

**GaussDB
24.3.0**

Kernel Release Notes

Issue 01
Date 2024-03-30



Copyright © Huawei Cloud Computing Technologies Co., Ltd. 2024. 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 3.x Versions.....	1
2 2.x Versions.....	27

1 3.x Versions

This section describes the kernel version updates of GaussDB.

Table 1-1 Versions 3.201, 3.207 and 3.208

Date	Feature	Description
2023-04	START WITH recursion performance optimized	The performance of the following statements is optimized. <ul style="list-style-type: none">• START WITH...CONNECT BY statements• START WITH...CONNECT BY statements with the connect_by_isleaf field
	Viewing running SQL statements in the current stored procedure	View SQL statements that are being executed in the current stored procedure.

Date	Feature	Description
	Session-level HA attributes	<ol style="list-style-type: none"> 1. Session-level standby node connection detection and strong synchronization of session-level logs are supported. <ol style="list-style-type: none"> a. Session-level standby node connection detection: After this function is enabled, the system checks whether the standby node can be connected. If the standby node cannot be connected, the primary node reports an error. b. Strong synchronization of session-level logs: After this function is enabled, sessions between the primary and standby nodes are strongly synchronized. 2. Single-node mode: <ol style="list-style-type: none"> a. Parameters for session-level standby node connection detection can be set, but the function does not take effect. After data is written to the primary node, a success message is returned. b. Parameters for strong synchronization of session-level logs can be set, but the function does not take effect. After data is written to the primary node, a success message is returned. 3. A full build of the standby node is supported. <ol style="list-style-type: none"> a. During a full build (the standby node cannot be connected): Parameters for session-level standby node connection detection can be set, but the function does not take effect. After data is written to the primary node, a success message is returned. After the parameters for strong synchronization of session-level logs are configured, the function takes effect. After most_available_sync is enabled and keep_sync_window is configured, new sessions for strong synchronization enter the maximum availability mode if the disconnection duration of the standby node exceeds the value of keep_sync_window. b. When the standby node is restoring the logs generated during the full build: Parameters for session-level standby node connection detection can be set, and the function can be enabled. Parameters for forced synchronization of session-level logs can be set, and the function can be enabled. 4. When the network is disconnected, the detection timeliness depends on parameters tcp_user_timeout

Date	Feature	Description
		<p>and wal_sender_timeout. When the primary and standby nodes are disconnected (including scenarios where the wal receiver process on the standby node exits normally or crashes), the connection status of the standby node is updated synchronously.</p> <p>Specifications:</p> <ol style="list-style-type: none"> 1. One primary node and one standby node are deployed. The data in the primary node is asynchronously synchronized to the standby node. 2. If the standby node cannot be connected, write request errors are reported during and the service needs to be retried.
	CURRENT OF CURSOR	<p>The CURRENT OF CURSOR syntax is supported.</p> <p>Restrictions for CURSOR statements are as follows:</p> <ol style="list-style-type: none"> 1. Only simple queries of a single table in the SELECT statement are supported. <ol style="list-style-type: none"> a. Subqueries are not supported. b. LIMIT/OFFSET is not supported. c. START WITH is not supported. d. WITH statements are not supported. 2. Multiple tables are not supported. 3. Partitioned tables are not supported. 4. It does not support row-store tables. 5. Hash bucket tables are not supported. 6. ORDER BY is not supported. <p>Restrictions for UPDATE statements are as follows:</p> <ol style="list-style-type: none"> 1. Only a single table can be updated. 2. UPDATE statements can be used only in stored procedures. 3. Multiple tables are not supported. 4. Multiple WHERE clauses are not supported. 5. WITH/USING and ORDER BY are not supported.
	Read-Only detection capability of CM disks enhanced	<p>The capability is enhanced in the following aspects:</p> <ol style="list-style-type: none"> 1. The read-only status is obtained from the database to ensure accuracy. 2. Read-only quorum arbitrates only the nodes that exceed the threshold. Other nodes are not affected. 3. When the primary node is read-only, a primary/standby switchover is automatically performed. An available standby node is promoted to primary to ensure that the instance can work properly.

Date	Feature	Description
	O&M database connection optimized	Peer authentication is not used. In addition, concurrency conflicts are resolved and O&M performance is improved.
	Unified name GaussDB	The name is changed to GaussDB.
	Automatic start of redistribution after lock timeout during scale-out	The automatic retry mechanism is provided for scale-out. If the lock times out for multiple times, the mechanism automatically exits. After a period of time, scale-out is retried to switch tables during off-peak hours, reducing the impact on user services.
	Software component lifecycle matching the product lifecycle	<p>The software component lifecycle can match the product lifecycle.</p> <ol style="list-style-type: none"> 1. Update underlying software components (such as open-source, third-party, and self-developed software). 2. Fix security vulnerabilities promptly within the product version lifecycle.
	Single-replica deployment enhanced (for non-production environments)	<ol style="list-style-type: none"> 1. Monitoring metrics are supported. 2. An alarm can be triggered when a node is set to read-only due to the fully occupied disk. <ol style="list-style-type: none"> a. O&M interface reporting adapts to the single-node mode. b. CM makes the read-only alarm information persistent on DCC.

Date	Feature	Description
	<p>Optimizer adaptation and robustness improved</p>	<p>The operator preference provides assurance if the cost model fails, and multi-version plan selection capability well as intelligent statistics are enhanced, improving the adaptability and robustness of the optimizer.</p> <ol style="list-style-type: none"> 1. In the cost convergence scenario, the operator preference mechanism is provided to prevent the optimizer from selecting a poor plan that is incorrectly estimated. In this way, the plan performance is improved by 50% when the statistics are invalid. 2. The robustness of intelligent statistics is enhanced, and statistics of multiple columns are integrated to create a statistics model, improving the evaluation performance of high-frequency values by 50%. The overall performance is improved by 20% to 30%. 3. Currently, many slow queries can be resolved only by using service hints. Multi-version plans are required to support more scenarios and provide an autonomous escape mechanism to automatically resolve more than 10% slow queries in specific scenarios. The performance deterioration in the TPC-C test model is less than 5%. This is available only to primary/standby instances.
	<p>Data restoration in extreme scenarios</p>	<p>In extreme scenarios, users can restore their data by outputting disk data files which need to be identified by the COPY command. Then, data can be imported to a new DB instance as needed.</p> <p>Generally, data restoration is based on existing redundancy technologies. This solution is used only in the following scenarios:</p> <ol style="list-style-type: none"> 1. Node recovery and primary/standby build 2. Disaster recovery 3. Backup and restoration 4. Hot patch 5. Other cases where the cluster needs to meet the expected requirements <p>Constraints:</p> <ol style="list-style-type: none"> 1. Only flushed data is parsed. 2. Data consistency cannot be ensured during offline data parsing. 3. Xlog replay cannot be guaranteed. 4. The database is offline. 5. Only heap table data in Astore and Ustore can be parsed.

