

Data Warehouse Service

Product Bulletin

Issue 02
Date 2024-04-29



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.

Contents

1 Product Notice	1
1.1 GaussDB(DWS) Hot Data Storage Resource Price Adjustment on May 27th, 2024	1
1.2 Checking the GaussDB (DWS) Version	2
1.3 Version 9.0.x (cloud native 3.0)	5
1.4 Version 8.2.0	9
1.5 Version 8.1.3	45
1.5.1 New Features in Version 8.1.3	46
1.5.2 New Features and Resolved Issues in 8.1.3.x	71
1.6 Version 8.1.1	78
1.7 Version 8.1.0	98
1.8 Version 8.0.1	114
1.9 Deleting a Database Interface	124
2 Version Support Notes	127

1 Product Notice

1.1 GaussDB(DWS) Hot Data Storage Resource Price Adjustment on May 27th, 2024

Notice on Adjusting the Price of GaussDB(DWS) Hot Data Storage Resources on May 27, 2024 00:00 (Beijing Time)

Dear customer:

Huawei Cloud will adjust the prices of GaussDB(DWS) hot data storage resources in the following regions. The new prices will take effect from 2024 05/27 00:00 (Beijing time). The price details are as follows:

Table 1-1 AP-Jakarta

	Original Price					New Price				
	Hourly (USD D/G B/Hour)	Monthly (USD D/G B/Month)	1 Year (USD D/G B/Year)	2 Years (USD D/G B/2 Years)	3 Years (USD D/G B/3 Years)	Hourly (USD /GB/ Hour)	Monthly (CNY per MB)	1 Year (USD /GB/ Year)	2 Years (USD /GB/2 Years)	3 Years (USD /GB/3 Years)
DWS 3.0 hot data storage	0.00125	0.6	6	10.08	12.96	0.0007	0.336	3.36	5.6448	7.2576

Table 1-2 EU-Dublin

	Original Price					New Price				
	Hourly (USD/G B/ Hour)	Monthly (USD/G B/ Month)	1 Year (USD/G B/ Year)	2 Years (USD/G B/2 Year s)	3 Years (USD/G B/3 Year s)	Hourly (USD /GB/ Hour)	Monthly (CNY per MB)	1 Year (USD /GB/ Year)	2 Years (USD /GB/ 2 Year s)	3 Years (USD /GB/ 3 Year s)
DWS 2.0 hot data storage	0.0015	0.72	7.2	12.096	15.552	0.000755	0.362	3.62	6.09	7.83

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.

1.2 Checking the GaussDB (DWS) Version

The cluster version of GaussDB(DWS) displayed on the management console is in the format of x.y.z.p, which is explained in [Figure 2-1](#).

Generally, new functions and features of GaussDB(DWS) are released based on the release number. For example, 8.1.3 and 8.2.0 are iteration versions. New features are added to new versions. For details about the version lifecycle, see [Version Lifecycle](#).

After each iteration version is released, a patch is released to fix problems, for example, patch 8.1.3.322. A patch only fixes problems and does not add new functions or features.

Figure 1-1 GaussDB(DWS) version description




Table 1-3 GaussDB (DWS) cluster version mapping

Released On	Cluster Version	Database Kernel Version	Description		
2024-02-28	8.2.1.119 (recommended version for whitelisted users)	8.2.1	-		
2023-09-25	8.2.0.107	8.2.0	Version 8.2.0		
2023-07-17	8.2.0.106				
May 25, 2023	8.2.0.103				
April 15, 2023	8.2.0.102				
March 13, 2023	8.2.0.101				
January 14, 2023	8.2.0.100				
2022-11-28	8.2.0				
2023-11-02	8.1.3.325 (recommended version)			8.1.3	New Features in Version 8.1.3
2023-10-16	8.1.3.323				
2023-08-04	8.1.3.322				
2023-06-26	8.1.3.321				
May 19, 2023	8.1.3.320				
March 9, 2023	8.1.3.310				
2022-12-19	8.1.3.300				
2022-10-31	8.1.3.200				
2022-08-23	8.1.3.110				
2022-06-20	8.1.3.100				
2022-04-15	8.1.3				
2022-06-20	8.1.1.500	8.1.1	Version 8.1.1		
2022-04-15	8.1.1.300				
2022-03-30	8.1.1.205				
2022-03-18	8.1.1.203				
2022-02-24	8.1.1.202				
2022-01-25	8.1.1.201				

Released On	Cluster Version	Database Kernel Version	Description
2021-12-09	8.1.1.200		
2021-07-30	8.1.1.100		
2021-05-15	8.1.0.100	8.1.0	Version 8.1.0
2021-02-08	8.0.1.500	8.0.1	Version 8.0.1
2020-07-31	8.0.1.100		
2020-04-28	1.7.2	8.0.0	-
2020-03-20	1.7.1		
2019-08-08	1.5.200	R8C10	-

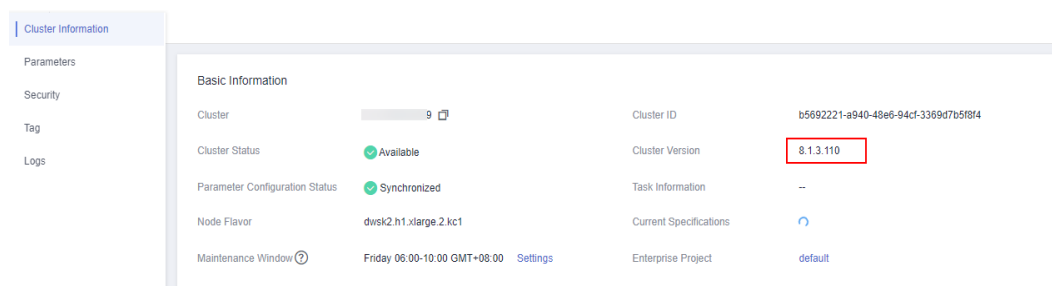
Querying the Cluster Version

Step 1 Log in to the GaussDB(DWS) console. In the navigation pane, Choose **Clusters**.



Step 2 In the cluster list, find the target cluster and click the cluster name. The **Basic Information** tab is displayed.

The cluster version is 8.1.3.110.



----End

Querying the Database Kernel Version

- Method 1: Query the cluster version by referring to [Querying the Cluster Version](#) and find the corresponding database kernel version in .
- Method 2: Connect to the DWS database. Run the following SQL statement:
select version();

```
gaussdb=> select version();
                version
-----
PostgreSQL 9.2.4 (GaussDB 8.1.1 build 9e73786c) compiled at 2021-10-29 21:14:06 commit 2597 last mr 4498 release
(1 row)
```

1.3 Version 9.0.x (cloud native 3.0)

[V9.0.3 Release Date] February 2024

[Release date of V9.0.2]: November 2023

Version 9.0.3

[Added new features]

Table 1-4 New features in 9.0.3

Scenario	Feature	Description	Reference
Data lakehouse	Readable and writable foreign tables	Readable and writable foreign tables are supported, allowing for read and write on a single foreign table.	-
	Foreign table query performance optimization	Predicate pushdown for in and not in enhances partition pruning.	-
	Complex data type support	Complex data types such as map, struct, array, and bytea/blob are supported.	-
	Parallel foreign table export	Parallel exporting of foreign tables is supported with enable_insert_ft_dop (default: disabled).	enable_insert_ft_dop
	INSERT OVERWRITE	The insert overwrite operation can be performed on either the entire table or individual partitions.	INSERT
	HiveMetaStore statistics accessibility	HiveMetaStore statistics of external schema tables can be read.	-

Scenario	Feature	Description	Reference
Decoupled storage and compute	Asynchronous read optimization	Asynchronous read is optimized.	-
	Data can be imported in parallel to tables with decoupled storage and compute.	Data can be imported to in parallel to OBS tables with decoupled storage and compute (enable_insert_dop enabled, which is disabled by default). Exporting data in parallel consumes more CPU and memory resources.	enable_insert_dop
	INSERT OVERWRITE for OBS tables with decoupled storage and compute	OBS tables with decoupled storage and compute support INSERT OVERWRITE.	INSERT
	Hot and cold local disk cache	The local disk cache supports hot and cold settings. Different cache policies can be configured for different partitions.	-
Enhanced engine	jsonb/json for column-storage	Column storage supports the jsonb/json type.	-

Version 9.0.2

[Added new features]

Table 1-5 New features in version 9.0.2

Scenario	Feature	Description	Reference
Data lakehouse	One-click lakehouse	create external schema is used to directly connect to HiveMetaStore metadata, avoiding complex create foreign table operations.	Cross-Cluster HiveMetaStore Access

Scenario	Feature	Description	Reference
	Parallel foreign table ANALYZE	This function is controlled by the enable_parallel_analyze parameter. By default, this function is enabled, which significantly improves the speed of collecting ANALYZE statistics. Note: Foreign tables do not have the AUTO ANALYZE capability. You are advised to manually perform ANALYZE to collect statistics after data import.	enable_parallel_analyze
	parquet/orc read and write for foreign tables	Overwriting, appending, and multi-level partition read/write are supported.	CREATE FOREIGN TABLE (SQL on OBS or Hadoop)
	Foreign table read acceleration by disk cache	Foreign tables can use the local disk cache for read acceleration.	-
	Foreign table data masking and row-level access control.	Foreign tables support datamasking and row-level access control.	-
Decoupled storage and compute	Decoupled storage and compute	You can specify COLVERSION=3.0 to create a column-store table as an OBS table with decoupled storage and compute. In this way, user data is stored in OBS and is charged on demand.	CREATE TABLE
	Asynchronous read and write for OBS tables with decoupled storage	OBS tables with decoupled storage supports asynchronous read and write.	-

Scenario	Feature	Description	Reference
	Parallel ANALYZE for OBS tables with decoupled storage	OBS tables with decoupled storage support parallel ANALYZE, which is controlled by the GUC parameter enable_parallel_analyze . This function is enabled by default.	enable_parallel_analyze
	Backup and restoration	OBS tables with separated storage and compute support only full backup and restoration.	-
	Elastic VW	Elastic Virtual Warehouse (VW) supports read/write acceleration. It takes 12 minutes to start an elastic VW, which is irrelevant to the current node scale, number of objects, and data volume. Integrating additional elastic VWs can bolster the system's read/write throughput, concurrent processing, and resource isolation.	Periodically Adding or Deleting DWS 3.0 Logical Clusters
	Elastic scale-out and ordinary scale-out.	Elastic scale-out does not redistribute data on OBS. Classic scaling redistributes all data. The system automatically determines the scaling mode based on whether the value of table_buckets divided by the total number of DNs exceeds the [min_table_buckets_per_dn, max_table_buckets_per_dn] range.	Impact on the System

Scenario	Feature	Description	Reference
	OBS monitoring view	Added the OBS monitoring view PGXC_OBS_IO_SCHEDULER_PERIODIC_STATS .	PGXC_OBS_IO_SCHEDULER_PERIODIC_STATS
	Disk cache view	Added the disk cache views PGXC_DISK_CACHE_STATS and PGXC_DISK_CACHE_ALL_STATS .	PGXC_DISK_CACHE_STATS PGXC_DISK_CACHE_ALL_STATS
	Clearing residual files	Added functions pgxc_scan_residualfiles and pgxc_rm_scan_residualfiles_archive to detect and clear residual files.	pgxc_scan_residualfiles() pgxc_rm_scan_residualfiles_archive()
Performance optimization	Memory sorting performance optimization	Optimized the memory sorting performance.	-

[Behavior Change]

1. **VACUUM FULL/ANALYZE/CLUSTER** are not supported for the entire database. No error is reported for the syntax, but the syntax is not executed actually. They are supported only for a single table.
2. OBS tables with decoupled storage and compute do not support delta tables. If **enable_delta** is set to **on**, no error is reported, but delta tables do not take effect.
3. Only scale-out is supported. During scale-out, tables that are being scaled out can only be queried.
4. The default sampling ratio is adjusted to one ten thousandth by setting **default_statistics_target** to **-0.01**.

1.4 Version 8.2.0

[Release date of V8.2.0.107] September 25, 2023

[Release date of V8.2.0.106] July 17, 2023

[Release date of V8.2.0.103] May 25, 2023

[Release date of V8.2.0.102] April 15, 2023

[Release date of V8.2.0.101] March 13, 2023

[Release date of V8.2.0.100] January 14, 2023

[Release date of V8.2.0]: November 28, 2022

New Feature

Table 1-6 New Functions in 8.2.0

Feature	Description	Reference
The pgcrypto encryption extension plug-in is added.	You can use pgcrypto to encrypt data.	Using pgcrypto to Encrypt Data
The object_name column in the audit log records all read and write objects in the SQL statement.	<ul style="list-style-type: none">The description about querying and auditing the records of multiple object names is added.Added the GUC parameter audit_object_name_format.	<ul style="list-style-type: none">Viewing Database Audit LogsAudit Switch
Certain user operations can be allowed in security mode.	The GUC parameter security_enable_options is added to control whether the grant_to_public , grant_with_grant_option , and foreign_table_options functions can be used in security mode.	Operation Audit
Session IDs can be recorded in audit logs.	<ul style="list-style-type: none">The configuration item normal_session_id is added to determine whether to generate session IDs in normal format.Modified the session_id field in pg_query_audit().	<ul style="list-style-type: none">behavior_compat_optionsSecurity Functions
MIXED AGG	The GUC parameter enable_mixedagg is added to control how the optimizer uses the Mixed Agg type.	Optimizer Method Configuration

Feature	Description	Reference
<p>The circuit breaker mechanism based on exception rules is optimized to prevent a slow SQL statement from affecting the entire cluster or resource pool.</p>	<ul style="list-style-type: none"> • Added the system catalog GS_BLOCKLIST_QUERY and system view GS_BLOCKLIST_QUERY for querying job blocklist information and exception information. • The query_exception_count_limit parameter is added to specify the maximum number of times that an exception rule can be triggered for a job. If this upper limit is exceeded, the job will be automatically added to the blocklist and cannot be executed. The job can be resumed only after it is removed from the blocklist. • Functions are added to record job exception information. <ul style="list-style-type: none"> <code>gs_increase_except_num(unique_sql_id int8)</code> <code>gs_increase_except_num(unique_sql_id int8, except_num int4)</code> <code>gs_increase_except_num(unique_sql_id int8, except_num int4, except_time int8)</code> <code>gs_update_blocklist_hash_info(unique_sql_id int8, is_remove boolean)</code> <code>gs_update_blocklist_hash_info()</code> <code>gs_append_blocklist(unique_sql_id int8)</code> 	<ul style="list-style-type: none"> • System View > GS_BLOCKLIST_QUERY • System View > GS_BLOCKLIST_QUERY • Load Management • Resource Management Functions

Feature	Description	Reference
	gs_remove_blocklist(unique_sql_id int8) gs_wlm_rebuild_exception_rule_hash()	

Feature	Description	Reference
<p>Full utilization of memory resources</p>	<ul style="list-style-type: none"> • The memory negative feedback mechanism is added to the PGXC_RESPOOL_RESOURCE_INFO view. • Changed the type of max_process_memory from postmaster to sighup. If a single DN is deployed on a server, max_process_memory = (Physical memory - vm.min_free_kbytes) x 0.6 Added the max_process_memory_auto_adjust parameter to control whether to enable the automatic adjustment function of the max_process_memory parameter. • Added the enable_wlm_internal_memory_limit parameter to specify whether to enable the built-in limit on estimated statement memory usage in load management. Added the enable_strict_memory_expansion parameter to determine whether to strictly control the increase of statement memory usage. Added the allow_zero_estimate_memory parameter to control whether the estimated statement memory usage can be 0. 	<ul style="list-style-type: none"> • PGXC_RESPOOL_RESOURCE_INFO • Memory • Resource Management

Feature	Description	Reference
<p>The exception rule interface is implemented using the kernel syntax.</p>	<ul style="list-style-type: none"> • The PG_EXCEPT_RULE system catalog is added to store information about exception rules. • Added syntaxes ALTER EXCEPT RULE, CREATE EXCEPT RULE, and DROP EXCEPT RULE. 	<ul style="list-style-type: none"> • PG_EXCEPT_RULESQL • ALTER EXCEPT RULE/CREATE EXCEPT RULE/DROP EXCEPT RULE
<p>Reconstructed resource management</p>	<ul style="list-style-type: none"> • The system catalog GS_WLM_USER_RESOURCE_HISTORY contains data on both CNs and DN. • The PGXC_TOTAL_USER_RESOURCE_INFO view is added to display real-time resource consumption information of users on all instances. The PGXC_WLM_USER_RESOURCE_HISTORY view is added to display historical information about resource consumption of all users on the corresponding instances. 	<ul style="list-style-type: none"> • GS_WLM_USER_RESOURCE_HISTORY • PGXC_TOTAL_USER_RESOURCE_INFO/PGXC_WLM_USER_RESOURCE_HISTORY
<p>A parameter is added in the optimizer phase. An error will be reported for the SQL statement run with too many stream threads.</p>	<p>The GUC parameter max_streams_per_query is added to control the number of stream nodes in a query plan.</p>	<p>Other Optimizer Options</p>
<p>If a GaussDB(DWS) cluster becomes read-only, you can perform TRUNCATE and DROP to quickly free up disk space and automatically cancel the read-only state.</p>	<p>In 8.2.0 and later versions, you can free up disk space by using DROP or TRUNCATE TABLE in a read-only cluster.</p>	<p>Removing the Read-only Status</p>

Feature	Description	Reference
<p>Hints can take effect in subqueries.</p>	<ul style="list-style-type: none">• Hint parameters can be used in DML statements, including INSERT, UPDATE, MERGE, and DELETE.• Hint enhancement in multi-table scenarios: Hints are enhanced for scenarios where the number of items in the FROM list exceeds the threshold, preventing hint failures and improving performance.• The enable_from_collapse_hint parameter is added to control whether to preferentially rewrite the FROM list with hints in effect.	<ul style="list-style-type: none">• Hint That Disables Subquery Pull-up• Other Optimizer Options

Feature	Description	Reference
Dynamic sampling	<ul style="list-style-type: none"> • The autoanalyze_mode parameter is added to control whether autoanalyze or autovacuum is enabled. • The autoanalyze_cache_number parameter is added to set the maximum number of tables whose statistics can be cached by the lightweight autoanalyze. • The comparison between common autoanalyze and lightweight autoanalyze is added. • The PV_RUNTIME_RELSTATS view is added to display table-level statistics generated by autoanalyze in the memory. • The PV_RUNTIME_ATTSTATS view is added to display column-level statistics generated by autoanalyze in the memory. • The pg_stat_get_runtime_relstats function is added to return table-level statistics generated by lightweight autoanalyze in the memory. • The pg_stat_get_runtime_attstats function is added to return 	<ul style="list-style-type: none"> • Automatic Cleanup • Analyzing a Table • PV_RUNTIME_RELSTATS • PV_RUNTIME_ATTSTATS • Statistics Information Functions

Feature	Description	Reference
	<p>column-level statistics generated by lightweight autoanalyze in the memory.</p> <ul style="list-style-type: none"> The pg_stat_set_last_data_changed_num function is added to set the number of historical modifications on the current node in the table. 	
<p>The on condition in the MERGE INTO statement supports system columns.</p>	<p>Added the syntax for modifying the status flag of an index.</p>	<p>ALTER TABLE</p>
<p>Sequence pushdown</p>	<p>Added the volatile_shipping_version parameter to control the pushdown scope of the volatile functions.</p>	<p>Optimizer Method Configuration</p>

Feature	Description	Reference
<p>Volatile temporary tables are supported.</p>	<ul style="list-style-type: none"> • The get_volatile_pg_class function is added to obtain the basic information about pg_class corresponding to the volatile temporary table, including table name parameters and the table list. • The get_volatile_pg_attribute function is added to obtain the basic information about pg_attribute corresponding to the volatile temporary table, including column parameters and the column list. • The VOLATILE keyword is added to CREATE TABLE and CREATE TABLE AS in the DDL syntax; and to SELECT INTO in the DQL syntax. • The DISCARD syntax is added to release internal resources related to database sessions in the current session. • VACUUM FULL does not support operations on volatile temporary tables. • The max_volatile_memory parameter is added to specify the maximum total memory occupied by contexts related to volatile temporary tables in all sessions. 	<ul style="list-style-type: none"> • Other Functions • CREATE TABLE • CREATE TABLE AS • SELECT INTO • DISCARD • VACUUM • Memory • Automatic Retry upon SQL Statement Errors

Feature	Description	Reference
	<ul style="list-style-type: none">• The max_volatile_tables parameter is added to specify the maximum number of volatile temporary tables created for each session.• The constraints on the volatile temporary table are added.	

