

TaurusDB

Product Bulletin

Issue 01
Date 2025-02-12



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2025. All rights reserved.

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

Trademarks and Permissions



HUAWEI and other Huawei trademarks are the property of Huawei Technologies Co., Ltd.

All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei Cloud and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Cloud Computing Technologies Co., Ltd.

Address: Huawei Cloud Data Center Jiaoxinggong Road
Qianzhong Avenue
Gui'an New District
Gui Zhou 550029
People's Republic of China

Website: <https://www.huaweicloud.com/intl/en-us/>

Contents

1 Product Bulletin.....	1
1.1 GaussDB(for MySQL) Will Be Renamed TaurusDB.....	1
1.2 TaurusDB Resource Display Pages Will Be Adjusted on the Console.....	1
1.3 Serverless DB Instances Will Be Released Commercially on Jan 19, 2024, 00:00 GMT+08:00.....	3
1.4 HTAP of Standard Edition Will Be Released Commercially on Jan 6, 2025, 00:00 GMT+08:00.....	3
2 Product Release Notes.....	4
2.1 TaurusDB Kernel Version Release History.....	4
2.2 TaurusDB Version Release Policy.....	17

1 Product Bulletin

1.1 GaussDB(for MySQL) Will Be Renamed TaurusDB

Change Content

Some time from October 2024 to December 2024, GaussDB(for MySQL) will be renamed TaurusDB.

Change Scope

We will update the product names in the portal, console, and Help Center documentation.

Change Impact

Please note that during this transition period, you may still see references to both GaussDB(for MySQL) and TaurusDB. Your services and product functions will not be affected.

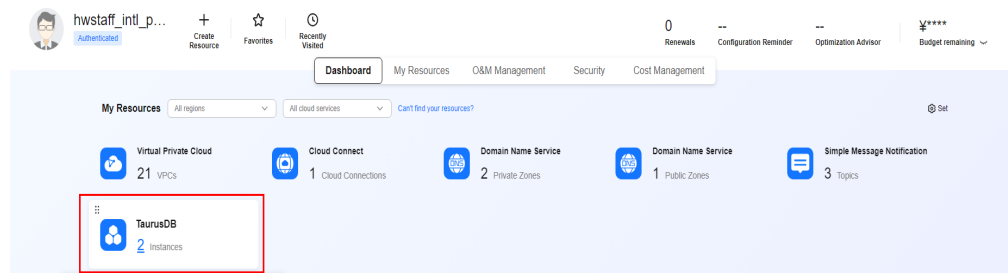
1.2 TaurusDB Resource Display Pages Will Be Adjusted on the Console

To improve user experience, Huawei Cloud plans to adjust the TaurusDB resource display pages on the console by region from April 2024.

Change Content

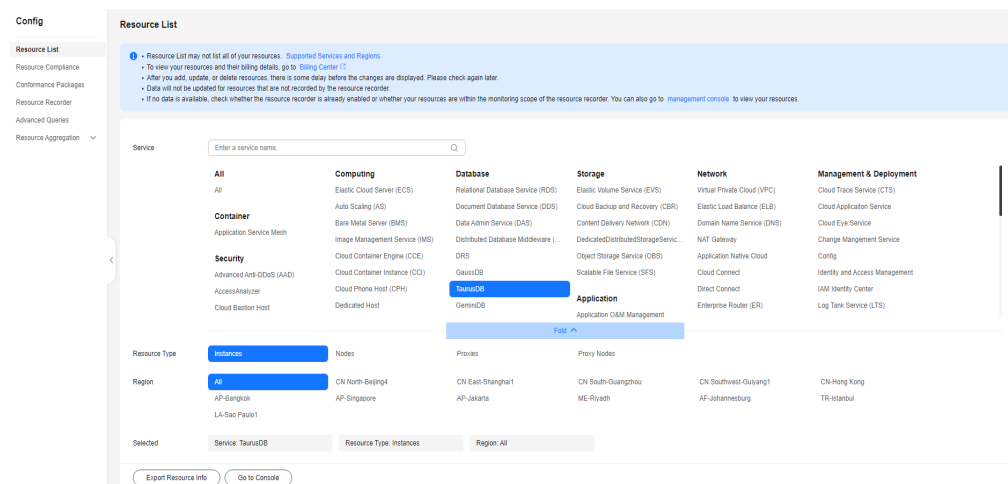
1. On the **Dashboard** page of the console, TaurusDB instance and node information is no longer displayed under GaussDB(for MySQL). Instead, it is displayed under TaurusDB. You can click **Go to Console** to access the **Instances** page.