Date	Feature	Description
	Undo files of Ustore managed by the segment-page mechanism	The undo files of Ustore are managed by the segment-page mechanism. This prevents file handle leakage when a large number of 1 MB undo files are generated.

Date	Feature	Description
	Global secondary indexes for distributed instances	<p>Global secondary indexes for distributed instances</p> <ol style="list-style-type: none"> Global unique indexes can be created on non-distribution columns. The performance of point queries on non-base table distribution keys is three times that of common index queries. <ol style="list-style-type: none"> Astore is supported. <p>Constraints:</p> <ul style="list-style-type: none"> Common UB-tree indexes can be created. Local indexes and expression indexes are not supported. GSIs whose distribution is the same as that of base tables cannot be created. The GSI distribution column in a base table cannot be updated. CREATE GSI CONCURRENTLY and PARALLEL are not supported. Hash-based GSIs can be created for row-store Astore tables and partitioned tables whose base tables are hash-based. Base tables cannot be replicated tables, bucketed tables, segment-page tables, list/range-based tables, Ustore tables, or column-store tables cannot be used as base tables. GSIs only support hash-based distribution. GSI can be only used in distributed instances. If _new\$\$ or _NEW\$\$ is added to a column name of a base table or an identifier (including ctid, xc_node_hash, xmin, xmax or tableoid), a GSI fails to be created for the base table whose column names are repeated. The GSI can be created only in GTM-Lite mode. IUD returning is not supported for base tables with GSIs. TABLE ACCESS BY INDEX ROWID is not supported. VACUUM FULL, COPY, GDS, and REINDEX TABLE/DATABASE are not supported. All GSIs will be skipped during execution. The following models are not supported: UPSERT, MERGE INTO, HOT UPDATE, max_datanode_for_plan (in PBE mode), logical replication, cluster, and For Update/share/nowait/skip Locked Compatibility syntax such as Start With is not supported. A maximum of 27 columns (including index keys and distribution keys) are supported.

Date	Feature	Description
		<ul style="list-style-type: none"> Other constraints are the same as those of hash distribution, UBTree, LP/FQS, and common indexes.
	CURRENT OF CURSOR	The CURRENT OF CURSOR syntax supports ORDER BY.
	SQL features supported by primary/standby Ustore instances	<p>SQL features supported by primary/standby Ustore instances</p> <ol style="list-style-type: none"> Global temporary tables Column constraints Parallel query Statistics enhancement: Ustore index active page estimation and cost adaptation.
	Out parameter overloading and capability enhancement for stored procedures	<ol style="list-style-type: none"> SQL PATCH scope: DML within a stored procedure Out function overloading scope: Functions with Out parameters can be created and invoked. Only Out parameters in direct value assignment can be returned. For example, in a = func(b), b is an out parameter and is returned.
	Overload escape in distributed scenarios	<p>Overload escape is supported in distributed scenarios.</p> <ol style="list-style-type: none"> After a certain proportion of slow SQL statements in a thread pool are executed, the CANCEL operation is triggered. Control or circuit breaker measures are taken for SQL statements that occupy a large number of resources to prevent normal SQL statements from being executed due to slow SQL statements.
	SQL PATCH in distributed instances	<p>SQL PATCH in distributed instances</p> <p>Constraints:</p> <ol style="list-style-type: none"> SQL PATCH operations cannot be synchronized between CNs. SQL PATCH cannot be used for DN. If the lightproxy plan delivered by a CN to a single node triggers the execution of SQL Patch, the plan is degraded to the FQS plan delivered by the CN to a single node, and SQL PATCH takes effect. Similar to HINT, SQL PATCH affects the capability and specifications of distributed plans. SQL PATCH cannot be used for DN. For distributed instances, if a stored procedure is pushed down, the DML in the stored procedure generates plans on DN and is not affected by SQL PATCH.

Date	Feature	Description
	Ultimate RTO when standby nodes can process read requests	Based on the ultimate RTO technology, standby nodes can process read requests. Constraints: <ol style="list-style-type: none"> 1. The DDL statement replay conflicts with the read-only mode of standby nodes. 2. When standby nodes process read requests, they consume CPU, memory, and disk space. The memory and disk space are configured based on GUC parameters. If used memory and disk space exceed the configured GUC values, the system will forcibly reclaim the overused resources.
	Syntax extension and cost model optimized	Syntax extension and cost model optimized <ol style="list-style-type: none"> 1. Auto analyze: <ol style="list-style-type: none"> a. Auto analyze is triggered when a large amount of data is inserted and updated in a short period of time to quickly process recycle task tables. b. By default, the sampling ratio and histogram precision are automatically increased based on the table size to improve the analysis accuracy of large tables. c. The adaptive estimation algorithm is provided to improve the NDV estimation accuracy in typical scenarios. 2. HINT can be used: <ol style="list-style-type: none"> a. In parameterized paths. b. In Bitmap Scan operations. c. Across query blocks. 1. The Plan Trace function can display the complete process of generating a query plan. This is available only to primary/standby instances.

Date	Feature	Description
	Index creation and rebuild online	<p>[Requirement description]</p> <ol style="list-style-type: none"> Indexes (including common indexes, global indexes, and local indexes) can be concurrently created and rebuild in Astore. Indexes (including common indexes, global indexes, and local indexes) can be concurrently created and rebuild in Ustore. <p>Specifications:</p> <ul style="list-style-type: none"> The online DML performance deteriorates by no more than 10%. Online index fields cannot be added, deleted, or modified. Only common tables and level-1 partitioned tables are supported.
	UstoreDFX and other features enhanced for primary/standby instances	<p>Primary/standby instances support the following Ustore features after related whitelist is enabled.</p> <ol style="list-style-type: none"> DFX capabilities are supplemented, and auxiliary verification for DML operations and REDO/UNDO functions is enhanced. VACUUM FULL is supported. UNDOMETA reduces the frequency and fields FSC structure is embedded and its performance is optimized.