Feature	Description	Reference
Enhanced MySQL compatibility	<ul style="list-style-type: none"> • The TRY_CAST keyword is added. • The CONV function is added to convert the given value or string into the value of a specific number system, and to output the result as a string. • The HEX function is added to return a specified value or a hexadecimal string. • The UNHEX function is added to perform the reverse operation of HEX(n). • The SPACE function is added to return a string consisting of a specified number of spaces. • The STRCMP function is added to compare the sizes of two strings. • The BIN function is added to convert the bigint type from decimal to binary, and to return the result as a string. • The substring_index function is added to search for delimiters in a case-sensitive manner, and to return the substring before the delimiter that appears for the <i>count</i>th time. • The rand function is added to return a random number in the range 0.0 to 1.0. • The truncate function is added to truncate a 	<ul style="list-style-type: none"> • Keyword • Character Processing Functions and Operators • Binary String Functions and Operators • Mathematical Functions and Operators • Date and Time Processing Functions and Operators • Type Conversion Functions • UUID Functions • Array Functions and Operators • CREATE TABLE • ALTER TABLE • PG_ATTRDEF • Syntax Compatibility Differences Among Oracle, Teradata, and MySQL

Feature	Description	Reference
	<p>number to specified decimal places.</p> <ul style="list-style-type: none"> ● The addtime function is added to return the result of a given datetime plus a time interval of a specified unit. ● The subtime function is added to return the result of a given datetime subtracted by a time interval of a specified unit. ● The timediff function is added to subtract one date parameter from another. ● The curdate function is added to return the current date. ● The curtime function is added to return the current time. ● The convert_tz function is added to convert the given date and time to those in the specified time zone. ● The try_cast function is added to convert a parameter to the value of a specified type. ● The cast(x, y) function is added to convert data types. ● The uuid function is added to generate a sequence number of the UUID type. ● The interval function is added to search for the last array index that is less than or equal to the target 	

Feature	Description	Reference
	<p>parameter n from the input integer array.</p> <ul style="list-style-type: none"> • The split function is added to separate strings by delimiter and return an array. • The ON UPDATE on_update_expr syntax is added to the CREATE TABLE syntax. • The MODIFY [COLUMN] ON UPDATE on_update_expr syntax is added to the ALTER TABLE syntax. • The adbin_on_update and adsrc_on_update columns are added to the PG_ATTRDEF system catalog. • Improved MySQL compatibility with the CURDATE, CURTIME(p), and backquote syntaxes. • In the ALTER FUNCTION, ALTER INDEX, ALTER TABLE and ALTER VIEW syntaxes, the new name can be prefixed with the schema name of the original table. The schema name cannot be changed at the same time. 	

Feature	Description	Reference
Value-based redistribution	<ul style="list-style-type: none"> • New GUC hint parameters are added. The enable_stream_ctescan parameter cannot be set at the subquery level, and the enable_value_redistribute parameter can be set at the subquery level. • The enable_value_redistribute parameter is added to specify whether to generate value-based redistribution plans. 	<ul style="list-style-type: none"> • Configuration Parameter Hints • Optimizer Method Configuration
Magic set enhancement	<ul style="list-style-type: none"> • The eager_magicset value is added to the rewrite_rule parameter. You can use eager_magicset to query rewriting rules (to push conditions from the main query to subqueries). • The windowagg_pushdown_enhancement parameter is added to specify whether to enable enhanced predicate pushdown for the window functions in aggregation scenarios. 	<ul style="list-style-type: none"> • Developer Options • Other Optimizer Options
SetOp supports two-level computing.	The setop_optmode parameter is added to control whether to perform deduplication on the query branch statements of a set operation (UNION/EXCEPT/INTERSECT) without the ALL option.	Optimizer Method Configuration

Feature	Description	Reference
<p>The performance of OBS foreign table import is optimized.</p>	<p>The file_split_threshold parameter is added, which is valid only for read only foreign tables in TEXT format. It is applicable to the import scenario where the number of files is smaller than the number of DNs.</p>	<p>CREATE FOREIGN TABLE (for OBS Import and Export)</p>
<p>GaussDB(DWS) can connect to OBS through an agency.</p>	<p>The security_token parameter is added to form a temporary security credential together with the temporary AK and SK.</p>	<ul style="list-style-type: none"> • CREATE SERVER • CREATE FOREIGN TABLE (for OBS Import and Export)
<p>The AK and SK can be specified using server when OBS foreign tables are exported in text format.</p>	<ul style="list-style-type: none"> • A description is added for the fdw_name parameter of FOREIGN DATA WRAPPER. GDS foreign table uses gsmpp_server and fdw_name is dist_fdw. • A constraint is added for the access_key, secret_access_key, and security_token parameters. FOREIGN DATA WRAPPER is set to dist_fdw. • The gsmpp_server parameter is changed to server_name. You can use the gsmpp_server created by the initial database by default, or use a custom server. 	<ul style="list-style-type: none"> • CREATE SERVER • CREATE FOREIGN TABLE (for OBS Import and Export)
<p>OBS read and write performance monitoring</p>	<p>The GS_OBS_LATENCY view is added to record the average latency of OBS.</p>	<p>GS_OBS_LATENCY</p>

Feature	Description	Reference
OBS read/write bandwidth views	The read/write bandwidth views GS_OBS_READ_TRAFFIC and GS_OBS_WRITE_TRAFFIC are added. Data is aggregated at an interval of 10 minutes. The logtime column is added to display time groups.	<ul style="list-style-type: none"> • GS_OBS_READ_TRAFFIC • GS_OBS_WRITE_TRAFFIC
Column storage supports ring buffer.	The GUC parameter enable_cstore_ring_buffer is added to control the column-store ring buffer.	enable_cstore_ring_buffer
Adaptation to GaussDB(DWS) hstore column-store vacuum	System functions are added for column-store vacuum. get_col_file_info(table_name) get_all_col_file_info() col_rebuild_file_relation(table_name)	Other Functions
Compressed text can be imported to OBS foreign tables dfs_fdw .	The compression parameter is added to specify the file compression mode for import and export.	CREATE FOREIGN TABLE (SQL on OBS or Hadoop)

Feature	Description	Reference
<p>The import with indexes is optimized to avoid generating too many Xlogs.</p>	<ul style="list-style-type: none"> • The PGXC_STAT_WAL view is added to show the WAL logs and data page traffic information of the current query. • The GUC parameter enable_wal_decelerate is added to set the WAL log rate limit. • The GUC parameter wal_decelerate_policy is added to control the action after rate limiting is triggered. • The GUC parameter wal_write_speed is added to specify the maximum WAL write speed allowed by each query on a single DN per second. • The GUC parameter wal_decelerate_trigger_threshold is added to specify the threshold for a query to trigger WAL write rate limiting on a single DN. • Two functions are added to the SQL syntax. • The pg_stat_wal_write() function is added to record the thread information of the current instance and collect statistics on the import volume and rate of WAL logs and data pages. • The pgxc_stat_wal_write() function is added to record information about threads 	<ul style="list-style-type: none"> • PGXC_STAT_WAL • Settings • Statistics Information Functions

Feature	Description	Reference
	interacting with the CN on each DN and collect statistics on the import volume and rate of WAL logs and data pages.	
The dist fdw foreign table in the LATIN1 database supports the dataencoding parameter.	The dataencoding parameter is added to convert between dataencoding and encoding during GDS import and export.	<ul style="list-style-type: none"> • CREATE FOREIGN TABLE (for CGS Import and Export) • CREATE FOREIGN TABLE (SQL on other GaussDB(DWS))
Bandwidth compression for GDS communication	The gds_compress is added to reduce the bandwidth usage for interconnection between clusters when users use GDS pipe files.	<ul style="list-style-type: none"> • CREATE FOREIGN TABLE (for CGS Import and Export) • CREATE FOREIGN TABLE (SQL on other GaussDB(DWS))
DDL lock timeout configuration	The GUC parameter ddl_lock_timeout is added to specify the DDL lock timeout interval.	ddl_lock_timeout
PG_LOCKS view enhancement	<ul style="list-style-type: none"> • Added the views PGXC_WAIT_DETAIL and PGXC_LOCKWAIT_DETAIL. • Columns including wait_on_pid and query_id are added. 	PGXC_WAIT_DETAIL PGXC_LOCKWAIT_DETAIL
The original pushdown plan is modified.	The implied_quality_optmode parameter is added to specify how to pass conditions for the equivalent columns in a statement.	implied_quality_optmode

Feature	Description	Reference
ANALYZE sampling adaptation (added in 8.2.0.100)	<ul style="list-style-type: none"> Added GUC parameters analyze_stats_mode and analyze_sample_mode. The default value of random_function_version is changed to 1. 	<ul style="list-style-type: none"> Automatic Cleanup Other Optimizer Options
Cost estimation enhancement (added in 8.2.0.100)	The default value of enable_extrapolation_stats is changed to on .	Other Optimizer Options
Column-store small CU view (added in 8.2.0.100)	Add the CU information function for column-store tables. <ul style="list-style-type: none"> <code>get_col_cu_info</code> <code>get_col_file_vacuum_info</code> <code>get_all_col_cu_info</code> 	Other Functions
Optimized dirty page monitoring view (added in 8.2.0.100)	Added view: PGXC_STAT_TABLE_DIRTY Added functions: pg_stat_get_tuple() and pgxc_stat_single_table	<ul style="list-style-type: none"> PGXC_STAT_TABLE_DIRTY Statistics Information Functions
SELECT does not block TRUNCATE (add in 8.2.0.100)	Added the GUC parameter ddl_select_concurrent_mode to control the concurrency of DDL and SELECT statements.	Lock Management
Support for HStore tables (added in 8.2.0.100)	Added the enable_hstore parameter to the CREATE TABLE syntax in the standalone hybrid data warehouse scenario.	<ul style="list-style-type: none"> Hybrid Data Warehouse CREATE TABLE

Feature	Description	Reference
<p>Fine-grained DR (added in 8.2.0.100)</p>	<ul style="list-style-type: none"> Added the function of backing up DR system catalogs and system views. Added the enable_disaster_restore and fine_disaster_table_role parameters to the CREATE TABLE syntax. Added the local coordinator only parameter to the LOCK syntax. Added the release and subscription SQL syntax. 	<ul style="list-style-type: none"> PG_FINE_DR_INFO/PG_STAT_LAST_OPERATION/PG_PUBLICATION/PG_PUBLICATION_RELEASE/PG_PUBLICATION_NAMESPACE PG_PUBLICATION_TABLES CREATE TABLE LOCK ALTER PUBLICATION/CREATE PUBLICATION/DROP PUBLICATION/ALTER SUBSCRIPTION/CREATE SUBSCRIPTION/DROP SUBSCRIPTION
<p>Display none of the non-aggregation function query columns after GROUP BY (added in 8.2.0.101)</p>	<p>The disable_full_group_by_mysql option is added to the GUC parameter behavior_compat_options. Users can set this parameter to display none of the non-aggregation function query columns after GROUP BY in a query.</p>	<p>Developer Guide > GUC Parameters > Miscellaneous Parameters > behavior_compat_options</p>
<p>Optimization of the returned number of limits on each DN in the distinct limit scenario (added in 8.2.0.101)</p>	<p>The enable_agg_limit_opt parameter is added to specify whether to optimize select distinct col from table limit N.</p>	<p>Developer Guide > GUC Parameters > Query Planning > Other Optimizer Options</p>
<p>The query performance of a single table is improved, and the restriction on inlist roughcheck is removed. (Added in 8.2.0.101)</p>	<p>The inlist_rough_check_threshold parameter is added to control the maximum number of values in the IN condition.</p>	<p>Developer Guide > GUC Parameters > Query Planning > Other Optimizer Options</p>

Feature	Description	Reference
Optimized hstore hot and cold table features (added in 8.2.0.101)	Hot and cold data can be managed in HStore tables.	Developer Guide > Hybrid Data Warehouse > Introduction to Hybrid Data Warehouse
Optimized DN Parallel Import in the PBE Scenario (added in 8.2.0.102)	Added the GUC parameter enable_parallel_batch_insert to control whether to enable concurrent data import to row-store and column-store tables.	Parallel Import

Behavior changes

Table 1-7 Behavior changes

Change Type	No.	Change	Description
Added	1	Added WindowAgg support for value partition execution plans	The window function is executed at the outermost layer of statements. If ORDER BY is not used, the data behaviors are in a default order. In version 8.2.0, after enable_value_redistribute is enabled by default, the data becomes unordered. You need to use the ORDER BY clause to specify the order.
Modified	2	Changed the exception rule configuration mode from tool configuration to syntax configuration.	In earlier versions, exception rules are configured for Cgroups using <code>gs_cgroup</code> . In 8.2.0, exception rules are configured using the CREATE/ALTER/DROP EXCEPT RULE syntax.
	3	Replaced the exception rule qualificationtime with elapsedtime .	In earlier versions, qualificationtime is used as a rule, but it can only be used together with cpuskepercent . Their effect is the same as that of elapsedtime . Therefore, this configuration item can be replaced with elapsedtime .

Change Type	No.	Change	Description
	4	Tid column name	Row storage supports the tid column names, which are not supported before.
	5	max_process_memory	Changed the number of CNs to half of that of DN.

SQL Syntax

Table 1-8 SQL syntax changes

Change Type	No.	Feature	Change Description
Added	1	Table parameter analyze_mode	Added the parameter analyze_mode , which specifies the ANALYZE modes supported by a table. If this parameter is not set, the previous mode remains unchanged.
	2	Hints for disabling subqueries	Added the hints for disabling subqueries. Format: no_merge (current query) and no_merge(releid) (specified object).
	3	DML hints	Added support for hints in INSERT, UPDATE, MERGE, and DELETE statements.
	4	AGG redistribution hints support specified column names.	Column names can be specified for AGG redistribution hints. The format is /+ redistribute ((*) (a b c d)) */ .
	5	Count(distinct) hints	Added hints for specifying the count(distinct) column, effective only in the double Hash Agg scenario.
	6	invisible	Added index OPTIONS to specify whether the index scan is visible in the execution plan.

Change Type	No.	Feature	Change Description
	7	Backquotes	Backquotes can be used to distinguish common characters from special characters.
	8	ON UPDATE in column definition	When the ON UPDATE attribute is specified in the column definition, the update timestamp is automatically filled when data is updated.
	9	CREATE/ALTER/DROP EXCEPT RULE	Added the syntax to add, create, and delete kernel exception rules.
	10	OPTIONS security_token in CREATE SERVER statements	Added OPTIONS security_token to transfer security tokens for using temporary AKs and SKs to access OBS.
	11	OPTIONS security_token in CREATE FOREIGN TABLE statements	Added OPTIONS security_token to transfer security tokens for using temporary AKs and SKs to access OBS.
	12	file_split_threshold in (dist_fdw) option in OBS foreign tables	Added the file_split_threshold to the OBS foreign table parameter (dist_fdw) option .
	13	gds_compress in GDS foreign tables	Added the gds_compress parameter, which is used for compressed transmission of data between DNs and GDSs during GDS foreign table interconnection. Currently, only the Snappy compression algorithm is supported.
	14	dataencoding in GDS foreign tables	Added the dataencoding parameter, which specifies the character set of the data stored in the latin1 database.

Change Type	No.	Feature	Change Description
	15	gds_compress in gc_fdw foreign tables (sql on other gaussdb foreign tables)	Added the gds_compress parameter, which is used for compressed transmission of data between DNs and GDSs during GDS foreign table interconnection. Currently, only the Snappy compression algorithm is supported.
	16	dataencoding in gc_fdw foreign tables (sql on other gaussdb).	Added the dataencoding parameter, which specifies the character set of the data stored in the latin1 database.
Modified	17	CREATE/ALTER REDACTION POLICY	The default value of deltarow_threshold is set to 10000 , which specifies the upper limit of rows when column-store tables are imported to the Delta table. This parameter is valid only when the table-level parameter enable_delta is enabled.
	18	Syntax of association between resource pools exception rules	Added the INHERIT option, which specifies whether a policy is inherited from other policies.
	19	single_node support for OBS foreign table syntax	In earlier versions, memory exception rules can be configured for resource pools. The exception rules in 8.2.0 are configured via the unified interface: CREATE/ALTER RESOURCE POOL 'pool' with (EXCEPT_RULE='rule1,2');
	20	Use server to specify the AK and SK of an OBS foreign table (dist_fdw).	Single_node supports OBS foreign tables (dfs_fdw).