Figure 1-1 Dashboard page



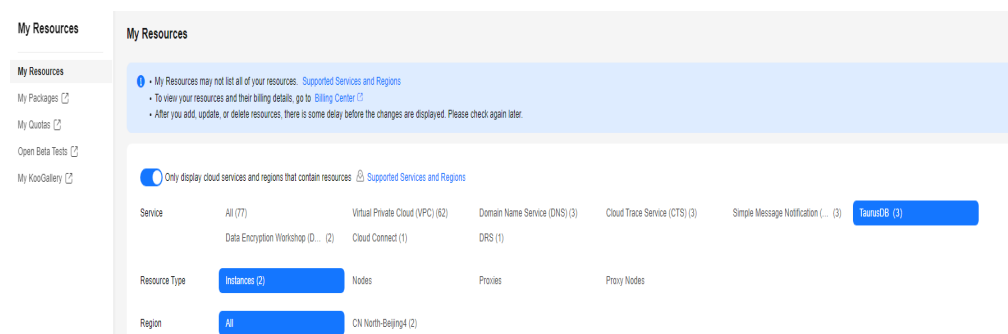
2. On the **Config > Resource List** page, TaurusDB instance and node information is no longer displayed under GaussDB(for MySQL). Instead, it is displayed under TaurusDB.

Figure 1-2 Resource List page



3. To view TaurusDB resources on the **Resources > My Resources** page, request the ***:*:list** permission.

Figure 1-3 My Resources page



Change Impact

After the adjustment, TaurusDB resources will be displayed under the TaurusDB service, instead of the GaussDB(for MySQL) service.

If you have any questions about the adjustment, [submit a service ticket](#) or contact us at +86-4000-955-988 or +86-950-808.

1.3 Serverless DB Instances Will Be Released Commercially on Jan 19, 2024, 00:00 GMT+08:00

Commercial Use Notice

TaurusDB instances billed on a serverless basis will be released commercially on Jan 19, 2024, 00:00 GMT+08:00.

The capacities of serverless DB instances automatically change based on application requirements, reducing costs. For billing details, see [Billing](#).

To avoid incurring service fees, if you do not want to continue using serverless DB instances, delete them on the console before Jan 19, 2024, 00:00 GMT+08:00.

Regions

CN North-Beijing4, CN East-Shanghai1, CN South-Guangzhou, AP-Singapore, and ME-Riyadh

Impacts

- Existing serverless DB instances
Serverless DB instances can be used free of charge before Jan 19, 2024, 00:00 GMT+08:00, but they will be billed normally after Jan 19, 2024, 00:00 GMT+08:00.
- New serverless DB instances
If you buy a serverless DB instance after Jan 19, 2024, 00:00 GMT+08:00, you will be billed for the instance. For details about how to buy such an instance, see [Buying a Serverless DB Instance](#).

1.4 HTAP of Standard Edition Will Be Released Commercially on Jan 6, 2025, 00:00 GMT+08:00

Commercial Use Notice

Huawei Cloud will commercially release TaurusDB HTAP of Standard Edition on Jan 6, 2025, 00:00 GMT+08:00. Since then, TaurusDB HTAP of Standard Edition will be billed.

To avoid unnecessary fees, if you do not want to continue using HTAP of Standard Edition, unsubscribe from or delete standard HTAP instances on the TaurusDB console before Jan 6, 2025, 00:00 GMT+08:00.