Date	Feature	Description
	Basic operator performance improved	<p>The performance of basic operators is improved.</p> <ol style="list-style-type: none"> 1. SRF 2. Expression framework flattened 3. AGG optimized 4. INNER UNIQUE optimized 5. Expression operator optimized 6. Parser optimized 7. Printtup optimized 8. 4. Communication performance optimized 9. Index scanning operators optimized 10. Predicates optimized 11. Subquery optimized 12. Noise floor elimination for basic operators 13. Indexes optimized <p>Performance improvement of basic operators: The end-to-end TPC-H performance is improved by 20%, and the TPC-H performance is improved by 40% (including the full table scan operators, expression operators, and aggregation operators).</p>

Date	Feature	Description
	SQL M* syntax compatibility	<p>GaussDB is compatible with some M* features and:</p> <ol style="list-style-type: none"> 1. Is compatible with MySQL RANGE, HASH, and LIST partition syntax. 2. Allows users to add columns before the first column of a table or after a specified column. 3. Adapts to the MySQL syntax for modifying column names or definitions. 4. Allows users to configure common character set and collation for schemas, tables, and columns. 5. Adopts to the modified CREATE, ALTER, and DROP EVENT syntax. 6. Is compatible with SELECT INTO. 7. Uses syntax tree of SET TRANSACTION to set the isolation level and read/write mode. <p>The following M* syntax is supported:</p> <ol style="list-style-type: none"> 1. GROUP_CONCAT function is supported. 2. User-defined variables in a session can be set through SET statements. 3. SET is enhanced in global variable configuration. 4. Prefix indexes are supported. 5. Delimiters are supported. 6. DELETE statements can be used for deleting multiple tables.
	Fully-encrypted database framework with software and hardware integrated	<p>Fully-encrypted database framework with software and hardware integrated: Based on the virtual TEE, the fully-encrypted framework can encrypt memory and ensure security isolation.</p> <ol style="list-style-type: none"> 1. Secure key transmission channel is supported. 2. The fully-encrypted execution framework integrates software and hardware and supports the virtual TEE execution. 3. Server-side encryption operator support range query and sorting operations. A whitelist is used to restrict syntax operations. 4. The JDBC client driver is supported. <p>This is available only to primary/standby instances.</p>

Date	Feature	Description
	Logic decoding capability enhanced	<ol style="list-style-type: none"> 1. Logical replication supports distributed strong consistency. <ol style="list-style-type: none"> a. Sequence preserving is ensured when distributed transactions are committed. b. Sequence preserving is ensured in the same transaction on a single DN. c. Sequence preserving is not supported between different DNs of the same transaction. <p>Constraints: This feature is only suitable for GTM-Lite mode.</p> 2. Logic decoding performance views are provided. <ol style="list-style-type: none"> a. Log sending speed (Mbit/s). b. Log read speed (Mbit/s). c. Log parsing speed (Mbit/s). <p>Constraints: For parallel decoding, only the parallel decoding performance views of the primary nodes for primary/standby instances are provided.</p> 3. Distributed logical decoding is supported. DRS can directly connect to each standby DN in the service cluster to perform logical decoding that is replayed in the cluster scale-out. 4. The following two monitoring metrics are added: <ol style="list-style-type: none"> a. Maximum number of transaction lines b. Number of large transactions (with more than 4096 lines) 5. A heartbeat message is added so that external tools can accurately calculate the synchronization latency. <ol style="list-style-type: none"> a. The latest WAL timestamp is provided for external tools to calculate the latency. b. Parallel decoding provides functions to query the latest WAL timestamp. Serial and parallel decoding provide heartbeat messages. 6. The logic decoding performance can reach 300 Mbit/s, which optimizes: <ol style="list-style-type: none"> a. The performance of saving data to the database is optimized. b. Single-slot decoding performance of primary nodes c. Concurrent replay performance of standby nodes d. Single-slot decoding performance of standby nodes

Date	Feature	Description
	Partitioned table functions and performance optimized	<ol style="list-style-type: none"> 1. LIST and RANGE multi-column partitioning are supported. <ol style="list-style-type: none"> a. The number of LIST partition keys can be increased from 1 to 16. b. The number of RANGE partition keys can be increased from 4 to 16. 1. Range/List/Hash (RLH) partitions: <ol style="list-style-type: none"> a. RLH partitioned tables mean LIST, RANGE, and HASH tables. b. LIST/RANGE tables support node group distribution. 1. The partition syntax is optimized as follows: <ol style="list-style-type: none"> a. LIST and level-2 partitions support SPLIT and MERGE functions. b. LIST, HASH, and level-2 partitions support CREATE TABLE ... LIKE statements. <p>Constraints: DDL and DML concurrency across partitions is not supported in distributed instances.</p> 1. Continuous improvement of partitioned table performance: <ol style="list-style-type: none"> a. The partitioned table performance in services is improved by 30%. b. Partition implicit type and function pruning are optimized. <p>This is available only to primary/standby instances.</p>

Date	Feature	Description
	MySQL syntax compatibility improved	<ol style="list-style-type: none"> 1. There are new 167 compatibility items for system views, system functions, and advanced packages. <ol style="list-style-type: none"> a. System functions: 39 b. System views: 124 c. Advanced package: 4 2. Distributed instances support the Go driver. 3. The dblink is supported. 4. PIVOT/UNPIVOT is supported. 5. The XML data type and related interfaces are supported. 6. Distributed instances support packages. 7. Primary/standby and distributed instances support ROWNUM. 8. The following 11 XMLDOM interfaces are supported: <ol style="list-style-type: none"> a. XMLDOM.GETFIRSTCHILD b. XMLDOM.ITEM c. XMLDOM.DOMNODE d. XMLDOM.MAKEELEMENT e. XMLDOM.GETCHILDNODES f. XMLDOM.DOMNODELIST g. XMLDOM.GETNODEVALUE h. XMLDOM.GETLENGTH i. XMLDOM.GETNODENAME j. XMLDOM.GETELEMENTSBYTAGNAME k. XMLDOM.GETTAGNAME l. XMLPARSER.NEWPARSER 9. The functions of tableof are enhanced. <ol style="list-style-type: none"> a. The array can be nested with tableof. b. Distributed instances support tableof.