Keywords

Table 1-9 Keywords

Change Type	No.	Keyword	Change Description
Added	1	TRY_CAST	Added the non-reserved keyword TRY_CAST for the compatibility with the TRY_CAST function in MySQL. If the type conversion fails, NULL is returned.

System Catalogs

Table 1-10 System catalogs

Change Type	No.	Name	Change Description
Added	1	gs_blocklist_query	Added the system catalog gs_blocklist_query , which is used to store blocklist statements and exception rule triggering information.
	2	pg_except_rule	Added the system catalog pg_except_rule , which is used to store information about exception rules created in the cluster, including the rule name, rule threshold, and action triggered by the rule.
Modified	3	pg_redaction_policy, pg_redaction_column	Added the inherited column, which specifies whether a policy is manually created or inherited.

System Functions

Table 1-11 System functions

Change Type	No.	Function	Description
Added	1	pg_stat_get_runtime_relstats	Queries table-level statistics generated by AUTOANALYZE.

Change Type	No.	Function	Description
	2	pg_stat_get_runtime_attstats	Queries column-level statistics generated by AUTOANALYZE.
	3	pg_stat_set_last_data_changed_num	Interface for manually setting the total size of historical modification for the remedy of the pgstat content loss.
	4	strcmp()	Compares two strings. If the first string equals the second string, 0 is returned. If the first string is less than the second string, -1 is returned. If the second string is less than the first string, 1 is returned. If the input parameters contain null, null is returned.
	5	hex()	Converts each character or decimal integer in the input parameter into a hexadecimal number.
	6	unhex()	Performs the reverse operation of hex() . This function interprets each pair of hexadecimal digits (in the argument) as a number and converts it into a character. The result character is returned in bytea format.
	7	space(n)	Returns a string consisting of n spaces.
	8	rand()	Obtains a random number ranging from 0.0 to 1.0.
	9	truncate(x,d)	Truncates x to d decimal places. If d is a negative number, then the number is truncated to the left side of the decimal point.
	10	addtime(expr1,expr2)	Adds time. expr1 is of the time or timestamp type, and expr2 is of the interval type.

Change Type	No.	Function	Description
	11	subtime(expr1,expr2)	Subtracts a time period expr2 from the time expr1 . expr1 is of the time or timestamp type. expr2 is of the interval type. The return type is the same as that of expr1 .
	12	timediff(expr1,expr2)	Calculates the time difference between two expr1 and expr2 (expr1 - expr2). Returns the time type. The arguments must be both of the time or timestamp type.
	13	curdate()	Returns the current date (date type).
	14	curtime([fsp])	Returns the current time (time type). fsp indicates the precision of the returned time.
	15	uuid()	Returns a universally unique identifier (UUID).
	16	convert_tz()	Convert the time zone.
	17	cast(expr,typename)	Converts expr to a specified type.
	18	digest()	Generates binary hash values of the given data based on different algorithms.
	19	hmac()	Calculates the MAC value for the data with the key.
	20	crypt () and gen_salt ()	The crypt() and gen_salt() functions are used for password hashing. crypt() executes hashes to encrypt data, and gen_salt() generates salted hashes.
	21	pgp_sym_encrypt()	Used for symmetric key encryption.
	22	pgp_sym_decrypt()	Decrypts a message encrypted using a PGP symmetric key.
	23	pgp_pub_encrypt()	Used for public key encryption.

Change Type	No.	Function	Description
	24	pgp_pub_decrypt()	Decrypts a message encrypted using a PGP public key.
	25	pgp_key_id()	Extracts the key ID of the PGP public or private key. If an encrypted message is used as the input, the ID of the key used to encrypt the message will be returned.
	26	armor()	Converts binary data into PGP ASCII-armor format by the CRC calculation and formatting of a Base64 string.
	27	dearmor()	Performs the conversion opposite to the armor() function.
	28	pgp_armor_headers()	Returns the armor header in data.
	29	encrypt	Original encryption function, which does not support any advanced functions of PGP encryption. The IV is 0.
	30	decrypt	Original decryption function, which does not support any advanced functions of PGP encryption. The IV is 0.
	31	encrypt_iv	Original decryption function, which does not support any advanced functions of PGP encryption. The IV can be set by users.
	32	decrypt_iv	Original decryption function, which does not support any advanced functions of PGP encryption. The IV can be set by users.
	33	gen_random_bytes	Generates cryptographically strong random bytes.
	34	gen_random_uuid()	Returns a random UUID of version 4.

Change Type	No.	Function	Description
	35	get_volatile_pg_class	Obtains the pg_class metadata related to all volatile temporary tables in memory (metadata in volatile temporary tables is not stored in system catalogs).
	36	get_volatile_pg_attribute	Obtains the pg_attribute metadata related to all volatile temporary tables in memory.
	37	gs_increase_except_num	Increases the number of query exceptions.
	38	gs_update_blocklist_hash_info	Updates the blocklist information in memory.
	39	gs_append_blocklist	Adds blocklist statements.
	40	gs_remove_blocklist	Removes statements from a blocklist.
	41	gs_wlm_rebuild_except_rule_hash	Triggers the building of an exception rule kernel hash table.
	42	pg_stat_wal_write	Queries the thread information on the current CN or DN, and the import volume and rate of WAL and data pages.
	43	pgxc_stat_wal_write	Queries the thread information about the interaction between CNs and DNs, and the import volume and rate of WALs and data pages.
	44	get_col_file_info	Obtains the number of empty CU files and the total number of CU files of a specified column-store table.
	45	get_all_col_file_info	Obtains the number of empty CU files and the total number of CU files of all column-store tables.

Change Type	No.	Function	Description
	46	col_rebuild_file_relation	Replaces the Relfilenode and reorganize the CU files of the specified column-store table. After the reorganization, the CU files become consecutive 1 GB files.
Modified	47	pg_query_audit/ pgxc_query_audit	Changed the thread_id field to the session id field to record the session where the statement is executed.
	48	login_audit_messages/ login_audit_messages_pid	Added the session id field to record the session where the statement is executed.

System Views

Table 1-12 System views

Change Type	No.	View	Change Description
Added	1	gs_obs_latency	Monitors OBS latency.
	2	pgxc_stat_wal	Queries the traffic information about WAL logs and data pages of the current query.
	3	pgxc_lockwait_detail	Displays detailed information about the lock wait chain on each node.
	4	pgxc_wait_detail	Displays details about the SQL waiting chains of all nodes, including the wait nodes and queries.
Modified	5	gs_obs_read_traffic	Data aggregation interval is changed to 10 minutes, and the logtime column is added to display time groups.
	6	gs_obs_write_traffic	Data aggregation interval is changed to 10 minutes, and the logtime column is added to display time groups.

GUC Parameters

Table 1-13 GUC parameters

Change Type	No.	Parameter	Change Description
Added	1	enable_wlm_internal_memory_limit	Controls whether to enable the internal restriction of the WLM. For example, the estimated memory of a job cannot exceed 80%, 90%, or 40%.
	2	enable_strict_memory_expansion	Controls whether job memory expansion is performed on DNs.
	3	allow_zero_estimate_memory	Controls whether the estimated memory of a job can be set to 0 when there is no statistics for the job.
	4	max_process_memory_auto_adjust	Controls whether to dynamically adjust max_process_memory during a primary/standby CM switchover.
	5	wlm_memory_feedback_adjust	Specifies whether to enable the memory negative feedback function in dynamic load management. The value can be on , off , or on(50,40) .
	6	autoanalyze_mode	A system-level AUTOANALYZE control parameter. By default, the normal mode is the same as before. If the mode is set to light , dynamic sampling is enabled.
	7	enable_redactcol_computable	This parameter specifies whether to enable the computable data redaction. By default, this function is disabled.
	8	windowagg_pushdown_enhancement	Controls whether to push down the intersection of PARTITION BY and GROUP BY when the window function and GROUP BY appear at the same time.

Change Type	No.	Parameter	Change Description
	9	enable_from_collapse_hint	Specifies whether the plans with hints take precedence in execution. By default, the parameter is enabled and statement-level hints are supported.
	10	enable_value_redistribute	Specifies whether to use value partition streams to optimize WindowAgg.
	11	setop_optmode	Specifies whether to perform deduplication on subqueries of the setop operation. The default value is disable , indicating that deduplication is not performed, which is the same as that in versions earlier than 8.2.0. The value force indicates that deduplication is forcibly performed. The value cost indicates that deduplication is performed based on the cost. (Note: Deduplication is performed only when the bottom layer of the setop branch is the stream operator.)
	12	security_enable_options	Indicates that some operations are enabled in security mode. By default, this parameter is left blank.
	13	audit_object_name_format	Specifies the displaying mode of objects in the object_name column in audit logs. The default value is single , indicating that a single object is displayed.
	14	enable_redactcol_equal_const	Restricts the constant obtaining behaviors in the new version. By default, this function is disabled, and equivalent comparison between anonymized columns and constants is prohibited.

Change Type	No.	Parameter	Change Description
	15	enable_mixedagg	Indicates whether to generate a mixed aggregate execution plan. This parameter is disabled by default.
	16	max_streams_per_query	Controls whether to report an error in the optimizer phase for a plan with a large number of streams to prevent streams from occupying too many connections. The default value is 10000 .
	17	agg_max_mem	Controls the maximum estimated memory when the number of aggregation columns exceeds 5 . This prevents CCN queuing caused by high estimated memory.
	18	max_volatile_memory	Maximum number of memory contexts related to the volatile temporary table. The unit is KB . The value ranges from 1024 to INT_MAX . The default value is 1 GB .
	19	max_volatile_tables	Maximum number of volatile temporary tables (including auxiliary tables such as TOAST and CUDESC) in a single session. The default value is 300 .
	20	enable_hashfilter	Determines whether to generate hash filters.
	21	volatile_shipping_version	Controls the sequence pushdown enhancing in 8.2.0. The default value is 1 .
	22	smp_thread_cost	Control the generated query_dop value. The default value is 1000 .
	23	enable_wal_decelerate	Specifies whether to limit the flow control rate when data is imported with indexes.
	24	wal_decelerate_policy	Specifies the behavior policy after rate limiting is triggered.

Change Type	No.	Parameter	Change Description
	25	wal_write_speed	Indicates the maximum WAL write speed (byte/s) of each query on a single DN per second.
	26	wal_decelerate_trigger_threshold	Specifies the WAL write rate limit triggered by each query on a single DN.
	27	enable_cstore_ring_buffer	Controls whether to enable the column-store ring buffer.
	28	enable_tsdbs_multi_temperature	Specifies whether to enable the OBS cold and hot storage for time series tables. By default, the OBS cold and hot storage is disabled.
	29	enable_tsdbs_multi_temperature	Specifies whether to enable the OBS cold and hot storage for time series tables. By default, the OBS cold and hot storage is disabled.
	30	ddl_lock_timeout	Specifies the lockwait time for DDL statements. If the lockwait times out, DDL statements are terminated. This parameter is valid only for level-8 locks. The default value is 0, indicating that this parameter does not take effect.
	31	build_backup_param	Specifies the minimum specifications for disk backup during incremental build.
Modified	32	comm_tcp_mode	The parameter type is changed from POSTMASTER to SIGHUP.
	33	max_process_memory	Indicates the maximum memory usage of a process. The type of the parameter is changed from POSTMASTER to SIGHUP. The minimum threshold is 2 GB. You are not advised to set it to the minimum threshold.

Change Type	No.	Parameter	Change Description
	34	rewrite_rule	The enumerated value eager_magicset is added to control the magic set enhancement.
	35	behavior_compat_options	The normal_session_id parameter is added to behavior_compat_options to determine whether to concatenate the session ID with node_name . By default, normal_session_id is disabled, and the session ID is concatenated with node_name .
	36	behavior_compat_options	Added behavior_compat_options to merge_into_with_trigger . After the parameter is enabled, MERGE INTO can be executed on tables with triggers. This parameter is disabled by default.
	37	behavior_compat_options	Added the option disable_jsonb_exact_match to control the operator matching. By default, this option is disabled. When this option is disabled, only exact match of JSONB types is supported. If exact match fails, the non-JSONB types are matched. When this option is enabled, and the operator matching does not distinguish types.

Change Type	No.	Parameter	Change Description
	38	hashjoin_spill_strategy	Policies 5 and 6 are added. The behavior is the same as that of policies 0 and 1. The original policies 0 and 1 are modified. Originally, if the memory of the internal table is large and the memory of the external table is small, exchange the internal and external tables. After the modification, if the memory of the internal table is large and the memory of the external table is small, continue to divide the internal table until the internal table cannot be divided. Then try to exchange internal and external tables
	39	sql_use_spacelimit	Modified the initial value of sql_use_spacelimit to 10% (upper limit) of the disk space where the current instance resides. If the value is less than 10 GB (lower limit), set the value to 10 GB.
	40	temp_file_limit	Modified the initial value of temp_file_limit to 10% (upper limit) of the disk space where the current instance resides. If the value is less than 10 GB (lower limit), set the value to 10 GB.
Deleted	41	enable_grant_public	Indicates that the grant to public syntax is not supported in security mode.
	42	enable_grant_option	Indicates that the grant with grant option syntax is not supported in security mode.

1.5 Version 8.1.3

1.5.1 New Features in Version 8.1.3

For details about new features in 8.1.3.x, see [New Features](#). For details about resolved issues in each patch, see [New Features and Resolved Issues in 8.1.3.x](#).

[Release date of V8.1.3.325]: November 2, 2023

[Release date of V8.1.3.323]: October 16, 2023

[Release date of V8.1.3.322]: August 4, 2023

[Release date of V8.1.3.321]: June 26, 2023

[Release date of V8.1.3.320]: May 19, 2023

[Release date of V8.1.3.310]: March 9, 2023

[Release date of V8.1.3.300]: December 19, 2022

[Release date of V8.1.3.200]: October 31, 2022

[Release date of V8.1.3.110]: August 23, 2022

[Release date of V8.1.3.100]: June 20, 2022

[Release date of V8.1.3]: April 15, 2022

New Features

Table 1-14 New functions in 8.1.3

Feature	Description	Reference
Partition addition and partition table management are supported.	PERIOD (interval type) and TTL (interval type) are now available in partition management.	CREATE TABLE ALTER TABLE CREATE TABLE PARTITION
JSON foreign tables are supported.	<ol style="list-style-type: none">1. READ ONLY HDFS foreign tables in JSON format are supported.2. READ ONLY OBS foreign tables in JSON file format are supported.	CREATE FOREIGN TABLE (SQL on Hadoop or OBS)

Feature	Description	Reference
Cross-Cluster interconnection	<ol style="list-style-type: none"> 1. Constraints are removed: <ol style="list-style-type: none"> a. The hang constraint of GDS is removed. b. The constraints on common aggregation functions, such as sum, count, min, max, and avg are removed. c. The constraints on scalar functions and analytical functions are removed. d. The strong verification of the column names and column types of foreign tables and source tables is added. e. The constraint that the column name cannot contain WHERE is removed. f. The limit constraint is removed. g. The temporary GDS foreign table policy is improved. 2. If the local cluster is the destination cluster, you can run the following statements: 	GDS-based Cross-Cluster Interconnection
Enhanced cold and hot table functions	Automatic migration of cold and hot tables	CREATE TABLE

Feature	Description	Reference
Enhanced stream data warehouse functions	The time series database stores stream data in real time.	Stream Data Warehouse GUC Parameters CREATE INDEX Functions and Expressions
The row-store ring buffer is controlled by a switch.	The ring buffer is controlled by a switch and its threshold is adjustable. Batch update of large tables no longer uses the ring buffer.	buffer_ring_ratio
Default storage mode of a table	It is controlled by the parameter default_orientation . If the storage mode is not specified during table creation, this parameter specifies the default storage mode.	default_orientation
GDS supports the euro sign.	The GBK character set supports the euro sign.	CREATE DATABASE
OBS foreign tables support XSKY S3.	OBS foreign tables are adapted to XSKY and the parameter eol is added.	CREATE FOREIGN TABLE (SQL on Hadoop or OBS)
SQL supports Roaring Bitmaps.	In the Internet, new retail, education, and gaming industries, target crowds need to be quickly selected based on their profiles. The SQL syntax now supports Roaring Bitmaps, which can be used for user profiling in Internet operations.	Bitmap Functions and Operators RoaringBitmap

Feature	Description	Reference
Case-insensitive is supported to improve Teradata compatibility.	In the data migration from Teradata to DWS, the source data contains a large amount of duplicate data with different letter cases. If DWS uses the distinct syntax to deduplicate the data, after deduplication, there is still a large amount of duplicate data remaining. If DWS uses the upper function to do that, the source data will be changed, and it requires a lot of procedures and complex logics. To reduce the workload, the case-insensitive feature is added to DWS.	Collation Support Collation Version Function
List partitioning	List partitioned tables are supported.	<ul style="list-style-type: none"> • Syntax: CREATE TABLE PARTITION ALTER TABLE PARTITION • System catalog: PG_PARTITION • Views: DBA_PART_INDEXES DBA_PART_TABLES DBA_TAB_PARTITIONS USER_IND_PARTITIONS USER_PART_INDEXES USER_TAB_PARTITIONS

Feature	Description	Reference
Enhanced MySQL compatibility	<p>ALTER TABLE can modify multiple columns at a time.</p> <p>ALTER TABLE supports COMMENT.</p> <p>ALTER TABLE supports CHANGE, DROP KEY, DROP PARTITION, and FIRST/AFTER.</p>	<p>ALTER INDEX, CREATE INDEX</p> <p>ALTER TABLE, CREATE TABLE</p> <p>ALTER TABLE PARTITION</p>
Projection pushdown	Unnecessary columns in a multi-column view are automatically deleted during reference to avoid redundant calculation.	Added the description for projection pushdown to rewrite_rule
Share scan	Share scan	<p>Stream plan CTEs support share scan.</p> <p>Added enable_stream_ctescan.</p> <p>The value STREAM_CTESCAN is added to logging_module.</p> <p>The optional [NOT] MATERIALIZED is added to the WITH clause in the SELECT syntax.</p>
Bloom Filter for column-store tables	In a HASH JOIN, if the thread of the foreign table contains HDFS tables or column-store tables, the Bloom filter is triggered.	enable_bloom_filter
Column-store replication tables support UPDATE .	DWS customers have a large number of MERGE INTO operations on replication tables. Like row-store replication tables, column-store replication tables now support UPDATE, DELETE, and MERGE INTO operations.	<p>MERGE INTO UPDATE UPSERT</p>

Feature	Description	Reference
Fine-grained permission management	<ol style="list-style-type: none"> Added the ALTER, DROP and VACUUM permissions at table level. Added the ALTER and DROP permissions at schema level. Add the preset roles role_signal_backend and role_read_all_stats. 	<ul style="list-style-type: none"> Preset Roles Added table-level permissions: VACUUM VACUUM DROP TABLE Added schema-level permissions: ALTER SCHEMA DROP SCHEMA ALTER DEFAULT PRIVILEGES GRANT REVOKE
The problem of the INSERT OVERWRITE deadlock on multiple CNs is fixed.	No deadlock occurs when INSERT OVERWRITE is triggered on multiple CNs at the same time.	-
Execution plan optimization	<ul style="list-style-type: none"> If the join conditions contain OR and the execution plan uses nestloop, the performance is poor. Therefore, the statement is changed to UNION ALL. To join column-store tables, if the join or filter conditions contain OR, the OR conditions are pruned before join. When large and small tables are joined across many DN, the small table is broadcast instead of being redistributed. 	<p>Adjusting Key Parameters During SQL Tuning</p> <p>Added projection_pushdown and or_conversion to rewrite_rule.</p>

Feature	Description	Reference
<p>Relfile-based space statistics and space control</p>	<ol style="list-style-type: none"> The query results of the pg_table_size, pg_relation_size and table_skewness views can be returned in seconds. Schema space control can collect statistics on the full space and the constraints are removed. The user/schema space calibration functions can be quickly executed, and the schema space calibration function can be automatically executed. The SQL interface allocates space quota based on the total schema space instead of the DN level. 	<ul style="list-style-type: none"> Workspace Management Added GUC parameters for load management: space_once_adjust_num space_readjust_schedule Added the system catalog: PG_RELFILENODE_SIZE Added functions: gs_table_distribution(schemaname text, tablename text) gs_table_distribution() pgxc_wlm_readjust_relfilenode_size_table() pgxc_wlm_readjust_relfilenode_size_table(integer)
<p>Enhanced performance fault locating and demarcation</p>	<ol style="list-style-type: none"> Enhanced communication exception monitoring logs Enhanced pgxc_stat_activity view Top SQL statements can be associated with unique SQL statements. Added the distributed view pg_session_wlmstat. 	<ul style="list-style-type: none"> Added the system view PG_COMM_QUERY_SPEED. Added the description for the lwtid and stmt_type column of PG_STAT_ACTIVITY. Added the unique_sql_id column information to GS_WLM_SESSION_HISTORY. Added the system view PGXC_SESSION_WLMSTAT.