Click [here](#) for more information about HTAP instances.

If you have any questions or suggestions, please [submit a service ticket](#) or call us at +86-4000-955-988 or +86-950-808.

Thank you for using Huawei Cloud.

2 Product Release Notes

2.1 TaurusDB Kernel Version Release History

This section describes the kernel version updates of TaurusDB.

2.0.60.241200

Table 2-1 Version 2.0.60.241200

Date	Description
2025-01-14	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Fast binlog positioning: If you pull binlogs from a TaurusDB instance using auto-positioning, it can be time-consuming to locate the correct binlog position when there are a large number of unread binlogs. For example, if there are 60,000 unread binlogs, it would take about 60 minutes to automatically locate the correct binlog position. With fast binlog positioning enabled, the time required can be reduced from 60 minutes to about 2 seconds. - Enhanced parallel query: You can use EXPLAIN statements to view parallel execution plans, including EXPLAIN FORMAT=TREE, EXPLAIN ANALYZE (supported in 2.0.60.241200 and later), and EXPLAIN FORMAT=JSON introduced in MySQL 8.0 and traditional EXPLAIN statements. ● Fixed issues <ul style="list-style-type: none"> - Fixed the issue that workloads are interrupted occasionally when Flink is used to pull binlogs from read replicas in some scenarios. - Fixed the issue in the open-source community that the results are incorrect due to duplicate JSON column index data (Bug #33334928). - Optimized the traffic distribution to OLAP. - Changed the ROW keyword from a reserved keyword to a non-reserved keyword.

2.0.57.240900

Table 2-2 Version 2.0.57.240900

Date	Description
2024-10-18	<ul style="list-style-type: none"> • New features and optimized features <ul style="list-style-type: none"> - Partition-level MDL: In MySQL Community Edition, you cannot perform both data manipulation language (DML) operations for accessing data of partitioned tables and data definition language (DDL) operations for maintaining partitions at the same time. This means that DDL operations can only be done during off-peak hours. To resolve such an issue, this version introduces partition-level metadata lock (MDL) to refine the lock granularity of a partitioned table from the table level to the partition level. After partition-level MDL is enabled, DML operations and specific DDL operations (such as adding and deleting partitions) on different partitions can be both performed, greatly improving concurrency between partitions. - Table recycle bin: After this function is enabled, the DROP TABLE statement that meets conditions does not directly delete a specified table. Instead, the table is temporarily stored in the recycle bin. When the maximum retention period expires, the table is automatically deleted in the backend. You can change the retention period of a deleted table in the recycle bin. You can also restore or permanently delete a table from the recycle bin at any time. • Fixed issues <ul style="list-style-type: none"> - Fixed the issue that CPU resources of each tenant are not strictly allocated based on the configured ratio in resource preemption scenarios. - Allowed Statement Outline to support views and EXPLAIN ANALYZE statements.

2.0.54.240600

Table 2-3 Version 2.0.54.240600

Date	Description
2024-07-19	<ul style="list-style-type: none"> • New features and optimized features <ul style="list-style-type: none"> - Optimized hot row update: Hot rows are frequently updated for flash sales, concert ticket reservations, and train ticket bookings for popular routes. This version enhances hot row update, which can be enabled automatically or manually. After hot row update is enabled, hot rows can be updated efficiently. - Non-blocking DDL: When you perform a DDL operation, if the target table has uncommitted long transactions or large queries, the DDL operation continuously waits for obtaining the MDL-X lock. As a result, service connections are stacked and blocked. This version supports non-blocking DDL, which allows new transactions to enter the target table even if the MDL-X lock cannot be obtained, ensuring the stability of the entire service system. - Multi-tenant management: This feature enables a database to serve multiple tenants, maximizing database resource utilization. - Binlog pull for read replicas: You can use read replicas as the data source to establish a binlog replication link and synchronize the binlogs in real time, which helps reduce the load on the primary node. - Column compression: TaurusDB introduces fine-grained column compression to reduce data page storage and save costs. Two compression algorithms, ZLIB and ZSTD, are provided. You can select either of them to compress infrequently accessed large columns based on the compression ratio and compression and decompression performance. - INTERVAL RANGE partitioned tables: In previous versions, if the data to be inserted into an existing RANGE partitioned table exceeds the range of existing partitions, the data cannot be inserted and an error is returned. With the support for INTERVAL RANGE partitioned tables in this version, the database can now create partitions based on rules specified by the INTERVAL clause when new data exceeds the range of existing partitions. - LIST DEFAULT HASH partitioned tables: This feature supports two types of partitions at the same level: LIST and HASH. Data is first inserted into LIST partitions. Data that does not comply with the LIST