Date	Feature	Description
	WDR optimization and DFX support	<ol style="list-style-type: none"> 1. WDR snapshots: <ol style="list-style-type: none"> a. A statement continues to be executed after it fails to be executed due to timeout, which does not affect the overall execution of a WDR snapshot. b. A WDR snapshot fails during the vacuum process. c. The retention period of WDR snapshots is not 8 days. 2. WDR reports: <ol style="list-style-type: none"> a. The WDR report generation speed is further accelerated. b. The WDR report format is optimized. c. WDR objects are sorted based on multiple dimensions (such as dead tuples, live tuples, vacuum, analyze, scanned rows). 3. There is an interface for querying GUC parameters such as search_path configured in a session. 4. WDR reports can be read by the standby nodes. 5. L0 FULL SQL plans can be viewed.
	Optimization on memory usage of stored procedures in Global PLSQL cache	<p>The shared cache based on the stored procedure addresses the high memory usage identified by the stored procedure compatibility. The existing PL/SQL session-level cache is optimized.</p> <ol style="list-style-type: none"> 1. Global PL/SQL Cache: All sessions share PL/SQL compilation intermediate products, and their execution status is maintained independently. 2. Syntax node optimization: The data structure of syntax nodes is optimized to reduce the memory usage of a single compilation product. <p>This is available only to primary/standby instances.</p>
	ODBC Gplan reuse	ODBC supports the JDBC Statement Pool function and Gplan reuse to improve SQL performance.

Date	Feature	Description
	Go drivers for fully-encrypted databases	<p>A fully-encrypted database:</p> <ol style="list-style-type: none"> 1. Can insert 60,000 rows per second. 2. Can query 90,000 rows per second. 3. Update temporary tables. 4. Support HIS key management. 5. Support server-side encryption and decryption functions. 6. Update CMK keys. 7. Support load balancing. <p>Go language drivers are available for fully-encrypted databases.</p>
	Minimum RTO due to improved primary/standby synchronization	The restriction that a standby node can reconnect to the primary node only after the replay is complete is removed, and the RTO performance is optimized.
	User-level audit settings	<p>The audit function is enhanced as follows:</p> <ol style="list-style-type: none"> 1. User-level audit is supported. The GUC parameter full_audit_users is added to set the list of audited users. Audit logs are recorded for all SQL statements executed by users in the list. Administrators' operations can be fully audited to audit administrators' access to user data. 2. The audit operation can be configured based on client information. The GUC parameter no_audit_client is added to configure the list of clients that do not need to be audited. Audit logs are not recorded for specified client apps and IP addresses (such as <code>cm_agent@127.0.0.1</code>). 3. Audit logs are recorded for high-risk system function calls.
	Strong security authentication supported by Roach	Roach supports SSL and Kerberos security certificate authentication.

Date	Feature	Description
	Security hardening	The following security vulnerabilities are fixed: CVE-2023-30608 CVE-2023-29469 CVE-2023-28484 CVE-2021-28235 CVE-2023-27538 CVE-2023-27536 CVE-2023-27535 CVE-2023-27534 CVE-2023-0466 CVE-2023-0465 CVE-2023-0464 CVE-2023-27533 CVE-2023-23931 CVE-2023-23915 CVE-2023-23914 CVE-2023-23916 CVE-2022-4304 CVE-2022-4450 CVE-2023-0215 CVE-2023-0286 CVE-2022-42898 CVE-2022-43552 CVE-2022-43551 CVE-2022-40303 CVE-2022-40304 CVE-2022-32221 CVE-2022-42916 CVE-2022-42915 CVE-2022-1941 CVE-2022-35252 CVE-2022-37434 CVE-2022-32205 CVE-2022-32208 CVE-2022-32206 CVE-2022-2097 CVE-2022-32207 CVE-2022-2068

Date	Feature	Description
		CVE-2022-27782 CVE-2022-27781 CVE-2021-22570 CVE-2022-29824 CVE-2022-1292 CVE-2022-27775 CVE-2022-27774 CVE-2022-27776 CVE-2022-22576 CVE-2021-3520 CVE-2021-37750

Table 1-2 3.103 (whitelist)

Date	Feature	Description
23.3	WDR optimization	Sort top SQL statements by the average duration of a single SQL statement execution. Full SQL statements can be sorted by avg based on the top SQL statements in the current WDR report.
	Online specifications change	Change instance specifications online with no downtime. DNs are rebooted in rolling mode, without the need to reboot the instance. This mitigates the impact on services.
	Full SQL tracing on standby nodes and statement-level wait events	1. Support full SQL tracing on standby nodes. 2. Measure time consumption statistics on statement-level wait events for full SQL tracing. Constraints: 1. Full SQL tracing on standby nodes is available only to primary/standby instances. 2. Time consumption statistics of wait events are measured for L0 full SQL tracing. The performance loss of wait events is less than 3%.
	Domain name configuration for gsql and gsloader	Domain name configuration for gsql and gsloader ensures that multiple domain names can be accessed in HA mode. 1. Multiple domain names can be configured for distributed instances. 2. Primary/Standby instances support automatic primary selection.

Date	Feature	Description
	Automatic creation of multi-column statistics	<p>Advantages:</p> <ol style="list-style-type: none"> 1. Solve the problem of composite index cost calculation. 2. Enhance multi-column statistics. 3. Solve the problem that an incorrect index is selected in typical scenarios.
	Backups created by standby nodes	Move some backup operations from primary nodes to standby nodes to reduce the usage of resources, such as the CPU and disk I/O, of primary nodes.
	SMP supported by global temporary tables	Global temporary tables support SMP.
	Backup and restoration for a single-node instance	<p>Backup and restoration for a single-node instance is supported. This feature meets the requirements of R&D and joint commissioning scenarios. It is used only for SQL development and function joint commissioning.</p> <p>NOTICE To use this feature, submit an application by choosing Service Tickets > Create Service Ticket in the upper right corner of the management console.</p> <p>Constraints:</p> <ol style="list-style-type: none"> 1. This is available only to primary/standby instances. 2. Service production environments are not supported. 3. Availability, reliability, and DR (the RTO cannot be ensured) are not committed. 4. O&M functions are not supported, including but not limited to node replacement, scale-out, upgrade, and specification change. 5. Fault injection tests and special tests on performance, pressure, and long-term stability are not supported. 6. There is no requirement on performance or impact on service performance.