Feature	Description	Reference
Resource load management optimization and reconstruction	<ol style="list-style-type: none"> Enhanced user resource monitoring: The user resource monitoring logic is modified. Enhanced resource monitoring: Job and user resource monitoring is optimized, and queue-level resource monitoring is added. 	<ul style="list-style-type: none"> User Resource Monitoring Enhanced resource monitoring: <ul style="list-style-type: none"> Added: Resource Pool Monitoring. Added system views: <ul style="list-style-type: none"> PGXC_RESPOOL_RESOURCE_INFO PGXC_RESPOOL_RUNTIME_INFO PGXC_RESPOOL_RESOURCE_HISTORY GS_RESPOOL_RESOURCE_INFO GS_RESPOOL_RUNTIME_INFO Added the system catalog: <ul style="list-style-type: none"> GS_RESPOOL_RESOURCE_HISTORY
Modified partition table features	<ol style="list-style-type: none"> Changed the limit on the number of boundary values: The number of partitions in a partition table cannot exceed 32767, and the number of boundary values of all partitions cannot exceed 32767. The default partition creation policy is modified. When a partition table is created, two default partitions are created. The partition time range of the two default partitions is PERIOD. 	CREATE TABLE PARTITION

Feature	Description	Reference
Redistribution column hint in the AGG process	During data migration from Oracle data warehouses, you can manually specify the distribution column using enhanced hints.	Plan Hint Optimization Stream Operation Hints
Added the function of returning NULL when the divisor is 0. (Added in 8.1.3.110)	To specify whether to report an error when the divisor is 0 in MySQL compatibility mode, the compatibility configuration item enable_division_by_zero_mysql is added to the behavior_compat_options parameter.	behavior_compat_options
Added the DDL lock timeout configuration. (Added in 8.1.3.200)	Added the GUC parameter ddl_lock_timeout , which specifies the DDL lock timeout interval.	ddl_lock_timeout
Enhanced the view PG_LOCKS . (Added in 8.1.3.200)	<ul style="list-style-type: none"> Added the views PGXC_WAIT_DETAIL and PGXC_LOCKWAIT_DETAIL. Add columns such as wait_on_pid and query_id. 	<ul style="list-style-type: none"> PGXC_WAIT_DETAIL PGXC_LOCKWAIT_DETAIL
Data masking compatibility configuration (Added in 8.1.3.310)	The GUC parameter redact_compat_options is added to configure the compatibility of the data masking techniques.	<ul style="list-style-type: none"> redact_compat_options

Feature	Description	Reference
<p>Concurrent truncate, exchange, and select operations (added in 8.1.3.320)</p>	<p>Added the GUC parameter ddl_select_concurrent_mode. This parameter is used to solve the problem that the query statement lasts for a long time and blocks DDL in the scenario where the data volume is large or the query is complex. The effect is the same as that of the Oracle database.</p> <p>Application scenarios:</p> <ul style="list-style-type: none"> • Concurrent truncate and select operations • Concurrent exchange and select operations <p>Scenario constraints:</p> <ul style="list-style-type: none"> • Concurrency is not supported when there are conflicts with locks of higher levels (higher than level 1). For example, autoanalyze is triggered by select when autoanalyze_mode is set to normal. • Concurrency is not supported when there are conflicts with locks in transaction blocks. <p>Connection constraints:</p> <ul style="list-style-type: none"> • When the JDBC, ODBC, or Python driver is used, if autocommit is set to false, this feature is not supported. • This feature is not supported by query statements delivered by Data Studio and 	<p>ddl_select_concurrent_mode</p>

Feature	Description	Reference
	auto-committed transactions.	
Enabling of the early stop optimization for Limit statements (added in 8.1.3.320)	Added the GUC parameter enable_limit_stop to specify whether to enable the early stop optimization for LIMIT statements.	enable_limit_stop
Early stop behavior change (Added in 8.1.3.322)	The default value of the GUC parameter enable_limit_stop is changed to on .	enable_limit_stop

SQL Syntax

Table 1-15 SQL syntax

Change Type	No.	Name	Change Description
Added	1	CREATE FOREIGN TABLE (SQL on Hadoop or OBS)	Added json for the parameter Format of OPTION .
	2	CREATE FOREIGN TABLE (SQL on Hadoop or OBS)	Added force_mapping for OPTION .
	3	Column type kvtype	Specifies the column type of time series tables in the CREATE TABLE and ALTER COLUMN syntax.
	4	Table-level parameter: sub_partition_count	Configures the number of level-2 partitions in a time sequence table.
	5	Table-level parameter ttd	Specifies the expiration time of auto-added partitions.
	6	Table-level parameter period	Specifies the interval for triggering partition adding and dropping tasks and the time range of a single partition.

Change Type	No.	Name	Change Description
	7	List partitioning support	The list partition operation syntax is added, including the CREATE and ALTER operations. The CREATE syntax is a separate syntax branch. In the AT syntax, except for the SPLIT syntax, other syntax are similar to those of the range partitioning.
	8	ALTER TABLE ADD INDEX/ DROP INDEX	The ALTER TABLE syntax branch is added to add and delete indexes.
	9	CREATE TABLE ... LIKE	The CREATE TABLE... LIKE syntax without brackets is added.
	10	ALTER TABLE ... ADD/ MODIFY COLUMN	The statements for adding and modifying columns are added to ALTER TABLE . COMMENT and DEFAULT can be set.
	11	Added the [NOT] MATERIALIZED syntax after with cte as .	Added the [NOT] MATERIALIZED syntax after with cte as .
	12	alter index ... comment	The COMMENT clause is now available in the ALTER INDEX syntax.
	13	ALTER TABLE ... add index comment	The COMMENT clause is now available in the alter table add index syntax.
	14	create index index_name on table_name comment	The COMMENT clause is now available in the CREATE INDEX syntax.
	15	copy to	Added server , bom , fileprefix , and maxrow to the OPTION of the COPY TO statement.
	16	CREATE FOREIGN TABLE (OBS import and export)	Added bom to the OPTION of the CREATE FOREIGN TABLE statement (OBS import and export).
	17	Operator =	If two Roaring Bitmaps are equal, true is returned. Otherwise, false is returned.

Change Type	No.	Name	Change Description
	18	Operator <>	If two Roaring Bitmaps are not equal, true is returned. Otherwise, false is returned.
	19	Operator &	Calculates the intersection of two Roaring Bitmaps.
	20	Operator	Calculates the union of two Roaring Bitmaps.
	21	Operator	Calculates the result of adding an ID to a Roaring Bitmap.
	22	Operator #	Calculates the XOR result of two Roaring Bitmaps.
	23	Operator -	Obtains the set of elements that are in a Roaring Bitmap but not in another Roaring Bitmap.
	24	Operator -	Removes a specified ID from a Roaring Bitmap.
	25	Operator @>	If the Roaring Bitmap before the operator contains the elements after the operator, true is returned. Otherwise, false is returned.
	26	Operator <@	If the Roaring Bitmap after the operator contains the elements before the operator, true is returned. Otherwise, false is returned.
	27	Operator &&	If two Roaring Bitmaps overlap, true is returned. Otherwise, false is returned.
	28	ALTER TABLE DROP KEY	Added the syntax for deleting indexes. The usage of this syntax is the same as that of ALTER TABLE DROP INDEX .
	29	ALTER TABLE CHANGE	Compatible with the CHANGE syntax in MySQL. Column names and column attributes can be modified at the same time.

Change Type	No.	Name	Change Description
	30	ALTER TABLE ... FIRST/ AFTER colname	The syntax layer supports the FIRST/AFTER syntax, but does not implement the actual semantics. It is applicable to the ALTER TABLE ADD/MODIFY/CHANGE COLUMN scenario and is controlled by the parameter skip_first_after_mysql .
	31	AGG HINT	Added Agg hints for performance optimization.
	32	ALTER TABLE ... ALTER COLUMN cstore_cu_sample_ratio	To reduce the work of CU decompression, samples are concentrated in randomly selected CUs. Therefore, it is difficult to obtain data features. This field attribute is used to select more CUs without increasing the sampling ratio to facilitate data feature obtaining.
Modified	33	Table-level parameter deltarow_threshold	The default value of deltarow_threshold is set to 10000 , which specifies the upper limit of rows when column-store tables are imported to the Delta table. This parameter is valid only when the table-level parameter enable_delta is enabled.

Change Type	No.	Name	Change Description
	34	More data types are supported by Btree indexes.	The data types supported by the B-tree indexes are extended to support row store tables, column-store tables, and time series tables. btree_gin provides data types int2, int4, int8, float4, float8, timestamp with time zone, timestamp without time zone, time with time zone, time without time zone, date, interval, oid, money, "char", varchar, text, bytea, bit, varbit, macaddr, macaddr8, inet, cidr, uuid, name, bool, bpchar, and all enum types.
	35	GRANT/REVOKE/ALTER DEFAULT PRIVILEGES	The ALTER/DROP/VACUUM permissions are added.
	36	ALTER TABLE ... DROP PARTITION	The DROP PARTITION syntax can be used to delete multiple partitions.
	37	DEFAULT clause in the CREATE/ALTER TABLE statement	The DEFAULT clause in the CREATE/ALTER TABLE statement does not support suffix operators (currently, only the factorial suffix operator "!" is used). For example, create table t (a int default 3!) is not allowed.

Keywords

Table 1-16 Keywords

Change Type	No.	Name	Change Description
Added	1	MATERIALIZED	Added the [NOT] MATERIALIZED syntax after with cte as . As a non-reserved keyword, it can be still used as an object name. When it is used as a column alias, AS must be added.

Change Type	No.	Name	Change Description
	2	time_fill	It is used as the keyword of a time filling expression to output the time_fill column. It cannot be used as a function name or user-defined data type name.
	3	fill_first/fill_last/fill_avg	It is used as the keyword of a time filling expression to output the filled column. It cannot be used as a function name or user-defined data type name.
	4	list	Specifies the type of a partitioned table. As a non-reserved keyword, it can still be used as an object name. When it is used as a column alias, AS must be added.
	5	tsfield/tstag/tstime	Specifies the type of a partitioned table. As a non-reserved keyword, it can still be used as an object name. When it is used as a column alias, AS must be added.

System Catalogs

Table 1-17 System catalogs

Change Type	No.	Name	Change Description
Added	1	rb_added	A value is added to RoaringBitmap.
	2	pg_partition	Added the boundexprs column to pg_partition .
	3	pg_relfilenode_size	New system catalog
	4	pg_attribute	The attkvttype column is added to pg_attribute to record the kvttype of columns.

Change Type	No.	Name	Change Description
	5	pg_collation	The case_insensitive record is added to support case-insensitive behaviors.

System Functions

Table 1-18 System functions

Change Type	No.	Name	Change Description
Added	1	rb_build	Converts an int array into a bitmap.
	2	rb_to_array	Reverse operation of rb_build . It converts a Roaring Bitmap into an int array.
	3	rb_and	Obtains the intersection of two Roaring Bitmaps.
	4	rb_or	Obtains the union of two Roaring Bitmaps.
	5	rb_xor	Obtains the XOR result of two Roaring Bitmaps.
	6	rb_andnot	Performs AND then negation operation on two Roaring Bitmaps.
	7	rb_cardinality	Calculates the cardinality of a Roaring Bitmap.
	8	rb_and_cardinality	Calculates the cardinality of the And result of two Roaring Bitmaps.
	9	rb_or_cardinality	Calculates the cardinality of the OR result of two Roaring Bitmaps.
	10	rb_xor_cardinality	Calculates the cardinality of the XOR result of two Roaring Bitmaps.
	11	rb_andnot_cardinality	Calculates the cardinality of the andNot result of two Roaring Bitmaps.

Change Type	No.	Name	Change Description
	12	rb_is_empty	Determines whether a Roaring Bitmap is empty.
	13	rb_equals	Determines whether two Roaring Bitmaps are equal.
	14	rb_intersect	Determines whether two Roaring Bitmaps intersect.
	15	rb_contain	Determines whether the first Roaring Bitmap contains a specified value.
	16	rb_add	Adds a value to a Roaring Bitmap.
	17	rb_remove	Deletes a value from a Roaring Bitmap.
	18	rb_flip	Reverses the Roaring Bitmap within a specified range.
	19	rb_min	Obtains the minimum value of a Roaring Bitmap.
	20	rb_max	Obtains the maximum value of a Roaring Bitmap.
	21	rb_rank	Returns the cardinality of values in a bitmap that are less than or equal to the specified Offset.
	22	rb_contain_rb	Check whether the first Roaring Bitmap contains the second Roaring Bitmap.
	23	rb_containedby_rb	Check whether the second Roaring Bitmap contains the first Roaring Bitmap.
	24	rb_containedby	Determines whether a specified value is included in a specified Roaring Bitmap.
	25	rb_iterate	Returns the ints corresponding to a Roaring Bitmap.
	26	rb_and_agg	Aggregates a Roaring Bitmap column based on the AND logic.
	27	rb_or_agg	Aggregates a Roaring Bitmap column based on the OR logic.

Change Type	No.	Name	Change Description
	28	rb_xor_agg	Aggregates a Roaring Bitmap column based on the XOR logic.
	29	rb_and_cardinality_agg	Obtains the cardinality of the AND result of a Roaring Bitmap column.
	30	rb_or_cardinality_agg	Obtains the cardinality of the OR result of a Roaring Bitmap column.
	31	rb_xor_cardinality_agg	Obtains the cardinality of the XOR result of a Roaring Bitmap column.
	32	rb_build_agg	Aggregates the int column into RoaringBitmap data.
	33	pgxc_wlm_readjust_relfilenode_size_table()	Space statistics calibration function. It does not recreate the PG_RELFILENODE_SIZE system catalog but recalibrates the user and schema space.
	34	gs_table_distribution()	Quickly queries a table size.
	35	pg_obs_cold_refresh_time	Modifies the time for automatic switchover of OBS multi-temperature tables.
	36	gs_clean_tag_relation	Deletes the useless data in the tagid row of the tag table. The input parameter of the function is the OID of the time series table. Each partition traverses the minimum value of the tagid column in the cu desc table to obtain the minimum tagid value in the entire time series table.
	37	proc_drop_partition	Drops a partition whose boundary time exceeds the TTL.
	38	proc_add_partition	Creates partitions for a partitioned table.
	39	pg_collation_actual_version	Returns the actual version number of an ICU collation.

Change Type	No.	Name	Change Description
	40	first	Indicates the first element in a group.
	41	last	Indicates the last element in a group.
	42	mode	Indicates the value with the highest occurrence frequency in a group.
	43	delta	Indicates the difference between two adjacent rows.
	44	percentile_of_value	Indicates the approximate percentile value.
	45	value_of_percentile	Indicates the approximate percentile.
	46	spread	Indicates the difference between the maximum value and minimum value in a group.
	47	pg_flush_buffers	Flushes all row-store dirty pages.
Modified	48	pg_stat_activity series views	The stmt_type and lwtid columns are added to the pg_stat_get_activity_with_conninfo , pg_stat_activity , pgxc_stat_activity and pg_stat_get_activity views.
	49	Added two rows to the pg_authid system catalog.	Added the preset roles pg_role_signal_backend and pg_role_read_all_stats .
	50	vac_fileclear_relation	Returns the number of files to be cleared in a specified table. If no files are to be cleared, 0 is returned.
	51	vac_fileclear_all_relation	Returns the number of files to be cleared in all column store tables. If no files are to be cleared, 0 is returned.

System Views

Table 1-19 System views

Change Type	No.	Name	Change Description
Added	1	pgxc_session_wlmstat	New view
	2	pg_comm_query_speed	New view. It is used to query sending information based on query_ids.
	3	pgxc_respool_resource_info pgxc_respool_runtime_info pgxc_respool_resource_history gs_respool_resource_info gs_respool_runtime_info	New views related to resource monitoring.
	4	pgxc_wait_detail pgxc_lockwait_detail	Added the enhanced view pg_locks .
Modified	5	DBA_PART_INDEXES DBA_PART_TABLES DBA_TAB_PARTITIONS USER_IND_PARTITIONS USER_PART_INDEXES USER_TAB_PARTITIONS	Added description for list partitions.
	6	pgxc_wlm_session_statistics	The logic for querying real-time information about Top SQL statements from all CNs is changed from serial to parallel to improve performance. The function remains unchanged.
	7	all_indexes	The definition of the view all_indexes in sys and pg_catalog is incorrect. If objects with the same name exist in different schemas, the result set expands.