Date	Description
	<p>partitioning rules is placed in the DEFAULT partition. If the DEFAULT partition has multiple partitions, HASH rules are used. LIST DEFAULT HASH partitioned tables are usually used in scenarios where LIST VALUES are unevenly distributed and cannot be fully enumerated.</p> <ul style="list-style-type: none"> ● Fixed issues <ul style="list-style-type: none"> - Optimized the table-level restoration performance. - Optimized the execution performance of read replicas of a high-spec instance in high-concurrency scenarios.

2.0.51.240300

Table 2-4 Version 2.0.51.240300

Date	Description
2024-03-30	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Added global consistency, which provides strongly consistent reads at the cluster level with low performance loss. - Added the SHOW BINARY LOGS NO BLOCK syntax, which prevents transaction commits from being blocked during the execution of SHOW BINARY LOGS. - Optimized the UNDO TRUNCATE capability, which solves the issue of undo space expansion caused by a large number of writes. - Enhanced the degree of parallelism for full restoration, which optimizes the backup and restoration efficiency. ● Fixed issues <ul style="list-style-type: none"> - The query results of window functions are incorrect, or errors occur when window functions are executed. - Database nodes break down when specific PREPARE statements are repeatedly executed after plan cache is enabled. - An error is reported due to inconsistent character sets when stored procedures are executed in sequence. - Query results do not meet the expectation when an on-disk hash join is performed after PQ is enabled. - An error is reported due to duplicate primary keys when a query involves performing a GROUP BY operation on temporary table fields.

2.0.48.231200

Table 2-5 Version 2.0.48.231200

Date	Description
2024-01-30	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Enhanced composite partitioning: In addition to RANGE-HASH and LIST-HASH of MySQL Community Edition, added RANGE-RANGE, RANGE-LIST, LIST-RANGE, LIST-LIST, HASH-HASH, HASH-KEY, HASH-RANGE, HASH-LIST, KEY-HASH, KEY-KEY, KEY-RANGE, and KEY-LIST. - Added the forward compatibility with GROUP BY implicit/explicit sorting in MySQL 5.7. - Added the forward compatibility with the max_length_for_sort_data parameter in MySQL 5.7, which optimizes the file sorting performance in specific scenarios. - Optimized the issue that accessing views in information_schema is slow due to incorrect execution plan selection. - Added the EXIST subquery in PQ. - Optimized restoration of database tables or instances to a specific point in time. ● Fixed issues <ul style="list-style-type: none"> - OpenSSL is upgraded. - The default value SYSTEM of the time_zone parameter impacts the efficiency of concurrent SQL statement execution in some scenarios. - SQL query results are incorrect when conditions are partially pushed down to a materialized derived table. - Performance suffers after PQ is enabled for on-disk hash joins in some scenarios. - The permissions page is not updated accordingly after a user is granted permissions on a database through the console and the database is later deleted in non-console mode.

