nonblocking DDL feature.</p> <p>Supported validate password plugin.</p> <p>Supported for storing historical deadlock information.</p> <p>Performance Optimization</p> <p>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.</p> <p>Bug Fixes</p> <p>Fixed the issue where creating a user with grant identified by failed, causing primary/standby interruption.</p> <p>Fixed the issue where GROUP_CONCAT did not correctly set USED_TABLES when the DERIVED_MERGE switch was enabled.</p> <p>Fixed the issue where the gtid_subset function failed to correctly handle null_value.</p> <p>Fixed the issue where dummy index cache failed to initialize system columns.</p> <p>Fixed the issue of instant add column in partition table exceeding the maximum number of columns.</p> <p>Fixed the issue where canal pulling binlog may cause OOM.</p>

		<p>Fixed the error in proxy when reusing connections with different users.</p> <p>Fixed the proxy's incorrect responses for row count, found rows, and db settings.</p> <p>Fixed the issue where the error messages during binlog sending and receiving were incomplete.</p> <p>Fixed anomalies in paging and pushdown calculations.</p> <p>Fixed potential invalidity of m_page after creating a subtree with Parallel DDL.</p> <p>Fixed the crash issue with instant modify under certain character sets.</p>
20220716	5.7.36	<p>New features</p> <p>Supported auto-increment column persistence for InnoDB.</p> <p>Supported precise memory statistics.</p> <p>Supported query-level memory monitoring.</p> <p>Supported recycle bin.</p> <p>Supported parallel DDL statements.</p> <p>Supported flashback query.</p> <p>Supported async rollback for internal XA transactions.</p> <p>Performance optimizations</p> <p>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.</p> <p>Bug fixes</p> <p>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.</p> <p>Fixed the issue where <code>buf/buf/pool</code> couldn't be viewed in PFS memory monitoring data.</p> <p>Fixed the issue where the returning statement might cause exceptions in some scenarios due to permission checks.</p> <p>Fixed the issue where an error was reported because the parser did not correctly handle semicolons in statements.</p> <p>Fixed the issue where single quotation marks in audit statements were not escaped.</p> <p>Fixed the issue of sudden memory usage increase on the ARM platform.</p> <p>Fixed the issue of source-replica inconsistency caused by modifying <code>`binlog_format`</code> after <code>writeset</code> was enabled.</p> <p>Fixed the issue of high CPU usage caused by exiting a large number of threads at the same time.</p> <p>Fixed bugs related to <code>`drop table partition force`</code>.</p> <p>Fixed the issue where binlog dump got stuck and caused the instance restart to become slow.</p> <p>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.</p> <p>Fixed the issue where <code>`show detail processlist`</code> displayed illegal characters.</p> <p>Fixed the issue where the <code>`thread_group`</code> lock was not released when the thread pool was closed in some cases.</p>

		<p>Fixed the issue where updating the parent table at the parallel table level caused the instance to run abnormally.</p> <p>Fixed the issue where virtual columns were calculated incorrectly on the replica.</p> <p>Fixed the issue where `gtid_subset` did not set `null_value` to `false` after executing a row.</p>
20211230	5.7.36	<p>New features</p> <p>Supported the official updates of MySQL 5.7.19–5.7.36.</p> <p>Supported source-replica buffer pool sync to speed up the performance recovery after HA switch (around 90 seconds faster than that in native mode).</p> <p>Added the backup lock feature to provide lightweight metadata locks to improve the service availability during backup.</p> <p>Performance optimizations</p> <p>Made functions related to `utf8/utf8mb4 my_charpos` inline to optimize the performance of UTF_8 functions in read_write scenarios.</p> <p>Upgraded jemalloc to v5.2.1.</p> <p>Optimized file number acquisition during binlog rotation.</p> <p>Optimized semi-sync replica I/O.</p> <p>Optimized hash scan aggregation.</p> <p>Accelerated the startup of crash recovery for large transactions.</p>
20211102	5.7.18	<p>New features</p> <p>Fixed the exception of the third-party data subscription tool caused by subscription to the comparison SQL for internal data consistency during tool usage.</p> <p>Note :</p> <p>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>
20211031	5.7.18	<p>New features</p> <p>Supported writeset replication.</p> <p>Performance optimizations</p> <p>Optimized the checkpoint mechanism to increase the backup success rate.</p> <p>Optimized the hash scan index selection.</p> <p>Optimized the hotspot update performance to support `insert on duplicate key update`.</p> <p>Bug fixes</p> <p>Fixed the error of unstable performance after hotspot update was enabled.</p> <p>Fixed the crash caused by rolling back the UPDATE operation after an instant DDL.</p>

		<p>Fixed the issue where the `CREATE TABLE AS SELECT` statement didn't inherit the compression attribute after column compression was enabled.</p> <p>Fixed the instance crash caused by the `show variables like 'tencent_root%'` statement after the `skip-grant-table` option was enabled.</p> <p>Fixed the crash of the Query Rewriter plugin in read-only mode.</p> <p>Fixed the error (error code: 1032) caused by hash scans in partitioned tables.</p> <p>Fixed the issue where the first large transaction's SBM was 0 in MTS mode.</p> <p>Fixed the crash of `stop slave` caused by `slave_preserve_commit_order=ON, slave_transaction_retries=0`.</p> <p>Fixed several XA transaction bugs.</p> <p>Fixed the issue where SQL splicing went wrong during `show create` after a JSON field with a default value was created.</p> <p>Fixed the issue where disconnected transactions could not be rolled back after transactions were blocked.</p> <p>Fixed the issue where there might be no statistics information for long records in InnoDB persistent mode.</p> <p>Ported 8.0 to fix the issue where `ANALYZE TABLE` might cause query retention.</p> <p>Fixed the issue where the InnoDB statistics couldn't be synced to the server layer in time after change.</p> <p>Fixed the issue where statistical sampling might block writes for too long and cause a crash (bug# 31889883).</p> <p>Fixed the possibility of reading zero rows during the InnoDB statistics update process (bug# 105224).</p> <p>Fixed the possible $O(N^2)$ behavior in MVCC (bug# 28825617).</p> <p>Fixed the crash caused by closing a temp table and triggering binlog rotation when a connection was released.</p>
20210630	5.7.18	<p>New features</p> <p>Added the new command SHOW SLAVE DETAIL [FOR CHANNEL channel] for displaying the binlog timestamp that the current replica has replayed.</p> <p>Supported transaction_read_only/transaction_isolation parameters.</p> <p>Performance optimizations</p> <p>Accelerated the application of hash scan on replicas by aggregating multiple identical binlog events.</p> <p>Bug fixes</p> <p>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.</p> <p>Fixed the issue where there might be no statistics information during the DDL process.</p> <p>Fixed the inaccurate undo log size in connection status statistics.</p> <p>Fixed the instance crash caused by querying the metadata_locks table.</p> <p>Modified `of` as a non-reserved keyword.</p>

		<p>Fixed the issue where the dynamically modified version number was not invalidly displayed in new connections.</p> <p>Fixed the issue where the wild pointer was accessed during page_cache cleaning.</p> <p>Fixed the issue where the execution of ALTER TABLE might report the "Incorrect key file for table" error.</p> <p>Fixed the excessive memory usage by partitioned tables.</p> <p>Fixed the issue where -1 was included in the TIME field in the result set returned by SHOW PROCESSLIST.</p> <p>Fixed the lock wait of XA transaction replication on replica nodes.</p> <p>Fixed the incorrect lock of partitioned tables in equal range query.</p>
20210331	5.7.18	<p>New features</p> <p>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.</p> <p>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 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.</p> <p>Supported querying the value of the `character_set_client_handshake` parameter.</p> <p>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.</p> <p>Performance optimizations</p> <p>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%.</p> <p>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.</p> <p>Bug fixes</p> <p>Fixed the memory leak issue.</p> <p>Implemented the JSON bug fixes provided in MySQL 8.0 to improve the stability of using JSON.</p> <p>Fixed the error (error code: 1032) caused by hash scans.</p> <p>Fixed concurrency security issues caused by hotspot update.</p> <p>Implemented the gcol bug fixes provided by MySQL in batches.</p> <p>Fixed the failure to compare DateTime data with String data in some cases.</p>

		<p>Fixed the bug where file handles cannot be released if source-replica buffer pool sync is enabled.</p> <p>Fixed the deadlocks caused by setting the `offline_mode` parameter and creating connections at the same time.</p> <p>Fixed the crashes caused by the `m_end_range` parameter incorrectly set in concurrent range queries.</p> <p>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.</p>
20201231	5.7.18	<p>New features</p> <p>Supported using `NOWAIT` and `SKIP LOCKED` in `SELECT FOR UPDATE/SHARE`.</p> <p>Supported dynamic setting of thread pooling mode or connection pooling mode by using the `thread_handling` parameter.</p> <p>Supported source-replica buffer pool sync.</p> <p>Supported monitoring of user connection status. Monitoring items include sync/async IO, memory, log size, CPU time, and lock duration.</p> <p>Performance optimizations</p> <p>Optimized the transaction subsystem to improve the high concurrency performance.</p> <p>Optimized the time to start crash recovery for large transactions.</p> <p>Optimized redo log flushing.</p> <p>Optimized the buffer pool initialization time.</p> <p>Optimized UTF8/UTF8MB4 string efficiency.</p> <p>Optimized audit performance.</p> <p>Revoked the restriction on the value of `gtid_purged` being empty.</p> <p>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.</p> <p>Optimized the `drop table` operations on big tables.</p> <p>Bug fixes</p> <p>Fixed the hang issue when querying `performance_schema`.</p> <p>Fixed the overflow issue of the `digest_add_token` function.</p> <p>Fixed the crash caused by ibuf access when the `TRUNCATE TABLE` command was executed.</p> <p>Fixed the query correctness issue caused by const propagation when `LEFT JOIN` statement is used.</p>
20200930	5.7.18	<p>Performance optimizations</p> <p>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.</p>

		<p>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.</p> <p>Bug fixes</p> <p>Fixed the crash caused by ibuf access when TRUNCATE TABLE was executed.</p> <p>Fixed cold backup failures when the quick column adding feature was enabled.</p> <p>Fixed performance degradation caused by frequently releasing InnoDB memory table objects.</p> <p>Fixed the query correctness issue caused by const propagation when `LEFT JOIN` statement is used.</p> <p>Fixed the core issue caused by rule class name conflict between SQL throttling and query rewrite.</p> <p>Fixed the concurrent update issue caused by the `INSERT ON DUPLICATE KEY UPDATE` statement in multiple sessions.</p> <p>Fixed the `duplicate key error` caused by concurrent INSERTs when `auto_increment_increment` is used.</p> <p>Fixed the crashes caused by evicting InnoDB memory objects.</p> <p>Fixed concurrency security issues caused by hotspot update.</p> <p>Fixed the coredump issue when enabling the thread pool after jemalloc was upgraded to v5.2.1.</p> <p>Fixed the incomplete audit log issue caused by fwrite error-free handling.</p> <p>Fixed the issue where `mysqld_safe` failed to print logs when it was started by a root user.</p> <p>Fixed the increase in the size of the DDL log file caused by `ALTER TABLE EXCHANGE PARTITION`.</p>
20200701	5.7.18	<p>Bug fixes</p> <p>Fixed the INNOBASE_SHARE index mapping error.</p>
20200630	5.7.18	<p>New features</p> <p>Supported using `NOWAIT` and `SKIP LOCKED` in `SELECT FOR UPDATE/SHARE` statements.</p> <p>Supported large transaction optimization, which can solve such problems as source-replica delay and backup failures caused by large transactions.</p> <p>Optimized audit performance to support async audit.</p> <p>Bug fixes</p> <p>Fixed the overflow of the `digest_add_token` function.</p> <p>Fixed the instance crash caused by `insert blob`.</p> <p>Fixed the source-replica replication interruption when a hash scan failed to find the record while updating the same row in an event.</p> <p>Fixed the hang issue when querying `performance_schema`.</p>
20200331	5.7.18	New features