Behavior Changes

Table 1-20 Behavior changes

Change Type	No.	Name	Change Description
Modified	1	In CREATE INDEX , the target table is a time series table.	Any index created in the time series table is converted to a double index in the tag table. The index column of the double index is the specified column that the index is created for.
	2	Secondary load management	Added secondary management to load management. Functions and multi-statements that contain complex queries may trigger multiple controls. You can set enable_transaction_parctl to off to disable secondary management. If you do so, transaction block statements and multi-statement management will also be disabled.
	3	Load management (autoanalyze management)	The control logic for queries triggering autoanalyze is changed from no control to control . You can set enable_transaction_parctl=off to disable the autoanalyze control.
	4	User monitoring view pg_total_user_resource_info	<ol style="list-style-type: none">1. Changed the CPU and memory resource usage and limit to the resource usage and limit in the cluster.2. The monitored CPU, I/O, and memory objects are changed from complex jobs to all jobs.3. The CPU monitoring logic is changed from Cgroup monitoring to job CPU monitoring summary.

Change Type	No.	Name	Change Description
	5	Audit logs	<p>1. Statements in a transaction are audited when the statement type is set to be audited, even if the transaction is not set to be audited.</p> <p>2. The DECLARE CURSOR statement is audited when the guc parameter audit_operation_exec is set to select.</p>

GUC Parameters

Table 1-21 GUC parameters

Change Type	No.	Name	Change Description
Added	1	space_readjust_schedule	Indicates whether to enable automatic calibration for users and schema spaces. The options are off , auto , and auto (<i>xxx</i> K/M/G). The default value is auto .
	2	space_once_adjust_num	Threshold of the number of files processed each time during slow build and fine-grained calibration in the space management and statistics functions. The default value is 300 .
	3	tag_cache_max_number	Indicates the maximum number of items cached in the tag cache of a global hash table. The default value is 10000000 .
	4	enable_tagbucket_auto_adapt	Indicates whether to enable tagbucket adaptive adjustment. The default value is on .
	5	cache_tag_value_num	Number of cached tag tuples during tag column laterread. The default value is 60000 .

Change Type	No.	Name	Change Description
	6	buffer_ring_ratio	Controls the threshold and switch of the ring buffer in a row-store table. The default value is 250 , that is, 1/4 (250/1000), which is the same as the previous logic.
	7	enable_stream_ctescan	This parameter is added. By default, this parameter is set to OFF after an upgrade and is set to ON after a new installation. Specifies whether to enable the share scan function for the stream plan.
	8	behavior_compat_options	The disable_including_all_mysql option is added to control whether the CREATE TABLE LIKE syntax is in INCLUDING ALL mode. By default, this option is not set and the CREATE TABLE LIKE syntax is in INCLUDING ALL mode.
	9	profile_logging_module	Added the GUC parameter profile_logging_module to configure the type of performance logs. By default, OBS, Hadoop, and REMOTE_DATANODE are enabled and MD is disabled after upgrade or installation. The method of setting this parameter is the same as that of setting logging_module .
	10	object_mtime_record_mode	Controls the behaviors of pg_object recording modification time. By default, the behavior of the old version is retained. The new option can be set to not record the TRUNCATE/GRANT/REVOKE operations on partitions.
	11	skew_option	Specifies whether to enable an optimization policy.
	12	ddl_lock_timeout	DDL lock timeout interval.

Change Type	No.	Name	Change Description
Modified	13	rewrite_rule	<ul style="list-style-type: none"> The orconversion option is added and is enabled by default. Indicates the OR optimization rules for equi-join. Change: The query plan converts nestloop to hashjoin. The projection_pushdown option is added and is enabled by default. Change: Unused columns in subqueries, CTEs, and views are deleted.
	14	default_orientation	The userset guc parameter default_orientation is added to control the default type of tables users can create (orientation is not set). The options are as follows: Row-store table (row), column-store table (column), and column-store table with delta enabled (column enabledelta)
	15	max_process_memory	It is an OM setting item configured during initial installation. The original formula is: Available memory x 0.7/(1 + Number of DNs). The coefficient 0.7 is changed to 0.8. It remains unchanged in scenarios such as upgrade and scale-out of the old cluster.
	16	enable_bloom_filter	Reused parameter. It is now used to control BloomFilter in a column-store table. The default value remains unchanged (Enabled).

Change Type	No.	Name	Change Description
	17	retry_ecode_list	Added the default value 45003 to solve the problem that wrong partitions may be queried when partitions are added and queried concurrently in a list partitioned table.
	18	auth_iteration_count	The default value is reduced from 50000 to 10000 . The change applies to both installation and upgrade. Indicates the number of hash iterations of the passwords of the client and server.

1.5.2 New Features and Resolved Issues in 8.1.3.x

8.1.3.325

Table 1-22 8.1.3.325 New Features/Resolved Issues

Type	Feature or Resolved Issue	Cause	Version	Handling Method
New Features	<ul style="list-style-type: none"> The GB18030_2022 character encoding is supported. The subquery hint function can be disabled. The GDS foreign table is compatible with invalid UTF-8 characters. 	-	-	-

Type	Feature or Resolved Issue	Cause	Version	Handling Method
Resolved Issues	Logs fail and the cluster is hung after DN restart.	When the stream thread error logs are printed, the stream thread responds to the cancel signal. When the logs are printed again, the stream thread is hung up.	8.1.3.320	Upgrade the version to 8.1.3.325.
	<ul style="list-style-type: none"> When a service table is created, the following error message is occasionally displayed: relation "xxx" already exists. The error could not read block xxx in file "xxx" is occasionally reported. 	When the OID usage surpasses 4.2 billion, the system reallocates OIDs. During this reallocation process, intermittent errors may occur.	8.1.3.323 and earlier versions	
	Checkpoint failures on the standby DN prevent the reclamation of xlogs, leading to their continuous accumulation.	During the redo checkpoint operation on the standby DN, the detection of incomplete btree splitting triggers an RM 11 error. This may indicate a btree splitting failure or that the process is still ongoing.	8.1.3.323 and earlier versions	
	If a lightweight update occurs on a column-store table while autovacuum runs in the background, the system may report an error Unsupported to update different rows with the same cuid and col in light update.	In lightweight update, the same cuid and col cannot be used to update different rows.	8.1.3.323 and earlier versions	

Type	Feature or Resolved Issue	Cause	Version	Handling Method
	The no_merge hint operation does not take effect in some scenarios.	The subquery hint function cannot be disabled.	8.1.3.323 and earlier versions	
	When a long-time query statement is executed on the background, if you press Ctrl+C to end the statement and then press Ctrl+Z to exit the gsql client, residual statements exist.	Upon receiving the CTRL+C signal, the CN thread enters the error jump process. In this case, the CN thread invokes EmitErrorReport, but EmitErrorReport does not respond to the signal.	8.1.3.323 and earlier versions	
	When a wide table is used with a real-time table for associated update, the memory usage is too high, affecting the performance.	Associating a wide table with a real-time table for updates introduces superfluous columns. This leads to excessive memory consumption during the optimizer's execution plan estimation, thereby deteriorating performance	8.1.3.323 and earlier versions	
	In the ShareScan plan, the memory of the tuplestore and batchstore for storing CTEs cannot be expanded. As a result, data may be written to disks.	cteMaxMem in the ShareScan execution plan is not copied, read, or written. As a result, the DN fails to obtain the value from the CTE execution plan, and the created tuplestore and batchstore for storing the CTE result set cannot be expanded in memory. As a result, data spills to disks.	8.1.3.323 and earlier versions	

Type	Feature or Resolved Issue	Cause	Version	Handling Method
	The error message could not determine data type of parameter is reported when a lightweight column-store update is performed.	In the remote query of the execution plan, the parame parameter is specified as its position in the base table, and the system column cannot match this parameter.	8.1.3.323 and earlier versions	

8.1.3.322

Table 1-23 8.1.3.325 new features/resolved issues

Type	Feature or Resolved Issue	Cause	Version	Handling Method
New Features	The default value of the GUC parameter enable_limit_stop is changed to on .	-	-	-
Resolved Issues	None	-	-	-

8.1.3.320

Table 1-24 8.1.3.320 new features/resolved issues

Type	Feature or Resolved Issue	Cause	Version	Handling Method
New Features	<p>Concurrent truncate, exchange, and select operations</p> <p>Added the GUC parameter ddl_select_concurrent_mode. This parameter is used to solve the problem that the query statement lasts for a long time and blocks DDL in the scenario where the data volume is large or the query is complex. The effect is the same as that of the Oracle database.</p> <p>Application scenarios:</p> <ul style="list-style-type: none"> • Concurrent truncate and select operations • Concurrent exchange and select operations <p>Scenario constraints</p> <ul style="list-style-type: none"> • Concurrency is not supported when there are conflicts with locks of higher levels (more than one level). For example, autoanalyze is triggered by select when autoanalyze_mode is set to normal. • Concurrency is not supported when there are conflicts with locks in transaction blocks. 	-	-	-

Type	Feature or Resolved Issue	Cause	Version	Handling Method
	<p>Connection constraints:</p> <ul style="list-style-type: none"> • When the JDBC, ODBC, or Python driver is used, if autocommit is set to false, this feature is not supported. • This feature is not supported by query statements delivered by Data Studio and auto-committed transactions. 			
	<p>Whether to enable early stop optimization for LIMIT statements can be controlled.</p> <p>Added the GUC parameter enable_limit_stop to specify whether to enable the early stop optimization for LIMIT statements.</p>	-	-	-
Resolved Issues	<p>The statement-level estimated memory exceeds the value of max_process_memory and causes ccn queuing.</p>	<p>The statement-level estimated memory exceeds the value of max_process_memory.</p>	8.1.3.310 and earlier versions	Upgrade the version to 8.1.3.320 or later.
	<p>When the MERGE operation is performed, an error message value xxx is out of range for type integer is displayed.</p>	<p>The pg_toast_get_baseid() function uses an integer interface for OID processing.</p>	8.1.3.310 and earlier versions	

Type	Feature or Resolved Issue	Cause	Version	Handling Method
	Querying the dirty page rate view of a large table containing 200,000 data rows results in significant memory consumption.	In the process of data querying, each node in the statistics-querying linked list requests memory to store tupStatus. Post-query, while the nodes' memory is freed, the allocated memory within the nodes for tupStatus remains unreleased.	8.1.1.x	
	When a Vacuum operation times out and is terminated, some DNs retain leftover threads, which then fail to respond to signals and obstruct the execution of subsequent statements.	The btvacuumscan thread does not respond to signals. As a result, after the Vacuum execution times out and is killed, some DNs have residual threads and cannot respond to signals.	8.1.1.x	
	Executing percentage sampling on partitioned and time series tables within the MySQL compatibility library triggers an error: unsupported feature with temporary/unlogged table for partitioned table.	In 8.1.3, the default inheritance mode is INCLUDING ALL to match the CREATE TABLE LIKE syntax in MySQL. However, this causes an error during ANALYZE percentage sampling because the command for automatically creating temporary replication tables is included the sampling.	8.1.3.310 and earlier versions	

Type	Feature or Resolved Issue	Cause	Version	Handling Method
	The SELECT query of the cold and hot tables triggers runtime autoanalyze , deteriorating the query performance.	If the select operation is performed After the TRUNCATE PARTITION, EXCHANGE PARTITION, DROP PARTITION, ALTER DISTRIBUTE, or ALTER COLUMN TYPE operation, AUTOANALYZE is triggered.	8.1.3.310 and earlier versions	
	The value of elapsed_time in pg_session_wlmstat is inconsistent with the actual execution time of the top SQL statement.	When querying the pg_session_wlmstat view, the elapsed_time value appears excessively high for statements in the RUNNING state.	8.0.x	
	When the SELECT permission is granted to a foreign table, the error message "has no distribute type" is displayed.	The default distribution mode is not specified for the OBS dfs server write-only foreign tables.	8.1.3.310 and earlier versions	

1.6 Version 8.1.1

V8.1.1.500 release date: June 20, 2022

[Release date of V8.1.1.300]: April 15, 2022

[Release date of V8.1.1.205]: March 30, 2022

[Release date of V8.1.1.203]: March 18, 2022

[Release date of V8.1.1.202]: February 24, 2022

[Release date of V8.1.1.201]: January 25, 2022

[Release date of V8.1.1.200]: December 9, 2021

[Release date of V8.1.1.100]: July 30, 2021

 NOTE

The database kernel functions of clusters 8.1.1.200, 8.1.1.201, 8.1.1.202, 8.1.1.203, 8.1.1.205, 8.1.1.300, 8.1.1.500, and 8.1.1.100 are the same. Some functions are adapted to the management console. For details, see [What's New](#).

New Features

Table 1-25 New Functions in 8.1.1

Feature	Description	Reference
2048 nodes in a cluster	Supported up to 2048 nodes in a cluster.	Technical Specifications
Resource monitoring performance tuning	<ol style="list-style-type: none"> Added functions pgxc_wlm_get_schema_space(cstring) and pgxc_wlm_analyze_schema_space(cstring). Added system views PGXC_TOTAL_SCHEMA_INFO, PGXC_TOTAL_SCHEMA_INFO_ANALYZE, and GS_WLM_SQL_ALLOW. 	<ul style="list-style-type: none"> Functions pgxc_wlm_get_schema_space(cstring) and pgxc_wlm_analyze_schema_space(cstring) System views PGXC_TOTAL_SCHEMA_INFO, PGXC_TOTAL_SCHEMA_INFO_ANALYZE, and GS_WLM_SQL_ALLOW
Lightweight UPDATE	<ol style="list-style-type: none"> Added descriptions about column-store tables and the lightweight UPDATE operation on column-store tables to "Precautions". Added the GUC parameter enable_light_colupdate. 	<ul style="list-style-type: none"> UPDATE enable_light_colupdate
Primary key CU of column-store tables	Supported primary key constraints of column-store tables.	<ul style="list-style-type: none"> Constraint Design ALTER TABLE CREATE INDEX
Space management	Added the GUC parameter bi_page_reuse_factor .	bi_page_reuse_factor

Feature	Description	Reference
GDS productization	Added system views PGXC_BULKLOAD_PROGRESS , PGXC_BULKLOAD_STATISTICS , and PG_BULKLOAD_STATISTICS .	<ul style="list-style-type: none"> PGXC_BULKLOAD_PROGRESS PGXC_BULKLOAD_STATISTICS PG_BULKLOAD_STATISTICS
Hot and cold data storage	<ol style="list-style-type: none"> Added the description of REFRESH STORAGE to "ALTER TABLE". Added the functions and descriptions of COLVERSION and STORAGE_POLICY to "CREATE TABLE". Added the function and description of the OBS tablespace to "CREATE TABLESPACE". 	<ul style="list-style-type: none"> ALTER TABLE CREATE TABLE
C function in SM4 CBC mode	<p>Added security functions gs_encrypt(encryptstr, keystr, cryptotype, cryptomode, hashmethod), gs_decrypt(decryptstr, keystr, cryptotype, cryptomode, hashmethod), and gs_hash(hashstr, hashmethod).</p> <p>Modified the description of gs_encrypt_aes128(encryptstr, keystr) and gs_decrypt_aes128(decryptstr, keystr).</p>	Security Functions
Support for the built-in MEDIAN function	Added the function median(expression) .	Aggregate Functions
Adjusting the cascade delete for tables with views	<ol style="list-style-type: none"> Added the description of rebuilding a view. Added the following description to the REBUILD parameter of ALTER VIEW: Only columns of the character, number, and time types in the base table can be modified. When view_independent is set to on, views can be automatically rebuilt. 	<ul style="list-style-type: none"> Creating and Managing Views ALTER VIEW

Feature	Description	Reference
Custom data redaction	<ol style="list-style-type: none"> Optimized data redaction. Provided data redaction functions in different forms. Updated the columns in the system catalog PG_REDACTION_COLUMN and system view REDACTION_COLUMNS. Added the syntax ALTER REDACTION POLICY, CREATE REDACTION POLICY, and DROP REDACTION POLICY. 	<ul style="list-style-type: none"> Data Redaction Data Redaction Functions PG_REDACTION_COLUMN and REDACTION_COLUMNS ALTER REDACTION POLICY CREATE REDACTION POLICY DROP REDACTION POLICY
After the password expires, a user can log in to the database only after changing the password.	<ol style="list-style-type: none"> "Keyword" in the <i>SQL Syntax</i> Added EXPIRATION (non-reserved). Added the security function gs_password_expiration, and the description that the system prompts users to change their passwords after the gs_password_deadline password expires. This is related to the GUC parameter password_effect_time. Added PASSWORD EXPIRATION period to the syntax ALTER ROLE, ALTER USER, and CREATE USER. Added the following description to CREATE ROLE: number of days before the login password of a role expires. A user needs to change the password before it expires. If the login password expires, the user cannot log in to the system. In this case, the user needs to ask the administrator to set a new login password and use it to log in to the system. Added the columns rolauthinfo, rolpwdexpire, and rolpwddtim to the system catalog PG_AUTHID. 	<ul style="list-style-type: none"> Keywords Security Functions ALTER ROLE, ALTER USER, and CREATE USER CREATE ROLE PG_AUTHID
Increasing the hash table size	Added the GUC parameter expand_hashtable_ratio .	expand_hashtable_ratio

Feature	Description	Reference
SMP adaptation enabled	Added "Suggestions for SMP Parameter Settings".	Suggestions for SMP Parameter Settings
PRETTY as the default value of EXPLAIN	Changed the default display format of EXPLAIN to PRETTY . Added the statistics of filtering and projection time to the new operator. The statistics and the operator execution time are displayed in the same row.	SQL Execution Plan Details
Enhanced concurrent SMP	<ol style="list-style-type: none"> Added the constraint of SMP: the short query scenario where the plan generation time accounts for a large proportion. Added the description of setting query_dop to 1 in the short query scenario. Added the GUC parameter query_dop_ratio. 	<ul style="list-style-type: none"> SMP Application Scenarios and Restrictions SMP Manual Optimization Suggestions query_dop_ratio
row2vec optimization	Added the GUC parameter enable_row_fast_numeric .	enable_row_fast_numeric