2.0.45.230900

Table 2-6 Version 2.0.45.230900

Date	Description
2023-11-24	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Added forward compatibility of datetime, timestamp, and time field behaviors. - Added on-disk hash joins in PQ. - Added INSERT and REPLACE SELECT functions in PQ. - Added log printing mechanism for connection and disconnection, which helps you locate connection-related issues quickly. - Added some useful information in slow query logs, which helps you locate slow SQL statements. - Allowed you to dynamically enable binlog. - Optimized the NDP bloom filter. - Allowed you to use the CAST (... AS INT) syntax. - Optimized the Nested Loop Join + Distinct performance. - Identified slice ID corresponding to the slow I/O quickly. - Added the sal_init log, which helps you to locate storage API timeout issues. ● Fixed issues <ul style="list-style-type: none"> - There are trx_id and cpu_time fields in full SQL statements. - Character strings can be converted into INT in WHERE conditions of PREPARE statements. - No crash issue occurs when DDL operations and queries are concurrently executed on a read replica. - The binlogs that are sharply generated in a short period of time can be cleared in a timely manner. - Execution results are consistent after PQ is enabled for multi-table JOIN SQL statements. - Backward Index Scan is compatible with ICP. - weight_string functions support LEVEL clauses. - The results of the same SQL statement using different indexes are consistent. - When NDP and PQ are enabled at the same time, recycle LSN is correct.

2.0.42.230600

Table 2-7 Version 2.0.42.230600

Date	Description
2023-08-31	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Added support for storing full and incremental backups on read replicas, which reduces the memory and CPU usage of the primary node. - Optimized UNDO damage location: When the undo damage occurs during startup, the undo damage log and the corresponding table name are printed. - Improved the query performance of read replicas. - Added the conversion of IN predicates to subqueries. - Supported large-scale commercial use of the NDP feature. - Optimized execution plans using statement outline. - Supported round functions in PQ. ● Fixed issues <ul style="list-style-type: none"> - The ORDER BY LIMIT and ORDER LIMIT result sets do not overlap when fast sorting and priority queue sorting algorithms are used. - Returned results are correct for PQ statements. - No errors are reported when PREPARE statements are executed. - No PQ assertion errors are reported on UNION queries. - The results of full-text index queries are correct after a read replica is promoted to the primary while a large amount of data is being inserted into the primary node. - When read replicas use the general_log and slow_log tables, warning logs will not be displayed. - After the value of the parameter innodb_lock_wait_timeout is changed, the actual timeout wait time is correct. - When a read replica is promoted to primary, there is no the error "Failed to find page in slice manager". - The percentage for the PWAL scanning progress in the SALSQL log cannot exceed 100%. - When the sqlsmith tool is executed, there is no the error "mysqld core dump" in the EXPLAIN phase of query statements. - In SELECT DISTINCT CAST functions, datetime can be converted to the float type correctly.

2.0.39.230300

Table 2-8 Version 2.0.39.230300

Date	Description
2023-05-11	<ul style="list-style-type: none"> ● New features and optimized features <ul style="list-style-type: none"> - Supported small-scale instances. - Optimized the solution when DDL statements on standby nodes fail. - Optimized the capacity calculation of salsql. - Supported the restriction on resources of a single SQL statement. - Supported the use of per thread for admin port and local socket. - Optimized the memory of pwalScanner. - Supported the modification of default_collation_for_utf8mb4 parameter. - Supported diagnosis on large transactions. - Supported the killing of idle transactions. - Accelerated incremental restoration. - Added database and account descriptions. - Supported the acceleration of buffer pool resize. ● Fixed issues <ul style="list-style-type: none"> - Ptrc does not lead to inconsistent execution results of Nestedloop join. - No crash issue occurs when subqueries are sorted using Windows functions. - When using rewrites view, tables are not evaluated to turn left joins into inner joins. - Execution results are returned from decimal data that meets specified filter criteria. - Memory is aligned. - Scan_row is correctly recorded in full logs.

2.0.28.18

Table 2-9 Version 2.0.28.18

Date	Description
2023-05-17	Errors of exceeded sorting memory are not reported for columns containing large JSON data.

2.0.28.17

Table 2-10 Version 2.0.28.17

Date	Description
2023-04-02	Character sets are not used in combination in prepared statements.