		<p>Added the official MySQL 5.7.22 JSON series functions.</p> <p>Supported the hotspot update feature as described in Real-Time Session for ecommerce flash sale scenarios.</p> <p>Supported the SQL throttling feature as described in Real-Time Session.</p> <p>Supported encryption with custom KMS keys.</p> <p>Bug fixes</p> <p>Fixed the crash caused by phrase search under multi-byte character sets in full-text index.</p> <p>Fixed the crash of the CATS lock scheduling module in high-concurrency scenarios.</p>
20190830	5.7.18	<p>New features</p> <p>Supported skipping the corrupted data and continuing to parse when a binlog is corrupted. If the source instance and binlog are both damaged, this feature helps restore data from the replica database for use as much as possible.</p> <p>Supported syncing data from non-GTID to GTID mode.</p> <p>Supported querying the "user thread memory usage" by executing the `SHOW FULL PROCESSLIST` statement.</p> <p>Supported quick column adding for tables as described in Overview. This feature does not replicate the data or use disk capacity/IO, and can implement changes in real time during peak hours.</p> <p>Supported persistent auto-increment values.</p> <p>Bug fixes</p> <p>Fixed the issue where replication would be interrupted if the column name in a `GRANT` statement contained reserved words.</p> <p>Fixed the issue where SQL execution efficiency dropped when reverse scan was performed on a partitioned table.</p> <p>Fixed the issue where the query result had an exception due to data inconsistency when using virtual column index and primary key.</p> <p>Fixed the issue where data was missing due to InnoDB primary key range queries.</p> <p>Fixed the issue where the system crashed when a DDL statement was executed for a table with spatial indexes.</p> <p>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.</p> <p>Fixed the issue where other events could not be executed as scheduled when an event was deleted.</p> <p>Fixed the issue where the aggregate query result was incorrect.</p>
20190615	5.7.18	<p>New features</p> <p>Supported transparent data encryption (TDE).</p>
20190430	5.7.18	<p>Bug fixes</p> <p>Fixed the issue where null pointer reference occurred when the LONGTEXT feature was used in subqueries.</p>

		<p>Fixed the issue where source-replica disconnection occurred due to hash scan.</p> <p>Fixed the issue where the replica I/O thread was interrupted due to source binlog switch.</p> <p>Fixed the crash caused by the use of `NAME_CONST`.</p> <p>Fixed the illegal mix of collation error caused by character set.</p>
20190203	5.7.18	<p>New features</p> <p>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, submit a ticket.</p> <p>Supported CATS lock scheduling.</p> <p>Supported creating and dropping temp tables and CTS syntax in transactions when GTID is enabled. To apply for this feature, submit a ticket.</p> <p>Supported implicit primary keys. To apply for this feature, submit a ticket.</p> <p>Supported users without super privileges to kill sessions of other users by configuring the `cdb_kill_user_extra` parameter (default value: `root@%`).</p> <p>Supported enterprise-grade encryption functions. To apply for this feature, submit a ticket.</p> <p>Bug fixes</p> <p>Fixed the issue where replication was interrupted when binlog cache file ran out of space.</p> <p>Fixed the hard error when `fsync` returned `EIO` and retries were made repeatedly.</p> <p>Fixed the issue where replication was interrupted and could not be recovered due to GTID holes.</p>
20180918	5.7.18	<p>New features</p> <p>Supported automatic killing of idle transactions to reduce resource conflicts. To apply for this feature, submit a ticket.</p> <p>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.</p> <p>Supported invisible indexes.</p> <p>Supported memory management with jemalloc, which can replace the jlibc memory management module to reduce memory usage and improve allocation efficiency.</p> <p>Performance optimizations</p> <p>Optimized binlog switch to reduce the `rotate` lock duration and improve system performance.</p> <p>Accelerated the crash recovery.</p> <p>Bug fixes</p> <p>Fixed the issue where an event became invalid due to source-replica switch.</p> <p>Fixed the crash caused by `REPLAY LOG RECORD`.</p> <p>Fixed the issue where the query result was incorrect due to loose index scans.</p>

20180530	5.7.18	<p>New features</p> <p>Supported SQL auditing.</p> <p>Supported table-level concurrent replication. To apply for this feature, submit a ticket.</p> <p>Performance optimizations</p> <p>Optimized replica instance locks to improve the sync performance of replica instances.</p> <p>Optimized the pushdown of the `SELECT ... LIMIT` statement.</p> <p>Bug fixes</p> <p>Fixed the issue where switch failed due to inconsistent positions between `relay_log_pos` and `master_log_pos`.</p> <p>Fixed the crash caused by `Crash on UPDATE ON DUPLICATE KEY`.</p> <p>Fixed the `Invalid escape character in string.` error when a JSON column was imported.</p>
20171130	5.7.18	<p>New features</p> <p>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, submit a ticket.</p> <p>Supported thread pool. To apply for this feature, submit a ticket.</p> <p>Bug fixes</p> <p>Fixed the error of `innodb_buffer_pool_pages_data` parameter overflow by calculating it based on `bytes_data`.</p> <p>Fixed the issue where speed limit plugin became unavailable in async mode.</p>

TXSQL Kernel Version	Description
20220303	<p>Bug fixes</p> <p>Fixed the abnormal release when the memory allocated by `mem_strdup` was used for `row_mysql_truncate_t::file_name` during async drop of big tables.</p>
20220302	<p>Bug fixes</p> <p>Fixed the memory leak issue in `sql_update.cc`.</p>
20220301	<p>New features</p> <p>Supported dynamically configuring the spin cycle (0-100) with the dynamic parameter `innodb_spin_wait_pause_multiplier`. This parameter is used for temporary adjustment and does not support fixing the change through the console.</p> <p>Supported printing deadlock loop information. After this feature is enabled through the parameter `innodb_print_dead_lock_loop_info`, when a deadlock occurs, you can run `show engine innodb status` to view the deadlock loop information.</p> <p>Bug fixes</p>

	<p>Fixed the issue where anonymous GTID transactions were generated in memory tables after replica restart.</p> <p>Fixed the issue where upgrade failed due to the missing `root@localhost` permission.</p> <p>Fixed the issue where the values of monitoring variables such as `innodb_row_lock_current_waits` were abnormal.</p> <p>Fixed the SQL type mapping error in the audit plugin.</p>
20211030	<p>New features</p> <p>Supported large transaction replication optimization.</p> <p>Performance optimizations</p> <p>Accelerated the application of hash scan.</p> <p>Bug fixes</p> <p>Fixed the OOM caused by a large number of table queries.</p> <p>Fixed the infinite loop error caused by setting `innodb_thread_concurrency` to 0.</p> <p>Fixed the issue where there were no statistics information for long records.</p> <p>Fixed the SBM jump error.</p> <p>Fixed the `LOCK_binlog_end_pos` hang` error.</p>
20210630	<p>New features</p> <p>Supported large transaction replication optimization.</p> <p>Bug fixes</p> <p>Fixed the incorrect copy when Index Merge was enabled.</p> <p>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.</p> <p>Fixed the bug that `max(id)` was greater than AUTO_INCREMENT in SHOW CREATE TABLE.</p>
20201231	<p>Bug fixes</p> <p>Fixed the error (error code: 1032) caused by hash scans.</p> <p>Fixed the issue where the source-replica auto-increment values were inconsistent due to the `REPLACE INTO` statement in `ROW` format.</p> <p>Fixed the memory leak caused by not freeing up the memory requested for parsing SQL statements.</p> <p>Fixed the issue where the sql_mode check is skipped when running `CREATE TABLE AS SELECT`.</p> <p>Fixed the issue where the `sql_mode` check was skipped when inserting default values.</p> <p>Fixed the issue where the `sql_mode` check was skipped when running UPDATE with bound parameters.</p>
20200915	<p>New features</p> <p>Supported the SQL throttling feature as described in Real-Time Session.</p> <p>Performance optimizations</p> <p>Optimized the initialization acceleration of buffer pool.</p> <p>Bug fixes</p> <p>Fixed the hang issue of `rename table` on both source and replica.</p>

	<p>Fixed the crash when `event_scheduler` was set to `disable` and `cdb_skip_event_scheduler` was changed from `on` to `off`.</p> <p>Fixed the `sync_wait_array` assertion failure when the maximum number of connections of `tencentroot` was not counted in `srv_max_n_threads`.</p> <p>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.</p> <p>Fixed the `INSERT ON DUPLICATE KEY UPDATE THE WRONG ROW` error.</p> <p>Fixed the `index_mapping` error.</p> <p>Fixed the MTR failure.</p> <p>Fixed the source-replica replication interruption when a hash scan failed to find the record while updating the same row in an event.</p>
20190930	<p>New features</p> <p>Supported querying the user thread memory usage by executing the `SHOW FULL PROCESSLIST` statement.</p> <p>Bug fixes</p> <p>Fixed GTID holes caused by the replication filter of the replica.</p> <p>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.</p> <p>Fixed the illegal mix of collation error caused by character set.</p> <p>Fixed the issue where the source-replica disconnection occurred due to hash scan.</p> <p>Fixed the crash caused by the use of `NAME_CONST`.</p> <p>Fixed the issue where the replica I/O thread was interrupted due to source binlog switch.</p> <p>Fixed the error of incompatible backups due to `innodb_log_checksum`.</p>
20190530	<p>Bug fixes</p> <p>Fixed the issue where dirty data might be read in RC mode.</p> <p>Fixed the issue where replica instance replay might fail due to the drop of temp table.</p> <p>Fixed the deadlock issue under high concurrency.</p>
20190203	<p>New features</p> <p>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, submit a ticket.</p> <p>Supported users without super privileges to kill sessions of other users by configuring the `cdb_kill_user_extra` parameter (default value: `root@%`).</p> <p>Supported creating and dropping temp tables and CTS syntax in transactions when GTID is enabled. To apply for this feature, submit a ticket.</p> <p>Performance optimizations</p> <p>Optimized the replication and replay of partitioned tables to improve efficiency.</p> <p>Bug fixes</p> <p>Fixed the source-replica data inconsistency issue caused by insufficient temporary space.</p> <p>Fixed the issue of suspended hot record updates.</p> <p>Fixed the issue where the value of `Seconds_Behind_Master` was abnormal during concurrent replication.</p>