Feature	Description	Reference
MySQL compatibility	<ol style="list-style-type: none"> 1. Added the negative processing of concat, log(x), left, and right, last_day and next_day return types of int + unknown operations, and compatibility differences of the operator ^. 2. Added the keywords IFNULL and TIMESTAMPDIFF. 3. Added the description of processing CASE, COALESCE, IF, and IFNULL in MySQL-compatible mode. 4. Added the MySQL-compatible schema and example of the behavior_compat_options option. 5. [OFFSET start [ROW ROWS]] LIMIT start, { count ALL } } supported by SELECT, SELECT INTO, and VALUES. 6. Added MySQL compatibility to the DBCMPATIBILITY of CREATE DATABASE. 7. Added the following content: compatibility differences between concat(str1, str2), left(str text, n int), and right(str text, n int) for character processing functions and operators; compatibility differences between log(x) and ^ for numeric operation functions and operators; time and date processing functions and operators, timestampdiff(field, timestamp1, timestamp2) / to_days(timestamp) / data_format; conditional expression functions, including if(bool_expr, expr1, expr2), ifnull(expr1, expr2), and isnull(expr). 	<ul style="list-style-type: none"> • Syntax Compatibility Differences Among Oracle, Teradata, and MySQL • Keywords • UNION, CASE, and Related Constructs • behavior_comp at_options • SELECT, SELECT INTO, and VALUES • CREATE DATABASE • Functions and Operators
Support for UPSERT	<p>Added the UPSERT syntax.</p> <p>Added IGNORE, AS, ON DUPLICATE KEY, and ON CONFLICT to the INSERT syntax.</p>	UPSERT

Feature	Description	Reference
Support for the INSERT , UPDATE , and DELETE operations of views	Added Updatable Views to "CREATE VIEW".	CREATE VIEW
One-click onsite information collection	Added the STATS boolean parameter to the EXPLAIN syntax.	EXPLAIN
One-click relationship analysis of a lock wait	<ol style="list-style-type: none"> Added the lock information function pgxc_get_lock_conflicts(). Added the system views PGXC_DEADLOCK and PGXC_LOCK_CONFLICTS. 	<ul style="list-style-type: none"> System Information Functions System views PGXC_DEADLOCK and PGXC_LOCK_CONFLICTS
Combination with an empty string of the CHAR type	Added the configuration item bpchar_text_without_rtrim to the parameter behavior_compat_option .	behavior_compat_options
Support for Python 3.x	Supported Python 3.	Using GDS to Import Data from a Remote Server
When to_date and to_timestamp process an empty string, 0001-01-01 is returned, and null is returned for TD.	<ol style="list-style-type: none"> Added "Syntax Compatibility Differences Among Oracle, Teradata, and MySQL". Added the configuration item convert_empty_str_to_null_td to behavior_compat_options. 	<ul style="list-style-type: none"> Syntax Compatibility Differences Among Oracle, Teradata, and MySQL behavior_compat_options
Support for INSERT OVERWRITE	Supported INSERT OVERWRITE .	INSERT
Obtaining all views queried by the current user	Added the view GS_VIEW_DEPENDENCY_PATH .	GS_VIEW_DEPENDENCY_PATH
Support for percentile_disc and percentile_cont	Added the following aggregate functions: percentile_disc(const) within group(order by expression) percentile_cont(const) within group(order by expression)	Aggregate Functions

SQL Syntax

Table 1-26 SQL syntax

Change Type	No.	Name	Change Description
Add	1	LIMIT offset,count	The LIMIT offset,count syntax is supported.
	2	EXPLAIN(STATS ON)...	The EXPLAIN(STATS ON) syntax is supported. It is used to export information for regenerating a plan.
	3	CREATE/ALTER/DROP REDACTION POLICY	Added the syntax for masking DDL statements.
	4	INSERT IGNORE INTO INSERT INTO ON DUPLICATE KEY UPDATE INSERT INTO ON CONFLICT DO UPDATE INSERT INTO ON CONFLICT DO NOTHING	UPSERT is supported.
	5	INSERT OVERWRITE INTO	INSERT OVERWRITE is supported.
	6	ALTER TABLE REFRESH STORAGE	Users can change hot data to cold data in tables.
	7	ALTER VIEW ONLY view_name REBUILD	ALTER VIEW ONLY view_name REBUILD is supported.
	8	ALTER SERVER REFRESH	You can update the configuration file of the HDFS server if the HDFS configuration was modified.
Modify	9	Operator	In MySQL compatibility mode, XOR is used. In ORA or TD mode, exponentiation is used.
	10	MERGE INTO	An error will be reported if the target table of the MERGE INTO statement contains triggers.
	11	CREATE/ALTER Table	Added options about cold_tablespace and storage_policy in Reloptions .

Keywords

Table 1-27 Keywords

Change Type	No.	Name	Change Description
Add	1	TIMESTAMPDIFF	The COL_NAME_KEYWORD keyword is added.
	2	IFNULL	The COL_NAME_KEYWORD keyword is added.
	3	REFRESH	Non-reserved keyword
Modify	4	IF	UNRESERVED_KEYWORD is changed to COL_NAME_KEYWORD .
	5	ISNULL	UNRESERVED_KEYWORD is changed to COL_NAME_KEYWORD .
	6	VERIFY	This keyword is changed from reserved to non-reserved.

System Catalogs

Table 1-28 System catalogs

Change Type	No.	Name	Change Description
Modify	1	PG_REDACTION_COLUMN	The func_expr column is added.
	2	PG_AUTHID	The rolpwdexpire and rolpwftime columns are added.

System Functions

Table 1-29 System functions

Change Type	No.	Name	Change Description
Add	1	pgxc_wlm_readjust_schema_space()	You can perform parallel calibration in schema space.

Change Type	No.	Name	Change Description
	2	pgxc_wlm_get_schema_space()	You can query the schema space information of the cluster.
	3	pgxc_wlm_analyze_schema_space()	You can analyze the schema space information of the cluster.
	4	gs_roach_enable_delay_ddl_recycle	You can enable DDL delay by specifying a backup set name.
	5	gs_roach_disable_delay_ddl_recycle	You can disable DDL delay by specifying a backup set name.
	6	gs_roach_stop_backup	You can disable row-store backup by specifying a backup set name.
	7	gs_roach_switch_xlog	You can configure whether to perform the checkpoint operation for xlog switch.
	8	pgxc_get_lock_conflicts()	This function returns information about conflicting locks on nodes.
	9	mask_none/mask_full/mask_partial	A built-in masking function is added.
	10	median	The median agg function is added.
	11	pgxc_node_stat_reset_time()	You can query the reset time of each node.
	12	pgxc_rel_iostat()	You can query the I/O statistics of each node.
	13	pgxc_redo_stat()	You can query the redo statistics of each node.
	14	pgxc_instance_time()	You can query the time statistics of each instance.
	15	pgxc_settings()	You can query the GUC settings of each node.
	16	pgxc_replication_slots()	You can query the replication slot statistics of each node.
	17	pgxc_stat_replication()	You can query the replication statistics of each node.

Change Type	No.	Name	Change Description
	18	pgxc_stat_bgwriter()	You can query the statistics of the bgwriter process on each node.
	19	pgxc_wait_events()	You can query statistics on the wait events of each node.
	20	pgxc_os_run_info()	You can query the OS performance statistics of the servers where instances are deployed.
	21	get_node_stat_reset_time()	You can query the time when the current node status is reset.
	22	get_local_rel_iostat()	You can query the I/O statistics of the current node.
	23	get_instr_wait_event()	You can query statistics on the wait events of the current node.
	24	pg_stat_get_db_total_blk_write_time()	You can query the total block write time.
	25	pg_stat_get_db_total_blk_read_time()	You can query the total block read time.
	26	pg_stat_get_db_total_temp_bytes()	You can query the total size of temporary files.
	27	pg_stat_get_db_total_temp_files()	You can query the total number of temporary files.
	28	pg_stat_get_db_total_deadlocks()	You can query the total number of deadlocks.
	29	pg_stat_get_db_total_conflict_all()	You can query the total number of conflicts.
	30	pg_stat_get_db_total_tuples_deleted()	You can query the total number of deleted tuples.
	31	pg_stat_get_db_total_tuples_updated()	You can query the total number of updated tuples.
	32	pg_stat_get_db_total_tuples_inserted()	You can query the total number of inserted tuples.
	33	pg_stat_get_db_total_tuples_fetched()	You can query the total number of fetched tuples.

Change Type	No.	Name	Change Description
	34	pg_stat_get_db_total_tuples_returned()	You can query the total number of returned tuples.
	35	pg_stat_get_db_total_blocks_hit()	You can query the total number of hit data blocks in the memory.
	36	pg_stat_get_db_total_blocks_fetched()	You can query the total number of read data blocks.
	37	pg_stat_get_db_total_xact_rollback()	You can query the total number of rollback transactions.
	38	pg_stat_get_db_total_xact_commit()	You can query the total number of commit transactions.
	39	pg_stat_get_db_total_numbackends()	You can query the total number of backends.
	40	gs_encrypt()	Encryption function
	41	gs_decrypt()	Decryption function
	42	gs_hash()	Hash function
	43	gs_password_expiration()	You can query the remaining password validity period (rolpwddtime and rolpwdexpire columns in the pg_authid system catalog) of the current user.
	44	to_char(date)	The to_char(date) function is added to be compatible with the return format of the DATE timestamp processed by to_char in TD mode.
	45	gs_wlm_set_queryband_action	You can configure the query_band load behavior.
	46	gs_wlm_set_queryband_order	You can configure the search priority of query_band.
	47	gs_wlm_get_queryband_action	You can query the query_band load behavior.
	48	CONCAT	A MYSQL compatibility rule is added. If the input parameter contains NULL, NULL will be returned.

Change Type	No.	Name	Change Description
	49	TIMESTAMPDIFF	The TIMESTAMPDIFF function can be used to return the time difference in a specified unit.
	50	TO_DAYS	The TO_DAYS function can return the number of days between the input time and 0.
	51	DATE_FORMAT	The DATE_FORMAT function can convert an input date and time into a string in a specified format. The string is in MYSQL format and starts with %.
	52	IF	The IF function can return the corresponding value based on the first input.
	53	IFNULL	The IFNULL function can return the first non-NULL value.
	54	ISNULL	The ISNULL function can check whether the input is NULL.
Modify	55	pg_resume_bkp_flag	You can obtain the delay DDL flag during backup and restoration.
	56	pg_query_audit	The begintime , operation_type , command_text , transaction_xid and query_id fields are added to audit logs. The time field is changed to endtime , and the type field is changed to audit_type . In addition, detail_info is used to store the command execution results instead of commands, which have been moved to command_text .
	57	pg_delete_audit	For security purposes, the interface for deleting audit logs cannot be provided. If this function is called, an error will be reported.

Change Type	No.	Name	Change Description
	58	log()	Logarithmic function. In ORA or TD mode, it represents a logarithm to base 10. In MySQL mode, it represents a natural logarithm.
	59	to_number	The to_number function instead of the numeric_in function is called. In TD mode, if the input is an empty string, null will be returned.
	60	left()	You can truncate a string. If the input parameter is a negative number, -n for example, all characters except the last <i>/n/</i> character will be returned in ORA or TD mode, and an empty string will be returned in MySQL mode.
	61	right()	You can truncate a string. If the input parameter is a negative number, -n for example, all characters except the first <i>/n/</i> character will be returned in ORA or TD mode, and an empty string will be returned in MySQL mode.
	62	last_day	Input parameters support the timestamptz type. In MySQL mode, the return type is date. In ORA or TD mode, the return type is timestamp.
	63	next_day	Input parameters support the timestamptz type. In MySQL mode, the return type is date. In ORA or TD mode, the return type is timestamp.
	64	add_months	Input parameters support the timestamptz type. In MySQL mode, the return type is timestamptz. In ORA or TD mode, the return type is timestamp.

Change Type	No.	Name	Change Description
	65	add_months	Input parameters support the date type. In MySQL mode, the return type is date. In ORA or TD mode, the return type is timestamp.
	66	pg_cbm_recycle_file	This parameter is added so that the concurrent backup and disaster recovery features can recycle CBM.
	67	pgxc_query_audit	The begintime , operation_type , command_text , transaction_xid and query_id fields are added to audit logs. The time field is changed to endtime , and the type field is changed to audit_type . In addition, detail_info is used to store the command execution results instead of commands, which have been moved to command_text .
	68	login_audit_messages	Enhanced the audit log feature.
Delete	69	add_policy/drop_policy/ alter_policy/enable_policy/ disable_policy	The old data making interface was deleted.

System Views

Table 1-30 System views

Change Type	No.	Name	Change Description
Add	1	GS_WLM_SQL_ALLOW	You can query the existing whitelist.
	2	PG_TOTAL_SCHEMA_INFO	You can query the schema space information of a node.
	3	PGXC_TOTAL_SCHEMA_IN FO	You can query the schema space information of a cluster.

Change Type	No.	Name	Change Description
	4	PGXC_TOTAL_SCHEMA_IN FO_ANALYZE	You can analyze the schema space information of the cluster.
	5	PGXC_LOCK_CONFLICTS	You can query the information about conflicting locks on each node.
	6	PGXC_DEADLOCK	You can query the information about lock waits in a distributed deadlock (including information about lock objects and lock holders).
	7	PGXC_NODE_STAT_RESET_ TIME	You can query the node reset time.
	8	GS_NODE_STAT_RESET_TI ME	You can check the reset time of the current node.
	9	GLOBAL_STAT_DATABASE	You can query global database statistics.
	10	GLOBAL_REL_IOSTAT	You can query global I/O statistics.
	11	PGXC_REL_IOSTAT	You can query node I/O statistics.
	12	GS_REL_IOSTAT	You can query the I/O statistics of the current node.
	13	GLOBAL_REDO_STAT	You can query global redo statistics.
	14	PGXC_REDO_STAT	You can query node redo statistics.
	15	GLOBAL_WORKLOAD_TRA NSACTION	You can query global workload transaction statistics.
	16	PGXC_INSTANCE_TIME	You can query instance time statistics.
	17	PGXC_SETTINGS	You can query node GUC settings.
	18	PGXC_REPLICATION_SLOT S	You can query the replication slot statistics of the nodes.
	19	PGXC_STAT_REPLICATION	You can query the replication statistics of the nodes.

Change Type	No.	Name	Change Description
	20	PGXC_STAT_BGWRITER	You can query node bgwriter statistics.
	21	PGXC_WAIT_EVENTS	You can query wait events on the nodes.
	22	GS_WAIT_EVENTS	You can query wait event statistics of the current node.
	23	PGXC_OS_RUN_INFO	You can query the OS performance statistics of the servers where instances are deployed.
	24	PG_LIFECYCLE_DATA_DISTRIBUTE	You can query the data distribution statistics of OBS cold and hot tables.
	25	PG_BULKLOAD_STATISTICS	You can encapsulate the <code>pg_stat_get_node_bulkload_statistics</code> function.
	26	PGXC_BULKLOAD_STATISTICS	You can encapsulate the <code>pgxc_stat_get_node_bulkload_statistics</code> function.
	27	PGXC_BULKLOAD_PROGRESS	This function can aggregate the query results of the pgxc_bulkload_statistics view and calculate the GDS service progress percentage.
	28	PG_QUERYBAND_ACTION	You can query all the query_band load behaviors.
	29	GS_VIEW_DEPENDENCY_PATH	You can query the dependency between cascading query views.
Modify	30	REDACTION_COLUMNS	The function_infos column is added to the system view definition.
	31	PGXC_GET_TABLE_SKEWNESS	Fixed the problem that the round-robin table is not displayed in the PGXC_GET_TABLE_SKEWNESS view.
	32	PGXC_STAT_DATABASE	You can query database statistics on each instance.

Behavior Changes

Table 1-31 Behavior changes

Change Type	No.	Name	Change Description
Modify	1	If the Not-null and CHECK constraints conflict, the details of the entire row are not printed.	If the Not-null and CHECK constraints conflict, the details of the entire row are not printed.
	2	Compatible with C80 behavior. By default, the implicit type conversion from interval to text is not matched.	If behavior_compat_options is set to enable_interval_to_text , this implicit type conversion is supported.
	3	Date type	In ORA mode, the date type is reversely parsed as pg_catalog.date .
	4	Agg plan	Required columns are generated when agg generates a stream plan.
	5	Median is used with collate.	Collate cannot be used with median, percentile_cont, or percentile_disc.
	6	Modifications on multi-column partitioning	The rule to check the boundary values of multiple columns during partition modification. If the boundary value of a column is the maximum value, the boundary values of other columns are ignored.

Change Type	No.	Name	Change Description
	7	Generate a plan with enable_nestloop is set to off and no association conditions.	Assume there are two tables that do not have equivalent join conditions with each other, but have equivalent join conditions with other tables. If the number of joined rows increases, the nestloop plan can be executed. After the rectification, the nestloop plan cannot be executed. For equivalent join columns that do not support hashjoin, if mergejoin is set to off , the nestloop cost will change. The execution plan will change from Nestloop + Indexscan to Nestloop + Materialize.

GUC Parameters

Table 1-32 GUC parameters

Change Type	No.	Name	Change Description
Add	1	behavior_compat_options	<ul style="list-style-type: none"> The convert_empty_str_to_null_td option is added to be compatible with the to_date, to_timestamp, and to_number functions that return null after processing an empty string in TD mode. The enable_interval_to_text option is added to determine whether to support the implicit conversion from interval to text.
	2	debug_group_dop	You can configure the DOP of each stream group for statement-level optimization during SMP adaptation. This value is left empty by default.

Change Type	No.	Name	Change Description
	3	enable_row_fast_numeric	Numeric data in row-store tables is flushed to disks in bigint format.
	4	expand_hashtable_ratio	Dynamic extension of the hash table
	5	query_dop_ratio	A logic DOP makes decisions based on the system resources and cost. If the decisions are inaccurate, a parameter is added for escape. The default value is 1.
	6	show_unshippable_warning	A parameter is added to determine whether to print logs that have not been pushed down to the client. By default, the function is disabled.
	7	assign_abort_xid	Session-level parameters are added. If data is deleted by mistake, incomplete recovery can be performed quickly. This parameter is left blank by default. The GUC parameter indicates rolling back the transaction submitted by a specific xid.
	8	bi_page_reuse_factor	Idle FSM factor during batch insert reuse (0 indicates returning to the old logic)
	9	check_cu_size_threshold	Userset level. When data is inserted to a column-store table, if the amount of data inserted to a CU is greater than check_cu_size_threshold , row size will be checked to avoid generating CUs greater than 1 GB.
	10	enable_light_colupdate	Userset level, indicating whether lightweight update is enabled

Change Type	No.	Name	Change Description
Modify	11	cost_model_version	The parameter value option 2 is added. The default value 1 remains unchanged. If the value is less than 2, the original function will be used for fixed value sampling. If the value is greater than 2, the function improved to enhance randomness is used for fixed value sampling.
	12	expected_computing_node_group	After an in-place upgrade, set expected_computing_nodegroup to bind .
	13	query_dop	For a newly installed cluster, the default value is changed to 1 . For an upgraded cluster, the default value remains unchanged.