2.0.28.16

Table 2-11 Version 2.0.28.16

Date	Description
2023-03-14	<ul style="list-style-type: none">• New features Reduced primary/standby latency.• Fixed issues<ul style="list-style-type: none">- No error occurs when JSON-related functions are used in prepare statements.- Query results are returned when filter criteria are specified.- No null pointer error is reported after Windows functions generate a temporary disk table.- The crash issue caused by the use of null pointers in Windows functions is resolved.- Prepared statements are executed successfully.

2.0.28.15

Table 2-12 Version 2.0.28.15

Date	Description
2023-01-11	<ul style="list-style-type: none"> ● New features <ul style="list-style-type: none"> - Supported SQL statement concurrency control. - Optimized read flow control. - Optimized the consistency of primary/standby execution plan. - Pre-created slices asynchronously. ● Fixed issues <ul style="list-style-type: none"> - No crash issue occurs when the system variable INNODB_VALIDATE_TABLESPACE_PATHS is disabled and the undo space truncate command is executed. - The query of information_schema.innodb_trx is fast. - The issue of inconsistent results is resolved: left joins now are turned into inner joins. - The crash issue caused by subquery optimization is resolved. - Values of the Instant field are correctly obtained under concurrent instant DDL and DML operations. - No OOM issue occurs when two InnoDB tables with FTS indexes are loaded. - No OOM issue occurs when the data dictionary of millions of tables is being updated.

2.0.28.12

Table 2-13 Version 2.0.28.12

Date	Description
2022-12-07	Scan errors triggered by Skip Scans are not displayed when a table with virtual columns is updated.

2.0.28.10

Table 2-14 Version 2.0.28.10

Date	Description
2022-11-16	During a primary/standby switchover, databases will not break down when connecting to the standby instance times out.

2.0.28.9

Table 2-15 Version 2.0.28.9

Date	Description
2022-09-23	<ul style="list-style-type: none">• The If(...) statement in Condition_pushdown::replace_columns_in_cond is modified.• The database does not break down when:<ul style="list-style-type: none">– Storage functions are invoked recursively.– Multiple tables are deleted or full-text search is performed.– SQL query statements of multiple window functions are executed.• Users with global permission can successfully run SHOW CREATE DATABASE.

2.0.28.7

Table 2-16 Version 2.0.28.7

Date	Description
2022-08-25	The ptrc crash problem in stored procedure is resolved.

2.0.28.4

Table 2-17 Version 2.0.28.4

Date	Description
2022-07-22	<ul style="list-style-type: none"> • Databases will not break down due to empty accounts. • When a temporary table used for aggregation is updated, BLOB points to the latest data.

2.0.28.1

Table 2-18 Version 2.0.28.1

Date	Description
2022-05-16	<ul style="list-style-type: none"> • New features <ul style="list-style-type: none"> - You can enable or disable orphaned definer check control. - TaurusDB supports transparent transmission of proxy IP addresses. - You can set the consistency level of your proxy instances to session consistency. • Fixed issues <ul style="list-style-type: none"> - The data dictionary on standby nodes is updated if DDL statements on the primary node are not submitted. - During a failover, the auto increment of the primary node is not rolled back. - The performance issue of standby nodes is resolved.

2.0.31.220700

Table 2-19 Version 2.0.31.220700

Date	Description
2022-08-12	<ul style="list-style-type: none"> ● New features and performance optimized <ul style="list-style-type: none"> - Supported SQL statement concurrency control. - Added a limit to concurrent numbers of Faster DDL. - Supported all Faster DDL operations in row format. - Extended full SQL fields. - Optimized flow control. - Supported the quick timeout of ALTER TABLE. - Supported the query of plan cache. - Optimized statistics on standby nodes. ● Fixed issues <ul style="list-style-type: none"> - Standby nodes do not break down after partition-table on the primary node is renamed. - The default buffer size of SQL tracer is modified. - When the truncate lsn of standby nodes lags behind, the standby nodes can start successfully. - The execution plan error is not displayed when SQL queries with the same range are executed. - The crash issue caused by empty accounts is resolved. - The crash issue caused by database dropping is resolved.