Date	Feature	Description
	SQL compatibility enhancement	<p>SQL compatibility is enhanced in the following aspects:</p> <ol style="list-style-type: none"> 1. Support the following advanced packages: <ol style="list-style-type: none"> a. UTL_FILE.FILE_TYPE b. DBMS_SQL.VARIABLE_VALUE c. UTL_FILE.FGETATTR d. DBMS_LOB.READ e. DBE_LOB f. DBE_SQL g. DBE_FILE h. DBE_RAW i. DBE_UTILITY 2. Cancel the limitation on concurrent sessions in an autonomous transaction. 3. Check the compilation integrity in stored procedures. <ol style="list-style-type: none"> a. Check the compilation integrity. b. View the dependency of compiled objects. Invalid objects can be recompiled. 4. Support the following custom types: <ol style="list-style-type: none"> a. Array b. Tableof c. TableofIndex d. Combination of the preceding three types 5. Support the following 33 system functions: TO_CHAR, SYSDATE, TO_NUMBER, UPPER, CHR, SYS_CONTEXT, TO_DATE, INSTRB, LPAD, ROUND, CAST, TRIM, SUM, ADD_MONTHS, LAST_DAY, MOD, LOWER, RTRIM, MONTHS_BETWEEN, TRUNC, RPAD, LEAST, GREATEST, REGEXP_SUBSTR, LTRIM, TRANSLATE, TREAT, EMPTY_CLOB, USER, REGEXP_REPLACE, AVG, NVL2, REGEXP_COUNT 6. Support the following system views: ALL_ARGUMENTS, ALL_IND_PARTITIONS, ALL_IND_SUBPARTITIONS, ALL_INDEXES, ALL_OBJECTS, ALL_SOURCE, ALL_TAB_COLUMNS, ALL_TAB_PARTITIONS, ALL_TABLES, DBA_ARGUMENTS, DBA_CONS_COLUMNS, DBA_CONSTRAINTS, DBA_DEPENDENCIES, DBA_DIRECTORIES, DBA_HIST_SNAPSHOT, DBA_HIST_SQL_PLAN, DBA_HIST_SQLSTAT, DBA_HIST_SQLTEXT, DBA_IND_COLUMNS, DBA_IND_EXPRESSIONS, DBA_IND_PARTITIONS, DBA_IND_SUBPARTITIONS, DBA_OBJECTS, DBA_PROCEDURES, DBA_SCHEDULER_JOBS,

Date	Feature	Description
		<p>DBA_SCHEDULER_PROGRAMS, DBA_SCHEDULER_RUNNING_JOBS, DBA_SEGMENTS, DBA_SOURCE, DBA_TAB_COLS, DBA_TAB_COLUMNS, DBA_TAB_PARTITIONS, DBA_TABLES, DBA_TRIGGERS, DBA_TYPE_ATTRS, GV\$INSTANCE, GV\$SESSION, PLAN_TABLE, USER_JOBS, USER_OBJECTS, USER_PROCEDURES, USER_TABLES, V\$INSTANCE, V\$MYSTAT, V\$SESSION, ALL_DEPENDENCES</p> <p>7. Enhance output parameters of stored procedures and functions:</p> <ol style="list-style-type: none"> a. Functions can return the record type. b. The stored procedure that contains the default value and out parameter can be invoked. <p>8. Improve the performance in stored procedures.</p> <ol style="list-style-type: none"> a. Large loop scenarios of stored procedures, for example, for i in 1 .. 100000 loop b. Autonomous transactions (10 packages, each containing 10 variables, and 100,000 calls) <p>9. Commit+exception performance optimization: Optimize performance bottlenecks in the commit+exception scenarios and optimize the resource release related to EXCEPTION</p>

Date	Feature	Description
	<p>Readable standby nodes in a distributed instance</p>	<p>Offload read requests from primary nodes to standby nodes in a distributed instance.</p> <p>Constraints:</p> <ul style="list-style-type: none"> • Weak read consistency is provided. However, only read consistency is guaranteed only for incremental snapshots within a session, not for those among sessions. • Session-level parameters are used to distinguish workloads on primary nodes from read workloads on standby nodes. If a standby node is faulty, its workloads cannot be switched to the primary node. • In abnormal scenarios such as primary/standby switchovers and standby node faults, an error will be reported for the ongoing read service on the concerned standby node, and a service retry is required. • After the resumption of a standby node or after adding replicas, an error will be reported for the ongoing read service on the concerned standby node, and a service retry is required. • If no proper standby node can be selected, an error is reported and a service retry is required. • After standby nodes process read requests, the RTO of the standby nodes is affected. If the RTO exceeds the specified threshold, traffic limiting or circuit breaker needs to be triggered after an error is reported on the service side. • In DDL playback scenarios, an API is provided for the service side to check DDL consistency on standby nodes. If DDL logs are consistent, standby nodes can process read requests. Otherwise, a circuit breaker is triggered.
	<p>Online index creation and ultimate RTO supported by primary/standby Ustore-based instances</p>	<p>The primary/standby Ustore-based instances support online index creation and ultimate RTO.</p> <ol style="list-style-type: none"> 1. Indexes can be created online without interrupting services. 2. An ultimate RTO improves the performance of primary/standby instances (The standby nodes in instances cannot process read requests).
	<p>Upsert and stored procedures supported by GPC</p>	<p>GPC supports upsert and stored procedures.</p> <ol style="list-style-type: none"> 1. In the pgxc framework, upsert can be executed through gplan. 2. Stored procedures, functions, and packages are supported.