20180915	<p>New features</p> <p>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.</p> <p>Supported automatic killing of idle transactions to reduce resource conflicts. To apply for this feature, submit a ticket.</p> <p>Bug fixes</p> <p>Fixed the crash caused by `REPLAY LOG RECORD`.</p> <p>Fixed the error of time data inconsistency between source and replica due to decimal precision issues.</p>
20180130	<p>New features</p> <p>Supported thread pool. To apply for this feature, submit a ticket.</p> <p>Supported dynamically modifying replication filtering rules for replica nodes.</p> <p>Performance optimizations</p> <p>Reduced performance fluctuation caused by `DROP TABLE`.</p> <p>Bug fixes</p> <p>Fixed the database crash caused by authentication password strings.</p>
20180122	<p>New features</p> <p>Supported SQL auditing.</p> <p>Bug fixes</p> <p>Fixed the integer overflow issue.</p> <p>Fixed the error caused by queries using full-text index.</p> <p>Fixed the issue where the replica crashed during replication.</p>
20170830	<p>Bug fixes</p> <p>Fixed the issue where binlog speed limit became invalid in async mode.</p> <p>Fixed the issue where the `buffer_pool` status was abnormal.</p> <p>Fixed the issue where `SEQUENCE` and implicit primary key conflicted.</p>
20170228	<p>Bug fixes</p> <p>Fixed the character encoding bug in `DROP TABLE`.</p> <p>Fixed the issue where a table contained symbols like decimal points or `replicate-wild-do-table` couldn't be used to filter databases correctly.</p> <p>Fixed the issue where SQL threads exited too early after the replica had a `rotate` event.</p>
20161130	<p>Performance optimizations</p> <p>Split the `lock_log` lock to reduce the time used by lock logs and improve the concurrency performance.</p> <p>Separated the ACK thread of the source to reduce the response time.</p> <p>Prohibited the user thread from being killed while waiting for ACK in order to prevent phantom reads.</p> <p>Fixed the unnecessary `lock_sync` lock when `sync_binlog != 1`.</p>

TXRocks Engine Kernel Version Release Notes

Last updated : 2025-04-01 17:11:55

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 Release Notes

MySQL 5.7 Kernel Version Release Notes

Minor Version	Note
20231030	New features Added virtual bulk load function. Introduced the "LOCK TABLES FOR BACKUP" syntax. Bug Fixes Resolve potential issues of GTID and binlog position mismatches when utilizing Xtrabackup.
20230401	New features By default, the TokuDB engine in the table creation statement is converted to the RocksDB engine.

Minor Version	Note
20240702	Bug fixes Fixed the issue that executing the SHOW ENGINE ROCKSDB STATUS command would cause a hang.
20231030	New features Added virtual bulk load function.
20230401	New features By default, the TokuDB engine in the table creation statement is converted to the RocksDB engine.

TXSQL Engine Kernel Version Release Notes

Last updated : 2025-05-09 11:51:17

This document introduces the kernel version updates for the LibraDB engine in TencentDB for MySQL. For detailed product features, see [LibraDB Engine Feature Highlights](#).

Version Numbering Rules

The kernel version of the LibraDB engine can be found in the instance details. It is represented by a version number composed of four segments separated by periods (.), each segment indicating a specific aspect of the version.

The first segment represents the primary version of the product. When the current version is incompatible with its predecessor, this number will be incremented by 1.

The second segment indicates the product's iterative version. When significant features are introduced, this version number is updated based on the current year and month.

The third segment represents the minor version of the product. When optimization features or bug fixes are implemented, this number will be incremented by 1.

The fourth segment represents the patch version of the product. When a critical bug is identified that requires immediate attention, a patch upgrade will be applied to this version.

Version Update Release Notes

Note:

TencentDB for MySQL is compatible with LibraDB kernel versions 1.2404.20.0 and later. Note that some of the historical version release information in this document is provided for reference only.

Version	Description
2.2410.1.0	<p>The latest version of the LibraDB engine kernel supports many new kernel features. This version is currently under Beta test. To apply for a trial, submit a ticket.</p> <p>Feature updates</p> <p>Storage features</p> <p>Added support for Columnar Storage Secondary Index. The supported index types include the Zonemap index, Bloom Filter index, and Bitmap index.</p> <p>Added support for concurrent FlushDMS operations to enhance data write efficiency, particularly improving update efficiency when managing large volumes of data.</p> <p>Added support for the Data Type and various JSON Function.</p> <p>Added support for loading range/list-partitioned tables into the LibraDB engine and loading partition DDL statements synchronously. For detailed syntax and function support related to</p>

partition operations (such as adding or dropping partitions or subpartitions), see [Data Loading Restrictions](#).

Added support for loading tables without primary keys in a columnar storage format, as well as for scenarios involving modifications to table primary keys. For details, see [Data Loading Restrictions](#).

Compute features

Added support for [Merge Join](#), improving query performance in primary key join scenarios. Added support for stream aggregation, improving GROUP BY performance in ordered data scenarios.

Added support for the automatic [Collecting Statistics](#) capability, removing the need for users to manually collect statistics regularly.

Added support for the random sampling feature, enhancing the efficiency of [Collecting Statistics](#) and reducing the resource consumption associated with statistics.

Added support for adaptive GROUP BY operators.

Added support for streaming CTEs. For details, see [CTE Syntax Use Instructions](#).

Added support for pushing Table Dual down to the storage layer for execution.

Added support for pushing the UNION ALL operator down to the storage layer for execution.

Added support for the Win Magic rewrite of the optimizer, which decouples correlated subqueries by rewriting them as window functions to improve query efficiency.

Added support for 14 MySQL-compatible functions. For details, see [Description of Supported String Functions](#).

Added support for string-type input values for the Hour, Minute, Second, and Microsecond functions.

Added support for loading tables with decimal-type primary keys (up to a maximum length of 256) into columnar storage.

Added support for NULL-aware anti-joins. When executing set operations, these joins intelligently assess whether a set is empty or contains null values, optimizing the execution efficiency of operations such as `NOT IN` and `<>ALL`. This enhancement significantly improves the performance of associated SQL queries.

System optimization

Optimized the error messages returned for some execution errors.

Optimized the performance of the grace hash join by employing a two-phase bucket acceleration mechanism, reducing the frequency of bucket scale-out.

Added support for right anti-joins/right semi-joins in the result set splitting feature, significantly alleviating long wait times for data reception caused by extremely large join result sets.

Added support for aggregation in the result set splitting feature, alleviating long wait times for data reception caused by extremely large aggregation output result sets.

Optimized the EXPLAIN feature by enabling the display of the table name and alias corresponding to the TableScan operator in the detailed information of the execution plan.

Optimized the EXPLAIN feature by enabling the display of replica information for the referenced table in the execution plan.

Added support for scenarios where the Unique Key column contains null values.

Optimized the metadata storage mechanism, significantly reducing excessive resource consumption associated with metadata.

	Optimized the execution method of the Apply operator, improving query performance in correlated subquery scenarios.
1.2404.21.0	<p>This release is a bugfixing and optimization update that primarily addresses defects identified in version 1.2404.20.0 and earlier.</p> <p>Feature updates</p> <p>Added support for DDL operations related to partitioned tables in the current version.</p> <p>Bug fixes</p> <p>Resolved the issue that data deletion failed when the primary key contained a timestamp type.</p> <p>Resolved the data synchronization error associated with SET type data.</p> <p>Resolved several errors in the table renaming scenario.</p>
1.2404.20.0	<p>This release is a bugfixing and optimization update that primarily addresses defects identified in version 1.2404.19.1 and earlier.</p> <p>Note:</p> <p>This is the first LibraDB kernel version supported by TencentDB for MySQL.</p> <p>Bug fixes</p> <p>Resolved the issue that caused data-loading objects to be inaccessible when the read-write instance was configured as case-sensitive.</p> <p>Resolved the issue that some existing instances could not support database proxies.</p> <p>Resolved the issue that Lazy Materialization statistics were displayed incorrectly.</p> <p>Resolved the issue that DMC was unable to retrieve user information for the read-only analysis engine.</p> <p>Optimized the default values for some kernel parameter settings.</p>
1.2404.19.1	<p>This release is a bugfix update that primarily addresses defects identified in versions earlier than 1.2404.19.0.</p> <p>Bug fixes</p> <p>Resolved the issue of inconsistent view definers when loading views into the read-only analysis engine.</p> <p>Resolved the issue that user synchronization failed for users with empty passwords in the mysql.user system table.</p> <p>Resolved the execution error encountered when executing DDL statements during SQL operations.</p> <p>Resolved the write exception issue in the ETL Writeback Acceleration feature when processing constant columns.</p>
1.2404.19.0	<p>Feature updates</p> <p>Added support for accessing the read-only analysis engine through the database proxy when using the HINT syntax.</p> <p>Added support for accessing the read-only analysis engine through DMC.</p> <p>System optimization</p> <p>Optimized the memory management methods for the SORT and AGGREGATION operators, enhancing memory usage efficiency. Following the optimization, memory used by</p>

	<p>the SORT and AGGREGATION operators can be released immediately upon completion of their execution, without waiting for the entire query execution to end.</p> <p>Optimized the memory usage method for join result output, reducing memory consumption when the JOIN operator produces a relatively large result set.</p> <p>Optimized the allocation policy for Lazy Materialization in primary key scenarios.</p> <p>Bug fixes</p> <p>Resolved the issue that Prepare Statement defaulted to using Plan Cache, resulting in multiple similar queries inadvertently matching a previously cached plan. This could lead to query errors when the types of constant parameters were inconsistent.</p> <p>Resolved the issue of query errors that occurred in some scenarios when output columns of VARCHAR type were used in a UNION ALL context.</p> <p>Resolved the issue that metadata retrieval errors were caused when DMC connected to the columnar storage engine due to insufficient permissions.</p> <p>Resolved the potential memory leak issue in scenarios involving frequent connection creation.</p> <p>Resolved the issue that temporary tables were prevented from being queried in the read-only analysis engine in scenarios involving changes to partitioned tables.</p> <p>Resolved the issue that the Lazy Materialization operator was not displayed in the execution plan under some scenarios.</p> <p>Resolved the issue that the read-only analysis engine was unable to delete fields from the join index.</p> <p>Resolved the out-of-bounds error in runtime filter assignment for multi-table join scenarios.</p> <p>Resolved the issue that the runtime filter occasionally returned empty data.</p> <p>Resolved the error caused by a mismatch between nullable and non-nullable types in UNION ALL scenarios.</p>
1.2404.17.0	<p>Feature updates</p> <p>Added support for the FROM_UNIXTIME and UNIX_TIMESTAMP functions.</p> <p>System optimization</p> <p>Optimized the permissions required to query system tables when database client tools, such as DBearer, access the read-only analysis engine.</p> <p>Optimized error messages for some unsupported functions.</p> <p>Bug fixes</p> <p>Resolved the issue that the read-only analysis engine incorrectly loaded Timestamp field values when the explicit_defaults_for_timestamp parameter differed across instances.</p>
1.2404.16.0	<p>Feature updates</p> <p>Added support for using the read-only analysis engine to accelerate the execution of query statements in <code>INSERT...SELECT...</code>, enabling faster writes to read-write instances. For details, see ETL Writeback Acceleration.</p> <p>Added support for querying and synchronizing data types, including BIT, ENUM, and SET.</p> <p>System optimization</p> <p>Optimized the naming conventions of some system tables for improved clarity, including the SLOW_LOG table.</p> <p>Adjusted the supported value size for regular columns to 16 MB, up from the previous limit of 5 MB.</p>