2.2 TaurusDB Version Release Policy

TaurusDB is an enterprise-grade cloud-native database fully compatible with MySQL. It decouples compute from storage and uses Huawei-developed Data Function Virtualization (DFV), which scales to up to 128 TB per instance. A failover can be complete within seconds. It provides the superior performance and high availability of a commercial database at the price of an open-source database.

This section describes the lifecycle and upgrade policy of each TaurusDB version.

Version Lifecycle

Table 2-20 Lifecycle of each TaurusDB version

TaurusDB Version Number	Status	Open-Source Compatibility	Released in Community	Adopted by Huawei Cloud	EOM	EOS
2.0.60	Commercial use	8.0.22	October 2020	January 2025	April 2025	January 2028
2.0.57	Commercial use	8.0.22	October 2020	October 2024	January 2025	October 2027
2.0.54	Commercial use	8.0.22	October 2020	July 2024	October 2024	July 2027
2.0.51	Commercial use	8.0.22	October 2020	April 2024	July 2024	April 2027
2.0.48	Commercial use	8.0.22	October 2020	January 2024	April 2024	January 2027
2.0.45	Commercial use	8.0.22	October 2020	November 2023	January 2024	November 2026
2.0.42	Commercial use	8.0.22	October 2020	September 2023	November 2023	September 2026
2.0.39	Commercial use	8.0.22	October 2020	June 2023	September 2023	June 2026
2.0.31	Commercial use	8.0.22	October 2020	September 2022	June 2023	September 2025
2.0.29	Commercial use	8.0.22	October 2020	July 2022	September 2022	July 2025
2.0.28	Commercial use	8.0.22	October 2020	June 2022	June 2023	June 2025
2.0.26	Commercial use	8.0.22	October 2020	April 2022	June 2022	April 2025
2.0.17	Commercial use	8.0.22	October 2020	February 2022	April 2022	February 2025
2.0.15	Commercial use	8.0.22	October 2020	November 2021	February 2022	November 2024
2.0.13	Commercial use	8.0.18	October 2019	September 2021	November 2021	September 2024


TaurusDB Version Number	Status	Open-Source Compatibility	Released in Community	Adopted by Huawei Cloud	EOM	EOS
2.0.12	Commercial use	8.0.18	October 2019	June 2021	September 2021	June 2024
2.0.11	Commercial use	8.0.18	October 2019	December 2020	June 2021	December 2023
2.0.8	Commercial use	8.0.18	October 2019	September 2020	December 2020	September 2023

- Adopted by Huawei Cloud: The commercial version of TaurusDB has been fully verified and is stable and reliable. You can use this version in production environments while enjoying the service-level agreement (SLA) provided by TaurusDB.
- End of Marketing (EOM): Once a TaurusDB version reaches its EOM, it is no longer possible to buy or create instances of that version.
- End of Service (EOS): Once a TaurusDB version reaches its EOS, it is no longer possible to create instances of that version. Additionally, technical support for existing instances is no longer available, including new feature updates, vulnerability or bug fixing, patch updates, service ticket guidance, and online check. The TaurusDB SLA will no longer be valid for that version.

DB Instance Versions

- DB instance major version: The format is $x.y$ (The current version TaurusDB 2.0 corresponds to MySQL 8.0 in the community.)
- Patch version: The format is $x.y.z.yy\text{mm}(n)$, where $x.y$ indicates the DB instance major version, z indicates the DB instance minor version, $yy\text{mm}$ indicates the year and month when the version is released, and (n) indicates the patch version.

Figure 2-1 DB instance versions



Upgrade Policy

To enjoy stable, reliable, and secure TaurusDB, you are advised to periodically upgrade your DB instance before EOS.

For details, see [Upgrading the Minor Kernel Version of a DB Instance](#).