Date	Feature	Description
	Performance improvement of basic operators	The performance of basic operators is improved. 1. The seqscan and PI operators are optimized. 2. The cost model is optimized and adjusted.
	Codegen supported by row-store expressions	The Codegen capability of row-store tables is enhanced: Codegen is enabled by default. Users can query whether Codegen is enabled the specified number of rows. By default, all queries are not split.
	Optimizer cost model optimization	The following basic capabilities are optimized: 1. cplan/gplan selection 2. coalesce selection rate estimation 3. nestloop/merge join in some scenarios 4. semi/anti join cost This is available only to primary/standby instances.
	Optimization on concurrent partitioned table query	Optimize concurrent query on partitioned tables in the following aspects: <ul style="list-style-type: none"> • DML/DDI concurrency across partitions is supported. • Static pruning and dynamic pruning are supported for expressions in partitioned tables. • Partitioned tables can be used as parameterized paths of internal tables. • MergeSort query plans are supported.
	MySQL syntax compatibility improved	Common MySQL syntax and APIs are supported. DELETE and UPDATE statements support the ORDER BY and LIMIT functions.
	Multiple IP addresses for Python drivers	Support multiple IP addresses for Python drivers, so that databases can be properly connected after a primary/standby switchover.
	PITR enhancement	Enhance PITR capabilities in the following aspects: <ul style="list-style-type: none"> • Performance deterioration caused by PITR backups is reduced for two-phase distributed transactions in distributed GTM-Lite mode. • PITR backups are not affected if the external storage device is faulty.
	Embedded C preprocessor	Embedded C preprocessor is supported.
	Overload escape in concurrency scenarios	After a certain proportion of slow SQL statements in a thread pool is executed, the CANCEL operation is triggered.

Date	Feature	Description
	Data page restoration	Data restoration is enhanced in the following aspects: <ol style="list-style-type: none"> 1. Data files on the standby node can be restored to prevent file deletion or damage. 2. Damaged pages can be automatically repaired during the creation of backup. There is no impact on the backup.
	GCM encryption and client sorting for fully-encrypted databases	Fully-Encrypted databases support GCM encryption and client sorting. <ol style="list-style-type: none"> 1. The sorting operation is implemented based on clients. 2. GCM encryption and bit transmission are supported. Constraints: Only gsql and JDBC clients are supported.
	SHA256 and SSL that uses SM series cryptographic algorithms	<ol style="list-style-type: none"> 1. The data channel between the client and server supports SSL-encrypted transmission using SM series cryptographic algorithms. Constraints: Only gsql clients are supported. 2. SHA256 encrypts and saves sensitive data in the Shanghai data base project.
	ANY permission management	Support the following 12 ANY permissions for five objects: <ul style="list-style-type: none"> • ALTER ANY TYPE, DROP ANY TYPE • ALTER ANY SEQUENCE, DROP ANY SEQUENCE, SELECT ANY SEQUENCE • ALTER ANY INDEX, DROP ANY INDEX • CREATE ANY TRIGGER, ALTER ANY TRIGGER, DROP ANY TRIGGER • CREATE ANY SYNONYM, DROP ANY SYNONYM
	Efficient data compression algorithm	With the efficient data compression algorithm, the read and write performance deterioration in the TPC-C scenario is less than 5%, and the data compression ratio of transactional databases reaches 2:1.

Date	Feature	Description
	ABO	<p>ABO supports the following service scenarios:</p> <ol style="list-style-type: none"> 1. Intelligent cardinality estimation: improves the cardinality estimation accuracy of multi-column equality queries. 2. Adaptive plan selection: provides cache multi-plan management and adaptive selection. <p>Specifications:</p> <ol style="list-style-type: none"> 1. Intelligent cardinality estimation: the ABO statistical algorithm enhanced (efficiency improved by 1x and performance improved by 50% in typical scenarios) 2. Adaptive plan selection: the efficiency improved by 1x in typical scenarios
	Display of query plans in running state	Display the execution plan specifications using views for specified slow SQL statements.
	Security hardening	<p>The following security vulnerabilities are fixed:</p> <ul style="list-style-type: none"> CVE-2022-29824 CVE-2022-27781 CVE-2022-27775 CVE-2022-35252 CVE-2021-37750 CVE-2022-32205 CVE-2022-27776 CVE-2022-32206 CVE-2022-37434 CVE-2022-22576 CVE-2022-27782 CVE-2022-27774 CVE-2022-32207 CVE-2022-2097 CVE-2022-32208 CVE-2022-32221 CVE-2022-42916 CVE-2022-42915

2 2.x Versions

This section describes the 2.x kernel version updates of GaussDB.

Table 2-1 2.7 Enterprise Edition

Date	Feature	Description
2022-11	Concurrent DML and TRUNCATE operations between different sessions of a global temporary table	Concurrent DML and TRUNCATE operations are supported between different sessions of a global temporary table.
	Snapshot backups for standby nodes	Standby nodes support snapshot backups. With typical public cloud resources, it takes less than 8 hours to fully back up and restore 12 TB of data. PITR is also supported. NOTICE Whitelist feature. To configure the whitelist permission, submit an application by choosing Service Tickets > Create Service Ticket in the upper right corner of the management console.
	Emergency measures for high-latency at the IaaS layer	If the latency increases due to underlying resources, many connections will be created on the service side. As a result, the DN thread pool and dynamic memory will be full. Emergency measures are provided to quickly rectify the fault, preventing long-time impact on services.

Date	Feature	Description
	No logic decoding for specific user operations	<p>Logical decoding logs are filtered by user. During the decoding, logical logs about database operations performed by blacklist users are not returned. This ensures that downstream data is not affected by logical replication of blacklist user operations.</p> <p>Username or user OID is required for blacklist user configuration. Logical log transactions contain user information, which provides the basis for downstream services to filter different users.</p>
	Read-only nodes in the cascaded standby server	<p>Read-only nodes provide inventory query services to offload the query load of the primary node.</p> <ol style="list-style-type: none"> 1. There is one primary node, two standbys nodes, and many read-only nodes in the HA architecture. The read-only nodes can be flexibly deleted. At least 3 read-only nodes are required. (One read-only node is connected to external services, and the other read-only nodes are for high availability.) 2. When network resources are normal, the log replication delay between the primary and read-only nodes is less than 5 seconds. 3. The query SQL statement can be executed for several hours. <p>NOTICE Whitelist feature. To configure the whitelist permission, submit an application by choosing Service Tickets > Create Service Ticket in the upper right corner of the management console.</p>
	Viewing SQL statements that are being executed in the current stored procedure	<p>The SQL statements that are being executed in the current stored procedure can be viewed, enhancing the DFX capability.</p>

Date	Feature	Description
	Security hardening	The following security vulnerabilities are fixed: CVE-2022-32221 CVE-2022-42916 CVE-2022-42915 CVE-2022-2097 CVE-2022-32205 CVE-2022-32206 CVE-2022-32207 CVE-2022-32208 CVE-2022-2068 CVE-2022-24302

Table 2-2 2.3 Enterprise Edition

Date	Feature	Description
2022-06	Logical replication capability enhanced	The logical replication capability is enhanced. Memory management is enhanced based on: <ol style="list-style-type: none">1. Parameter that controls the memory threshold of transactions. When the threshold is exceeded, logical logs are automatically flushed to disks.2. Parameter that controls the memory threshold of logical replication slots. When the threshold is exceeded, logical logs are automatically flushed to disks.3. Views or functions that display memory structure of parallel logic decoding.

