

Data Warehouse Service

Product Bulletin

Issue 02
Date 2024-04-29



Copyright © Huawei 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 Technologies Co., Ltd.

Trademarks and Permissions



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

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

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: <https://www.huawei.com>

Email: support@huawei.com

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:

<https://www.huawei.com/en/psirt/vul-response-process>

For vulnerability information, enterprise customers can visit the following web page:

<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

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 8.1.3.....	4
1.3.1 New Features in Version 8.1.3.....	4
1.3.2 New Features and Resolved Issues in 8.1.3.x.....	30
1.4 Version 8.1.1.....	37
2 Version Support Notes.....	58

1 Product Notice

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

[1.2 Checking the GaussDB \(DWS\) Version](#)

[1.3 Version 8.1.3](#)

[1.4 Version 8.1.1](#)

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				
Hot data storage resource	Hourly (US D/G B/Hour)	Monthly (US D/G Month)	1 Year (US D/G B/Year)	2 Year (US D/G B/2 Years)	3 Year (US 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)

	Original Price					New Price				
DWS 3.0 hot data storage	0.00 125	0.6	6	10.0 8	12.9 6	0.000 7	0.336	3.36	5.644 8	7.257 6

Table 1-2 EU-Dublin

	Original Price					New Price				
Hot data storage resource	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 Year s)	3 Years (USD D/G B/3 Year s)	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 2.0 hot data storage	0.00 15	0.72	7.2	12.0 96	15.5 52	0.000 755	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
2023-11-02	8.1.3.325 (recommended version)	8.1.3	1.3.1 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		
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		
2021-12-09	8.1.1.200		
2021-07-30	8.1.1.100		

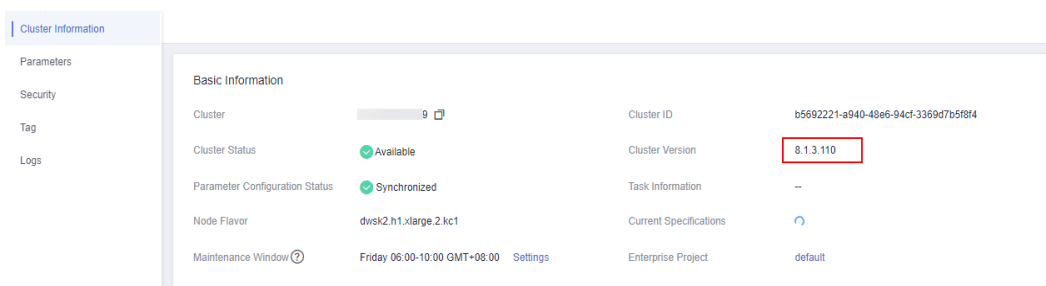
Querying the Cluster Version

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 [Table 1-3](#).
- 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 8.1.3

1.3.1 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 [1.3.2 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 1-4 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</p> <p>UPDATE</p> <p>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 permissions to tables: VACUUM VACUUM DROP TABLE Added permissions to schemas: 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 the following 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(schemaname text, tablename text) pgxc_wlm_readjust_relfilenode_size_table() pgxc_wlm_readjust_relfilenode_size_table()
<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 columns 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_RESPPOOL_RESOURCE_INFO PGXC_RESPPOOL_RESOURCE_INFO PGXC_RESPPOOL_RESOURCE_INFO PGXC_RESPPOOL_RESOURCE_INFO PGXC_RESPPOOL_RESOURCE_INFO - Added the system catalog: <ul style="list-style-type: none"> GS_RESPPOOL_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 1-5 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 1-6 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 1-7 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 1-8 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 cuDESC 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 1-9 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 1-10 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 1-11 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 1000000 .
	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.

1.3.2 New Features and Resolved Issues in 8.1.3.x

8.1.3.325

Table 1-12 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 1-13 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 1-14 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	

1.4 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 .

New Features

Table 1-15 New Functions in 8.1.1

Feature	Description	Reference
2048 nodes in a cluster	Supported up to 2048 nodes in a cluster.	
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. 	
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. 	
Primary key CU of column-store tables	Supported primary key constraints of column-store tables.	
Space management	Added the GUC parameter bi_page_reuse_factor .	
GDS productization	Added system views PGXC_BULKLOAD_PROGRESS , PGXC_BULKLOAD_STATISTICS , and PG_BULKLOAD_STATISTICS .	

Feature	Description	Reference
Hot and cold data storage	<ol style="list-style-type: none"> 1. Added the description of REFRESH STORAGE to "ALTER TABLE". 2. Added the functions and descriptions of COLVERSION and STORAGE_POLICY to "CREATE TABLE". 3. Added the function and description of the OBS tablespace to "CREATE TABLESPACE". 	
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>	
Support for the built-in MEDIAN function	<p>Added the function median(expression).</p>	
Adjusting the cascade delete for tables with views	<ol style="list-style-type: none"> 1. Added the description of rebuilding a view. 2. 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. 	

Feature	Description	Reference
Custom data redaction	<ol style="list-style-type: none"> 1. Optimized data redaction. 2. Provided data redaction functions in different forms. 3. Updated the columns in the system catalog PG_REDACTION_COLUMN and system view REDACTION_COLUMNS. 4. Added the syntax ALTER REDACTION POLICY, CREATE REDACTION POLICY, and 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"> 1. "Keyword" in the <i>SQL Syntax</i> Added EXPIRATION (non-reserved). 2. 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. 3. Added PASSWORD EXPIRATION period to the syntax ALTER ROLE, ALTER USER, and CREATE USER. 4. 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. 5. Added the columns rolauthinfo, rolpwdexpire, and rolpwddtim to the system catalog PG_AUTHID. 	
Increasing the hash table size	Added the GUC parameter expand_hashtable_ratio .	
SMP adaptation enabled	Added "Suggestions for SMP Parameter Settings".	

Feature	Description	Reference
PRETTY as the default value of EXPLAIN	<p>Changed the default display format of EXPLAIN to PRETTY.</p> <p>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.</p>	
Enhanced concurrent SMP	<ol style="list-style-type: none"> 1. Added the constraint of SMP: the short query scenario where the plan generation time accounts for a large proportion. 2. Added the description of setting query_dop to 1 in the short query scenario. 3. Added the GUC parameter query_dop_ratio. 	
row2vec optimization	Added the GUC parameter 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 DBCOMPATIBILITY 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). 	
Support for UPSERT	<p>Added the UPSERT syntax.</p> <p>Added IGNORE, AS, ON DUPLICATE KEY, and ON CONFLICT to the INSERT syntax.</p>	

Feature	Description	Reference
Support for the INSERT , UPDATE , and DELETE operations of views	Added Updatable Views to "CREATE VIEW".	
One-click onsite information collection	Added the STATS boolean parameter to the EXPLAIN syntax.	
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. 	
Combination with an empty string of the CHAR type	Added the configuration item bpchar_text_without_rtrim to the parameter behavior_compat_option .	
Support for Python 3.x	Supported Python 3.	
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. 	
Support for INSERT OVERWRITE	Supported INSERT OVERWRITE .	
Obtaining all views queried by the current user	Added the view 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)	

SQL Syntax

Table 1-16 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 1-17 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 1-18 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 1-19 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 1-20 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 pg_stat_get_node_bulkload_statistics function.
	26	PGXC_BULKLOAD_STATISTICS	You can encapsulate the pgxc_stat_get_node_bulkload_statistics 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 1-21 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 1-22 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.

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)
8.1.3.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

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.
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.

Source Version	EOS	Suggestion	Upgrade Plan
8.1.1.x~8.1.3.3 23	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 .