	<p>Optimized the logic of the StorageTableReader operator for retrieving column names when the query plan is empty.</p> <p>Bug fixes</p> <p>Resolved the issue that, in subquery scenarios, the Apply operator logic was still applied when the primary query result was empty, leading to a query error.</p> <p>Resolved the issue that the Apply operator caused a crash during multi-threaded competition.</p>
1.2404.15.2	<p>This release is a bugfixing update that primarily addresses defects identified in versions earlier than 1.2404.15.1.</p> <p>Bug fixes</p> <p>Resolved the issue that the instance crashed during data queries when high-concurrency <code>UPDATE</code> operations were performed on the read-write instance after it was loaded as a column-stored table.</p>
1.2404.15.1	<p>System optimization</p> <p>Optimized the default settings for database log retention. Specifically, if an instance's disk size is 500 GB or less, the default maximum log retention is set to 5 GB; for instances with a disk size over 500 GB, it is set to 10 GB.</p> <p>Enabled Lazy Materialization in the read-only analysis engine by default. For details on its capabilities and tuning options, see Lazy Materialization.</p> <p>Added support for loading tables whose names begin with an underscore (_) into the read-only analysis engine.</p> <p>Improved the restart efficiency of the read-only analysis engine.</p> <p>Bug fixes</p> <p>Resolved the issue that some versions of Navicate displayed an insufficient permissions error when connecting to the read-only analysis engine.</p> <p>Resolved the issue that the read-only analysis engine failed to load data when DDL statements were executed consecutively.</p> <p>Resolved the syntax compatibility issue related to DEFAULT (0) in the <code>CREATE TABLE</code> statement.</p>
1.2404.10.0	<p>Feature updates</p> <p>Added support for the DATE_SUB function.</p> <p>Added support for users to manually collect table statistics. For details, see Collecting Statistics.</p> <p>Added support for synchronous loading of ENUM data in special scenarios.</p> <p>Note:</p> <p>Although ENUM data can be loaded into the read-only analysis engine in special scenarios, querying data of the ENUM field type is still not supported.</p> <p>System optimization</p> <p>Adjusted the kernel to manage the number of versions retained for expired data in storage, preventing excessive disk space usage caused by retaining an excessive number of data snapshots.</p> <p>Optimized kernel resource allocation during data clearing, preventing degradation of user SQL execution performance resulting from excessive consumption of system resources.</p>

	<p>Adjusted the logging level for kernel information to prevent excessive log output from consuming too much disk space.</p> <p>Adjusted the maximum value of the storage field. The previously supported default limit of 1 MB has been increased to 5 MB.</p> <p>Optimized the error message for data querying when the data is not fully loaded into the read-only analysis engine. It now clearly indicates that you need to wait until all data is fully loaded before performing a query.</p> <p>Optimized the error message for unsupported data types and functions during execution. You can now clearly identify the unsupported items in the error message.</p> <p>Optimized the internal trigger frequency and selection logic for background tasks, ensuring they do not impact the execution performance of users' online SQL queries.</p> <p>Improved the version number information displayed when accessing the read-only analysis engine instance through the MySQL client to prevent potential misunderstandings.</p>
1.2404.7	<p>Feature updates</p> <p>Added support for real-time loading of hyperscale data into the read-only analysis engine.</p> <p>Added support for ultra-high performance in complex data queries.</p>

Database Proxy Version Release Notes

Last updated : 2025-05-27 11:53:36

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 as instructed in [Upgrading Kernel Minor Version](#).

Regarding the database proxy versions, please be advised that versions beginning with 1.4 are beta, versions starting with 1.3 are deemed stable, and versions below 1.3 are categorized as public beta versions and will no longer receive updates.

The instance architectures supporting database proxy’s momentary disconnection prevention capabilities include MySQL dual-node (high availability edition) and MySQL triple-node.

Release Notes for the Stable Version

Release Notes for Beta Version

Database Proxy Version	Release Date	MySQL Kernel Version Requirements	Description
1.3.17	2025-5	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue where errors occasionally occurred during read-only transaction initiating after transaction splitting is enabled. Addressed the issue where character set errors may occur during connection reuse after the connection pool is enabled.
1.3.16	2025-3	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that executing prepared statements at a read-only address could occasionally result in errors. Addressed the problem that enabling connection pooling may lead to exceptions in Capabilities Flags. Fixed the issue that the connection might be interrupted when the replication delay of a read-only instance changed frequently after delayed read-only instance removal was enabled.
1.3.15	2024-12	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Performance optimizations Optimized memory usage under high load conditions, addressing the issue of prolonged high memory utilization. Fixes

			Resolved the issue where the anti-disconnection feature occasionally failed.
1.3.14	2024-11	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	<p>Feature updates</p> <p>The statements SHOW BINLOG EVENTS and SHOW BINARY LOGS are now routed to the root instance.</p> <p>Optimized the issue that querying large messages may result in excessive memory consumption.</p> <p>Fixes</p> <p>Fixed the occasional inconsistency between the character set sorting of the handshake messages and the backend databases.</p> <p>Fixed the issue that the primary-secondary switch in the primary database may cause interruption in the read-only address connection.</p> <p>Fixed the issue of the potential authentication failure when a specified IP address matches multiple accounts simultaneously.</p>
1.3.13	2024-8	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	<p>Feature updates</p> <p>System variables related to SELECT binlog will be sent to the primary database for execution.</p> <p>Added new high-availability detection statements to reduce the migration caused by false fault detections.</p> <p>Optimized the memory usage of the database proxy kernel in multi-threading scenarios.</p> <p>Fixes</p> <p>Fixed the issue that an error was reported when a precompiled statement was executed in a read-only transaction.</p> <p>Fixed the performance issue caused by frequent connection establishment with an incorrect username.</p> <p>Fixed the issue that uncommitted transactions existed on read-only nodes because COMMIT and ROLLBACK statements were routed to all nodes after transaction split was enabled.</p>
1.3.12	2024-6	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	<p>Feature updates</p> <p>The SHOW VARIABLES statement will be sent to the primary database for execution.</p> <p>Fixes</p> <p>Fixed the issue that a large number of connections triggering momentary disconnection prevention could cause some connections to be stuck for too long.</p> <p>Fixed the issue that a large number of connections triggering re-load balancing could cause some</p>

			connections to fail to be broken as expected.
1.3.11	2024-5	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that too many accounts could lead to connection failures. Note: Version 1.3.12 already includes this fix, so it is recommended to upgrade to 1.3.12 directly.
1.3.10	2024-3	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that encountering specific messages during message processing could cause anomalies.
1.3.9	2024-1	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that annotations in multi-statements caused incorrect parsing.
1.3.8	2023-9	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Feature updates Supports more functions updated in MySQL 5.7 and 8.0. Optimized the parser cache to reduce OOM probability under a large number of complex SQL statements. Added traffic monitoring metrics for the database proxy. Supports the dynamic CLB feature. Fixes Fixed the issue that the momentary disconnection prevention feature might fail for some connections under a high connection count.
1.3.7	2023-4	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue of route errors that occurred when the SELECT...FOR UPDATE statement was executed in some cases. The SELECT @@read_only statement will be routed to the primary database to avoid using the read_only tag under certain frameworks, causing the system to mistakenly determine that the database proxy was unwritable. Fixed the issue that the database proxy node was switched due to database instance errors in some scenarios.
1.3.5	2023-1	MySQL 5.7 20211030 and later	Performance optimizations

		MySQL 8.0 20211202 and later	Optimized the performance when a large number of read statements were executed concurrently.
1.3.4	2022-12	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that the returned data of SHOW PROCESSLIST was incomplete.
1.3.3	2022-11	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that an error may occur during statement precompilation under a session-level connection pool.
1.3.2	2022-10	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Fixes Fixed the issue that an error may occur during statement precompilation.
1.3.1	2022-10	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Feature updates Allowed instances with a weight of 0 to sustain read requests when the weight of all valid instances under the database proxy is 0. Supported the multi-AZ deployment architecture where read-only instances can be mounted across AZs. Provided the read-only mode. Supported transaction split. 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.
1.2.1	2022-6	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Feature updates Supported the lower_case_table_names parameter. By default, the case is not verified.
1.1.3	2021-6	MySQL 5.7 20211030 and later MySQL 8.0 20211202 and later	Feature updates Supported using a hint that indicates the route information in precompiled MySQL statements. When a hint is used in a PREPARE statement to specify the route destination, subsequent EXECUTE messages will be sent to the specified backend node. Fixes Fixed the issue that the frontend connection was reset immediately after source-replica switch of the source instance on the database proxy was performed.

			<p>Fixed the issue that 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.</p> <p>Fixed the issue that MySQL 8.0 might return incorrect handshake information and cause connection failures.</p>
1.1.2	2021-6	<p>MySQL 5.7 20211030 and later</p> <p>MySQL 8.0 20211130 and later</p>	<p>Feature updates</p> <p>Supported MySQL 8.0.</p> <p>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.</p> <p>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.</p> <p>Updated the internal memory management mechanism to reduce the memory usage.</p> <p>Fixes</p> <p>Fixed the issue that the client connection persisted after the backend connection was closed due to timeout.</p> <p>Fixed the issue that the internal cache might cause excessively rapid increase of the memory utilization.</p> <p>Fixed the occasional issue that messages in an incorrect format were returned.</p>
1.0.1	2022-2	<p>MySQL 5.7 20201230 and later</p>	<p>Feature updates</p> <p>Supported MySQL 5.7.</p> <p>Supported read/write separation.</p> <p>Supported read weight assignment in read/write separation.</p> <p>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.</p> <p>Supported the configuration of the least number of 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>

			<p>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.</p> <p>Supported using a hint to specify the destination node for routing.</p>
--	--	--	---

Database Proxy Version	Release Date	MySQL Kernel Version Requirements	Description
1.4.4	2025-3	MySQL 5.7 20211230 and later MySQL 8.0 20221215 and later Read-only analysis engine 2.2410.4.0 and later	Feature updates Supported participating in database proxy read-write separation and removal upon faults of LibrADB instances with the read-only analysis engine. Supported using the hint /* to ap */ in SQL statements to route SQL statements to a read-only analysis engine instance.
1.4.2	2023-8	MySQL 5.7 20211230 and later MySQL 8.0 20221215 and later	Fixes Fixed the issue of excessive memory usage after the transaction-level connection pool is enabled.
1.4.1	2023-7	MySQL 5.7 20211230 and later MySQL 8.0 20221215 and later	Feature updates Supports the transaction-level connection pool feature. Fixes Fixed the issue that DDL statements would get stuck in the secondary database.

Functionality Features

Killing Idle Transactions Automatically

Last updated : 2025-02-19 16:09:56

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 above.

Kernel version: MySQL 5.7 20180918 and above.

Kernel version: MySQL 8.0 20200630 and above.

Use Cases

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.

Instructions

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	Type	Default Value	Valid Values/Value Range	Description
<code>cdb_kill_idle_trans_timeout</code>	Yes	ulong	0	[0, 31536000]	If it is <code>0</code> , the feature is disabled; otherwise, a transaction idle for <code>cdb_kill_idle_trans_timeout</code> or <code>wait_timeout</code> seconds will be killed.