Date	Feature	Description
	Ustore DFX capability enhanced	<ol style="list-style-type: none"> 1. Online verification: Important data and fields are verified during service running. 2. Offline verification: For onsite problems and data exceptions, if the processes do not exit, analyze verify can be used to obtain and output error page information. 3. Module verification: The undo, redo, upage, or ubtree module is classified into three levels: fast, skip recovery, and complete. Key verification analysis can be provided by module. 4. Fault location information enhanced: There is more detailed fault location information in logs (pg_log) and views. The uheap, undo, ubtree, and flashback modules support views. 5. Error message "Snapshot too old" resolved: The message is displayed in a new transaction due to forcible reclamation of undo space.
	SQL PATCH	<p>SQL patches are designed for database administrators, O&M personnel, and other personnel that need to optimize SQL statements. If the execution plan or execution mode of a query statement does not meet expectations, a SQL patch can adjust the query execution plan or mode without directly modifying the statement, to meet the expected effect. In addition, you can flexibly create, delete, or change a SQL patch as needed. In actual GaussDB O&M scenarios, unique SQL IDs are used to locate query statements (normal-stmt and prepared-stmt) and hints are used to optimize query performance. The scenarios for adjusting the statement execution mode are as follows:</p> <p>Scenario 1: Adding hints (such as hint-index and hint-GUC) to statements to adjust the query compilation mode of the optimizer.</p> <p>In scenario 2: Performing plain-text rewriting on query statements in special emergency cases.</p> <p>NOTICE SQL-PATCH is not supported for distributed instances.</p>

Date	Feature	Description
	Security hardening	The following security vulnerabilities are fixed: CVE-2022-31799 CVE-2022-27781 CVE-2022-27782 CVE-2022-1292 CVE-2022-29824 CVE-2022-27774 CVE-2022-27776 CVE-2022-22576 CVE-2022-27775 CVE-2021-3520

Table 2-3 2.2 Enterprise Edition

Date	Feature	Description
2022-04	Go driver supported by primary/standby instances	<ol style="list-style-type: none"> Go driver supports SSL connections to the database. Go driver supports database handles. Go driver supports Stmt. Go driver supports database transactions. Go driver supports query results (Rows and Row). Go driver supports column types.
	SMP for stored procedures	SMP parallel execution plans can be generated and executed for query statements in stored procedures.
	Ustore hardening and optimization	The Ustore adapts to the framework of the parallel index creation.
	WDR report enhanced	The ASP capability is enhanced. The WDR capability is enhanced.
	View O&M capacity enhanced	You can view disk usage details by pg_ls_waldir() for wal diagnosis and pg_ls_tmpdir() for tmp diagnosis.

Table 2-4 2.1 Enterprise Edition

Date	Feature	Description
2022-02	Ustore commercial use capability improved	<ol style="list-style-type: none"> 1. The Ustore resources can be better managed. The undo memory is optimized and zones are loaded as required. When Ustore is not used, the memory overhead and undo file space are less than 1 MB. 2. Ustore supports WDR reports which include: <ol style="list-style-type: none"> a. Number of accesses to UHeap pages (including QUERY, MODIFY, DELETE, and PRUNE operations). b. Number of accesses to undo records c. Average access length of undo chains d. Rate at which undo and translot files are flushed to disks e. Rate at which undo files are generated in a specified period. f. Rate at which undo files are reclaimed in a specified period g. Index-related information (number of layers and page visits) 3. Ustore flashback: DROP operations, TRUNCATE operations, and partitioned tables can be flashed back. The baseline for flashback query performance is provided.
	Driver and syntax compatibility enhanced	<ol style="list-style-type: none"> 1. The select for update wait n syntax is supported. 2. The conversion rule of the Decode type is compatible with O*. 3. New syntax needs to be connected to GK Smith (a pressure test tool). 4. The following regular expressions are supported: REGEXP_COUNT, REGEXP_INSTR, REGEXP_SUBSTR, and REGEXP_REPLACE. 5. execute direct on in query statements supports precompilation. 6. The criteria for updating or inserting data can be specified.
	Online database maintenance without downtime	<ol style="list-style-type: none"> 1. Indexes can be created online without downtime. 2. Adding a standby node does not restart the primary node.

Date	Feature	Description
	Row- and column-store engine capability improved	<ol style="list-style-type: none"> 1. Row-store table query execution is converted to vectorized execution, improving performance in the case of hybrid loads. Under the same resource environments and workloads, the TPC-H performance of vectorized execution is 30% higher than that of the row-store engine and row-store execution framework. 2. Column storage supports MERGE INTO subqueries.
	Logic decoding performance and functions optimized	<ol style="list-style-type: none"> 1. The logic decoding performance is optimized. The DML decoding performance reaches 100 Mbit/s (DDL decoding is not supported). 2. Logical decoding supports parsing DML operations on a specified table. In JDBC, a new parameter is added to support log parsing. (You can call the existing libpq logical walsender API to parse logs).
	Performance improvement for enterprise-class application scenarios	<ol style="list-style-type: none"> 1. When the same amount of data is updated concurrently, the average update latency and latency difference on a single DN are almost the same as those on MySQL. 2. There is no obvious periodic fluctuation within 2 hours in the TPC-C benchmark with 8,000 warehouses on an Arm-powered single-node instance with 128 vCPUs and 512 GB SAS SSD or equivalent specifications. 3. walwriter optimization: The performance is optimized under heavy load (the vCPU usage is 60% or higher).
	Enterprise-level capability optimized	<ol style="list-style-type: none"> 1. The IP address, port number, and host name of a dynamic server can be changed. 2. A client driver can be configured with multiple IP addresses of the primary and standby nodes, but is only connected to the primary node. In libpq, after an primary/standby switchover occurs, a standby node can be automatically promoted to primary. 3. The compatibility of gsql with SQLPlus is enhanced. 4. The JDBC interface supports uppercase and lowercase letters.