1.7 Version 8.1.0

Release date: May 15, 2021

Cluster version: 8.1.0.100

Feature Changes

Table 1-33 8.1.0 feature changes

Feature	Description
Enhanced ANALYZE for temporary tables and single-transaction operations.	Transactions and stored procedures support the ANALYZE command. No error is reported when ANALYZE and ALTER TABLE are executed concurrently. The automatic analyze function supports temporary tables.
Improved performance of the vectorized executor.	<ul style="list-style-type: none"> Enhanced vectorized execution (HASH FULL JOIN) string_agg/listagg supports vectorization.
Partition pruning with non-constant partition column filter criteria	Partition pruning is optimized. Currently, partition pruning can be performed only for constants. Variables that need to be calculated cannot be pruned and need to be optimized.

Feature	Description
Schema-level full backup + Table- and schema-level fine-grained restoration	Schema-level full backup + Table- and schema-level fine-grained restoration
Supported third-party user-defined data conversion of GDS.	Supported third-party user-defined data conversion.
Optimized RTO.	RPO = 0 and RTO < 60s in intra-cluster HA scenarios; RTO < 30s in POC scenarios
SQL on HDFS	Supported ORC data export using OBS foreign tables.
Supported XML function parsing.	Supported XML function parsing.
Modified distribution columns.	Modified distribution columns.
Supported automatic job migration when a CN fails.	If the current CN is faulty, its jobs can be automatically migrated.
Supported smooth upgrade of C functions.	Supported smooth upgrade of C functions.
Supported Teradata DSQL compatibility of gsql.	<ol style="list-style-type: none"> Supported equivalent comparison of dynamic variable strings (excluding other logical operations, AND/NOR logic, and nesting). GOTO label The THEN module supports only SQL and GOTO statements.
Synchronized the time zone with the IANA Time Zone Database (tz).	Synchronized the time zone with the IANA Time Zone Database (tz).

SQL Syntax

Table 1-34 SQL Syntax

Change Type	No.	Name	Change Description
Add	1	Column-store partitioned tables support the split operation.	Column-store partitioned tables support the split operation.
	2	CREATE TABLE tbl(LIKE tbl2 INCLUDING DROPCOLUMNS)	The LIKE DROPCOLUMNS syntax is supported.

Change Type	No.	Name	Change Description
	3	ALTER TABLE tbl DISTRIBUTE BY	Added the syntax for modifying the distribution column of a table.
	4	ALTER TABLE set version	You can switch between the old and new formats of column-store tables.
	5	CREATE/ALTER FOREIGN TABLE Write Only	Data can be written in OBS ORC format.
	6	CREATE/ALTER FOREIGN TABLE Write Only	Data can be written in OBS HDFS format.
	7	XML	The XML data type can be used in tables and functions.
	8	ALTER VIEW	alter view viewname REBUILD;
Modify	9	VACUUM FULL	A deadlock occurs when concurrent VACUUM FULL operations are performed on the same table.
	10	DROP TEXT SEARCH CONFIGURATION	If the text search configuration to be deleted is the current default_text_search_config , an error is reported and the deletion fails.
	11	CREATE/ALTER ROLE/ USER AUTHINFO 'ldap'	The syntax for specifying the user authentication type as LDAP is modified to remove the 64 character length limit.
	12	CREATE USER/ROLE with authinfo 'ldap' password disable	Modified the syntax for creating an LDAP user.
	13	ALTER USER/ROLE with authinfo 'ldap' password disable	Modified the syntax for creating an LDAP user.
Delete	14	alter table <i>Partitioned_table</i> exchange with unlogg <i>Table</i>	Forbid partitioned table exchange with unlogged tables.
	15	Inheritance can be modified using ALTER TABLE ... inherit .	The optimizer is not adapted to table inheritance. The alter table inherit interface needs to be disabled.

Keywords

Table 1-35 Keywords

Change Type	No.	Name	Change Description
Add	1	xmltable	The COL_NAME_KEYWORD keyword has been adapted for inspection.
	2	xmlnamespaces	The COL_NAME_KEYWORD keyword has been adapted for inspection.

System Catalogs

Table 1-36 System catalogs

Change Type	No.	Name	Change Description
Add	1	dbms_om.gs_wlm_session_info	New system catalog
	2	dbms_om.gs_wlm_operator_info	New system catalog
Modify	3	pg_jobs	Same as the old system catalog pg_job , the metadata relationship is associated using OIDs. If the shared system catalog is full, and you cannot add a transaction after deleting one, pg_job will be renamed as pg_job_proc .
	4	pg_proc	The textanycat and anytextcat functions are changed to non-STRICT functions.
	5	pg_proc	The database_to_xml, database_to_xmlschema, and database_to_xml_and_xmlschema functions are changed to non-STRICT functions.
	6	pg_catalog.gs_wlm_session_info	It is changed to a view.

Change Type	No.	Name	Change Description
	7	pg_catalog.gs_wlm_operator_info	1. It is changed to a view. 2. The nodename column is added.
	8	pg_catalog.gs_wlm_ec_operator_info	It is changed to a view.
	9	pg_catalog.gs_wlm_ec_operator_info	1. It is changed to a view. 2. The nodename , plan_node_name , ec_operator , and ec_fetch_count columns are added.

System Functions

Table 1-37 System functions

Change Type	No.	Name	Change Description
Add	1	pgxc_get_residualfiles	Unified CN query function of pg_get_residualfiles.
	2	pgxc_rm_residualfiles	Unified CN query function of pg_rm_residualfiles.
	3	pgxc_verify_residualfiles	Unified CN query function of pg_verify_residualfiles.
	4	query_to_xmlschema	Maps the query content to an XML schema document.
	5	query_to_xml	Maps the query result to an XML file.
	6	query_to_xml_and_xmlschema	Maps the contents of a query to XML values and schema documents.
	7	table_to_xmlschema	Maps the contents of a relational table to an XML schema document.
	8	table_to_xml	Map the contents of the relation table to XML values.
	9	table_to_xml_and_xmlschema	Maps the contents of a relational table to XML values and schema documents.

Change Type	No.	Name	Change Description
	10	cursor_to_xmlschema	Maps the contents of a cursor query to an XML schema document.
	11	cursor_to_xml	Maps the contents of a cursor query to an XML document.
	12	schema_to_xmlschema	Maps the contents of the entire schema into an XML schema document.
	13	schema_to_xml	Maps the contents of the entire schema to an XML document.
	14	schema_to_xml_and_xmlschema	Maps the contents of the entire schema to XML values and schema documents.
	15	database_to_xmlschema	Maps the contents of the entire database into XML schema documents.
	16	database_to_xml	Maps the contents of the entire database to XML documents.
	17	database_to_xml_and_xmlschema	Maps the contents of the entire database to XML values and schema documents.
	18	xmlpi	Creates an XML processing instruction.
	19	xmlcomment	Creates an XML comment with the specified text as the content.
	20	xmlserialize	Generates a string from a value of the xml type.
	21	xmlparse	Generates a value of the xml type from character data.
	22	xpath	Returns an array of XML values corresponding to the node set generated by the XPath expression.
	23	get_instr_unique_sql_remote_cns	Obtain the unique sql data on all CNs except the current CN.
	24	xml_is_well_formed	Check whether the text string is a well-formed XML value.

Change Type	No.	Name	Change Description
	25	xml_is_well_formed_content	Check whether the text string is well-formed.
	26	xml_is_well_formed_document	Checks whether the text string is a well-formed document.
	27	xmlconcat	Concatenates a list of XML values into a single value.
	28	xmllagg	Aggregate function that concatenates input values.
	29	IS DOCUMENT	Determines whether the parameter value is a correct XML file.
	30	IS NOT DOCUMENT	Determines whether the parameter value is an incorrect XML file.
	31	xmlexists	Determines whether the XPath expression in the parameter returns any node.
	32	xpath_exists	Determines whether the XPath expression in the parameter returns any node.
	33	xmlelement	Produces an XML element with the given name, attributes, and content.
	34	xmlforest	Produces an XML forest (sequence) of an element using the given name and content.
	35	pg_xlog_replay_completion	You can query the Xlog redo progress of the current DN.
	36	pg_data_sync_from_dummy_completion	You can query the progress of transferring data pages from the dummy standby node to the current DN.
	37	pg_stat_xlog_space	You can query the Xlog space usage on the current DN.
	38	pgxc_stat_xlog_space	Show the Xlog space usage on all primary DNs.
	39	xmlroot	Modifies the attributes of the root node of an XML value.

Change Type	No.	Name	Change Description
	40	pg_get_residualfiles	This command is used to query the list of all residual files on the current node.
	41	pg_get_running_jobs	This view queries jobs that are running on the current node.
	42	pg_is_residualfiles	You can query whether a specified relfilenode is a residual file in the current database.
	43	pg_rm_residualfiles	Deletes files from a specified residual file list.
	44	pg_verify_residualfiles	This command is used to check whether the residual file list recorded in a specified file is a residual file.
	45	table_distribution	Add a table_distribution function whose parameter type is regclass.
	46	XMLTABLE	The XMLTABLE function can be used to parse XML data.
Modify	47	pgxc_get_senders_catchup_time	The catchup process information field is added.
	48	pg_stat_get_data_senders	The catchup process information field is added.
	49	pgxc_get_thread_wait_status	The num_node_display parameter is added.
	50	pg_stat_get_status	The num_node_display parameter is added.
	51	pg_catalog.pgxc_get_wlm_session_info_bytime	The first parameter is changed from case-sensitive to case-insensitive.

System Views

Table 1-38 System views

Change Type	No.	Name	Change Description
Add	1	pg_job_single	Job information of a single node.
	2	pg_job	Used for forward compatibility with the pg_job system catalog and distributed pg_job_single collection.
	3	pgxc_get_stat_all_partitions	Obtain the number of insert/update/delete/live/dead tuples and dirty page rate of each partition in all partitioned tables.
	4	gs_view_dependency	The dependency of the cascading query view is the union of the preceding two functions.
Modify	5	pgxc_get_instr_unique_sql	Run the get_instr_unique_sql_remote_cns statement to obtain the unique SQL data on other CNs.
	6	pg_stat_replication	Debugged the function.
	7	pg_get_senders_catchup_time	The catchup process information field is added.
	8	pg_catalog.gs_wlm_operator_history	The nodename field is added.
	9	pg_catalog.pgxc_wlm_operator_history	The nodename field is added.
	10	pg_catalog.pgxc_wlm_operator_info	The nodename field is added.

Behavior Changes

Table 1-39 Behavior changes

Change Type	No.	Name	Change Description
Modify	1	CN retry supports the stored procedure that affects real-time printing.	The output of the stored procedure is printed in real time. If cn retry occurs, a notice is displayed, indicating that the output message is repeated. If cn retry occurs before printing, no notice is output.

GUC Parameters

Table 1-40 GUC parameters

Change Type	No.	Name	Change Description
Add	1	wal_compression_level	Zlib compression level of the PFI log compression function. The default value is 9.
	2	wal_compression	Specifies whether to enable the PFI log compression function. This function is disabled by default.
	3	max_xlog_backup_size	Size of the Xlog backup file. When the size of a backup Xlog file exceeds the value of this parameter, the earliest backup Xlog file is automatically deleted until the size of the backup Xlog file is less than 90% of the value of this parameter.
	4	max_cache_partition_num	Specifies the maximum number of partitions reserved in the memory during redistribution. If the number of partitions exceeds the value of this parameter, the earliest partition will be written to disks in CU format.

Change Type	No.	Name	Change Description
	5	password_encryption_type	Specifies the encryption type of user passwords. The value 2 is added, indicating that the password is encrypted using SHA256. If the current cluster is upgraded from 8.0.0 or an earlier version, the default value of this parameter is 1 for forward compatibility. The default value is 2 for a newly installed cluster.
	6	join_num_distinct	Specifies the default distinct value of the join column. The default value is -20.
	7	cost_model_version	Indicates whether the cost estimation optimization takes effect. The default value is 1, indicating that the optimization takes effect.
	8	qual_num_distinct	Default distinct value of the filter column. The default value is 200.
	9	behavior_compat_options	The varray_verification option is added to roll back the new verification.
	10	behavior_compat_options	The check_function_conflicts option is added to check whether non-IMUMTABLE behavior exists in the IMMUTABLE function.
	11	auto_process_residualfile	Switch for automatically recording residual files. The default value is true, indicating that the function is enabled.
	12	default_colversion	Specifies the default format for creating a column-store table. The default format is 1.0.
	13	enable_partition_dynamic_pruning	Whether to support dynamic pruning during partition table scanning. Dynamic pruning is enabled by default.

Change Type	No.	Name	Change Description
	14	enable_join_pseudoconst	Indicates whether to generate a join expression on the join expression that is equal to a constant. Similar to the t1 inner join t2 on t1.a=t2.a where t1.a = 1 scenario, the join expression can be generated based on t1.a=t2.a. (In the earlier version, such join conditions cannot be used for join cond.)
	15	view_independent	Parameter function: Specifies whether to enable the view dependency decoupling function. Default value: off
	16	enable_view_update	The update of a single table view is supported.
Modify	17	enable_index_nestloop	If the system is upgraded from C80 or an earlier version to the latest patch version, set enable_index_nestloop to off. enable_index_nestloop is upgraded from 651 650 and is forward compatible. The default value of enable_index_nestloop is on during installation.
	18	archive_mode	The default value of the xlog archive switch is changed from off to on. In some scenarios, the performance is slightly affected. In POC scenarios such as performance comparison, you are advised to manually avoid this problem.
	19	cost_param	The default value is changed to 16, which corresponds to the optimized cost estimation.
	20	rewrite_rule	The partialpush option is removed.

Change Type	No.	Name	Change Description
	21	behavior_compat_options	Added the strict_concat_functions parameter for forward compatibility of the textanycat and anytextcat functions.
	22	behavior_compat_options	The 'strict_text_concat_td' option has been added to change textcat, textanycat, and anytextcat functions to strict functions in TD mode.
	23	behavior_compat_options	The strict_text_concat_td option is added to be compatible with the null stitching behavior in T/D mode.
	24	behavior_compat_options	The bpchar_text_without_rtrim option is added to be compatible with the style of processing spaces at the end of the bpchar string in TD mode.

Time Zone

Table 1-41 Time zone

Change Type	No.	Name	Change Description
Modify	1	Asia/Beijing: The definition of the Beijing time zone is changed to be the same as that of Asia/Shanghai.	<ol style="list-style-type: none"> The time zone offset before 1901 is changed from 08:00:00 to 08:05:43. The DST rules from 1940 to 1949 are added. <ul style="list-style-type: none"> The DST is used from June 1 to October 12, 1940. The DST is used from March 15, 1941 to November 1, 1941. The daylight saving time (DST) was used from January 31, 1942 to September 1, 1945. The DST was used in the three years because the government did not adjust the time zone rules nationwide during World War II. The DST is used from May 15, 1946 to September 31, 1946. The DST is used from April 15, 1947 to October 31, 1947. The DST is used from May 1, 1948 to September 30, 1948. DST is used from May 1, 1949 to September 30, 1949. DST rules were also adjusted between 1986 and 1991. <ul style="list-style-type: none"> From 1986 to 1991, the DST was changed from 00:00 to 02:00. From 1987 to 1991, the DST was adjusted to the

Change Type	No.	Name	Change Description
			standard time zone, and the first Sunday after September 10 was changed to the first Sunday after September 11.

Change Type	No.	Name	Change Description
	2	Modified the definition of Asia/Shanghai.	<ol style="list-style-type: none"> 1. The time zone offset before 1901 is changed from 08:05:57 to 08:05:43. 2. Adjust the DST rules from 1940 to 1941. <ul style="list-style-type: none"> • The DST time in 1940 is changed from June 3 to October 1 to June 1 to October 12. • The DST time in 1941 is changed from March 16 to October 1 to March 15 to November 1. 3. The DST rules from 1942 to 1949 are added. <ul style="list-style-type: none"> • The daylight saving time (DST) was used from January 31, 1942 to September 1, 1945. The DST was used in the three years because the government did not adjust the time zone rules nationwide during World War II. • The DST is used from May 15, 1946 to September 31, 1946. • The DST is used from April 15, 1947 to October 31, 1947. • The DST is used from May 1, 1948 to September 30, 1948. • DST is used from May 1, 1949 to September 30, 1949. 4. DST rules were also adjusted between 1986 and 1991. <ul style="list-style-type: none"> • From 1986 to 1991, the DST was changed from 00:00 to 02:00. • From 1987 to 1991, the DST was adjusted to the

Change Type	No.	Name	Change Description
			standard time zone, and the first Sunday after September 10 was changed to the first Sunday after September 11.
	3	Asia/Urumqi	After May 1980, the time zone offset is changed from GMT+8 to GMT+6.
	4	The five time zones defined based on the physical time zone are changed to two time zones.	<ol style="list-style-type: none"> 1. The time zone names Asia/Harbin (long white time zone), Asia/Chongqing (longshu time zone), and Asia/Shanghai (original standard time zone) are reserved. All the time zone names point to the newly defined Asia/Shanghai (standard time zone of China). It is used in the whole country. 2. The Asia/Kashgar (Kunlun time zone) and Asia/Urumqi (New Tibet time zone) time zone names are reserved, and the definitions point to the newly defined Asia/Urumqi (Urumqi time zone). Two time zone definitions are used in the local area.
	5	The initial default time zone may change.	During database initialization, a time zone that best matches the operating system time zone is used as the default time zone. The time zone definition supported by the database is changed. As a result, the initial default time zone may be changed.