Parallel Replication

Last updated : 2024-07-22 11:30:52

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.

Use Cases

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.

Instructions

For MySQL 5.6 and 5.7, you can enable this feature by setting the `slave_parallel_type` parameter to the newly added value `TABLE` . 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 parameter description](#)

[MySQL 5.7 parameter description](#)

[MySQL 8.0 parameter description](#)

Parameter	Effective Immediately	Type	Default Value	Valid Values/Value Range	Description

slave_parallel_type	Yes	char*	SCHEMA	SCHEMA/TABLE	The level of parallel replication on the replica node. SCHEMA: Replication events of different schemas can be executed in parallel. TABLE: Replication events of different tables can be executed in parallel.
---------------------	-----	-------	--------	--------------	--

Parameter	Effective Immediately	Type	Default Value	Valid Values/Value Range
slave_parallel_type	Yes	char*	LOGICAL_CLOCK	DATABASE/TABLE/LOGICAL_CLOCK

Parameter	Effective Immediately	Type	Default Value	Valid Values/Value Range	Description
slave_parallel_type	Yes	char*	LOGICAL_CLOCK	DATABASE/LOGICAL_CLOCK	The replication type of the replica node. DAT: Data replication. Repl: Replication of data.

						data exec LOG Repl of th clock can l para
--	--	--	--	--	--	--

Dynamic Thread Pool

Last updated : 2024-02-18 11:24:04

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 above.

Kernel version: MySQL 5.7 20201230 and above.

Use Cases

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.

Performance Impact

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.

Instructions

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:

Valid 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 `show full processlist` command displays the following new status:

`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	Type	Default Value	Valid Values/Value Range
-----------	-----------------------	------	---------------	--------------------------

thread_pool_idle_timeout	Yes	uint	60	[1, UINT_MAX]
thread_pool_oversubscribe	Yes	uint	High Stability Parameter Template: 10 High Performance Parameter Template: 16	[3,32]
thread_pool_size	Yes	uint	The number of CPUs on the current machine	[1,1000]
thread_pool_stall_limit	Yes	uint	500	[10, UINT_MAX]
thread_pool_max_threads	Yes	uint	100000	[1,100000]
thread_pool_high_prio_mode	Yes, session	enum	transactions	transactions\\statement\\none

thread_pool_high_prio_tickets	Yes, session	uint	UINT_MAX	[0, UINT_MAX]
threadpool_workaround_epoll_bug	Yes	bool	false	true/false

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

Status	Description
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
max_req_latency_us	The maximum time in milliseconds for a user connection to wait in the queue in the thread group
conns_timeout_killed	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 (net_wait_timeout)
connections_moved_in	The total number of connections migrated from other thread groups to this thread group
connections_moved_out	The total number of connections migrated from this thread group to other thread groups
connections_moved_from_per_thread	The total number of connections switched from the one-thread-per-connection mode to this thread group
connections_moved_to_per_thread	The total number of connections switched from this thread group to the one-thread-per-connection mode
events_consumed	The total number of events processed by the thread group
average_wait_usecs_in_queue	The average waiting time of all events in the queue in the thread group

NOWAIT

Last updated : 2024-07-22 11:31:50

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).

Use Cases

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.

Instructions

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 : 2024-07-22 11:32:19

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, TSQL 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.

Use Cases

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.

Notes:

1. For DELETE ... RETURNING, the returned data rows are pre-images, while for INSERT/REPLACE ... RETURNING, they are post-images.
2. Currently, UPDATE ... RETURNING is not supported.
3. For INSERT/REPLACE ... RETURNING, columns in the outer table are currently invisible to the subquery in the RETURNING clause.
4. 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.

Instructions

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, 'EEEE'), (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	EEEE
7	GGGGGGG

```
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 : 2024-12-30 18:05:42

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.

The column compression feature enables the compression of infrequently accessed large fields without compressing those small fields that are frequently accessed. This approach not only reduces the storage space required for the entire row of fields but also enhances the efficiency of read and write operations.

For instance, consider an employee table `create table employee(id int, age int, gender boolean, other varchar(1000) primary key (id))`. When there is frequent access to the smaller fields such as `id`, `age`, `gender`, and less frequent access to the larger field `other`, it is advisable to designate the `other` column as a compressed column. Generally, only read and write operations on the `other` column will trigger its compression and decompression, while access to other columns will not. This strategy further reduces the storage size of row data, enabling swifter access to frequently accessed smaller fields and significantly reducing the storage space required for less frequently accessed larger fields.

Note :

The parameter ``cdb_column_compression_enabled`` serves as the toggle for the column compression feature.

The column compression feature in MySQL 5.7 is disabled by default. Should you wish to utilize it, please [submit a ticket](#) to enable this function.

The column compression feature is enabled by default in MySQL 8.0 kernel versions 20221215 and above.

Owing to the adoption of the open-source collaborative version's column compression capability in the MySQL 8.0 version, its implementation diverges from that of the MySQL 5.7 version. Below, we will delineate the usage instructions for the column compression capabilities of both versions. You may click on "**MySQL 5.7 Column Compression**" or "**MySQL 8.0 Column Compression**" to toggle and learn more.

[MySQL 5.7 Column Compression](#)

[MySQL 8.0 Column Compression](#)

Supported Versions

Kernel version: MySQL 5.7 20210330 and above.

Use Cases

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

Instructions

Supported data types

1. `BLOB` (including `TINYBLOB`, `MEDIUMBLOB`, and `LONGBLOB`)
2. `TEXT` (including `TINYTEXT`, `MEDIUMTEXT`, and `LONGTEXT`)
3. `VARCHAR`
4. `VARBINARY`

Note :

Here, the maximum length of `LONGBLOB` and `LONGTEXT` is $2^{32}-2$ bytes, which is one byte less than $2^{32}-1$ supported by native MySQL as described in [String Type Storage Requirements](#).

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 i
    [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>	Yes
<code>CREATE TABLE t2 SELECT * FROM t1;</code>	Yes
<code>CREATE TABLE t2(a BLOB) SELECT * FROM t1;</code>	No

ALTER TABLE:

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.

Parameter Description

Parameter	Effective Immediately	Type	Default Value	Valid Values/Value Range	Description
<code>cdb_column_compression_enabled</code>	Yes	bool	FALSE	TRUE/FALSE	The column compression feature is on or off. When creating a table, compression is prohibited for columns that are already compressed and remain unchanged.
<code>innodb_column_compression_zlib_wrap</code>	Yes	bool	TRUE	TRUE/FALSE	If enabled, the zlib header and footer for performance are added to the compressed data for checksum verification.

innodb_column_compression_zlib_strategy	Yes	Integer	0	[0,4]	Column compression policy. Values: Z_DEFLATE, Z_FILT, Z_HUFF, Z_RLE; 4. Generally, Z_DEFLATE is the best for text data, while image data is better for Z_FILT.
innodb_column_compression_zlib_level	Yes	Integer	6	[0,9]	Column compression level. Value range: 0-9. 0 indicates no compression. The higher the value, the smaller the compressed data size, but the longer the compression duration.
innodb_column_compression_threshold	Yes	Integer	256	[0, 0xffffffff]	Column compression threshold. Value range: 1-4294967295. 1 indicates that data whose length is greater than or equal to this threshold will be compressed. If the compressed data is larger than the original data, no compression is added.
innodb_column_compression_pct	Yes	Integer	100	[1, 100]	Column compression ratio. Value range: 1-100. 1 indicates that the compressed data size after compression is less than or equal to 1% of the original data size. If the compressed data size is greater than the original data size, no compression 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	Type	Description
<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	Scope	Description
<code>Compressed column '%-.192s' 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' 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 0–9 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.

Notes

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.

3. Column compression employs the zlib compression algorithm to condense data. However, applying column compression to data that has already been compressed typically does not significantly enhance the compression outcome. Taking JPEG images as an example, JPEG is a highly optimized lossy compression format that has already reduced the image data to a smaller file size. Therefore, in practical applications, the effect of column compression on JPEG data is not optimal.

Supported Versions

Kernel version: MySQL 8.0 20221215 and above.

Use Cases

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

Instructions

Supported data types

1. `BLOB` (including `TINYBLOB` , `MEDIUMBLOB` , and `LONGBLOB`)
2. `TEXT` (including `TINYTEXT` , `MEDIUMTEXT` , and `LONGTEXT`)
3. `VARCHAR`
4. `VARBINARY`
5. `JSON`

The syntax is as follows:

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

or

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

Compression Threshold

The compression threshold is governed by the parameter `innodb_min_column_compress_length` , which defaults to 256. Should the original size of a column exceed the value of this parameter, compression will be applied; otherwise, only a compression header is added without actual data compression being performed.

The supported DDL syntax is summarized as follows:

CREATE TABLE:

DDL	Whether the Compression Attribute is Inherited
<code>CREATE TABLE t2 LIKE t1;</code>	Yes
<code>CREATE TABLE t2 SELECT * FROM t1;</code>	No
<code>CREATE TABLE t2(a BLOB) SELECT * FROM t1;</code>	No

ALTER TABLE:

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

<pre>ALTER TABLE t1 MODIFY COLUMN a BLOB COMPRESSED;</pre>	Alters a non-compressed column into a compressed one.
--	---

Parameter Description

Parameter	Effective Immediately	Type	Default Value	Valid Values/Value Range	Description
innodb_zlib_column_compression_level	Yes	UINT	6	[0-9]	Zlib compressi where 0 denotes nc compressi 1 signifies the fastest compressi and 9 represents the highes level of compressi achievable As the sca progresse: from 1 to 9 the the compressi speed decreases yet the compressi ratio increases.
innodb_zstd_column_compression_level	Yes	UINT	3	[1-22]	Zstd compressi employs a scale wher 1 signifies the fastest compressi speed, and 22 denotes the highes

					level of compressi As the sca progresses from 1 to 2 the compressi speed decreases yet the compressi ratio increases.
innodb_min_column_compress_length	Yes	UINT	256	[1, UINT_MAX32]	Adjust the compressi threshold, measured bytes. Should the original length of a column be greater tha or equal to the value c this parameter compressi will be applied. Otherwise only a compressi header wil be added, without actual data compressi being performed

Multiple Algorithms

TencentDB for MySQL 8.0 version supports three compression algorithms: ZLIB, LZ4, and ZSTD. It is also permissible to omit specifying the algorithm, in which case ZLIB will be selected as the default algorithm.

The syntax is as follows:

```
CREATE TABLE t1(  
    id INT PRIMARY KEY,  
    b BLOB COMPRESSED ALGORITHM = [ZLIB|LZ4|ZSTD]  
);
```

or

```
CREATE TABLE t1(  
    id INT PRIMARY KEY,  
    b BLOB COLUMN_FORMAT COMPRESSED ALGORITHM = [ZLIB|LZ4|ZSTD]  
);
```

Compression Algorithms and Compression Levels

1. ZLIB: Currently, multiple compression levels of ZLIB are offered, with the parameter designated as `innodb_zlib_column_compression_level`, ranging from 0 to 9. Here, 0 signifies no compression, 1 denotes the fastest compression, and 9 represents the highest degree of compression, with the default set to 6.
2. LZ4: It is consistent with MySQL's Page compression and does not support multi-level compression of LZ4. When using the LZ4 compression algorithm, you need to pay attention to the fact that the maximum original length of LZ4 compression is $2^{31}-1$, while the maximum length of LONGBLOB is $2^{32}-1$. When the compressed original data When the length is greater than or equal to 2^{31} , ZLIB will be used implicitly for compression.
3. ZSTD: ZSTD, short for ZStandard, offers three modes of compression. This section details the support for non-streaming, dictionary-less standard compression, providing multiple levels of compression. The parameter for this is `innodb_zstd_column_compression_level`, with values ranging from 1 to 22. A setting of 1 denotes the fastest compression, while 22 indicates the highest degree of compression, with the default setting being 3.

Display of Compression Attributes

This section presents the default compression algorithm: `ALGORITHM = ZLIB`.