Date	Feature	Description
	4 GB for CLOB/BLOB	CLOBs/BLOBs can store 4 GB of data. The performance of the CLOBs/BLOBs smaller than 1 GB is not inferior to that of the previous version. Constraints: More than 1 GB of data can be read and processed only by advanced package functions. An error is reported when more than 1 GB of data is transferred by system functions. Operator and string functions do not support more than 1 GB of data. The maximum buffer size in a stored procedure is 32 KB. DISTINCT, GROUP BY, or ORDER BY operations are not allowed in LOB columns. An advanced package supports a maximum of 4 GB data. To update tables, you need to add update in lob_write .
	Deletion and update of specified partitions for primary/standby instances	Primary/standby instances support deletion and update of specified partitions (level-1 partitions for level-1 partitioned tables and level-2 partitions for level-2 partition tables). For the deletion and update of the same amount of data (including point-to-point deletion and batch deletion scenarios), the performance of partitioned tables is the same as that of non-partition tables. After a partition is specified, the global index scan cannot be selected.
	Global cache of syscache system tables	This feature reduces the high memory usage in high-concurrency and complex query scenarios. In the case of any concurrent load, the total memory usage of the GSC or LSC cannot exceed the upper limit specified by the corresponding GUC parameters (GSC: <code>global_syscache_threshold</code> ; LSC: <code>local_syscache_threshold</code>). After GSC/LSC is enabled, the loss in the standard benchmark scenario (tpcc/sysbench) does not exceed 5%.

Date	Feature	Description
	Module function and reliability hardening	<ol style="list-style-type: none"> 1. Capacity expansion hardening: <ol style="list-style-type: none"> a. Redistribution adaptive lock timeout and lock levels are optimized. b. Redistribution escape method: The kernel can trigger fast service failure. c. Redistribution supports different scenarios, and the mode can be dynamically configured. d. Abnormal residual data can be cleared during redistribution. e. Resumable data transfer is supported. 2. Data reliability <ol style="list-style-type: none"> a. The primary node restores damaged pages or files from the standby node: main fork is used for common row-store tables and init fork is used for unlogged tables. After the clog and cslog files are damaged, a function is called. Currently, the file-level does not support the segment-page format. b. The standby node repairs damaged pages or files from the primary node: Table data files can be automatically repaired during replay. Currently, the file-level does not support the segment-page format. c. If the verification page and file are damaged during the backup, the damaged page and file can be automatically recovered on the standby node by calling a function. d. Logical decoding of a specified LSN segment is supported when the meta information is not changed. 3. Ultimate RTO enabled by default: <ol style="list-style-type: none"> a. PITR adaptation b. The ultimate RTO flow control is configurable, and the replay speed views are provided. 4. DFX and quality hardening: <ol style="list-style-type: none"> a. Data consistency mechanism between pgxc_node and cn/dn-handler is hardened. 5. UUID authentication: A verification mechanism is used when the primary and standby nodes are connected. The UUID is generated based on the dbname character string for verification.
	PITR for distributed instances	Distributed instances support PITR to further: <ol style="list-style-type: none"> 1. Enhance backup and restoration capability. 2. Realize resumable recovery.

Date	Feature	Description
	Stored procedures	<ol style="list-style-type: none"> Autonomous transactions support global variables. Cursors can be closed in exceptions. When the JDBC executes stored procedures, the input and output parameters do not need to be registered.
	Permissions	<p>Fine-grained permissions and ANY permissions are supported. The following ANY permissions and syntax are supported:</p> <ul style="list-style-type: none"> CREATE ANY TABLE ALTER ANY TABLE DROP ANY TABLE SELECT ANY TABLE INSERT ANY TABLE UPDATE ANY TABLE DELETE ANY TABLE CREATE ANY FUNCTION EXECUTE ANY FUNCTION CREATE ANY PACKAGE EXECUTE ANY PACKAGE CREATE ANY TYPE
	Security	<ol style="list-style-type: none"> When an audit file is damaged, audit logs generated after the damage can be queried. An error in an SSL certificate revocation list (CRL) does not affect normal authentication. Audit logs contain transaction IDs, which are used to associate data changes with audit operations. The password of the O&M account is encrypted and flushed to disks.
	Software component lifecycle matches the product lifecycle	<p>New lower-layer software components (such as open-source software, third-party software, and self-developed platforms) are used and will be updated in a timely manner within their lifecycle to match the product lifecycle. Security vulnerabilities can be fixed in a timely manner within the lifecycle of product versions.</p>
	ODBC commercial use	<ol style="list-style-type: none"> Distributed load balancing is supported. The primary/standby switchover is supported.

Date	Feature	Description
	CM enterprise-level capability improved	<ol style="list-style-type: none"> 1. CMS records key arbitration event logs. 2. CMS depends on Huawei-developed DCC component, instead of ETCD, to select the primary node. 3. The arbitration logic is reconstructed to expand the arbitration capability.
	Performance improved	The performance of a full-table scan based on linetitem 1X is improved by 30%, and the response latency is shorter than 1s.
	O&M monitoring capability improved	<ol style="list-style-type: none"> 1. Normalized SQL processing is added to session statistics for diagnosing historical slow SQL statements. 2. The mem trace performance is enhanced to enrich memory fault locating methods. 3. Full-link trace from JDBC to the kernel is supported. 4. Kernel metric collection is optimized to support single-node collection and reporting. 5. The get_node_disk_and_log_status interface can be used in primary/standby instances.
	Paxos enterprise-level capability improved	<ol style="list-style-type: none"> 1. The DCF protocol can prevent frequent switching. 2. You can configure standby Paxos nodes are promoted to primary by priority.
	Security hardening	The following security vulnerabilities are fixed: CVE-2022-23308 CVE-2020-28473 CVE-2021-22947 CVE-2021-22946 CVE-2021-22945 CVE-2021-37750 CVE-2021-36222