1.8 Version 8.0.1

V8.0.1.500 release date: February 8, 2021

V8.0.1.100 release date: July 31, 2020

SQL Syntax

Table 1-42 SQL syntax

Change Type	No.	Name	Change Description
Add	1	Setting the temporary tablespace quota	The CREATE USER/ROLE user1... TEMP SPACE '5GB' syntax is added to set the temporary space quota.
	2	Changing the temporary tablespace quota	The ALTER USER user1 TEMP SPACE '5GB'; syntax is added to modify the temporary space quota.
	3	Setting the space limit for the intermediate calculation result set to be flushed to disks	The CREATE USER/ROLE user1... SPILL SPACE'unlimited'; syntax is added to set the temporary space quota of the intermediate result set.
	4	Modifying the upper limit of the space for storing intermediate calculation results	The ALTER USER user1 SPILL SPACE '5GB'; syntax is added to set the temporary space quota for the intermediate result set.
	5	ALTER TABLE	The alter table modify not null [enable] / alter table modify null syntax is added.
	6	CREATE INDEX	The index name can contain schema when an index is created.
	7	ALTER SEQUENCE	The ALTER SEQUENCE syntax is added.
	8	CREATE SYNONYM	The CREATE [OR REPLACE] SYNONYM syn_name FOR obj_name syntax is added.
	9	DROP SYNONYM	The DROP SYNONYM IF EXISTS syn_name syntax is added.
	10	ALTER SYNONYM	The ALTER SYNONYM syn_name OWNER TO newowner syntax is added.
	11	CREATE USER/ROLE with ldap password disable	AD authentication is supported. When creating a user, set the user authentication type to LDAP.

Change Type	No.	Name	Change Description
	12	ALTER USER/ROLEwith ldap password disable	AD authentication is supported, and the user authentication type is changed to LDAP.
	13	Adding parameter options to the copy fault tolerance table	The log errors data option is added to support user-defined selection. Fill in the rawrecord field.
Modify	14	GIN index	The GIN index is disabled in the 300 scenario.
Delete	15	alter table <i>Partitioned_table</i> exchange with unlogg <i>Table</i>	Do not exchange partitioned tables with unlogged tables.

Keywords

Table 1-43 Keywords

Change Type	No.	Name	Change Description
Add	1	SYNONYM	The non-reserved keyword SYNONYM is added.
	2	VERIFY	The reserved keyword VERIFY is added.
	3	FAST	The non-reserved keyword FAST is added.
	4	COMPLETE	The non-reserved keyword COMPLETE was added.
Modify	5	FUNCTION	The keyword FUNCTION is changed to a non-reserved keyword U.
	6	OVER	The keyword OVER is changed to a non-reserved keyword U.
	7	SPLIT	The SPLIT keyword is changed to a non-reserved keyword U.
	8	RETURN	The type of keyword RETURN is changed to a non-reserved keyword U.

Change Type	No.	Name	Change Description
	9	ISNULL	The ISNULL type is changed to a non-reserved keyword U.

System Catalogs

Table 1-44 System catalogs

Change Type	No.	Name	Change Description
Add	1	PG_SYNONYM	Added the pg_synonym system catalog for storing the mapping between synonym objects and associated object names.
	2	PG_REDACTION_POLICY PG_REDACTION_COLUMN	pg_redaction_policy stores object masking policies and status expressions. pg_redaction_column stores object column information, including which columns need to be anonymized and the anonymization mode.
	3	PG_OBJECT	Stores the creation time, creation user, and last modification time of a table/foreign table/view/index/sequence/function.
Modify	4	PG_AUTHID	The authinfo column is added to identify the authentication type of a user. The default value is empty. If the user is an AD user, the value is ldap.

System Functions

Table 1-45 System functions

Change Type	No.	Name	Change Description
Add	1	dbms_redact.add_policy dbms_redact.alter_policy dbms_redact.enable_policy dbms_redact.disable_policy dbms_redact.drop_policy	<ul style="list-style-type: none"> • Add a data masking policy. • Modify the masking policy information, including adding a masking column to a table object and modifying the masking mode of the column. • Enable or disable a data masking policy. • Deleting a data masking policy
	2	pgxc_get_workload_sql_count	Provides statistics on the number of SELECT/UPDATE/INSERT/DELETE statements executed in all workload Cgroup on all CNs of the current cluster and the number of DDL, DML, and DCL statements.
	3	pgxc_get_workload_sql_elapsed_time	Provides statistics on response time of SELECT/UPDATE/INSERT/DELETE statements executed in all workload Cgroup on all CNs of the current cluster.
	4	get_instr_unique_sql	Provides unique SQL information collected on the current node.
	5	reset_instr_unique_sql(cstring, cstring, INT8)	Deletes the collected Unique SQL information.
	6	pgxc_get_instr_unique_sql	Provides complete Unique SQL information collected on all CNs in a cluster.
	7	pg_check_authid	Checks whether the user exists based on the user OID.

Change Type	No.	Name	Change Description
Modify	8	pg_stat_get_sql_count	Provides statistics on the results of SELECT/UPDATE/INSERT/DELETE/MERGE INTO statements executed by all users on all the nodes, their response time, and the number of DDL, DML, and DCL statements.
	9	pgxc_get_sql_count	Provides statistics on the results of SELECT/UPDATE/INSERT/DELETE/MERGE INTO statements executed by all users on the current node, their response time, and the number of DDL, DML, and DCL statements.

System Views

Table 1-46 System views

Change Type	No.	Name	Change Description
Add	1	ALL_SYNONYMS DBA_SYNONYMS USER_SYNONYMS	<ul style="list-style-type: none"> ALL_SYNONYMS displays all synonyms accessible to the current user. DBA_SYNONYMS displays all synonyms in the database. Only users with system administrator permissions can access this view. USER_SYNONYMS displays synonyms accessible to the current user.
	2	REDACTION_POLICIES REDACTION_COLUMNS	The two tables correspond to two system catalogs. The OID is converted to the name to display the policy information.

Change Type	No.	Name	Change Description
	3	GS_WORKLOAD_SQL_COUNT	Displays statistics on the number of SQL statements executed in workload Cgroups on the current node, including the number of SELECT , UPDATE , INSERT , and DELETE statements and the number of DDL, DML, and DCL statements.
	4	PGXC_WORKLOAD_SQL_COUNT	Displays statistics on the number of SQL statements executed in workload Cgroups on all CNs in a cluster, including the number of SELECT , UPDATE , INSERT , and DELETE statements and the number of DDL, DML, and DCL statements.
	5	GLOBAL_WORKLOAD_SQL_COUNT	Displays statistics on the number of SQL statements executed in all workload Cgroups in a cluster, including the number of SELECT , UPDATE , INSERT , and DELETE statements and the number of DDL, DML, and DCL statements.
	6	GS_WORKLOAD_SQL_ELAPSE_TIME	Displays statistics on the response time of SQL statements in workload Cgroups on the current node, including the maximum, minimum, average, and total response time of SELECT , UPDATE , INSERT , and DELETE statements. The unit is microsecond.
	7	PGXC_WORKLOAD_SQL_ELAPSE_TIME	Displays statistics on the response time of SQL statements in workload Cgroups on all CNs in a cluster, including the maximum, minimum, average, and total response time of SELECT , UPDATE , INSERT , and DELETE statements. The unit is microsecond.

Change Type	No.	Name	Change Description
	8	GLOBAL_WORKLOAD_SQL_ELAPSE_TIME	Displays statistics on the response time of SQL statements in all workload Cgroups in a cluster, including the maximum, minimum, average, and total response time of SELECT , UPDATE , INSERT , and DELETE statements. The unit is microsecond.
	9	PGXC_INSTR_UNIQUE_SQL	Displays the complete Unique SQL statistics of all CN nodes in the cluster.
	10	GS_INSTR_UNIQUE_SQL	You can query the Unique SQL execution information collected by the current node.
Modify	11	ALL_OBJECTS DBA_OBJECTS USER_OBJECTS	Modify the view definition and extend the SYNONYM object.
	12	PGXC_THREAD_WAIT_STATUS	Modify the view definition.
	13	GS_SQL_COUNT	Provides statistics on the number of SELECT/UPDATE/INSERT/DELETE/MERGE INTO statements executed by all users on the current node, response time, and the number of DDL, DML, and DCL statements.
	14	PGXC_SQL_COUNT	Provides statistics on the number of SELECT/UPDATE/INSERT/DELETE/MERGE INTO statements executed by all users in the cluster, response time, and the number of DDL, DML, and DCL statements.
	15	ALL_OBJECTS DBA_OBJECTS USER_OBJECTS	Modify the view definition and add the creation time and last modification time of the object.

Change Type	No.	Name	Change Description
	16	PG_TABLES	Modify the view definition and add the creation user, creation time, and last modification time of the table.

Behavior Changes

Table 1-47 Behavior changes

Change Type	No.	Name	Change Description
Add	1	CN retry is disabled for certain COPY FROM statements.	<ul style="list-style-type: none"> The COPY FROM STDIN statement is not supported. The gsql \copy from meta-command is not supported. JDBC CopyManager copyIn cannot be used to import data.

GUC Parameters

Table 1-48 GUC Parameters

Change Type	No.	Name	Change Description
Add	1	sql_use_spacelimit	Specifies the maximum storage space for a single SQL statement on a single DN.
	2	remote_read_mode	Sets the remote read mode.
	3	enable_incremental_catchup	Sets whether to enable incremental catchup.
	4	wait_dummy_time	Maximum time for waiting for the secondary server to start or send back the scan list during incremental catchup.
	5	enable_xlog_group_insert	It specifies whether the xlog enables the group insertion mode on the ARM platform.

Change Type	No.	Name	Change Description
	6	topsql_retention_time	Number of days for storing data in the gs_wlm_session_info and gs_wlm_operator_info tables.
	7	user_metric_retention_time	Indicates the number of days for storing historical resource monitoring data.
	8	instance_metric_retention_time	Indicates the number of days for storing historical resource monitoring data of an instance.
	9	enable_user_metric_persistent	Indicates whether to enable the function of monitoring and dumping historical user resources.
	10	enable_instance_metric_persistent	Indicates whether to enable the function of dumping historical instance resource monitoring data.
	11	COMM_IPC	Switch for locating communication performance problems. This parameter specifies whether to print the information about the packets sent and received by each communication node.
	12	COMM_PARAM	Communication performance problem locating switch. This parameter specifies whether to print the session parameter settings during node communication.
	13	plan_cache_mode	Controls the behavior of the cached plan. The custom plan, generic plan, or auto mode can be forcibly used (the default value is auto).
Modify	14	audit_system_object	The 18th digit is added to indicate whether to audit SYNONYM objects. The maximum value is changed to 524287.

Change Type	No.	Name	Change Description
	15	comm_max_stream	In the FusionCube environment, the default value of this parameter is changed to 10000 after the cluster is installed.
	16	enable_dynamic_workload	The default value is changed from off to on.
	17	behavior_compat_options	The convert_interval_to_text member is added to control whether the implicit type conversion from interval to text is supported. By default, the implicit type conversion is not supported.

1.9 Deleting a Database Interface

To ensure product evolution and maintenance, GaussDB (DWS) provides the metadata sunset function for obsolete GUC parameters, system functions, and system views. To ensure historical compatibility, sunset functions and views are stored in SUNSET EXTENSION, which is created by users and does not affect subsequent version upgrades.

CREATE EXTENSION SUNSET

When you query sunset system functions and views, an error message is displayed indicating that the functions or views do not exist. GaussDB (DWS) allows you to create sunset system functions and views in CREATE EXTENSION SUNSET mode. The EXTENSION contains sunset functions and views, its actual content is the definition before sunset.

The creation method is as follows:

The administrator runs the CREATE EXTENSION SUNSET statement to create the user. The user needs to be created only once. The upgrade of later versions is not affected. The permissions on related functions and views are the same as those in the version before sunset.

```
CREATE EXTENSION sunset;
CREATE EXTENSION
```

```
\dx+ sunset
```

```
Objects in extension "sunset"
Object
```

```
Description
```

```
-----
function pg_delete_audit(timestamp with time zone,timestamp with time zone)
```

```

function pg_log_comm_status()
function pg_stat_get_realtime_info_internal(oid,oid,bigint,cstring,oid)
function pg_stat_get_wlm_session_info_internal(oid,oid,bigint,oid)
function pg_stat_get_wlm_session_info(oid)
function pg_stat_get_wlm_session_iostat_info(integer)
function pg_stat_get_wlm_statistics(integer)
function pg_sync_cstore_delta()
function pg_sync_cstore_delta(text)
function pg_user_iostat(cstring)
function pgxc_log_comm_status()
function pgxc_pool_check()
function signal_backend(bigint,integer)
function update_pgjob(bigint,"char",bigint,timestamp without time zone,timestamp without time zone,timestamp without time zone,timestamp without time zone,timestamp without time zone,smallint)
view gs_wlm_session_info_all
view pg_session_iostat
view pg_wlm_statistics
(17 rows)

```

Deleting a Table Through the Database Interface of the 8.3.0 Cluster Version

Table 1-49 Sunset System Functions and Views

Type	Change	Constraint
Function	pg_sync_cstore_delta(text)	Merge delta table data to the primary table in column-store mode. This parameter has been discarded. This syntax is replaced by the VACUUM DELTAMERGE xxx syntax.
	pg_sync_cstore_delta	Merge delta table data to the primary table in column-store mode. This parameter has been discarded. This syntax is replaced by the VACUUM DELTAMERGE xxx syntax.
	pgxc_pool_check	Check whether the connection data cached in the connection pool is consistent with pgxc_node. This check item has been discarded.
	pg_delete_audit	This function is used to delete audit logs and has been discarded. You are not allowed to manually delete audit logs.
	pg_log_comm_status	Log recording status of the internal structure of the SCTP communication layer on the current node. This parameter is discarded.
	pgxc_log_comm_status	Indicates the log recording status of the internal structure of the SCTP communication layer on all nodes. This parameter is discarded.
	signal_backend	This parameter has been discarded.
	pg_stat_get_realtime_info_internal	This interface is used to obtain real-time top SQL information at the query level. This interface is discarded. This command is replaced by the PGXC_WLM_SESSION_STATISTICS view.

Type	Change	Constraint
	pg_stat_get_wlm_session_info_internal	Obtains top SQL information at the query level. This parameter is discarded. This function is replaced by the pg_stat_get_wlm_realtime_session_info function.
	pg_stat_get_wlm_session_info	Obtains session information, which has been discarded. This command is replaced by the PGXC_WLM_SESSION_INFO view.
	pg_stat_get_wlm_statistics	Obtains session information, which has been discarded. This command is replaced by the PGXC_WLM_SESSION_INFO view.
	pg_user_iostat	This interface is used to obtain the I/O information of a user. This interface is discarded. This command is replaced by the PGXC_TOTAL_USER_RESOURCE_INFO view.
	pg_stat_get_wlm_session_iostat_info	Obtains the I/O information of a query. This parameter has been discarded. This view is replaced by the PGXC_WLM_SESSION_STATISTICS or PGXC_WLM_SESSION_INFO view.
System views	gs_wlm_session_info_all	Obtains session information, which has been discarded. It is replaced by PGXC_WLM_SESSION_INFO.
	pg_wlm_statistics	Displays load management information after a job is complete or an exception has been handled. This parameter has been discarded. This view is replaced by the PGXC_WLM_SESSION_INFO view.
	pg_session_iostat	Obtains the I/O information of a session. This parameter is discarded. This view is replaced by the PGXC_WLM_SESSION_STATISTICS view.


2 Version Support Notes

The cluster version of GaussDB(DWS) displayed on the management console is in the format of x.y.z.p, which is explained in [Figure 2-1](#).

Generally, new functions and features of GaussDB(DWS) are released by release number, for example, 8.1.3 and 8.2.0.

After each iteration version is released, a patch is released to fix problems, for example, patch 8.1.3.322. A patch only fixes problems and does not add new functions or features.

Figure 2-1 GaussDB(DWS) version description



Version Lifecycle

[Table 2-1](#) provides the lifecycle of GaussDB(DWS) versions (release number), helping you plan your version update.

For versions that have reached EOS or are not [Recommended Version](#), you are advised to upgrade them to the latest stable version by referring to [Version Upgrade Policies](#).

Table 2-1 GaussDB(DWS) version lifecycle

Version	Running status	Released On	EOM Date (End of Marketing)	EOS Date (End of Service)
8.2.0.x	Released	2020-11-30	2025-09-30	2027-03-30
8.1.3.x	Released	2022-04-15	2024-06-30	2025-12-31
8.1.1.x	EOM	2021-07-30	2023-06-30	2025-12-31
8.1.0.x	EOM	2021-05-15	2023-06-30	2025-12-31
8.0.1.x	EOM	2020-07-31	2023-06-30	2024-12-31

Version Upgrade Policies

Table 2-2 Version incorporation

Source Version	EOS	Suggestion	Upgrade Plan
1.5.x/1.6.x	Yes	Direct upgrade is not supported. You are advised to migrate data to 8.1.3.325 (recommended version).	You are advised to contact technical support engineers for evaluation before performing this operation.
1.7.x	Yes	You are advised to upgrade the version to 8.1.3.325 (recommended version).	A cross-version upgrade to 8.1.3.325 is not supported. A multi-hop upgrade is required. You are advised to contact technical support for evaluation before the upgrade.
8.0.1.x	No	You are advised to upgrade the version to 8.1.3.325 (recommended version).	A cross-version upgrade to 8.1.3.325 is not supported. A multi-hop upgrade is required. You are advised to contact technical support for evaluation before the upgrade.

Source Version	EOS	Suggestion	Upgrade Plan
8.1.0.x	No	You are advised to upgrade the version to 8.1.3.325 (recommended version).	A cross-version upgrade to 8.1.3.325 is not supported. A multi-hop upgrade is required. You are advised to contact technical support for evaluation before the upgrade.
8.1.1.x~8.1.3.323	No	You are advised to upgrade the version to 8.1.3.325 (recommended version).	It can be directly upgraded to 8.1.3.325. You are advised to contact technical support for evaluation before the upgrade. It can be upgraded on the console. For details, see Upgrading a Cluster .
8.2.0.x (whitelist user version)	No	You are advised to upgrade the version to 8.2.1.119 (recommended version for whitelisted users).	It can be directly upgraded to 8.2.1.119. You are advised to contact technical support for evaluation before the upgrade. It can be upgraded on the console. For details, see Upgrading a Cluster .
8.2.1.1xx (whitelist user version)	No	You are advised to upgrade the version to 8.2.1.119 (recommended version for whitelisted users).	It can be directly upgraded to 8.2.1.119. You are advised to contact technical support for evaluation before the upgrade. It can be upgraded on the console. For details, see Upgrading a Cluster .