```
CREATE TABLE t2 (a VARCHAR(100) COMPRESSED) ENGINE=InnoDB;  
  
SHOW CREATE TABLE t2;
```

Notes

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 8.0.22, data can be imported directly.

If the official MySQL version is 8.0.22 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.
3. In terms of physical backups, since the fields are already compressed within InnoDB at the time of backup, the version used for the backup must also support column compression.
4. Physical upgrades from MySQL 5.7 to MySQL 8.0 are not supported.
5. Column compression utilizes ZLIB, LZ4, and ZSTD compression algorithms to condense data. However, applying column compression to data that has already been compressed typically does not significantly enhance the compression outcome. Taking JPEG images as an example, JPEG is a highly optimized lossy compression format that has already reduced the image data to a smaller file size. Therefore, in practical applications, the effect of applying column compression to JPEG data is not optimal.

Flashback Query

Last updated : 2024-07-22 11:35:35

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](#).

Use Cases

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%'`).

Instructions

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 the 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)
```

Example of creating a table from 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';
```

Parameters

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

Parameter	Scope	Type	Default Value	Value Range/Valid Values	Restart Required
Innodb_backquery_enable	Global	Boolean	OFF	ON/OFF	No
Innodb_backquery_window	Global	Integer	900	1-86400	No
Innodb_backquery_history_limit	Global	Integer	8000000	1-9223372036854476000	No

--	--	--	--	--	--

Performance Features

Parallel Query

Overview

Last updated : 2024-07-22 12:34:17

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: You have many parameters at hand to control the conditions of enabling or disabling parallel query. This helps you make queries smarter and more adaptable to your business scenarios with no transformation needed.

Supported Statement Scenarios and Restricted Scenarios

Last updated : 2024-07-22 12:34:38

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.

Restriction	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 `index_merge` 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 : 2024-07-22 12:34:53

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

Prerequisites

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

Parameters

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

Parameter	Variable Type	Scope	Default Value	Value Range
txsql_max_parallel_worker_threads	Integer	Global	{MIN(DBInitCpu,0)}	0– {MAX(DBInitCpu-2,2)}

txsql_parallel_degree	Integer	Global/session	4	0-64

Suggested 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:

--	--	--	--

Parameter	Variable Type	Scope	Default Value
innodb_txsql_parallel_partitions_per_worker	Integer	Global/Session	13
txsql_optimizer_context_max_mem_size	Integer	Global/Session	{MIN(DBInitMemory*52429,83886144)}
txsql_parallel_cost_threshold	Integer	Global/Session	50000
txsql_parallel_exchange_buffer_size	Integer	Global/Session	1048576
txsql_parallel_table_record_threshold	Integer	Global/Session	5000

--	--	--	--

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](#).

Database List Parameter Settings Account Management				
Batch Modify Parameters	Default Template <small>NEW</small>	Custom Template	Import Parameters	Export Parameters
Save as Template	Intelligent Parameter Tuning <small>NEW</small>	txsql_max_parallel_worker_threads 🔍 Recent Modifications		
Parameter Name	Instance...	Default Value ^①	Current Value	Acceptable Values
Search by parameter name "txsql_max_parallel_worker_threads", 1 result in total Back to List				
txsql_max_parallel_worker_threads ^①	No	(MIN(DBInitCpu,0))	🔄 0	[0-(MAX(DBInitCpu-2,2))] ^①

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

TencentDB for MySQL allows you to specify the parallel execution mode of a SQL statement through the HINT statement. For detailed directions, see [HINT Statement Control](#).

References

- [Viewing Parallel Query](#)
- [HINT Statement Control](#)

HINT Statement Control

Last updated : 2024-07-22 12:35:02

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

Feature	Command Line	Description
Enable parallel query	<code>SELECT /*+PARALLEL (x) */ ... FROM ...;</code>	<code>x</code> indicates the parallelism for the SQL statement, which should be greater than <code>0</code> .
Disable parallel query	<code>SELECT /*+PARALLEL (x) */ ... FROM ...;</code>	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> <p>Specify the table to be included in the plan through</p> <code>PARALLEL . SELECT /*+PARALLEL (t) */ ... FROM ...;</code> <p>Specify the table to be excluded from the plan through</p> <code>NO_PARALLEL . SELECT /*+NO_PARALLEL (t) */ ... FROM ...;</code>	<code>t</code> is the table name.
Specify both the parallel table and parallelism	<code>SELECT /*+PARALLEL (t x) */ * ... FROM ...;</code>	<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	<code>SELECT /*+SET_VAR (var=n) */ * ... FROM ...;</code>	<code>var</code> is the parallel query parameter in the

the HINT statement, which takes effect only for the specified SQL statement

session 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.

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

Set `txsql_optimizer_context_max_mem_size=500000` for a statement, which means to adjust the maximum memory size it can apply for in the parallel query plan environment to `500000`.

References

[Enabling/Disabling Parallel Query](#)

Viewing Parallel Query

Viewing Parallel Query

Last updated : 2024-07-22 12:35:17

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.

EXPLAIN statement:

```
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)
```

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      |
+-----+-----+-----+-----+-----+-----+-----+
| 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.

References

[Enabling/Disabling Parallel Query](#)

[HINT Statement Control](#)

Large Transaction Replication

Last updated : 2024-07-22 12:35:36

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.

Use Cases

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 Data

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

Instructions

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

Parameter	Status	Type	Default Value	Description
<code>cdb_optimize_large_trans_binlog</code>	true	bool	false	Switch of transaction optimization
<code>cdb_optimize_large_trans_binlog_last_affected_rows_threshold</code>	true	ulonglong	10000	Large transaction replication condition number (time)
<code>cdb_optimize_large_trans_binlog_aver_affected_rows_threshold</code>	true	ulonglong	10000	Large transaction replication condition average rows
<code>cdb_sql_statistics</code>	true	bool	false	Switch of execution collector
<code>cdb_sql_statistics_info_threshold</code>	true	ulonglong	10000	Maximum statement of <code>CDB_SQL_STATISTICS</code>

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	Type	Description
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_LONGLONG	The number of executions of this SQL statement.
TOTAL_AFFECTED_ROWS	MYSQL_TYPE_LONGLONG	Total number of affected rows.
AVER_AFFECTED_ROWS	MYSQL_TYPE_LONGLONG	Average number of affected rows.
LAST_AFFECTED_ROWS	MYSQL_TYPE_LONGLONG	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 : 2024-07-22 12:35:46

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.

Use Cases

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.

Performance Impact

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.

Instructions

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	Effective Immediately	Type	Default Value	Valid Values/Value Range	Description
cdb_plan_cache	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 parameter. If needed, [submit a ticket](#) for assistance.

You can run the `show cdb_plan_cache` command to query information about execution plan cache hits. The command will return the following fields:

Field	Description
sql	A SQL statement with the question mark (?) which represents that the execution plan of this statement has been cached.
mode	SQL cache mode. Currently, only the <code>prepare</code> mode is supported.
hit	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 : 2024-07-22 12:35:56

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.

Use Cases

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

Performance Data

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.

Instructions

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	Type	Default Value	Valid Values/Value Range	Description
<code>innodb_flush_redo_using_fdatasync</code>	Yes	bool	false	true/false	Whether to call <code>fdatasync()</code>

					to flush redo logs
--	--	--	--	--	--------------------

Auto-Increment Column Persistence

Last updated : 2024-07-22 12:36:06

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.

Use Cases

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

Instructions

This feature is enabled in the kernel by default.

Buffer Pool Initialization

Last updated : 2024-07-22 12:36:19

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.

Use Cases

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 (%)
50 GB	2.55 s	0.13 s	1,962%
200 GB	10.28 s	0.52 s	1,977%
500 GB	25.72 s	1.32 s	1,948%

Instructions

This feature is enabled in the kernel by default.

FAST DDL

Last updated : 2024-07-22 12:36:36

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 and above.

Kernel version: MySQL 5.7 20210331 and above.

Use Cases

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 Data

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.

Instructions

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

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

Note:

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

Invisible Index

Last updated : 2024-07-22 12:36:47

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.

Use Cases

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

Instructions

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 : 2024-07-22 12:37:00

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.

Use Cases

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

Performance Data

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

Thread Count	FCFS (FIFO)	CATS	Performance Improvement
128	11,999	12,005	0%
256	6,609	10,137	53%
512	3,453	9,365	171%
1,024	2,196	7,015	219%

Instructions

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	Type	Default Value	Valid Values/Value Range	Description
<code>innodb_trx_schedule_algorithm</code>	Yes	string	<code>auto</code>	<code>[auto,fcfs,cats]</code>	Specify the transaction scheduling algorithm

Note:

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

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

Computation Pushdown

Last updated : 2024-07-22 12:37:15

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.

SUM optimization applies to kernel version MySQL 5.7 20180918.

Use Cases

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 Data

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.

Instructions

During the execution of an 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	Type	Default Value	Valid Values/Value Range	Description
<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 : 2024-07-22 12:38:10

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 later.

Kernel version: MySQL 8.0 20200630 and later.

Use Cases

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.

Instructions

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

Audit

Last updated : 2025-05-28 16:10:52

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.

Use Cases

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.

Performance Impact

Database audit supports synchronous audit mode. Synchronous audit records all audit logs synchronously, and the average performance impact is less than 6%.

Instructions

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 : 2024-11-27 11:20:24

Overview

The quick column addition 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 later.

Kernel version: MySQL 8.0 20200630 and later.

Use Cases

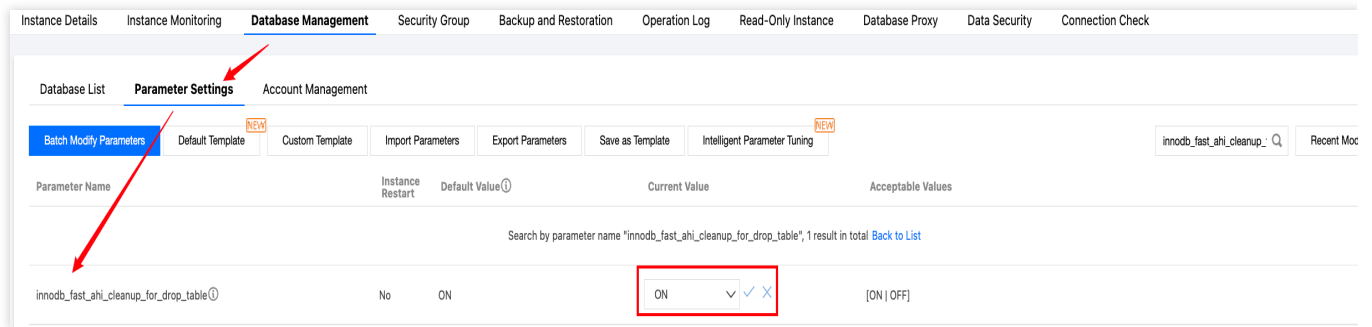
This feature is suitable for adding columns to a table with a high volume of data.

Performance Data

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

Instructions

Before the feature of adding columns in seconds is used, it is necessary to set the parameter `innodb_fast_ahi_cleanup_for_drop_table` to ON.



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.

Note :

The setting of the parameter `innodb_alter_table_default_algorithm` is only applicable to MySQL version 5.7. For MySQL version 8.0, the default algorithm for executing Alter Table operations in kernel subversion 20230630 and later is `instant`, with this parameter unable to be set.

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 `COMPRESSED` row format.

`INSTANT ADD COLUMN` is not supported in tables with a full-text index.

`INSTANT ADD COLUMN` is not supported for temp tables.

Second-Level Column Modification

Last updated : 2024-12-30 17:52:01

Description of the Feature

Second-level Column Modification operations are only recorded in the Data Dictionary Table to log Column Modification Information, avoiding the data copying that was previously necessary for column modification operations. This significantly reduces the time required for Large Table Column Modifications, minimizing the impact on Application Systems and resource consumption.

Supported Versions

Kernel version MySQL 8.0 20230702 and above.

Applicable Scenario

This feature is suitable for modifying columns in Tables with Large Data Volumes.

Test Results

Number of table rows	Modification Duration without using Second-level Column Modification Feature	Modification Duration with Second-level Column Modification Feature
1 million	22.9 seconds	0.01 seconds
10 Million	13 minutes and 39.72 seconds	0.01 seconds
100 million	3 hours, 51 minutes, 16.40 seconds	0.01 seconds

Use Instructions

Second-level Column Modification Syntax ALTER TABLE add algorithm = instant Clause, column modification can be performed with the following statement:

```
ALTER TABLE modify_tab_col MODIFY COLUMN c1 BIGINT,ALGORITHM=INSTANT;
```

Add the parameter `cdb_instant_modify_column_enabled` to control the Second-level Column Modification feature, which can be set to ON/OFF.

Parameter Name	Status	Type	Default Value	Valid Values/Value Range	Description
<code>cdb_instant_modify_column_enabled</code>	yes	bool	OFF	ON/OFF	Feature Toggle, determining whether to enable the Second-level Column Modification feature.

Note:

Users cannot directly modify the parameter values mentioned above. To make changes, [submit a ticket](#).

Second-level Column Modification Restrictions

Support is limited to modifying column types only; changes to the field's nullable, unsigned/signed, charset are not supported, but modifying the default property is allowed.

Only certain type modifications are supported, and only length can be increased. Currently, conversions are only supported between char and varchar, binary and varbinary, and amongst tinyint/smallint/mediumint/int/bigint.

For example:

`char(10) → char(100)`

`char(10) → varchar(100)`

`varchar(10) → char(100)`

`varchar(10) → varchar(100)`

`binary(10) → binary(100)`

`binary(10) → varbinary(100)`

`varbinary(10) → binary(100)`

`varbinary(10) → varbinary(100)`

`tinyint/smallint/mediumint/int → bigint`

`tinyint/smallint/mediumint → int`

`tinyint/smallint → mediumint`

`tinyint → smallint`

Note:

Modifications between char/varchar and binary/varbinary, as well as integer types, are not supported. For example, changing from char to binary, binary to varchar, or reducing the range of integer types, such as from bigint to int or int

to smallint, is not allowed.

A single column can only be modified using `INSTANT MODIFY COLUMN` once, but multiple columns can be modified with it simultaneously.

After a single column is added/modified using `INSTANT ADD/MODIFY COLUMN` for the first time, any subsequent modifications to that column must be done in a non-instant manner.

[INSTANT ADD COLUMNS](#) and `INSTANT MODIFY COLUMNS` operations must be executed separately. You may first execute `INSTANT ADD COLUMNS`, then `INSTANT MODIFY COLUMNS`, or vice versa. You cannot perform `INSTANT MODIFY COLUMN` on a column that has been added with `INSTANT ADD COLUMN`.

You cannot modify the column name and column type at the same time. Instead, you can modify the column name first and then column type.

Import/export is not supported.

Encryption and compression are not supported.

Async Deletion of Big Tables

Last updated : 2025-04-16 17:29:32

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 20220715 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

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_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	Dynamic	Type	Default	Value Range	Description
<code>innodb_adaptive_hash_index</code>	yes	string	ON	ON/OFF	Whether to enable InnoDB adaptive hash index. ON: Enable. OFF: Disable.
<code>innodb_async_truncate_work_enabled</code>	yes	string	OFF	ON/OFF	Whether to enable asynchronous big table deletion. ON: Enable. OFF: Disable.

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
<code>innodb_fast_ahi_cleanup_for_drop_table</code>	yes	string	ON	ON/OFF

innodb_adaptive_hash_index	yes	string	ON	ON/OFF
innodb_table_drop_mode	yes	string	ASYNC_DROP	SYNC_DROP/ASYNC_D

innodb_async_truncate_size	yes	int	128	128 - 168
----------------------------	-----	-----	-----	-----------

Hotspot Update

Last updated : 2024-07-22 12:42:32

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.

Use Cases

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 Data

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



Instructions

For more information, see [Hotspot Update Protection](#).

SQL Throttling

Last updated : 2024-06-18 17:24:18

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 8.0 20211202 and Later

Kernel Version MySQL 5.7 20200331 and Later

Kernel Version MySQL 5.6 20200915 and Later

Use Cases

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

Instructions

For more information, see [SQL Throttling](#).

Statement Outline

Last updated : 2024-07-22 12:42:52

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.

Use Cases

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.

Performance Impact

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.

Instructions

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	Type	Default Value	Valid Values	Description
cdb_opt_outline_enabled	Yes	bool	false	true/false	Whether to enable the outline feature.

Note:

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

Binding outline

To bind an outline, replace the SQL statement with another where the SQL syntax is not changed but some hint information is added to tell the optimizer how to execute the statement.

The syntax is in the format of `outline "sql" set outline_info "OUTLINE";`. Note that the string after `outline_info` must start with `"OUTLINE:"`, which is followed by the SQL statement with the hint information added. For example, you can add the index in column `a` to table `t2` in the SQL statement `select *from t1, t2 where t1.a = t2.a` as follows:

```
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 make the feature more flexible, TXSQL allows you to add optimizer 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 `"OPT:"`, which is followed by the optimizer hint information to be added. For example, you can specify `SEMIJOIN` of `MATERIALIZATION/DUPSWEEDOUT` for the SQL statement `select *from t1 where t1.a in (select b from t2)` as follows:

```
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 `test.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 Name	Description
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 Name	Description
origin	Original SQL statement fingerprint.
outline	SQL statement fingerprint after the outline is bound.

TXRocks Engine Overview

Last updated : 2024-07-22 12:48:36

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.

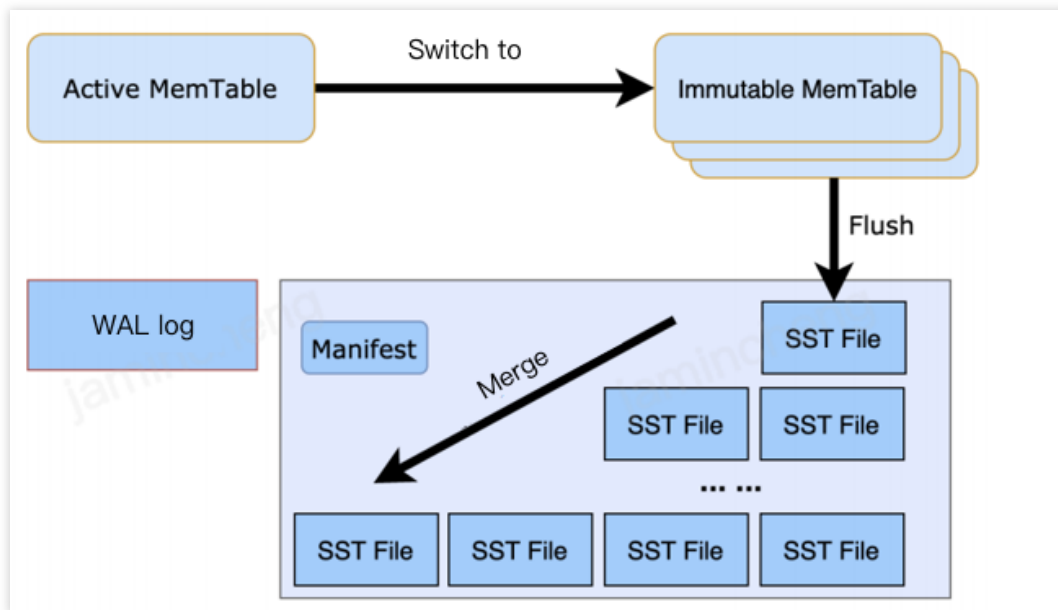
For a write request, the new version of the record is first written to an active MemTable, and then WAL is written for data durability. After the MemTable and WAL file are written for the request, a response can be returned.

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.

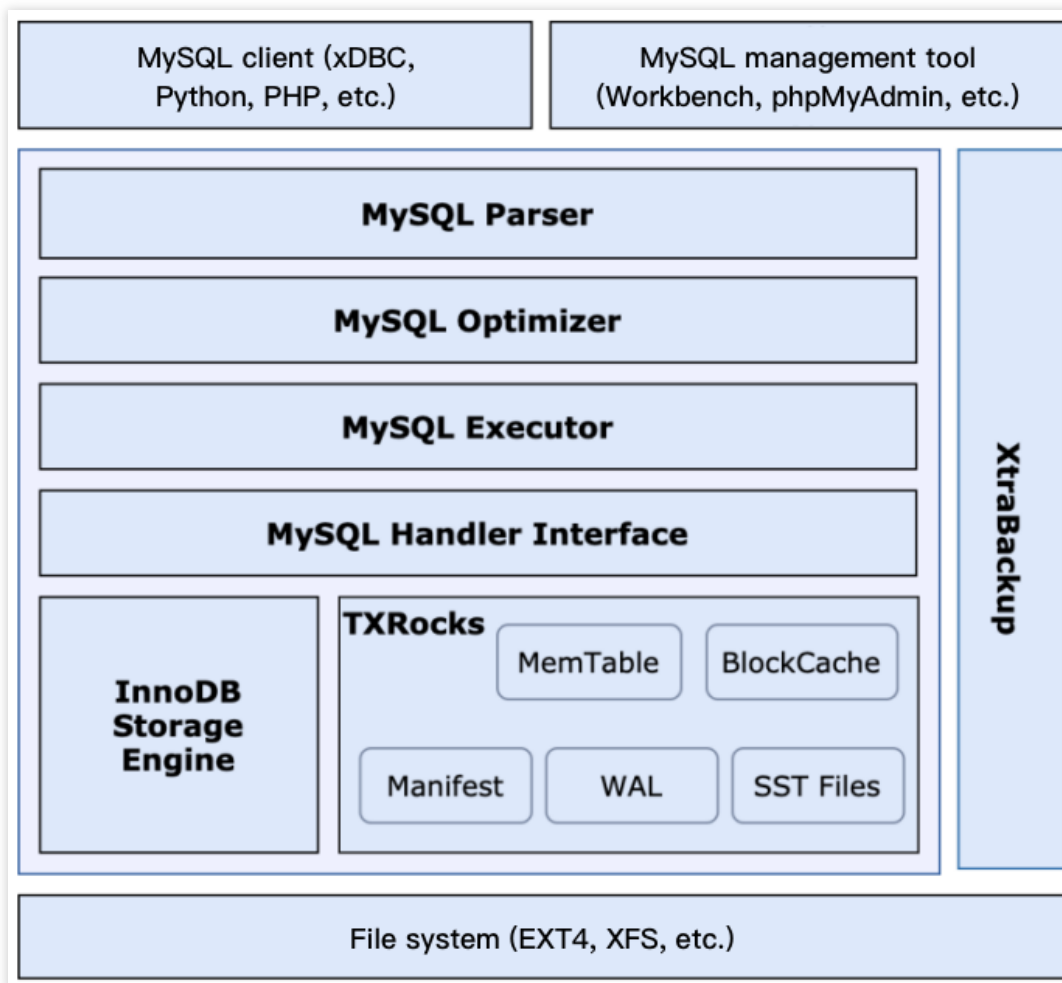
When a record is read, it will be searched for in the active MemTable, immutable MemTable, SST files at L0, and SST files at L1–6 in sequence. Once the record is found in any component, the latest version of the record is found and will be returned immediately.

When a range scan is performed, an iterator will be generated at each level of data containing MemTables. The iterators perform a merge sort to find the next record. As shown in the read process, if the LSM tree has too many levels, the read performance, especially the performance of range scan will drop significantly. Therefore, to maintain a

better LSM tree shape, the backend continuously performs compaction operations to merge low-level data to a higher level and thus reduce the number of levels.



TXRocks Architecture



TXRocks Strengths

Less 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.

Use Cases

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 : 2024-07-22 12:50:43

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.

Product Overview

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.

Prerequisites

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	MySQL5.5	MySQL5.6	MySQL5.7	MySQL8.0
Engine	InnoDB	RocksDB NEW		

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:

Category	Feature	TXRocks Limit
DDL	Online DDL	Not supported. For example, `ALTER TABLE ... ALOGRITHM=INSTANT` is not supported. Only the `COPY` algorithm is supported for partition management operations.
SQL	Foreign key	Not supported.

	Partitioned 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"</code> reported.
	Encrypted table	Not supported.
Index	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.
	Multi-valued index	Not supported.
Replication	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.
Transaction and lock	LOCK NOWAIT and SKIP LOCKED	Not supported.
	Gap lock	Not supported.
	Savepoint	Not supported.
	Partial LOB field update	Not supported.
	XA transaction	Not recommended.

Parameter Description

Note:

When creating a TencentDB for MySQL instance, you can 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	Restart Required	Default Value	Value Range/Valid Values	Description
rocksdb_use_direct_io_for_flush_and_compaction	Yes	ON	ON/OFF	Whether to use D
rocksdb_flush_log_at_trx_commit	No	1	0/1/2	Controls when to It is similar to <code>innodb_flush</code> and indicates whe be synced when k 0: Transactions a committed. 1: Transactions a committed. 2: Transactions a second.
rocksdb_lock_wait_timeout	No	1	1– 1073741824	Lock wait timeout
rocksdb_deadlock_detect	No	ON	ON/OFF	Whether to enable it is enabled, all d recorded in <code>mys</code>
rocksdb_manual_wal_flush	Yes	ON	ON/OFF	If the total size of <code>rocksdb_max_</code> RocksDB will forc to the disk to ensi can be deleted.

MySQL 8.0 parameter list

Parameter	Restart Required	Default Value	Value Range/Valid Values	Descriptio
rocksdb_flush_log_at_trx_commit	No	1	0/1/2	Controls v It is simila <code>innodb</code> and indica be syncce 0: Transa committec 1: Transa committec

				2: Transa second.
rocksdb_lock_wait_timeout	No	1	1-1073741824	Lock wait
rocksdb_merge_buf_size	No	524288(=512K)	100- 18446744073709551615	Size of the secondary
rocksdb_merge_combine_read_size	No	8388608 (=8M)	524288(=512K)- 18446744073709551615	Size of the during sec
rocksdb_deadlock_detect	No	ON	ON/OFF	Whether t
rocksdb_manual_wal_flush	Yes	ON	ON/OFF	If the total rocksdb RocksDB to the disk can be de

RocksDB Monitoring Metrics

RocksDB monitoring metrics are as listed below:

Metric	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 : 2024-07-22 12:50:55

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

Background

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](#).

Practical Tutorial of TXRocks

Last updated : 2024-07-31 10:00:32

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

Background

Scenario: The import of massive amounts of data to a database with the TXRocks engine needs to be accelerated.

Impact: The `Rows inserted during bulk load must not overlap existing rows` error may be reported when massive amounts of data are imported.

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;
...
Import the 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;
...
Import the 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.

LibraDB Engine

LibraDB Engine Feature Highlights

Last updated : 2025-05-09 11:51:17

Feature Introduction

The LibraDB engine focuses on efficient analysis queries, offering a read-only extension for real-time and high-performance complex SQL processing. With its columnar storage capability, vectorized parallel execution engine, and distributed query optimizer, users can quickly achieve high-performance analysis directly within the database. Additionally, LibraDB's columnar storage has been optimized for high-QPS updates and ACID transactions, ensuring real-time data consistency.

Principles

The LibraDB engine kernel automatically converts data from the primary node into columnar storage while synchronizing the binlog in real time, ensuring data remains consistently aligned with the primary node. Additionally, by leveraging parallel processing capabilities and a vectorized execution engine, it efficiently processes complex user SQL queries and accelerates query execution.

Supported Features

The LibraDB engine kernel boasts a variety of outstanding features. Below is a brief introduction to the features supported by the product.

1. Parallel Computing Capability

In the LibraDB engine, the computing resources of nodes are fully leveraged to execute user SQL queries. When an SQL query is processed within the LibraDB engine, it can be divided into multiple subtasks that run concurrently across multiple processors. This multi-threaded execution effectively enhances CPU utilization and optimizes I/O resource usage, achieving accelerated processing.

1. Vectorized Execution Engine

Regarding the LibraDB engine, data is not only stored in columns but is also processed on a column basis. In traditional OLTP engines like TXSQL, computations are typically performed on row-based storage since transactions primarily involve point queries, reads, and writes. However, with the LibraDB engine focused on analytical processing,

the computational workload of a single SQL query can be immense. To address this, we have implemented a vectorized execution mode in LibraDB. This mode processes in-memory columnar data in batches, invoking SIMD instructions once per batch, which reduces function call overhead and minimizes cache misses. Additionally, it fully leverages the parallel capabilities of SIMD instructions, significantly shortening computation time.

1. Support for Columnar Storage in High-Speed Change Scenarios

The primary instance of TencentDB for MySQL supports online data operations with a QPS exceeding one million. As a LibraDB engine designed to support real-time data analysis, it should ensure data consistency even under scenarios of rapid data changes. Traditional columnar storage offers advantages for high-volume data writes, but its performance is inadequate when dealing with large-scale deletions and updates. In conventional real-time data warehouse settings, a good practice is to convert updates into delete-and-insert operations and implement batch data processing at the data synchronization layer. However, columnar storage still faces performance limitations in deletion scenarios. Considering these factors, traditional columnar storage inevitably leads to high data latency and fails to achieve the desired real-time data analysis outcomes.

The LibraDB engine ensures data consistency for changes in high-concurrency scenarios by leveraging optimizations and support at the storage layer. This prevents missed analysis opportunities caused by data latency from frequent updates in read-write instances.

1. Targeted Data Loading Capability

In TencentDB for MySQL, not all data possesses analytical value, so it is not required to load every object as columnar storage. The LibraDB engine allows you to specify which objects to load. You can configure this either through the data loading console or by using SQL commands in the command line.

ETL Writeback Acceleration

Last updated : 2025-05-09 11:51:17

In databases, the `INSERT INTO <Table> SELECT...` statement is often executed. If the execution time of `SELECT` is too long, the overall execution efficiency becomes low. The read-only analysis engine can handle complex queries in online SQL statements. By using this feature, the `SELECT` query in an `INSERT...SELECT...` statement can be executed in the read-only analysis engine to accelerate the `SELECT` query. The data is automatically written back to the read-write instance, thereby greatly improving the SQL execution efficiency.

Technical Principles

As shown in the above figure, you can execute the `INSERT...SELECT...` statement in the read-only analysis engine. The `SELECT` query will be accelerated by the read-only analysis engine, and the query results will be directly written to the read-write instance, thereby improving the overall execution efficiency.

Application Scenario

Note:

This feature applies only to scenarios that are not sensitive to data latency. The read-only analysis engine adopts the asynchronous replication mode. Therefore, a certain time delay may exist between the query results and those in the read-write instance when there is a latency in the read-only analysis engine.

This feature is recommended in scenarios where the query conditions are complex, the SQL statement execution time is long, and the data volume of the query result set is small. In such scenarios, this feature can significantly improve performance because the read-only analysis engine will accelerate `SELECT` queries.

This feature does not bring performance benefits in all scenarios. Instead, the performance may deteriorate in some scenarios. For example:

When the `SELECT` query in an `INSERT...SELECT...` statement is relatively simple, reading data from the read-only analysis engine and writing the data back to the read-write instance will incur additional network overhead. The strengths are not obvious compared with the method of directly reading data from the read-write instance.

When the result set of the `SELECT` query in an `INSERT...SELECT...` statement contains a large volume of data, the main performance bottleneck lies in the process of transmitting the result set over the network and writing it to the read-write instance. In this scenario, this feature cannot be used to improve performance.

Ordered Pagination Feature

Last updated : 2025-05-09 11:51:17

In MySQL, the data on the second page may be duplicated with the data on the first page if an order by clause with duplicate field values is used together with a limit clause for pagination queries. This is because MySQL uses heapsort, which is an unstable sorting method. When duplicate values exist, the output result changes each time.

Therefore, when an order by...limit... statement is used, it is suggested that indexes are set for fields in the order by clause or other fields are added to the order by clause to maintain data uniqueness.

To solve this problem, the read-only analysis engine provides the ordered pagination feature. By using this feature, you do not need to worry about the database logic, and no duplicate data exists in the result output by order by...limit..., ensuring the sequence consistency.

Implementation Principles

In SQL execution scenarios involving sorting and pagination, the read-only analysis engine performs implicit sorting on the full data by default after it determines the business sorting logic. This avoids the problem of data duplication caused by limit operations.

Ordered Pagination Scenarios

Only the limit operator is used, and no order by operator is used.

Duplicate order by keys exist, and only some fields in the output result are sorted.

The subquery involves sorting, but sorting is not performed in the outer query.

Impact on Performance

If the ordered pagination feature is enabled, additional sorting operations will be performed. Therefore, some queries may experience performance degradation. Please choose whether to use the ordered pagination feature according to the actual business situation.

Usage Instructions

For details on enabling and using the ordered pagination feature, see [Paging Order Preservation Feature](#).